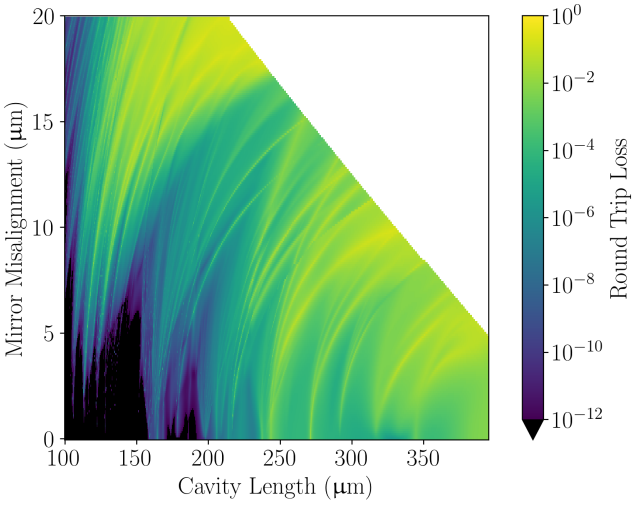
Towards Scalable Ion Trap Nodes For Cavity-Enhanced Quantum Networking



William Hughes

Keble College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2022

Abstract

Ion trap quantum information processors have shown impressive performance in demonstrations of elementary quantum operations and simple algorithms. However, the realisation of quantum computers with 'real-world' utility will require a significant increase in the number of qubits controlled. A promising approach to scaling is a modular architecture, where interactions between modules are mediated by single photons collected from ions in each module and distributed over a photonic network. However, the entanglement of remote ions at useful rates requires extremely high photon collection efficiency, which is best achieved by placing the ion at the centre of an optical cavity. This work investigates different aspects of cavity-enhanced remote entanglement and determines characteristics of both the cavities and the ion-trap module required for scalable quantum computation.

The limitations on high-probability photon production in short times are identified, leading to an intuitive method for the optimisation of cavities with residual transverse misalignment that encodes practical trade-offs. The round-trip loss and resonant mode of cavities with non-spherical mirrors and residual transverse misalignment is studied by expanding the resonant field in a Hermite Gauss basis, uncovering complex and unwanted behaviour for cavities with mirrors that are not highly spherical. The vulnerability to cavity birefringence of remote entanglement schemes using polarisation-encoded photonic information is considered, and options for avoiding these concerns by encoding in the time degree of freedom are presented. Finally, a miniature atomic source suitable for scalable ion trap systems is developed and characterised.

The combined results of this thesis provide a theoretical and technical foundation for a new generation of cavity-enhanced quantum network experiments to be built over the coming years.

Acknowledgements

My first thanks go to my supervisors Prof. Andrew Steane for his guidance and especially to Dr. Joe Goodwin for his support, formidable enthusiasm, and teaching me all that I know (but probably not all that he knows) about cavity physics. I would also like to thank the other members of the cavity team I have worked with during my project: Dr. Shaobo Gao for his relaxed comradeship in the early days, Dr. Ezra Kassa for his attention to detail, the current team of Dr. Jake Blackmore and Dr. Thomas Doherty for their patience and assistance in helping my research and the writing of this document, Master's student Sebastian Saner, MPhys project students Emil Østergaard, Oliver Lowe, and David O'Shea, summer students James Bate and Tim Wohlers-Reichel, and new DPhil student Lorenzo Versini. Thanks to all for making the cavity team such a welcoming place to work, and I look forward to seeing the continuation of all we have achieved. Thanks also to the wider Ion Trap Quantum Computing group, now far too numerous to drop individual names, for your support and friendship. It is certainly true that not everything has transpired exactly to plan set out at the start, but I have always found the ion trap group and the cavity team to be supportive, and I feel privileged to have been a part of them.

I would also like to thank those who have been part of my life outside of work. It is impossible to name everyone, but special thanks to 'The Chain' from Cambridge for being the best friends I could ever hope to have, to Dani and Amanda for their company as the world locked down, and to Ashley for her kindness and friendship. Finally, I must thank my family: My parents and sister for supporting me regardless of whether they should, and finally, to Michael for being there throughout the DPhil, and confusing quite a few others in the process.

List of Publications

Publications

An optically heated atomic source for compact ion trap vacuum systems. S. Gao, W.J. Hughes, D.M. Lucas, T.G. Ballance, J.F. Goodwin. *Rev. Sci. Instrum.* **92**, 033205, (2021)

Effects of cavity birefringence in polarisation-encoded quantum networks. E. Kassa, W.J. Hughes, S. Gao, J.F. Goodwin. *New J. Phys.* **25**, 013004, (2023)

Optimisation of scalable ion-cavity interfaces for quantum photonic networks. S. Gao, J.A. Blackmore, W.J. Hughes, J.F. Goodwin. *Phys. Rev. Applied* **19**, 014033, (2023)

Losses of misaligned microcavities with non-spherical mirrors. W.J. Hughes, T.H. Doherty, J.A. Blackmore, D.V. Karpov, P. Horak, J.F. Goodwin. *in preparation*

Mode mixing calculations for cavities with transverse misalignment. W.J. Hughes, T.H. Doherty, J.A. Blackmore, D.V. Karpov, P. Horak, J.F. Goodwin. *in preparation*

Rapid mode-matching algorithm for monolithic optical cavities. S. Saner, W. J. Hughes, J. A. Blackmore, T.H. Doherty, D.M. Lucas, J. F. Goodwin. *in preparation*

Patents

Patent Pending: "Oven assembly for producing spatially propagating neutral atoms". United Kingdom Patent Application No. GB 2208026.1

Contents

List of Figures **x**

List of Tables **xii**

1 Introduction **1**
1.1 Quantum Computing . . . 1
1.2 Ion Trap Quantum Computing . . . 2
1.2.1 Ion Traps . . . 2
1.2.2 Ion Traps for Quantum Computing . . . 3
1.3 Architectures for Scaling Ion Trap Quantum Computation . . . 4
1.3.1 Scaling Problems with Long Ion Chains . . . 4
1.3.2 QCCD Approach . . . 4
1.3.3 Modular Quantum Computation with Photonic Interconnects 5
1.3.4 Comparison of Scaling Methods . . . 6
1.3.5 Rationale for Ion-Cavity Integration . . . 7
1.3.6 Development of Ion Traps for Scalable Quantum Computation 8
1.4 Thesis Context . . . 8
1.5 Thesis Outline . . . 10

2 Cavity Theory **12**
2.1 Introduction to Cavities . . . 12
2.2 Introduction to Cavity Modes . . . 13
2.2.1 Maxwell's Equations . . . 13
2.2.2 Cavity Types and Boundary Conditions . . . 13
2.2.3 The Paraxial Equation . . . 14
2.2.4 Conditions for Cavity Resonance and Mode Structure . . . 16
2.2.5 Misaligned Cavities . . . 21
2.3 Cavity Coupling . . . 22
2.3.1 Coherent Interaction . . . 22
2.3.2 Incoherent Interaction . . . 25
2.3.3 Cooperativity and Cavity-enhanced Photon Collection . . . 25
2.3.4 Cavity Geometry . . . 26
2.4 Cavities for Remote Entanglement . . . 27
2.4.1 Remote Entanglement Schemes . . . 27

2.4.2 Choice of Cavity Transition . . . 29
2.5 Cavity Manufacture . . . 31
2.5.1 Introduction . . . 31
2.5.2 Mirror Manufacture . . . 31
2.5.3 Cavity Assembly and Mechanical Structure . . . 32
2.5.4 Considerations for the Future Experiment . . . 33
2.6 Conclusion . . . 33

3 Optimising Photon Extraction **35**
3.1 Introduction . . . 35
3.2 Typical Photon Extraction Schemes . . . 35
3.2.1 Overview . . . 35
3.2.2 Direct Excitation . . . 37
3.2.3 Raman Process . . . 37
3.2.4 vSTIRAP . . . 38
3.3 Photon Driving Simulations and Temporal Impurity . . . 38
3.4 Finite Time Driving . . . 39
3.5 Three-Level System: Bounded Approach . . . 40
3.5.1 The System and Problem . . . 40
3.5.2 Calculating the Ideal Photon Shape . . . 41
3.6 Numerical Method for More General Emitter-Cavity Systems . . . 44
3.6.1 Setting up the System . . . 44
3.6.2 Calculating Probabilities . . . 46
3.6.3 Enforcing the Initial Condition . . . 47
3.6.4 Normalising Probabilities . . . 48
3.6.5 Optimising Probabilities . . . 48
3.6.6 Calculating the Required Driving . . . 49
3.6.7 Results . . . 50
3.7 Conclusions . . . 55

4 Cavity Optimisation **58**
4.1 Introduction . . . 58
4.2 Optimisation Scenario . . . 59
4.2.1 System . . . 59
4.2.2 Metric . . . 60
4.3 Optimisation Overview . . . 64
4.3.1 Separation of Geometric and Optical Properties . . . 64
4.3.2 Bounding the Geometrical Optimisation . . . 65
4.3.3 Summary of Optimisation Procedure . . . 68
4.4 Expected Cavity Performance . . . 69
4.5 Robustness to Parameter Errors . . . 72
4.6 Conclusion . . . 73

5 Mode Mixing **75**
5.1 Introduction . . . 75
5.2 Geometric Theory of Cavity Modes . . . 77
5.3 Mode mixing by Mirror Transformation . . . 79
5.3.1 Mode Mixing Introduction . . . 81
5.3.2 Replacing Coordinates with Operators . . . 83
5.3.3 Calculating Polynomial Mirror Surface Profiles . . . 84
5.3.4 Calculating the Gaussian Surface Profile . . . 85
5.3.5 Taking the Exponent of the Surface Profile . . . 86
5.3.6 Translating the Mirror . . . 87
5.3.7 Spherical Cap Mirror . . . 89
5.3.8 Summary . . . 89
5.4 Results . . . 90
5.4.1 Simulation Method . . . 90
5.4.2 Spherical cap Mirror . . . 94
5.4.3 Gaussian Mirror . . . 102
5.5 Conclusion . . . 114

6 Polarisation Encoding and Polarisation Oscillations **117**
6.1 Introduction . . . 117
6.2 Cavity Birefringence . . . 118
6.3 Modelling Cavity Birefringence in Remote Entanglement . . . 120
6.3.1 Cavity Geometry . . . 120
6.3.2 Simulating Production of Photon Wavepackets . . . 121
6.3.3 Photonic Measurement . . . 126
6.4 Correcting for Birefringent Effects . . . 129
6.4.1 Pre-herald Photonic Qubit Correction . . . 129
6.4.2 Windowed Post-Selection . . . 132
6.4.3 Spin-Qubit Rotation . . . 132
6.5 Effect of Cavity Birefringence on Remote Entanglement Fidelity . . . 133
6.5.1 Analytical suggestions . . . 133
6.5.2 Structure of Numerical Investigations . . . 134
6.5.3 Illustrative Examples . . . 135
6.5.4 Dependence on Birefringence Orientation . . . 140
6.5.5 Dependence on Magnitude of Birefringence . . . 142
6.5.6 Summary of Results . . . 142
6.5.7 Deliberately Applied Birefringence . . . 144
6.6 Further Issues with Polarisation Encoding . . . 145
6.7 Conclusions . . . 146

7 Time-Bin Encoding **147**
7.1 Introduction . . . 147
7.2 Principle of Time-Bin Encoding . . . 148
7.3 Advantages of Time-Bin Encoding Over Polarisation Encoding . . . 148
7.4 Disadvantages of Time-Bin Encoding Compared to Polarisation Encoding 150
7.5 Detailed Protocol . . . 152
7.5.1 Introduction . . . 152
7.5.2 Manifolds . . . 152
7.5.3 Pulse Shaping . . . 153
7.5.4 Geometry . . . 154
7.5.5 Driving Protocol and Temporal Mixing . . . 155
7.5.6 Specific Schemes . . . 157
7.6 Conclusion . . . 161

8 Atomic Sources **164**
8.1 Background and Theory of Atomic Sources . . . 164
8.1.1 Introduction . . . 164
8.1.2 Approaches to Sources . . . 165
8.1.3 Miniature Optically-Heated Atomic Sources . . . 166
8.1.4 Theory of Thermal Sources . . . 167
8.1.5 Theory of Measuring Fluorescence . . . 169
8.1.6 Thermal Models for Atomic Sources . . . 171
8.2 MITAS . . . 171
8.2.1 Introduction and Construction . . . 171
8.2.2 Collection of Fluorescence . . . 172
8.2.3 Thermal Performance . . . 176
8.2.4 Temporal Response . . . 177
8.2.5 Predicted Lifetime . . . 178
8.3 'Pico Oven' . . . 178
8.3.1 Source Design . . . 178
8.3.2 Source Preparation . . . 180
8.3.3 Oven mounting and Beam Geometry . . . 181
8.3.4 Heating and Oven Imaging System . . . 182
8.3.5 Excitation-Collection System . . . 184
8.3.6 Initial Alignment . . . 185
8.3.7 Experimental Control . . . 186
8.3.8 Cracking . . . 186
8.3.9 Spectroscopy . . . 187
8.3.10 Summary of Characterisation . . . 188
8.3.11 Results . . . 189
8.4 Discussion . . . 193
8.4.1 Evaluation of Sources . . . 193
8.4.2 Improvements to Experiment . . . 194

9 Conclusion and Outlook **198**
9.1 Criteria for Successful Cavity-Enhanced Remote Entanglement 198
9.1.1 Criterion 1: Negligible Temporal Mixing 198
9.1.2 Criterion 2: Cavity Design Tolerates Residual Misalignment . 199
9.1.3 Criterion 3: Protocol Tolerates Residual Cavity Birefringence . 203
9.1.4 Criterion 4: Suitable Driving Scheme 204
9.2 Atomic Sources 206
9.3 Areas for Further Investigation 207
9.3.1 Simulation of Remote Entanglement Scheme 207
9.3.2 Mode Mixing 208
9.3.3 Experimental Development 209

References **211**

A Calculating Probability Matrices for Multi-Level System Optimisation **226**

B Choosing the Basis Size for Multi-Level System Optimisation **228**

C Derivation of Operators in the Hermite Gauss Basis **229**

D Finding the Gaussian Profile Mirror Matrix in the Hermite Gauss Basis **231**

E Mode Transformations in the Hermite Gauss Basis **233**
E.0.1 Changing the Mode Waist 233
E.0.2 Changing the Mode Angle 233
E.0.3 Determining the Fundamental Mode 234
E.0.4 Removing Rotation from Mode Occupations 235

F Uniqueness of Post-Measurement Corrections to a Single Qubit **236**

G Evaluating the Relative Emission Intensity to a Detector of Finite Size **238**

List of Figures

2.1 Cavity stability diagram 17
2.2 Mode of an optical resonator 18
2.3 Schematic cavity spectrum 20
2.4 Mode of a misaligned cavity 22
2.5 Schematic of ion-cavity coupling 23
2.6 Level diagram of $^{88}Sr^{+}$ 30

3.1 Three-level system diagram 36
3.2 Multi-level system diagram 45
3.3 Comparison of numerical and analytical performance bounds 52
3.4 Multi-level system targeted optimisations 53
3.5 Impact of Zeeman splitting on entanglement success probability . . . 55

4.1 Ion-cavity optimisation scenario 60
4.2 Comparison of photon extraction as function of output time window . 63
4.3 Dependence of clipping loss on cavity geometry 67
4.4 Bounded optimisation of mode divergence angle 69
4.5 Optimised photon extraction probability and cavity transmission . . . 70
4.6 Optimised cavity parameters and photon extraction performance . . . 71
4.7 Robustness of cavity performance to parameter errors 72

5.1 Modes of cavities with Gaussian-shaped mirrors in geometrical approach 80
5.2 Schematic of mode mixing method 83
5.3 Schematic of transverse translations in mode mixing method 88
5.4 Comparison of developed mode mixing methods and results from literature 91
5.5 Example of low loss mode with complicated transverse structure . . . 93
5.6 Round-trip loss and overlap with geometric expectation for cavity with spherical mirrors 96
5.7 Comparison of mode mixing and clipping calculations of round-trip loss 98
5.8 Specific example configurations for cavities with spherical mirrors . . 100
5.9 Round-trip loss and overlap with geometric expectation for cavity with Gaussian mirrors 103
5.10 Higher order transverse mode occupations for large Gaussian mirror . 104

5.11 Higher order transverse mode occupations for Gaussian mirrors of varying width . . . 109
5.12 Cavity configurations with prominent mixing of seventh order mode . 111
5.13 Mode shapes and round-trip losses near for near-confocal configurations 112
5.14 Dependence of round-trip loss features on wavelength . . . 114

6.1 Schematic diagram of polarisation encoded remote entanglement . . . 118
6.2 Possible photon production geometries . . . 122
6.3 Poincare sphere schematic . . . 124
6.4 Polarisation encoded remote entanglement with birefringent cavities . 128
6.5 Schematic of schemes to mitigate impact of birefringence . . . 130
6.6 Case study of birefringent system: photon wavepackets and correction dynamics . . . 136
6.7 Case study of birefringent system: fidelity . . . 138
6.8 Dependence of fidelity on birefringence orientation . . . 141
6.9 Dependence of fidelity on birefringence angle . . . 143

7.1 Schematic of time-bin encoded remote entanglement . . . 148
7.2 Level diagram for $^{88}Sr^{+}$ . . . 152
7.3 Example time-bin encoded scheme . . . 155
7.4 Temporal mixing in time-bin encoding . . . 157
7.5 Relative strengths of cavity decays in $P_{\frac{3}{2}} \rightarrow D_{\frac{5}{2}}$ transition . . . 159
7.6 Suggested scheme: magnetic field aligned to cavity axis . . . 160
7.7 Suggested schemes: magnetic field perpendicular to cavity axis . . . 161

8.1 Picture of MITAS cell . . . 172
8.2 Schematic diagram of MITAS apparatus . . . 173
8.3 Beam geometry for MITAS test . . . 175
8.4 Thermal model for MITAS system . . . 176
8.5 Temporal response of MITAS system . . . 177
8.6 Designs of 'Pico-Oven' system . . . 179
8.7 Images of 'Pico-Oven' parts . . . 180
8.8 Loading of 'Pico-Oven' system . . . 181
8.9 Chamber and beams for 'Pico-Oven' test . . . 182
8.10 Optical heating system for 'Pico-Oven' test . . . 183
8.11 Imaging and fluorescence collection system for 'Pico-Oven' test . . . 185
8.12 Observed ^{44}Ca isotope peak . . . 188
8.13 Measured beam profile for 'Pico-Oven' test . . . 190
8.14 Thermal model for 'Pico-Oven' system . . . 191
8.15 Image of gold coating for 'Pico-Oven' system . . . 192

G.1 Collection efficiency for detection of finite solid angle . . . 238

List of Tables

6.1 Summary of basis states for simulation of ion-birefringent cavity systems 124
6.2 Output wavepacket components for ion-birefringent cavity system . . 126

8.1 Properties of 'Pico-Oven' test system 188
8.2 Summary of measured and projected properties of atomic sources . . 195

1 Introduction

1.1 Quantum Computing

Quantum physics is the most accurate known theory for describing the behaviour of the world at small scales, and yet its signature phenomena, such as entanglement or wave-particle duality, are disturbingly unintuitive and defy everyday explanation [1]. In this theory, the state of a system is specified as a vector in a complex Hilbert space whose dimension is the product of the dimensions of the constituent systems. This makes it impossible to fully simulate even modest quantum systems on the most powerful supercomputers, as the full Hilbert space is too large to capture [2].

However, the inherent complexity of many-body quantum systems also lends them unique abilities for information processing [3]. Quantum systems could be used to directly simulate problems that are classically challenging, the most obvious examples being many-body quantum systems themselves, a technique known as quantum simulation [4]. But, much as classical computation in contemporary times mainly utilises digital classical computers rather than analogue classical simulators, there exists the possibility of creating a general-purpose programmable quantum information processor: a quantum computer [5]. In 1985, Deutsch showed that a quantum Turing machine could simulate everything a classical Turing machine could, but also solve certain problems more efficiently [6], laying the formal foundation for a quantum computer. Great encouragement for the practical possibility of such

a machine was taken when quantum error correction was shown to be possible [7, 8], dramatically reducing the (still demanding) requirements on the accuracy of elementary quantum logic gates.

The possible applications of quantum computers are still a subject of active research, but quantum algorithms are more efficient than the best known classical algorithms for certain problems such as prime number factorisation [9], unstructured searches [10], or finding properties of solutions to systems of linear equations [11]. The promise of solving hitherto intractable problems has inspired efforts to realise quantum computers experimentally, and a number of qubit types have been proposed, including trapped ions [12], superconducting circuits [13], NV centres [14], linear photonics [15], electronic spins in silicon [16], rare-earth ions [17], and neutral atoms [18].

1.2 Ion Trap Quantum Computing

1.2.1 Ion Traps

Ion traps are devices designed to confine ions to a specific region of space. To do so, they need to create a potential minimum, however, as expressed in Earnshaw's theorem, creating a confining electrostatic potential in all three spatial directions simultaneously is impossible. One solution to this problem is to trap particles via a combination of static electric and magnetic fields, known as a Penning trap [19, 20]. However, for quantum computing applications, the most common approach is to use a radiofrequency or Paul trap [21].

In a linear Paul trap, a static confining potential along the trap axis is combined with an oscillating saddle potential in the transverse directions. Charged particle motion in the non-axial directions is described by the Mathieu equations. These imply that for certain combinations of the saddle oscillation frequency and charge-to-mass ratio the ions are stably confined in a 'pseudopotential' [22, 23] with the addition of a smaller oscillatory 'micromotion' component [24]. Once confined, the trapped particles can be laser-cooled [25, 26], even to the ground state of the potential [27]. These traps were used to trap [28], perform spectroscopy [29], and explore quantum phenomena with single particles [30] prior to the interest in ion traps for quantum computing.

1.2.2 Ion Traps for Quantum Computing

The ability to trap a desired number of ions in a highly controlled environment led to proposals for ion traps to be used for quantum computing [31, 32]. In these systems, the quantum bits (qubits) are typically encoded in two electronic states of the ion which, for the case of singly-ionised Group II elements discussed in this thesis, has a single outer shell electron. Different types of qubit can be defined in the level structure [33], with distinct properties and strengths. For clock qubits [34] defined between certain levels within the hyperfine structure, the qubit frequency is first-order field insensitive, while optical qubits [35] offer a qubit transition at optical frequencies. This work will mostly be concerned with 'Zeeman' qubits within the atomic fine structure [36], which benefit from a long decay lifetime (T_1) and simple electronic structure, but are sensitive to magnetic field noise.

Operations on an individual qubits may be effected by manipulating the state with external radiation. Qubit-qubit interactions typically use the Coulomb interaction in a variety of ways [31, 37, 38], although Rydberg interactions can also be used [39]. These single and two-qubit gates can be performed using both lasers [40] and microwaves [41, 42] with outstanding fidelity [43–48] and two qubit gates have been demonstrated between different ionic species [49, 50]. Gates are typically much shorter than the coherence times of the qubits [51], which can extend to minutes [52]. However, operations do tend to be slower than in other architectures such as superconducting circuits, and increasing operation speed is an active area of research [39, 53].

Running a quantum algorithm would entail the trapping and cooling of ions, preparation of the initial electronic state through optical pumping [54, 55], the execution of the quantum gates, and then the readout of the final qubit states. This readout is usually achieved through state-dependent fluorescence [56–58], which determines the qubit state by observing whether the ion fluoresces strongly when illuminated with near-resonant laser light. Recently, the state-of-the-art in practical ion trap quantum computing has enabled error-correction steps to improve gate fidelity [59–61], although a demonstration of fault tolerant computation has not been made. More details on the application of trapped ions to quantum information processing can be found in [12, 62, 63].

1.3 Architectures for Scaling Ion Trap Quantum Computation

1.3.1 Scaling Problems with Long Ion Chains

Despite the impressive demonstrations of elementary quantum operations in ion traps, scaling to large numbers of ion qubits remains an outstanding challenge [12, 64]. High fidelity entangling gate demonstrations have used ions arranged in a linear chain along the trap axis interacting through their mutual coulomb repulsion. While it is possible to perform gates on arbitrary ions in a long chain through careful control of the interaction [65], such an arrangement is not easily scalable because extending to more qubits leads to spectral crowding of the motional modes of the chain that carry the interaction [66]. Though high fidelity, high rate entangling gates can be maintained for longer chains through careful control of the amplitude [67], frequency [68], and phase [69] of the driving pulse, such fine adjustment is not a practical solution for arbitrary ion numbers. The control of long-chain systems would also bring additional challenges: Harmonic traps do not produce uniform ion spacing [70] (extra control of the axial potential is required to do so [71, 72]), the probability that collisions with background gas and thermal reordering disrupt the chain increases as the chain becomes longer [73, 74], and optically-addressing individual ions with low crosstalk in these scenarios would be challenging [75].

1.3.2 QCCD Approach

One solution to the issues presented by a single, long chain of ions is to use a trap chip that can harbour multiple shorter ion chains in separate trapping potentials with ions individually shuttled throughout the chip using control voltages. This means that entangling gates can be performed on pairs or small chains of ions while controlling the quantum state of a much larger system of ions. Such a device is called a quantum charge coupled device (QCCD) [76] and requires traps with multiple zones and many individually controlled electrodes. Particular zones on these trap chips could be customised for specific purposes, for example gates or readout [77].

Traps designed according to the QCCD principle have already been used to perform interesting experiments, generating multi-qubit entangled states [78], performing teleported gates [79] and mid-circuit measurements [80], as well as realising topological phases [81] and simulating quantum dynamics [82]. These successes have established the QCCD approach as the leading route for scaling ion trap quantum processors in the near term, however, the approach is not without its drawbacks. The first concern is that increasing the qubit capacity would involve repeatedly developing traps with larger capacity. Secondly, and more fundamentally, in a QCCD approach, quantum information is always encoded in ions; this means the connectivity of the system is determined by the physical layout of the chip, and the speed of operations between qubits distant on the chip is determined by the maximum speed of shuttling. Improving shuttling speed is an active area of research [83–85], but these speeds are, and are likely to remain, far below relativistic levels.

1.3.3 Modular Quantum Computation with Photonic Interconnects

An alternative scaling approach — the approach considered for this thesis — forms a large-scale quantum computer from a collection of simpler 'nodes' joined by photonic interconnects. Each node contains a small QCCD, which at its simplest could consist of just a few qubits and a few trapping zones. Within each node, local entangling gates proceed through the coulomb interaction, while entanglement between nodes is achieved via transmission of single photons entangled with the trapped ion qubits that emitted them. These photons are routed through the photonic interconnects to perform remote ion-ion entanglement through an entanglement swapping measurement [86]. The combination of remote and local entangling operations with local single qubit gates and readout permits the modular system to host universal quantum computation [87].

To enact this proposal practically, further aspects must be considered. Firstly, the ideal properties of the ion that emits the single photons (the 'network ion') are convenient photon emission wavelengths and simple level structure for fast manipulation. These differ from, and at points oppose, the ideal properties of ions

for storing and manipulating quantum information ('computational ions'), which are long coherence times and high-fidelity gate operations, metrics generally aided by access to complex hyperfine structure in which to define the qubit. Secondly, remote entanglement fidelities may never reach the level of local entangling gates. Therefore, to establish the high-fidelity remote entanglement necessary for quantum computation, entanglement purification procedures [88, 89], which probabilistically convert multiple low quality Bell pairs into fewer, higher quality Bell pairs, may be required [90, 91], and there has been much work on understanding how networked quantum computation can achieve high fidelities with imperfect remote operations [92, 93]. These considerations mean that practical proposals for modular quantum computation usually contain more than one ion and more than one species per node [94]. The network ion is used to generate remote entanglement, which is swapped to the computational ion for more accurate information processing. This leaves the network ion free to generate another Bell pair for any purification steps. These multispecies proposals additionally allow for the computation ions to be cooled continuously through sympathetic cooling [95].

1.3.4 Comparison of Scaling Methods

Compared to the QCCD architecture, modular quantum computation brings additional challenges. Firstly, the state-of-the-art remote entanglement rate and fidelities ($182\ \mathrm{s}^{-1}$ at 94.0(5)% fidelity [96]) are far below what can be achieved with local gates[1]. In particular, the rate will have to improve dramatically or remote entanglement operations will be a serious bottleneck in a modular quantum computer. Secondly, the modular architecture requires a large number of nodes, which each include a trap, a vacuum system, control hardware and external optics. These nodes will have to be made much simpler and more reliable or a many-node quantum system would be prohibitively complex to commission and maintain. For this reason, the QCCD architecture is likely to maintain its advantage in the near term.

[1] Higher remote entanglement rates have been achieved in other platforms such as superconducting circuits (50 kHz at 79% [97], or 7 kHz at 77% [98]), but the fidelity of these operations is also much lower.

However, modular quantum computation does offer persuasive advantages. Firstly, while scaling in the QCCD approach requires the trap chip to be modified, in the networked approach the node itself does not need to change as the system size increases. This means that the network nodes could be made simple, reliable and easily maintainable, with the design reaching a greater level of refinement than a QCCD device. Secondly, the fibre network, rather than the physical layout of the nodes, sets the connectivity of the system, and the potential for higher connectivity affords the possibility of more efficient algorithms and error correction [99]. Finally, information propagates through the network via photons travelling close to the speed of light in free space, considerably faster than reasonable transport speeds in a QCCD chip. These advantages in the simplicity of scaling, connectivity, and information transport speed become more significant as the system size increases, making networked quantum computation a compelling long-term option.

1.3.5 Rationale for Ion-Cavity Integration

For useful networked quantum computation, the rate of remote entanglement must be increased. The rate achieved in state of the art demonstrations is fundamentally constrained by the limited photon collection efficiency afforded by refractive optics [96]. Optical cavities can increase the efficiency of photon collection, in principle towards unity [100], which could provide orders of magnitude increase in remote entanglement rate.

Though less developed than atom-cavity QED [101–103], single trapped ions have been coupled to optical cavities and used to map out the optical field [104], reach the 'strong coupling' regime [105], and produce ion-photon entanglement [106, 107]. Remote ion-ion cavity-enhanced entanglement has recently been demonstrated [108], although neither the rate nor the fidelity demonstrated equalled those obtained with refractive optics. The aim of this thesis is to understand the physical systems and experimental protocols necessary to perform cavity-enhanced high-rate high-fidelity remote entanglement.

1.3.6 Development of Ion Traps for Scalable Quantum Computation

Both QCCD and network approaches to scaling require traps with many controllable electrodes to shuttle the ions and ideally a small footprint [109]. Much of the recent work in the field has thus been directed towards the development of surface electrode traps [110, 111], where a planar electrode geometry produces a potential minimum above the trap surface [112, 113]. These traps offer compelling advantages in the ease of fabrication [114] with electrode geometries defined by standard lithography [115].

However, three dimensional traps, offer advantages in that they can provide deeper, more harmonic traps with higher ion lifetimes [116, 117] and shield the trapping potential more effectively from the disturbance induced by the proximity of the dielectric mirror surfaces required when using microcavities for photon collection [118]. Due to difficulties in manufacturing the necessary structures, microfabricated 3D traps have typically involved multiple stacked layers [44, 119] or complex fabrication processes [120, 121]. However, recent advances in fabrication have permitted the production of monolithic 3D microtraps through laser-enhanced etching [122].

Desires to improve stability and simplicity have led to increasing integration of functional components into or around the trap structure [123, 124], including waveguides [125, 126], detectors [127], fibres [128], mirrors [129–131], current-carrying wires [132] and voltage controllers [133]. In this spirit, the envisaged cavity-integrated trap will be taken to be a microfabricated, monolithic 3D structure, with a cavity integrated into the trap itself. This broad specification is the motivation behind the work in this thesis, although an exact physical construction is not crucial to the research presented.

1.4 Thesis Context

I began my DPhil in October 2018 as one of the founding members of the Ion Trap Group's Cavity Team with fellow DPhil student Shaobo Gao and supervisor Dr. Joseph Goodwin. The aim of this project was (and remains) to build a scalable cavity-enhanced quantum network with trapped ions. During the first year, the dedicated project funding was not yet available and my efforts were mainly devoted

to understanding the physics of ion-cavity interactions, assisting Shaobo on work concerning temporal impurities in photon wavepackets that would feature in his thesis, and laying the intellectual foundation for our method of optimising cavities [134]. I also worked with the other members of the team and Dr. Tim Ballance of ColdQuanta UK on a project to test miniature atomic sources [135], taking some preliminary data, analysing the published data and revising the published manuscript. In the second year, the team was able to purchase its own capital equipment, and as a preliminary project we began establishing apparatus to automatically mode match and characterise cavities. In March 2020, I helped Master's student Sebastian Saner set up the apparatus with which he and Shaobo were able to characterise our first microcavities, with the results a key part of a future publication on automated mode matching that is currently in preparation[2].

However, in March 2020, the disruption due to the coronavirus pandemic began. In the immediate term, this prevented the continuation of lab work, and we switched our focus to theoretical studies that could be conducted remotely. During this time, I began to look closely at the impact cavity birefringence could have on the remote entanglement protocols we were planning. Working with Dr. Ezra Kassa (who had joined the team for a short postdoc role before a planned move to Okinawa), I modelled the effect of birefringence, finding that realistic levels of birefringence could be problematic for the experiments we were considering. These results led to a publication [136], for which I conducted the numerical modelling, produced the data figures, and wrote the analysis text with Dr. Joseph Goodwin.

The initial lockdown began to ease, but disruption to experimental efforts continued, as we were unable to access many specialist facilities and key collaborations were impacted. During the months towards to the end of 2020, I continued to work on the papers that we had started, while inspecting prototypes for our initial trap designs. At the beginning of 2021, with pandemic lockdowns drawing to a close, the decision was made that, along with completing my various strands of theoretical and simulation work, I should pivot towards atomic source development as the pre-pandemic plan of developing the trap system was no longer realistic within

[2]This work will imminently be completed with data from a long-awaited batch of high-finesse microcavities, the production of which fell victim to pandemic related disruptions.

the DPhil. I developed the apparatus to test an improved design of miniature atomic source along with Dr. Jake Blackmore, observing neutral atom flux in the Autumn of 2021. Tests of these sources continued for the remainder of the DPhil, first with the help of Dr. Jake Blackmore and later Dr. Thomas Doherty as the designs of Dr. Joseph Goodwin evolved.

Towards the end of my DPhil, I returned to study further some promising ideas that I had begun to develop in my first year regarding the impact of mirror profile variations on cavity loss. Different methods of producing micromirrors result in substrates that may deviate from an ideal spherical shape, and we were particularly interested in whether non-spherical mirrors could be used in cavities with residual transverse misalignment, where the cavity mode would not intersect the mirror at the nominal centre. To this end, I led a collaboration with Dr. Peter Horak at the University of Southampton to efficiently model mirrors with transverse misalignment in an extension to the 'mode mixing' method. The physics uncovered in this investigation was considerably richer than anticipated, and the methods I developed provide a very powerful way of exploring this. These methods, and the results that emerged through their application, form the basis of two publications currently in preparation.

The work presented, while largely my own, was conducted in a collaborative group setting. More detailed descriptions of my own contribution, and the contribution of others, are given in the 'Contribution and Acknowledgements' section at the start of each chapter. The author would also like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out the work presented in the thesis. doi:10.5281/zenodo.22558

1.5 Thesis Outline

Chapter 2 will begin by covering the theory of the ion-cavity interaction and how this can lead to an enhanced collection rate. The basic mechanism of different remote entanglement methods will be presented and the choice to focus on two-photon remote entanglement schemes justified. Chapter 3 will discuss different atomic processes that can result in cavity-enhanced photon emission, and present

a novel approach for optimising the photonic output. Chapter 4 will present a simplified model developed in our group that encodes some of the compromises inherent in cavity design, resulting in a simple procedure to optimise cavity geometry given manufacturing tolerances. In Chapter 5, Developments to the mode mixing method for calculating cavity eigenmodes, particularly in the case of transverse mirror misalignment, will be presented. These are used to study the impacts of transversely misaligned and Gaussian-shaped cavity mirrors to understand which misalignments and deformations to the cavity mirror structure could be tolerated in a remote entanglement experiment. Chapter 6 will focus on a remote entanglement protocol using polarisation-encoded photonic qubits, discussing the impact that cavity birefringence can have on this type of entanglement scheme. In Chapter 7, an alternative approach, using time-bin encoded photons, is introduced and specific schemes for the remote entanglement protocol are suggested. Chapter 8 details another strand of work which concerns efforts to developed a miniature optically-heated atomic source for integration into ion trap systems. Finally, Chapter 9 presents the conclusions of the study and identifies areas for further research.

2
Cavity Theory

2.1 Introduction to Cavities

The light-matter interaction is a fundamental part of trapped ion quantum information processing, finding application in cooling and quantum logic operations. Spontaneous decay of excited electronic states can also be understood through the same interaction, but with quantum fluctuations of vacant vacuum modes in place of the real photons in coherent electric fields. These vacuum fluctuations are not immutable; adjustments to the boundary conditions of the electric field, such as those provided by the mirrors of optical cavities, can shape these fluctuations [137]. These changes in the vacuum field can cause radical changes to emission properties including, under the right conditions, near-unit probability emission into a cavity mode of modest numerical aperture. In this way, cavities can overcome inherent limitations of refractive optics, potentially taking remote entanglement to a new speed regime. The physics of the ion cavity interaction and enhanced collection of photons has been covered in great detail elsewhere [100, 138, 139], and this chapter will provide a summary of the key results that form the foundation for the work presented in the subsequent chapters.

Contribution and Acknowledgements

No original research is presented in this chapter.

2.2 Introduction to Cavity Modes

2.2.1 Maxwell's Equations

To begin, the resonant modes of optical cavities will be introduced. For the purposes of this document, a rigorous description of the quantised cavity mode, for example through expansion in modes of the universe [140, 141] or the input-output formalism [142], is unnecessary, and the field inside the cavity will be analysed classically on the understanding that the dynamics of the quantised field should not contradict classical theory. The quantised nature of the field only features when single photons occupying the cavity modes are measured [143].

In classical optics, the electric field $\boldsymbol{E}$ with temporal frequency ω inside a cavity should satisfy Maxwell's equations in free space

$$\begin{aligned} \boldsymbol{E}(x,y,z,t) &= \boldsymbol{\epsilon}(x,y,z)e^{i\omega t}, \\ \left(\nabla^2 + k^2\right)\boldsymbol{\epsilon} &= \mathbf{0}, \\ k &= \frac{\omega}{c}, \\ \boldsymbol{\nabla}\cdot\boldsymbol{\epsilon} &= 0, \end{aligned} \tag{2.1}$$

using the notations and conventions in [144], where $\boldsymbol{\epsilon}$ is the spatial dependence of the electric field and $\boldsymbol{k}$ the wavevector.

2.2.2 Cavity Types and Boundary Conditions

An optical cavity confines light oscillating at optical frequencies to a restricted region of space by imposing boundary conditions of the solutions to Maxwell's equations. While there are many different types of cavity such as ring [145] or photonic crystal resonators [146], for the purposes of integration with a trapped ion, the type of cavity utilised is a concave Fabry-Perot resonator, consisting of two concave mirrors facing each other. Such cavities offer high optical access [147] and do not require the emitter to be near the cavity structure. The latter is particularly important for trapped ions as the dielectric mirrors that offer the highest optical reflectivities distort the trapping potential [118], harbour stray charges [148, 149], and introduce heating [150, 151]. The cavity could be asymmetric, which can confer advantages [152], but to keep

the ion distant from the cavity structure, symmetric concave Fabry-Perot cavities are preferred, using two mirrors of equal curvature with the ion in the centre of the cavity.

The boundary conditions imposed by the mirrors of the Fabry-Perot cavity are complex in their microscopic detail [153], but for simplicity, we will assume that the mirrors are perfectly reflective. In this case, the electric field parallel to the mirror surface for the incoming wave should cancel that of the outgoing wave. It should be noted that cavity mirrors are not perfectly reflective, either due to inevitable scattering or absorption, or by transmission of the cavity mode through a partially transmissive mirror, which is crucial for many cavity applications including ours. This means that the concept of an internal, isolated cavity mode is not strictly correct as the cavity mode must couple to the environment [154]. In this chapter, the cavity mode is calculated on the assumption of perfect reflection, and loss then introduced to mimic coupling to external modes [142]. This approach is appropriate for resonators with high quality factors such as those considered here, but is not suitable for lossier resonators [155].

2.2.3 The Paraxial Equation

In light of the boundary conditions imposed by the Fabry-Perot cavity, approximate solutions to the intracavity field can be derived. First, given that the expected solution should propagate between the mirrors of the Fabry-Perot cavity, a plane wave dependence will be separated from the solutions of Maxwell's equations. This is achieved substituting into Eq. 2.1 the solution $\boldsymbol{\epsilon} = \boldsymbol{G}e^{\mp ikz}$, where $\boldsymbol{G}$ collects the non-plane wave behaviour, with z defined by the axis of the Fabry-Perot cavity and the sign $\mp$ indicating propagation towards positive and negative z respectively. This leads to

$$\left(\nabla_T^2 \mp 2ik\frac{\partial}{\partial z} + \frac{\partial^2}{\partial z^2}\right)\boldsymbol{G} = 0, \tag{2.2}$$

$$\boldsymbol{\nabla}_T \cdot \boldsymbol{G}_T \mp ikG_L + \frac{\partial G_L}{\partial z} = 0, \tag{2.3}$$

$$\boldsymbol{G} = (G_x, G_y, G_L),$$

$$\boldsymbol{G}_T = (G_x, G_y),$$

with $\boldsymbol{\nabla}_T = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}\right)$ the transverse derivative operator. The paraxial approximation takes Eq. 2.2 and assumes that, because the beam is directed along z, transverse variation to the beam function $\boldsymbol{G}$ will be much stronger than axial, and thus the second derivative with respect to z will be small compared to the other terms. Note that this does not guarantee the fulfilment of the divergence condition (Eq. 2.3), which will become important when discussing cavity birefringence in Chapter 6. This approximation yields the paraxial equation

$$\left(\nabla_T^2 \mp 2ik\frac{\partial}{\partial z}\right)\boldsymbol{G}_T = 0. \tag{2.4}$$

The electric field may be separated into a polarisation and a scalar part $\boldsymbol{G}_T = \hat{\boldsymbol{\epsilon}}u$, where $\hat{\boldsymbol{\epsilon}}$ is a fixed polarisation direction that lies orthogonal to the beam axis z. The scalar mode function u relates to the electric field through

$$\boldsymbol{E}(x, y, z, t) = \hat{\boldsymbol{\epsilon}}u(x, y, z)e^{-ikz}e^{i\omega t}. \tag{2.5}$$

The majority of cavity mode physics can be understood from the solutions of the paraxial equation in the scalar cavity function u:

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \mp 2ik\frac{\partial}{\partial z}\right)u = 0. \tag{2.6}$$

The advantage of separating the plane wave dependence and then taking the paraxial approximation is that Eq. 2.6 can be solved conveniently with the boundary conditions imposed by the Fabry-Perot cavity. This condition is that the scalar function u should have an equiphase surface along the mirror. The solutions of interest are therefore the 'Gaussian beams', which are specific transverse field configurations with constant curvature equiphase surfaces to match the constant curvature of the concave mirror. The transverse extent of the beam is characterised by a 'waist' $w(z)$, which is the $1/e$ extent of the amplitude envelope. This envelope narrows to and expands from a central waist w_0, chosen without loss of generality to lie at $z = 0$. There is some freedom in the functions used to describe the Gaussian

beam solutions to the paraxial approximation, but throughout this document the functions used will be the Hermite-Gauss solutions[1].

Cavity function:

$$u^{(\pm)}_{n_x,n_y}(x,y) = aH_{n_x}\left(\frac{\sqrt{2}x}{w(z)}\right)H_{n_y}\left(\frac{\sqrt{2}y}{w(z)}\right)e^{-\frac{x^2+y^2}{w(z)^2}}e^{\mp ik\frac{x^2+y^2}{2R_u(z)}}e^{\pm i(n_x+n_y+1)\Psi_G},$$

Amplitude normalisation: $a = \frac{1}{w(z)}\sqrt{\frac{2}{\pi}\frac{1}{2^{n_x+n_y}n_x!n_y!}}$,

Beam waist: $w(z) = w_0\sqrt{1+\left(\frac{z}{z_0}\right)^2}$, (2.7)

Rayleigh range: $z_0 = \frac{\pi w_0^2}{\lambda}$,

Wavefront radius of curvature: $R_u(z) = z\left(1+\left(\frac{z_0}{z}\right)^2\right)$,

Gouy phase: $\Psi_G(z) = \arctan\left(\frac{z}{z_0}\right)$,

where the wavelength $\lambda = 2\pi/k$, and H_i are the Hermite polynomials with $n_x, n_y \in \mathbb{N}$ the x and y the indices of the transverse profile and z_0 the Rayleigh range of the beam.

2.2.4 Conditions for Cavity Resonance and Mode Structure

For a Fabry-perot cavity with mirrors of radius of curvature R and axial separation L, the Gaussian beam solution of Eq. 2.7 must satisfy $R_u(L/2) = R$, leading to

$$w_0 = \sqrt{\frac{\lambda L}{2\pi}\sqrt{\frac{2R}{L}-1}}, \tag{2.8}$$

$$z_0 = \frac{L}{2}\sqrt{\frac{2R}{L}-1}. \tag{2.9}$$

Therefore the cavity length and mirror curvature set the central waist w_0 of the mode, and, in turn the length scale z_0 over which this waist expands along the cavity axis. The condition that the cavity has a resonant Gaussian mode can be expressed in terms of the stability parameters $g_1^{(s)} = 1 - L/R_1$ and $g_2^{(s)} = 1 - L/R_2$, where R_1 and R_2 are

[1]It is also common to express these solutions in terms of Laguerre-Gauss polynomials [156], but the choice of Hermite-Gauss will be convenient for Chapter 5.

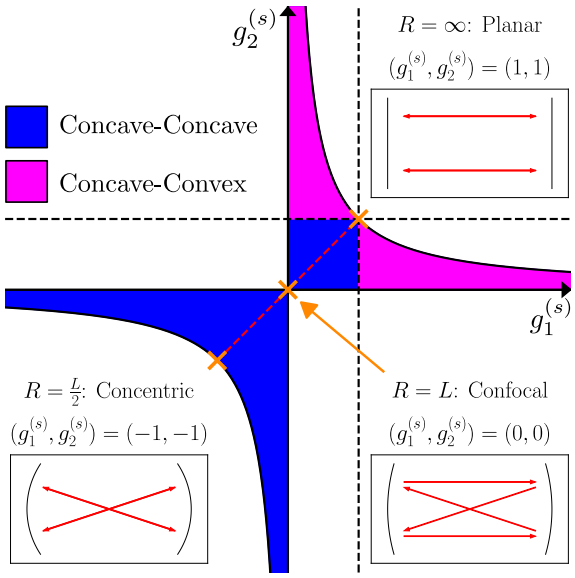


Figure 2.1: Stability diagram for Fabry-Perot resonators in terms of the stability parameters $g_1^{(s)}$ and $g_2^{(s)}$. Stable configurations, which satisfy $0 \leq g_1^{(s)} g_2^{(s)} \leq 1$, are highlighted, in blue for the case that both mirrors are concave, and purple where one mirror is concave and the other convex. Restricting to symmetric cavities means that $g_1^{(s)} = g_2^{(s)}$, and a stable cavity must occupy a configuration on the red line. At one end of this line is the planar configuration, where both mirrors are perfect planar surfaces. On the other end is the concentric configuration, where the mirror surfaces are opposite sides of the same sphere. In the middle is the confocal configuration, where the centre of curvature of one mirror lies on the surface of the other.

the radii of curvature of the two mirrors. The condition that there is a resonant mode is $0 \leq g_1^{(s)} g_2^{(s)} \leq 1$, which, for the case of a symmetric cavity, leads to $0 < L/R < 2$. A planar cavity has $L/R = 0$ (the radii of curvature of the mirrors are infinite on the scale of the cavity), while $L/R = 2$ is a 'concentric' cavity, where the mirror profiles are two diametrically opposed caps of a single common sphere. The cavity stability criterion and important cavity configurations are shown schematically in Fig. 2.1.

The mode divergence is described by Eq. 2.7. At distances far beyond z_0 either side of the centre, the beam waist expands according to

$$\begin{aligned} \lim_{z \to \infty} w(z) &= \theta z, \\ \theta &= \frac{\lambda}{\pi w_0}, \end{aligned} \tag{2.10}$$

where θ is the divergence angle of the mode. In highly planar configurations, the waist is very wide, and the beam diverges slowly; in near-concentric configurations, the waist is small and the divergence angle is large.

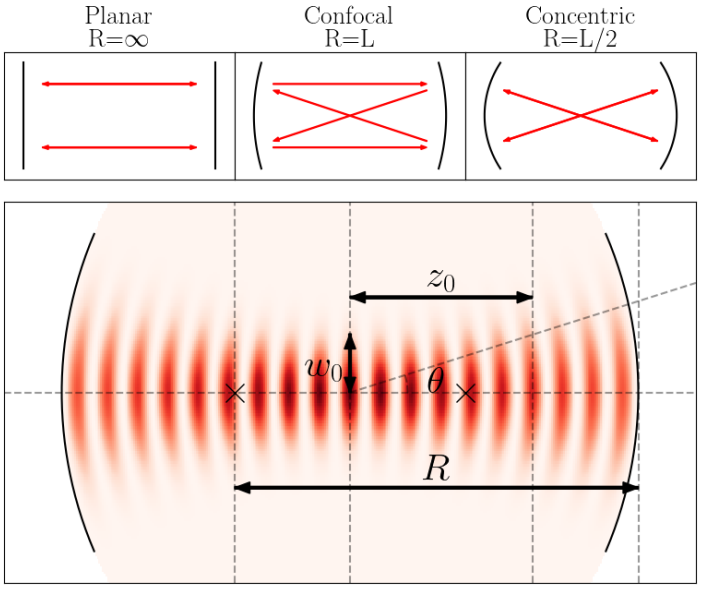


Figure 2.2: Diagram of a symmetric Fabry-Perot cavity and the resonant fundamental mode. Top: Ray diagrams indicating the range of configurations of a Fabry-Perot cavity. Configurations exist between planar, where both mirrors are flat, and concentric, where the mirrors share a common centre. Bottom: The fundamental mode of a Fabry-Perot resonator, showing the standing wave with central waist w_0, Rayleigh range z_0 and divergence angle θ. The centres of the spheres defining the two mirrors of radius of curvature R are marked by crosses.

The cavity parameters L and R, the mode parameters w_0, z_0, and θ, and the valid geometries for a stable optical cavity are shown schematically in Fig. 2.2.

However, there is a second requirement on the cavity mode, which is that optical fields from successive round trips must constructively interfere. This is ensured if the phase φ_{RT} accrued by the travelling wave solution over one round trip of the cavity is a multiple of 2π. This phase depends upon the cavity length, the wavelength of light and the transverse mode indices

$$\varphi_{\mathrm{RT}} = 4\pi\frac{L}{\lambda} - 4(n_x + n_y + 1)\arctan\left(\sqrt{\frac{L}{2R-L}}\right). \tag{2.11}$$

Rearranging for frequency, the resonant frequencies of the cavity are

$$
\begin{aligned}
f_{N_{\mathrm{an}},n_x,n_y} &= \nu_{\mathrm{FSR}} N_{\mathrm{an}} + \nu_T (n_x + n_y + 1), \\
\nu_{\mathrm{FSR}} &= \frac{c}{2L}, \\
\nu_T &= \nu_{\mathrm{FSR}} \frac{2}{\pi} \Psi_G = \nu_{\mathrm{FSR}} \frac{2}{\pi} \arctan\left(\sqrt{\frac{L}{2R-L}}\right), \\
&= \nu_{\mathrm{FSR}} \frac{1}{\pi} \arccos\left(\pm\sqrt{g_1^{(s)} g_2^{(s)}}\right),
\end{aligned}
\tag{2.12}
$$

with $N_{\mathrm{an}} \geq 1$ the integer number of antinodes in the cavity standing wave. This expression defines two important frequencies: The *free spectral range* ν_{FSR} is the frequency increase for each additional node in the longitudinal standing wave, and the *transverse mode splitting* ν_T is the frequency increase for adding an excitation to either transverse direction, meaning modes of the same excitation index $\mathcal{I} = n_x + n_y$ are degenerate. The expression for this frequency can be written in terms of the Gouy phase Ψ_G, cavity parameters L and R, or in terms of the stability parameters. This last expression is not restricted to symmetric cavities, and the sign $\pm$ should match the sign of $g_1^{(s)}$ and $g_2^{(s)}$.

All cavities resonances are lossy. These losses derive from the mirror roughness (leading to scattering), absorption in the mirror coatings or other intra-cavity element, as well as the transmission that is necessary if light is to be coupled into or extracted from the cavity. The round-trip loss $\mathcal{L}_{\mathrm{RT}}$ is proportional to the cavity power decay rate 2κ through

$$
2\kappa = \nu_{FSR} \mathcal{L}_{\mathrm{RT}}. \tag{2.13}
$$

This lends the resonant peaks a FWHM of $2\kappa/2\pi$. An oft-quoted characteristic of cavities is the finesse

$$
F \approx \left(\frac{2\pi}{2\kappa}\right) \nu_{FSR} \approx \frac{2\pi}{\mathcal{L}_{\mathrm{RT}}}, \tag{2.14}
$$

which, for high finesse cavities, is the ratio of free spectral range to resonant linewidth. Cavity mode structure can be measured spectroscopically, by coupling monochromatic laser light into the cavity and monitoring transmitted or reflected optical power while scanning the laser frequency or cavity length. A schematic cavity spectrum showing cavity mode structure and linewidths is presented in Fig. 2.3.

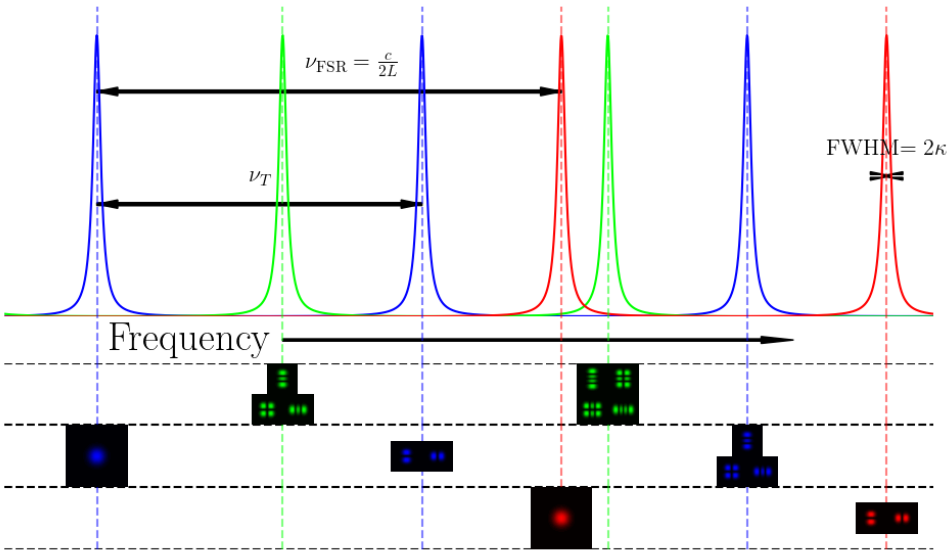


Figure 2.3: A schematic diagram of the transmission spectrum of a Fabry-Perot cavity. Note that the diagram does not show a complete set of cavity modes. Top: Cavity enhancement of selected modes from an example section of a cavity spectrum. The spectrum is composed of a series of distinct resonances, with each resonance approximately Lorentzian with a FWHM equal to the cavity power decay rate 2κ. The modes of the same longitudinal order are in the same colour and form a series of states separated by frequency ν_T, where increasing frequency increases the transverse order. Different colours represent different longitudinal orders, separated by frequency ν_{FSR}, with the green resonances containing one fewer longitudinal node than the blue, and the red containing one more. Bottom: Transverse mode structure corresponding to each resonance. The lowest transverse order for each longitudinal order is the fundamental, with a single mode, the next is the first order excitation with two degenerate modes, and this degeneracy increases by one for every successive order. Different orders of longitudinal and transverse modes can overlap in frequency if ν_T and ν_{FSR} are related by an integer fraction, which occurs at particular lengths. Note that seeing a mode on a transmission spectrum requires the probe laser to couple spatially into the pertinent mode in addition to matching its frequency, and therefore measured cavity spectra do not typically have uniform peak heights across transverse orders.

Cavities for enhancement of ion-ion remote entanglement are designed to operate at wavelengths in or near the visible, but geometric considerations, such as keeping the cavity mirrors further from the ion than the trap electrodes, mean cavities are typically several hundred micrometres long. Therefore, the cavity length will be much longer than the wavelength, and there will be many nodes and antinodes in the longitudinal standing wave. For the case discussed in this work, where the enhanced transition is a dipole transition and no coupling to the ion motion is intended, the ion should be placed at an antinode to maximise coupling. In the

transverse direction, the mode of choice is the simplest, 'fundamental' mode, with $n_x = n_y = 0$. This is preferred for coupling to an emitter over higher order transverse modes because it is the most transversely compact. Firstly, this provides the highest possible concentration of field at the centre and, secondly, the mode interacts with a smaller region of the mirrors, and is therefore generally less lossy for mirrors of finite size. Typically, cavity QED experiments will couple an emitter to just this transverse mode, which is isolated in frequency as long as the transverse splitting ν_T is not an integer fraction q/p of the free spectral range. For a given radius of curvature these frequency coincidences occur at particular lengths

$$L_{p,q} = 2R \left[\frac{\tan^2\left(\frac{\pi q}{2p}\right)}{1 + \tan^2\left(\frac{\pi q}{2p}\right)} \right], \tag{2.15}$$

for which all transverse modes where $\mathcal{I} = n_x + n_y$ is a multiple of p have overlapping frequencies. Here, p is the difference in transverse excitation index $\mathcal{I}$ between the co-resonant modes, and q the difference in longitudinal index N_{an}

For cavities with a fixed radius of curvature and length, the frequency of the spectroscopy laser can be adjusted to hit a resonance (as in Eq. 2.12). However, when coupling to an emitter in the cavity, the frequency will be set by the electronic transition of interest. To make the cavity resonant with the target frequency, the length is scanned, with a change in length of $\lambda/4$ required to tune the cavity to resonance with the fundamental transverse mode.

2.2.5 Misaligned Cavities

When the mirrors of a cavity are transversely misaligned, in order for the mode to encounter both mirrors at normal incidence, the mode axis must tilt relative to the optical axis of the mirrors, following the line defined by the centres of curvature of the two mirrors, as shown in Fig. 2.4. The mode angle relative to the optical axis of the mirrors,

$$\phi = \frac{\Delta x}{2R - L}, \tag{2.16}$$

for transverse mirror misalignment Δ_x thus increases as the cavity approaches the concentric configuration, where the centres of the two spheres coincide. As mode tilt

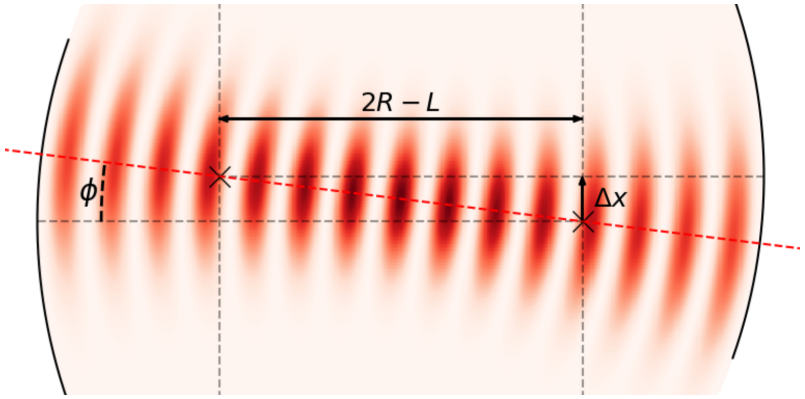


Figure 2.4: Diagram of the mode of a cavity whose mirrors are misaligned transversely by Δx. The mode axis tilts by an angle ϕ so that the mode is propagates normal to both mirrors at intersection. This means that the mode axis passes through the centres of both spheres that define the mirrors, marked by crosses.

increases, more of the mode will escape the confines of the finite diameter mirrors, which increases mode loss. The impact of residual misalignment on cavity performance will be discussed in Chapter 4 for spherical mirrors using a simple model, and then in Chapter 5 for more general mirror shapes using a more sophisticated approach.

2.3 Cavity Coupling

2.3.1 Coherent Interaction

The cavity modes described in Sec. 2.2 can be coupled to an emitter within the cavity. The simplest system to understand features an emitter with just two levels, a ground state $|g\rangle$ and an excited state $|e\rangle$, which are coupled by a dipole-allowed transition and a single cavity mode with annihilation operator a as depicted in Fig. 2.5. The ion and cavity mode vacuum fluctuations interact through the Hamiltonian

$$H_{\text{ic}} = \hat{\boldsymbol{d}} \cdot \hat{\boldsymbol{\epsilon}}^{\boldsymbol{V}}, \tag{2.17}$$

where $\hat{\boldsymbol{d}}$ is the dipole operator with transition matrix element $\boldsymbol{d_{eg}} = -\langle e| e\boldsymbol{r} |g\rangle$ where $\boldsymbol{r}$ is position of the outer electron relative to the nucleus, and $\hat{\boldsymbol{\epsilon}}^{\boldsymbol{V}}$ is the electric field operator at the position of the emitter, with $\boldsymbol{\epsilon}^{\boldsymbol{V}}$ corresponding electric field.

Ion trap quantum information experiments typically use a fixed magnetic field to split the degeneracy of electronic states with different angular momenta. The isotopes of interest for networking will be assumed to have no nuclear spin to simplify

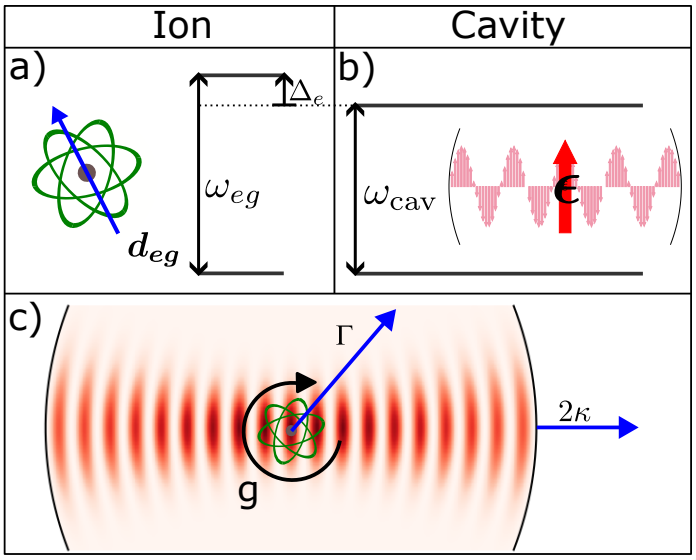


Figure 2.5: Schematic diagram of ion-cavity coupling. a) The ion has an excited state $|e\rangle$ and ground state $|g\rangle$, where the transition between the states has an angular frequency ω_{eg} and associated dipole moment $\boldsymbol{d_{eg}}$. b) The cavity has a resonant mode with angular frequency ω_{cav} and vacuum fluctuation electric field $\boldsymbol{\epsilon}$ in the centre of the cavity. The cavity mode is detuned by Δ_e from the atomic transition frequency. c) When the ion is located within the cavity mode, the electronic transition couples to the cavity mode, even if the mode is vacant. The cavity QED system can be described by three rates: Γ, the decay rate of state $|e\rangle$, 2κ, the power decay rate of the cavity mode, and g, the coherent coupling strength between the electronic and cavity mode excitation.

the level structure. Thus, the cavity can address a single transition $|e\rangle \rightarrow |g\rangle$ in the fine structure of the ion with state $|e\rangle$ ($|g\rangle$) having electronic angular momentum $J = J_e$ (J_g) with projection $m_J = m_{J_e}$ (m_{J_g}) onto the quantisation axis. The dipole strength of this fine transition is

$$\boldsymbol{d_{eg}} = \langle J_e || e\boldsymbol{r} || J_g \rangle \left\langle J_e, m_{J_e} \middle| J_g, 1, m_{J_g}, q \right\rangle \tag{2.18}$$

where $q = m_{J_e} - m_{J_g}$ and $\langle J_e || e\boldsymbol{r} || J_g \rangle$ is a reduced matrix element, a result of the Wigner-Eckart theorem [157].

The vacuum fluctuation strength is typically evaluated by setting the energy in the vacuum field equal to $\hbar\omega_{\mathrm{cav}}/2$, but can be derived more rigorously [153]. The

vacuum fluctuations can be expressed as components

$$\begin{aligned}\boldsymbol{\epsilon^V} &= \epsilon^V_{-1}\hat{\boldsymbol{e}}_{\boldsymbol{-1}} + \epsilon^V_0\hat{\boldsymbol{e}}_{\boldsymbol{0}} + \epsilon^V_{+1}\hat{\boldsymbol{e}}_{\boldsymbol{+1}} \\ \hat{\boldsymbol{e}}_{\boldsymbol{-1}} &= \frac{1}{\sqrt{2}}(\hat{\boldsymbol{e}}_{\boldsymbol{x_B}} - i\hat{\boldsymbol{e}}_{\boldsymbol{y_B}}) \\ \hat{\boldsymbol{e}}_{\boldsymbol{0}} &= \hat{\boldsymbol{e}}_{\boldsymbol{z_B}} \\ \hat{\boldsymbol{e}}_{\boldsymbol{+1}} &= \frac{1}{\sqrt{2}}(-\hat{\boldsymbol{e}}_{\boldsymbol{x_B}} - i\hat{\boldsymbol{e}}_{\boldsymbol{y_B}})\end{aligned} \tag{2.19}$$

in a spherical basis defined by the magnetic field [158], where $(\hat{\boldsymbol{e}_{\boldsymbol{x_B}}}, \hat{\boldsymbol{e}}_{\boldsymbol{y_B}}, \hat{\boldsymbol{e}}_{\boldsymbol{z_B}})$ are a right-handed coordinate system with $\hat{\boldsymbol{e}}_{\boldsymbol{z_B}}$ aligned along the magnetic field. Conservation of angular momentum dictates that the electric field will only couple to the electronic transition if the projection of momentum in the quantisation direction is conserved. This means that a specific dipole transition will only couple to one component of vacuum fluctuations $\epsilon^V_q\hat{\boldsymbol{e}}_{\boldsymbol{q}}$.

The cavity coupling

$$\begin{aligned}g &= \sqrt{\frac{3\Gamma_{J_e,J_g}\lambda^2 c}{8\pi V}}\left\langle J_e, m_{J_e}\middle| J_g, 1, m_{J_g}, q\right\rangle \frac{\epsilon^V_q}{|\boldsymbol{\epsilon^V}|} \\ V &= \frac{\pi L w_0^2}{4}\end{aligned} \tag{2.20}$$

can then be expressed, where Γ_{J_e,J_g} is the partial linewidth of spontaneous decay from the manifold J_e to J_g [159]. This form splits the cavity coupling into three parts. The first depends on the intrinsic strength both of the transition between manifolds and the cavity field, the second encodes the geometric strength of the specific angular momentum transition and the third weights the alignment of the cavity field with the electronic dipole. This leads to the coherent Hamiltonian for the entire ion-cavity system

$$H = g(a^\dagger |g\rangle\langle e| + a|e\rangle\langle g|) + \Delta_e|e\rangle\langle e|, \tag{2.21}$$

where a is the annihilation operator of the cavity mode, and Δ_e is the detuning of ω_{eg} above cavity resonance (see Fig. 2.5). This Hamiltonian has been expressed in a frame rotating with the cavity mode.

2.3.2 Incoherent Interaction

As integral to the emitter-cavity interaction as the coherent coupling are incoherent processes, specifically cavity loss and spontaneous emission. The evolution of the coupled ion-cavity system can be simulated with the master equation

$$\begin{aligned} \frac{d}{dt}\rho &= -i[H,\rho] + \hat{L}(\rho), \\ \hat{L}(\rho) &= \sum_n (2C_n\rho C_n^\dagger - \rho C_n^\dagger C_n - C_n^\dagger C_n \rho), \\ C_0 &= \sqrt{\Gamma}\,|g\rangle\langle e|\,, \\ C_1 &= \sqrt{2\kappa}a, \end{aligned} \tag{2.22}$$

where ρ is the density matrix of the system and Γ is the spontaneous emission rate from $|e\rangle$.

However, on the assumption that after an incoherent process the system state never returns to the simulated Hilbert space, the quantum state can instead be calculated through a non-Hermitian Hamiltonian approach, which is simpler to model and understand:

$$\begin{aligned} H^{\mathrm{NH}} &= H - i\gamma\,|e\rangle\langle e| - i\kappa a^\dagger a, \\ i\frac{\partial\psi}{\partial t} &= H^{\mathrm{NH}}\psi, \end{aligned} \tag{2.23}$$

where H^{NH} is the non-Hermitian Hamiltonian, $\gamma = \Gamma/2$ is the amplitude decay rate of the excited state, and ψ the state of the ion-cavity system. The assumption is usually reasonable for cavity loss processes if $|g\rangle$ is long-lived, but for spontaneous emission this is always an assumption which is only reasonable under specific circumstances. More discussion on this topic is given in Chapter 3.

2.3.3 Cooperativity and Cavity-enhanced Photon Collection

Different regimes of cavity physics can be accessed through different ratios of the three rates g, κ and γ. A useful dimensionless parameter for comparing the coherent and incoherent processes in an emitter cavity system is the cooperativity [160]

$$C = \frac{g^2}{2\kappa\gamma}. \tag{2.24}$$

This parameter has many applications, but in the context of cavity-enhanced networking, its major role is in determining the long-time limit of photon generation in a cavity as $2C/(2C+1)$ [161], which will be discussed further in Chapter 3. It has an alternative formula [134]

$$\begin{aligned} C &= 6\alpha_B \left(\frac{\epsilon_q^V}{|\boldsymbol{\epsilon}^V|}\right)^2 \theta^2 \langle n_{\mathrm{RT}} \rangle , \\ \alpha_B &= \frac{\Gamma_{J_e,J_g} \left| \left\langle J_e, m_{J_e} \middle| J_g, 1, m_{J_g}, m_{J_e} - m_{J_g} \right\rangle \right|^2}{\Gamma}, \\ \langle n_{\mathrm{RT}} \rangle &= \frac{1}{\mathcal{L}_{\mathrm{RT}}}, \end{aligned} \tag{2.25}$$

where α_B is the branching ratio of all excited state decays into the coupling state and $\langle n_{\mathrm{RT}} \rangle$ is the average number of round trips a photon completes in the cavity before it is lost. This form emphasises that the cooperativity is proportional to the atomic branching ratio, the geometric overlap of the electronic decay with the cavity mode, the total solid angle covered by the cavity mode (proportional to θ^2) and the average number of round trips of a photon in the cavity. It makes explicit the crucial point that the cavity allows for high collection probabilities, even if the solid angle subtended by the cavity is not large, provided the cavity coupling is enhanced by a large number of photon round trips. The exact techniques for using cavities to extract photons will be discussed in Chapter 3, but this fact is the basis for the interest in enhancing remote entanglement performance using cavities.

2.3.4 Cavity Geometry

Hitherto, the ion cavity interaction has been discussed in the context of two-level emitters and single cavity modes, but this picture is too simple for understanding remote entanglement protocols for two reasons. Firstly, because the cavity typically supports two degenerate polarisation modes, and secondly because the electronic structure of the ion is much richer.

The scalar mode theory presented in Sec. 2.3.1 factors out the polarisation of the mode, requiring only that the polarisation is transverse to the cavity axis. This means that the cavity actually supports two degenerate vector modes for each scalar

mode, with polarisations aligned with basis vectors $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_1}}$ and $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_2}}$, that tile the plane perpendicular to the cavity axis $\hat{\boldsymbol{e}}_{\boldsymbol{z}}$. These vector modes are most conveniently expressed as either linearly or circularly polarised, and an emitter can couple to both subject to the geometrical considerations of Sec. 2.3.1. Note that due to birefringence of the cavity mirrors these vector modes might not be perfectly degenerate. This phenomenon will be discussed in depth in Chapter 6, but is not considered here.

Regarding ion structure, the states $|e\rangle$ and $|g\rangle$ belong to larger manifolds containing states of different angular momentum projections. The transitions between these additional states can also couple to the cavity modes (even if they may be off-resonant due to Zeeman splittings). The orientation of the cavity with respect to the magnetic field is typically chosen to strengthen desired couplings and forbid unintended ones, a process best illustrated with the two example configurations that will be discussed in the context of remote entanglement protocols.

One configuration aligns the cavity axis $\hat{\boldsymbol{e}}_{\boldsymbol{z}}$ with the magnetic field $\hat{\boldsymbol{e}}_{\boldsymbol{z_B}}$. In this case, the two transverse modes of the cavity correspond to $\sigma_\pm$ emission, and π-polarised decay into the cavity is forbidden because this polarisation is completely parallel to the cavity axis.

The second configuration has the magnetic field orthogonal to the cavity axis (for concreteness it will be aligned along $\hat{\boldsymbol{e}}_{\boldsymbol{x}}$). In this case $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_1}} = \hat{\boldsymbol{e}}_{\boldsymbol{y_B}} = (1/\sqrt{2})i(\hat{\boldsymbol{e}}_{\boldsymbol{-1}} + \hat{\boldsymbol{e}}_{\boldsymbol{+1}})$ and $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_2}} = \hat{\boldsymbol{e}}_{\boldsymbol{0}}$. Here $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_2}}$ is (by convention) V-polarised and couples only to π decays, whereas $\hat{\boldsymbol{e}}_{\boldsymbol{\epsilon_1}}$ is H-polarised and couples to both σ_- and σ_+ transitions. Usually only one of these σ-polarised transitions will be resonantly enhanced.

2.4 Cavities for Remote Entanglement

2.4.1 Remote Entanglement Schemes

Optical cavities can strengthen the vacuum fluctuations in resonant cavity modes, which can be used to enhance the rate of spontaneous emission of a emitter into these cavity modes. This allows for more efficient extraction of photons from emitters, and these photons can be used to perform remote entanglement schemes via a number of schemes, which will be discussed in turn.

Perhaps the most physically direct photon-mediated remote entanglement scheme involves emitting a photon whose state is entangled with one emitter, and arranging for this photon to be absorbed by a second emitter, in such a way that, upon absorption of the photon, the second emitter will be entangled with the first [162], as demonstrated with atoms in [101]. The major problem with this scheme is that it is sensitive to photon losses; if the photon was not absorbed by the second emitter (because it was never collected from the first, or passed through the second, or was lost in transit), the emitters will not be entangled, and there is no signal that the entanglement procedure has failed. Practical ion-cavity systems do not presently have light-matter coupling efficiencies sufficient for photon loss to be tolerated.

Thus, heralded schemes, where the success of the entanglement process is signalled by a measurement, are preferred. The simplest is a single photon entanglement scheme presented in [163] and used in many experiments [164, 165]. In this scheme, both emitters are weakly excited and, upon detection of a single photon, the emitters are projected into an entangled state. An advantage of this scheme is that, as only one photon must be collected, the success rate is linearly proportional to the photon extraction probability. There are two major disadvantages. First, the requirement of weak excitation means very high success rates are incompatible with high fidelity. Secondly, the phase of the entangled state is sensitive to photon path length differences on the scale of an optical wavelength, which would require careful stabilisation in a practical device. More details of these issues can be found in [166].

Because of the problems with single photon schemes, this work considers two-photon probabilistic remote entanglement schemes of the type proposed in [86]. In these schemes, each emitter emits a single photon through a process that leaves the emitter state entangled with a binary photonic degree of freedom (a photonic qubit). These photons are routed through a non-polarising beamsplitter (NPBS) to erase the information of which emitter produced each photon. If the photonic qubit states cannot be easily resolved (e.g. when using frequency qubits), the ions are projected to an entangled state upon measuring one photon output on both sides of the beamspliter through the Hong-Ou-Mandel effect [167]. A faster rate can be achieved by resolving the photonic qubit and projecting to an entangled

state upon detection of two photons, one for each photon qubit state, with no restriction on the sides of the NPBS at which the photons are detected [168]. Because the protocol relies on interference, the photons emitted from each node should be indistinguishable [169] (at least from the perspective of measurement [170, 171]) for the entanglement protocol to function. The photonic degree of freedom could be photon frequency [172], or an easily resolved degree of freedom like polarisation (as has been utilised in previous ion-ion remote entanglement with [108] and without [173] cavity enhancement), or wavepacket time bin [168][2].

For two-photon schemes with a resolvable photonic qubit encoding, the probability of entanglement success for any given attempt (assuming that each emitter has equal probability of emitting either photonic qubit state) is

$$P_{\text{ent}} = \frac{1}{2}\left(P_{\text{ext}}\eta_n\eta_d\right)^2, \tag{2.26}$$

where P_{ext} is the probability that a photon is extracted from the cavity, η_n the efficiency of routing this photon to a detector and η_d the efficiency of the detector. The factor of 1/2 arises because in only 50% of trials will the detected photons have different qubit states (this factor is 1/4 when the photon qubit states cannot be resolved). This equation highlights that the rate of remote entanglement is quadratically dependent on the probability that a photon is collected from the emitter. It is for this reason that cavities, which can increase P_{out} towards unity, could be very beneficial for use in such remote entanglement protocols.

2.4.2 Choice of Cavity Transition

A key consideration for a cavity-enhanced quantum network is the transition used to generate network photons. Here, this will be discussed in the context of $^{88}Sr^+$, whose level diagram is shown in Fig. 2.6. The most important considerations for the transition will be analysed in turn.

First, the final state should be long-lived enough to store information, at least until it can be mapped elsewhere. This limits the protocol to transitions decaying to the $S_{\frac{1}{2}}$, $D_{\frac{3}{2}}$ and $D_{\frac{5}{2}}$ manifolds.

[2]The merits of polarisation and time bin encoding will be discussed in Chapters 6 and 7 respectively.

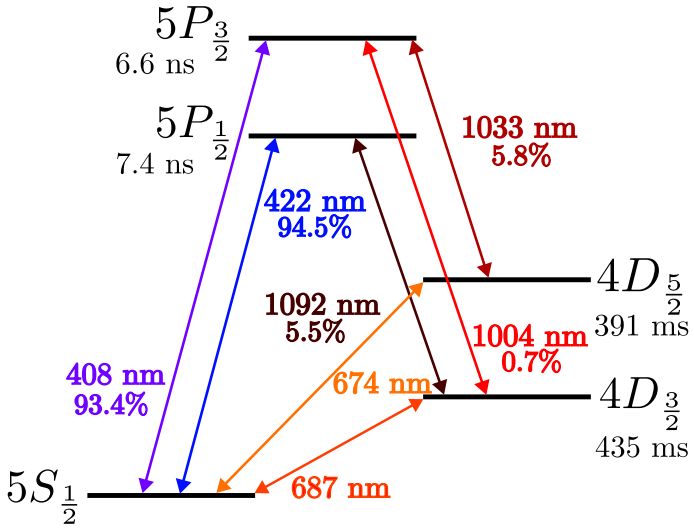


Figure 2.6: Level diagram for the $^{88}Sr^{+}$ ion. Each transition is indicated by an arrow with a label denoting the wavelength and branching ratio of the decay α_B. Manifolds (with the exception of the ground state $S_{\frac{1}{2}}$) are labelled with lifetime. The ion features S and D manifolds with sufficiently long lifetimes to host qubits. Dipole allowed transitions are either short-wavelength high-branching ratio transitions $P_{\frac{1}{2}} \rightarrow S_{\frac{1}{2}}$ and $P_{\frac{3}{2}} \rightarrow S_{\frac{1}{2}}$, or infrared low-branching ratio transitions $P_{\frac{1}{2}} \rightarrow D_{\frac{3}{2}}$, $P_{\frac{3}{2}} \rightarrow D_{\frac{3}{2}}$ and $P_{\frac{3}{2}} \rightarrow D_{\frac{5}{2}}$.

Secondly, the decay should naturally be strong (i.e. electric dipole-allowed) so that the photons can be extracted quickly. While the extraction rate will also depend on the characteristics of the cavity, it would, for instance, be challenging to fabricate a fast interface on the $D \rightarrow S$ quadrupole transitions even though a high probability of extraction is possible [174]. Therefore, only dipole-allowed transitions will be considered, and lower wavelength transitions are favoured by this consideration as the transitions are generally stronger [107].

Thirdly, the transition should allow for efficient collection (so that a photon is likely to be collected in a reasonable number of attempts). As summarised in Eq. 2.25, this favours transitions with strong branching ratios and wavelengths of low round-trip loss. Branching ratios are generally higher for lower wavelengths. However, scattering loss due to mirror roughness σ scales with λ^{-2} [175], meaning that cavities at short operating wavelengths tend to have low finesse [176]. Therefore, while free-space implementations tend to favour short wavelengths, the preferred option is not so clear cut for cavity-enhanced networking.

Fourthly, there should be low transmission losses within the network. For convenient fibre-based networks, this favours telecom band photons, but even near-

infrared photons generally have much lower loss than visible and UV photons. This factor is a much bigger concern for long-range networks (where frequency conversion has been used to reduce fibre losses [177, 178]), but would become progressively more important in a modular quantum computer as the number of modules increases and the fibre length between distant modules grows.

Finally, blue or ultraviolet radiation is problematic as the low wavelength leads to surface charging from the laser used for spectroscopic locking of the cavity length [148, 179] and can cause finesse loss unless care is taken [180].

On balance, the combination of factors would appear to favour an infrared $P \rightarrow D$ dipole transition, of which there are 3 possible options in $^{88}Sr^{+}$. The process of extracting photons on these transitions will be discussed in Chapter 3, before a presentation of a reasonable extraction protocol is made in Chapter 7.

2.5 Cavity Manufacture

2.5.1 Introduction

This thesis will not primarily be concerned with cavity manufacture, and as such the introduction given here will be brief. However, trends and techniques in the physical construction of optical cavities are important motivators for the physics investigated.

2.5.2 Mirror Manufacture

Demonstrations of ultralow loss [181], strong cavity coupling [182], and impressive photon extraction [183] have been made using cavities with distributed Bragg reflector coatings on superpolished mirror substrates. However, applications of remote entanglement require not just high extraction efficiency but high rate. This means the cavity should facilitate a high number of photon round trips in a short time, and therefore must be short. However, the cavity must be long enough that there is space for the trap structure to shield the ion from the effects of the dielectric mirror coating. The combination means that the mirrors must have a tight radius of curvature to provide the necessary extraction probability and rate while remaining compatible with trapping. Superpolished mirrors have not been fabricated at these tight curvatures [184], but several manufacturing techniques have been developed to do so.

Laser Ablation

Perhaps the most popular method for forming such mirrors is laser ablation, where a focussed CO_2 laser both vaporises and melts part of the substrate to leave a low-roughness depression [185, 186]. Such mirrors can enable very high finesse [187], but the mirror formed usually inherits the Gaussian shape of the ablating laser beam [188], and may also be elliptical [189], although recent developments have reduced these fabrication defects [190, 191].

Laser ablated mirrors are often formed on the end of fibres to yield fibre Fabry-Perot cavities [185], with the supposed advantage that the extracted light is already coupled into a fibre [192]. However, the cavity mode is often poorly matched into the fibre; integrated matching lenses [193] and collection fibres with larger modes [191] have been used to alleviate these issues.

Focussed Ion Beam Milling

Another option for forming depressions with tight curvature is focussed ion beam (FIB) milling, where a precise beam of ions mills the desired structure [194]. This process allows for a high degree of control over the mirror shape [195], but has traditionally been used for smaller cavities and tighter curvatures than would suit integration with an ion trap [196].

Reactive Ion Etching

Reactive ion etching can remove material from a surface, forming a smooth depression that can be coated to form a mirror [197]. Etching procedures have been used for silicon [198] and fused silica [199] substrates to form high finesse cavities, although the shape is not necessarily perfectly spherical [200].

2.5.3 Cavity Assembly and Mechanical Structure

To form an optical cavity requires two mirrors be aligned facing each other, and the length set by the intended resonant frequency. To hit resonance, the cavity length needs to be accurate to at least $\lambda/(2\mathcal{F})$. It is impractical to manufacture devices to this precision, so cavities typically have a method of length tuning

(usually piezoelectric [201] but thermal tuning is possible [202]). The transverse and angular alignment are not required to meet such exacting precision requirements. Traditionally, cavity mirrors have been positioned carefully before being glued in place on their length adjusters [203], although microadjusters have been used for active transverse manipulation in some experiments [192], with similar systems of this type proposed [204, 205].

Another consideration is mechanical vibration. Fibre-based cavities have traditionally suffered more severely from mechanical noise [187] than more monolithic designs [206], due to bending modes of bare fibres and the distance of the mirrors from a common base [203]. Much effort has been devoted to improving the stability, either through a more intrinsically stable structure [202] or sophisticated stabilisation [201]. The cavity in a practical quantum network node would ideally be highly passively stable.

2.5.4 Considerations for the Future Experiment

The remainder of the document will aim to understand the physics of remote entanglement relevant for a practical system. To do this, a specific set of characteristics for the device will be imagined. In order to make the cavity passively stable and simple, the conceived device will be quasimonolithic, with mirrors held rigidly in the transverse directions without prior alignment, and adjustment only available in the cavity length. Due to the lack of active transverse alignment, the mode may form at a tilted angle relative to the mirror axes (as discussed in Sec. 2.2.5). Therefore it will be assumed that light is extracted through free space fibre coupling, so that the coupling can be adjusted to suit the tilt-angle of the mode. The possibility of non-zero mode tilt also means the shape of the mirror away from its central depression will play an important role in the analysis.

2.6 Conclusion

The theory of cavity modes, treatment of ion-cavity interactions, and details of recent cavity construction presented in this chapter form the basis of the results presented in this thesis, which asks what design choices are appropriate for an ion-cavity

interface for remote entanglement. Despite the apparent simplicity of a single ion interacting with a single cavity mode, there are many aspects of physics that must be considered in the design of these future experiments.

For completeness, it should also be noted that, while in this thesis the purpose of a cavity is principally to enhance the photon collection efficiency from an emitter, a cavity is an excellent general tool to interface light and matter. As such, there are many other possible uses of cavities, such as detecting photons without destruction [102], or performing photon-photon [103] or ion-ion [207] gates, and so motivations for integrating a high-cooperativity cavity in an ion trap are not necessarily limited to remote entanglement schemes.

3 Optimising Photon Extraction

3.1 Introduction

To implement a cavity-enhanced quantum network, single photons must be extracted from an emitter-cavity system. The theory of single-mode cavity QED was discussed in Chapter 2, and will be developed in this chapter to overview photon extraction and understand how it may be optimised.

Contribution and Acknowledgements

The chapter begins with a summary of photon extraction methods discussed in the literature. From Sec. 3.5 onwards (inclusive), original work concerning the optimisation of photon extraction is presented, which was conducted solely by the author.

3.2 Typical Photon Extraction Schemes

3.2.1 Overview

For the cavity-enhanced quantum network, the ion trapped inside the cavity must emit a photon into the cavity mode, with the state of the photon entangled with the final electronic state of the ion, and this photon must then be collected. There are several different techniques for producing the photon in the cavity mode, which will be discussed in the context of the generic three-level system depicted in Fig. 3.1.

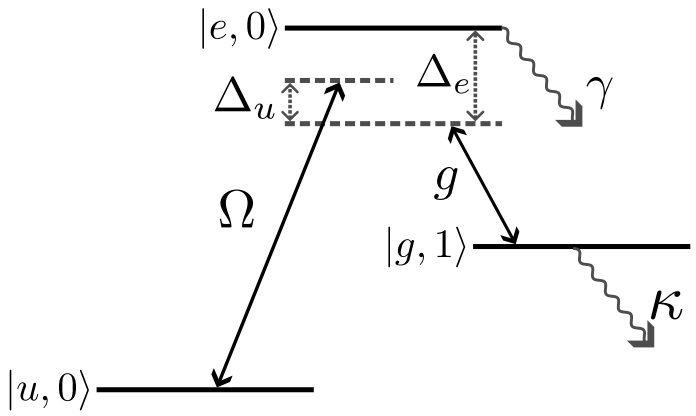


Figure 3.1: The level scheme for the three-level system. The Hilbert space is formed of three levels that are the tensor product of the ion's electronic state and the cavity mode Fock state. These levels are the initial ion state with no cavity photon $|u,0\rangle$, an excited ion state with no cavity photon $|e,0\rangle$, and a final ion state with a cavity photon $|g,1\rangle$. The excited state and final state are coupled by the cavity coupling g, with the excited state detuned from resonance with the cavity by Δ_e. The initial and excited state are coupled by a laser field with time-dependent coupling Ω. The detuning of the laser field from Raman resonance with the cavity is Δ_u, although this is only nominal as the laser field can be adjusted to include an arbitrary detuning in its time dependence. The excited state amplitude decays at a rate γ due to spontaneous emission, and the cavity state amplitude at a rate κ due to field loss from the cavity.

The Hamiltonian for the three-level model is

$$H = \Delta_u \left|u,0\right\rangle\!\left\langle u,0\right| + \Delta_e \left|e,0\right\rangle\!\left\langle e,0\right| + \left[-i\Omega^* \left|u,0\right\rangle\!\left\langle e,0\right| + \text{h.c.}\right] + \left[ig \left|e,0\right\rangle\!\left\langle g,1\right| + \text{h.c.}\right], \tag{3.1}$$

which is a straightforward extension of the two-level Hamiltonian from Chapter 2, where Ω is a time dependent drive strength from initial state $|u,0\rangle$ to excited state $|e,0\rangle$, and the factor of i before g is made for algebraic convenience. The detuning of the initial state from Raman resonance with $|g,1\rangle$ is Δ_u, which is typically set to zero. This leads to a non-Hermitian Hamiltonian

$$H^{\text{NH}} = H - i\gamma \left|e,0\right\rangle\!\left\langle e,0\right| - i\kappa \left|g,1\right\rangle\!\left\langle g,1\right|. \tag{3.2}$$

As discussed in Chapter 2, this system is most correctly simulated using the master equation, and to perform this, the Hilbert space would need to be extended to contain all of the collapse operators that properly account for all the partial decay rates in the real system. On the assumption that no decay process can be followed by subsequent cavity decay, the simpler non-Hermitian Hamiltonian is valid and can be integrated according to the Schrödinger equation. However, the non-Hermitian

approach cannot capture an important effect known as temporal mixing (discussed in Sec. 3.3 of this chapter), and therefore care must be taken before using it to ensure that the underlying assumptions are appropriate.

3.2.2 Direct Excitation

A conceptually straightforward way of generating photons in the cavity mode (demonstrated in [208]) is to directly excite the emitter to the excited state $|e,0\rangle$ far faster than the cavity dynamics, and then let the atomic excitation decay into the cavity mode before leaking out through a partially transmissive mirror . This method is most directly analogous to the fast excitation used in previous free-space remote entanglement experiments [96, 167]. This method has a total photon generation probability through the cavity

$$P_{\text{gen}}^{\text{direct}} = \frac{\kappa}{\kappa+\gamma}\frac{2C}{2C+1}, \tag{3.3}$$

where C is the cooperativity [209]. This scheme has an advantage in its conceptual simplicity and speed, but requires a laser capable of fast excitation and does not control the output shape of the photon. In addition, the probability of photon extraction is reduced by a factor of $\kappa/(\kappa+\gamma)$ from the maximum that could be extracted from the given three-level system (discussed later in Sec. 3.4 of this chapter). This attenuation, a feature of the coherent mechanism of cavity enhancement and not seen in free-space implementations, can be severe, particularly for longer cavities with high finesse, where γ may exceed κ by a large factor.

3.2.3 Raman Process

Another possible paradigm is to use a two-photon transition, applying a laser drive in Raman resonance with the cavity transition ($\Delta_u = 0$), but with a substantial detuning Δ_e of the excited state $|e,0\rangle$ [210]. In this two-photon resonant process, the ion absorbs a photon from the drive pulse, and emits a photon into the cavity mode with effective two-photon coupling $g\Omega/\Delta_e$, and this population then leaks out to be collected [211]. The intensity profile of the applied laser determines the temporal profile of the emitted photons [212]. This scheme can be extended to systems with

multiple cavity decays to produce ion photon entanglement tuned by the relative phases and intensities of the drive fields [106] and is used in current experiments [183].

3.2.4 vSTIRAP

A widely used method for producing controlled photons from emitter-cavity systems is vSTIRAP. The method takes its name from the adiabatic transfer procedure STIRAP [213, 214], where the 'v' signifies that the role of one of the control fields is taken by the vacuum mode of the cavity [215]. The principle can be understood from the treatment of [139] considering the three-level system on Raman resonance ($\Delta_u = 0$) and ignoring decay effects. One of the eigenstates (the 'dark state') of this Hermitian system is

$$\begin{aligned} |\Psi_0\rangle &= \cos(\theta)\,|u,0\rangle - \sin(\theta)\,|g,1\rangle\,, \\ \tan(\theta) &= \frac{\Omega}{g}, \end{aligned} \tag{3.4}$$

where Ω is the drive strength defined in Eq. 3.1. The system is initialised in $|u,0\rangle$, and with no drive field applied, this is an eigenstate of the Hamiltonian. The drive intensity is then gradually increased until $\Omega \gg g$. Provided the parameters of the Hamiltonian are changed sufficiently slowly to ensure adiabatic evolution, the state will have remained in $|\Psi_0\rangle$ throughout the process, and the final state of the system will be $|g,1\rangle$ without the excited state ever being occupied. In a more detailed treatment, the cavity decay of $|g,1\rangle$ disrupts the adiabatic evolution, and means perfect efficiency cannot be achieved [160]. This procedure can be used simply to generate photons [216], but the control vSTIRAP allows over the evolution of the dark state means the shape [217] and phase [218] of output photons can be controlled, permitting the production of entangled photon pairs [219] or the encoding of quantum information in the time domain of single photons [220].

3.3 Photon Driving Simulations and Temporal Impurity

Photon production schemes can be simulated through integrating a non-Hermitian Hamiltonian or using a master equation approach. However, only the master equation can capture a subtle error process resulting from spontaneous emission to $|u,0\rangle$

followed by cavity emission. After spontaneous emission to the initial state, the emitter-cavity system is projected back to $|u, 0\rangle$, and the photon wavepacket produced from that time. This means that the photon output wavepacket depends on the stochastic spontaneous emission. In any given remote entanglement attempt, such an error could introduce distinguishability between the photons from the two nodes, reducing the fidelity of remote entanglement [108, 221]. In the master equation approach, this phenomenon causes the cavity to emit a probabilistic mixture of photon wavepackets [222], and thus this process will be termed 'temporal mixing'.

The options for mitigating this error are

1. Use direct excitation (Sec. 3.2.2) so that all photons are generated at the same time.

2. Improve the cooperativity of the system so that spontaneous emission is unlikely.

3. Use a transition for which the decay to the initial state has a low branching ratio. Spontaneous emission is therefore well-approximated as loss from the system, and subsequent cavity emission is unlikely [223].

The first solution suffers the drawbacks discussed in Sec. 3.2.2, namely that the probability of extraction is suboptimal, and control over the photon generation process is relinquished. The second solution would be ideal, but the cooperativity of state-of-the-art ion-cavity interfaces is not sufficient to achieve this for lambda-systems where the branching to the initial state is strong. The third option allows for temporal mixing to be made negligible, even using gradual drive processes and realistic cavity parameters. For the remainder of the document, it will be assumed that the third option is implemented by ensuring the decay to the initial state has a small branching ratio.

3.4 Finite Time Driving

The direct excitation approach to producing photons for cavity-enhanced remote entanglement has disadvantages in that the photon extraction probability does not reach the limit of an infinitely slow photon generation process, and the photon output

shape cannot be controlled. This would suggest the use of a more 'gradual' method of photon production that transfers the photon into the cavity over a period of time, reducing the probability of spontaneous emission. It can be shown that the maximum photon extraction probability from the three-level system is $2C/(2C+1)$, whatever method is chosen, be it vSTIRAP, Raman transfer, or any other type of drive [161]. This limit is only achieved in infinite time.

While this result puts useful limits on the extraction, a practical experiment must produce finite-length photons. The natural question to ask is what level of photon extraction can be obtained from these systems in finite time and how this would be achieved. One approach would be numerical optimisation of the driving pulse. With this approach, practical constraints on the driving pulse can also be encoded, but the major drawback is that it is practically impossible to guarantee the methods accounts for all possible pulses. It would be helpful to have a method to calculate the optimum finite time photon extraction probability, without involving the vagaries of numerical optimisation of the driving pulse.

The results presented below show such a scheme. First, an analytical method for the three-level system that bounds the output is presented, and then numerical method optimising the output photon shape is developed, with applications to more general systems.

It is important to remember throughout that optimising the photonic output is by no means optimal; a high rate is desirable for most quantum protocols, but simplicity and fidelity in the presence of noise and defects are often more important. The approach to maximising output should be viewed as a tool to understand the physical limits of these systems, and therefore understand whether it is worth attempting to extract more probability from a system in practice.

3.5 Three-Level System: Bounded Approach

3.5.1 The System and Problem

The first approach investigated will bound the possible output probability from the three-level system (Sec. 3.2.1) in time T using analytical methods. This will be done using a non-Hermitian framework under the assumption (discussed in Sec. 3.3)

that no cavity photon can be produced after a decay process (spontaneous emission or cavity decay) has already occurred.

3.5.2 Calculating the Ideal Photon Shape

The analysis initially follows Goto (2019) [161] to link the probabilities of cavity emission and spontaneous emission to the photon shape, and then optimises this shape to calculate the optimum output. The analysis is very similar also to Vasilev (2010), [224], which discusses how to achieve a particular photon shape, and Utsugi (2022), [225], which restricts the output photons to Gaussian wavepackets.

The starting point is the set of equations of motion for the three-level system, taken directly from the Schrödinger equation in the non-Hermitian approximation

$$\begin{aligned} |\Psi(t)\rangle &= \alpha_u |u,0\rangle + \alpha_e |e,0\rangle + \alpha_g |g,1\rangle , \\ \dot{\alpha_u} &= -i\Delta_u \alpha_u - \Omega^* \alpha_e, \\ \dot{\alpha_e} &= -(\gamma + i\Delta_e)\alpha_e + \Omega \alpha_u + g\alpha_g, \\ \dot{\alpha_g} &= -\kappa \alpha_g - g\alpha_e, \end{aligned} \tag{3.5}$$

where $|\Psi(t)\rangle$ is the wavefunction of the ion-cavity system. From these equations, the integral relation

$$\begin{aligned} I_e = \frac{\kappa^2}{g^2} I_g + \frac{\kappa}{g^2}(I_\cdot + I_\cdot^*) + \frac{1}{g^2} I_{g'}, \\ I_e = \int_0^T |\alpha_e(t)|^2 \, dt, \quad I_g = \int_0^T |\alpha_g(t)|^2 \, dt, \\ I_\cdot = \int_0^T \dot{\alpha_g}(t)\alpha_g^*(t) \, dt, \quad I_{g'} = \int_0^T |\dot{\alpha_g}(t)|^2 \, dt, \end{aligned} \tag{3.6}$$

can be derived. These integrals can be related to physical probabilities: P_E, the probability of spontaneous emission, P_R, the probability of decay through the cavity, and P_C, the occupation of the cavity at time T. The probability condition corresponding to the integral relation is

$$\begin{aligned} P_E = \frac{\gamma\kappa}{g^2} P_R + \frac{2\kappa\gamma}{g^2} P_C + \frac{2\gamma}{g^2} I_{g'} \\ P_E = 2\gamma I_e, \quad P_R = 2\kappa I_g, \quad P_C = I_\cdot + I_\cdot^* \end{aligned} \tag{3.7}$$

where the form for P_C assumes the condition (that will later be enforced) that the cavity is vacant at the start of the process.

A further constraint on the probabilities is that the total probability of all outcomes at all times must add to 1. The set of total probabilities includes P_R, P_E and P_C, but also includes the probability that the initial state (excited state) is populated, P_u (P_e). The assumption of arbitrary driving means that, provided the probability required to produce α_g remains across $|u,0\rangle$ and $|e,1\rangle$, the driving to produce the requisite α_e can be performed. Therefore, the limiting case occurs where the probability outside the initial state reaches unity at one time during the process. However, in order to obtain an upper bound to the performance, it is assumed the probabilities at time T are distributed across P_R, P_E and P_C only, and that these probabilities sum to unity, a condition which is in all cases less or equally restrictive than the true probability condition. This assumption leads to

$$1 = (1 + \frac{\gamma\kappa}{g^2})P_R^u + (1 + \frac{2\kappa\gamma}{g^2})P_C^u + \frac{2\gamma}{g^2}I_{g'}, \tag{3.8}$$

where the subscript u indicates that the probabilities are not physical but pertain to the upper bound. The analysis of [224], aimed at producing photons of desired shape, assumes the cavity to be unoccupied at the end of the process, so derives a less restrictive, but very similar, upper bound. The output probability can then be written

$$\begin{aligned} P_R^u &= \frac{1}{1 + \frac{1}{2C}(1+m)}, \\ C &= \frac{g^2}{2\kappa\gamma}, \\ m &= \frac{I_{g'} + \kappa(1+C)(I_{\cdot} + I_{\cdot}^*)}{\kappa^2 I_g} \equiv \frac{F}{G}, \end{aligned} \tag{3.9}$$

in terms of the cooperativity C, and a parameter m that is purely a function of the shape of the output photon (rather than its scale), and should be minimised to maximise the output. The functionals F and G are defined solely for notational convenience. To minimise m, assume that the function α_g changes by $\delta\alpha_g$, causing a change in F (G) of δF (δG). The condition that m is an extremum gives $\delta F/\delta G = F/G$. Now define a shape functional

$$S_q = I_{g'} + \kappa(1+C)(I_{\cdot} + I_{\cdot}^*) - \kappa^2 q I_g = F - qG, \tag{3.10}$$

for a general scalar q. Extrema of S_q satisfy $\delta F/\delta G = q$. Therefore, the procedure to find extrema of m with respect to changes in the photon shape α_g is a two-stage process: first, find the extrema of S_q for all q and second, check that $q = F/G$. Minimising S_q proceeds according to

$$\begin{aligned} S_q &= \int_0^T L_q \, dt, \\ L_q &= \dot{\alpha}_g \dot{\alpha}_g^* + \kappa(1+C)\left(\dot{\alpha}_g \alpha_g^* + \dot{\alpha}_g^* \alpha_g\right) - \kappa^2 q \alpha_g \alpha_g^*, \\ \frac{\partial L_q}{\partial \alpha_g} &= \frac{d}{dt}\left(\frac{\partial L_q}{\partial \dot{\alpha}_g}\right), \end{aligned} \tag{3.11}$$

by analogy with the Euler Lagrange equations and produces a differential equation

$$\ddot{\alpha_g} + \kappa^2 q \alpha_g = 0, \tag{3.12}$$

for the wavefunction of the cavity mode. Setting the boundary condition that the cavity is vacant at the start of the process gives

$$\begin{aligned} \alpha_g(t) &= A \sin(\omega_q t), \\ \omega_q^2 &= \kappa^2 q, \end{aligned} \tag{3.13}$$

where A is a free parameter that can be set to adjust the sum of probabilities. The second stage of the procedure demands

$$\cos(2\omega_q T) - \frac{2\omega_q}{2\kappa(1+C)} \sin(2\omega_q T) = 1 \tag{3.14}$$

is satisfied. This equation has many solutions, however, when satisfied, $q = F/G = m$, and therefore the solution of interest is the minimum q. There is a solution at $q_f = \frac{\pi^2}{(\kappa T)^2}$, corresponding to $\alpha_{g_f} = A \sin\left(\frac{\pi t}{T}\right)$, which is a photon of sinusoidal amplitude whose wavefunction begins at zero, and reaches zero again at T. This solution 'closes' the photon at the final time, resulting in no remaining probability in the cavity. However, there is a better solution at a slightly lower q, which retains the sinusoidal amplitude, but with a slower temporal frequency of the shape so that the photon does not fully close.

Thus the upper bound can be summarised

$$\begin{aligned} 1 &= \cos(2\omega_m T) - \frac{2\omega_m}{2\kappa(1+C)}\sin(2\omega_m T), \\ m &= \left(\frac{\omega_m}{\kappa}\right)^2, \\ P_R^u &= \frac{1}{1+\frac{1}{2C}(1+m)}, \end{aligned} \tag{3.15}$$

where m is taken as the smallest value to satisfy the first equation.

The upper bound is an optimised output derived under the condition that the populations remaining in the cavity P_C, emitted spontaneously P_E and decayed through the cavity P_R sum to unity at time T. This is less restrictive than the true condition that these probabilities, in addition to the probability in the excited state P_e reaches unity at some point in the process. A lower bound to the output probability P_R^l can be obtained by assuming the photon retains the same amplitude profile, but satisfies the true probability constraint. This is a lower bound because the photon has not be optimised in shape for this more correct condition.

The upper and lower bounds can be readily calculated, but before those results are presented, the numerical method will be developed.

3.6 Numerical Method for More General Emitter-Cavity Systems

3.6.1 Setting up the System

The ideas presented in Sec. 3.5 inspire a numerical method applicable to more general systems. These systems again contain a single initial state $|u,0\rangle$ and excited state $|e,0\rangle$, but now many distinct states with occupied cavity modes $|g_j,1_j\rangle$ for $1 \leq j \leq j_M$, each coupled to the excited state with coupling rate g_j. Note that the occupied photon modes (whose single-occupancy states are denoted $|1_j\rangle$) are not necessarily orthogonal to each other, but the combined ion-photon states $|g_j,1_j\rangle$ should be mutually orthogonal. For cavity applications there would typically be two (for magnetic field aligned with cavity axis) or three states (for magnetic field perpendicular to cavity axis). The level scheme is shown in Fig. 3.2 for an example of two levels. The equations of motion of the system can be written

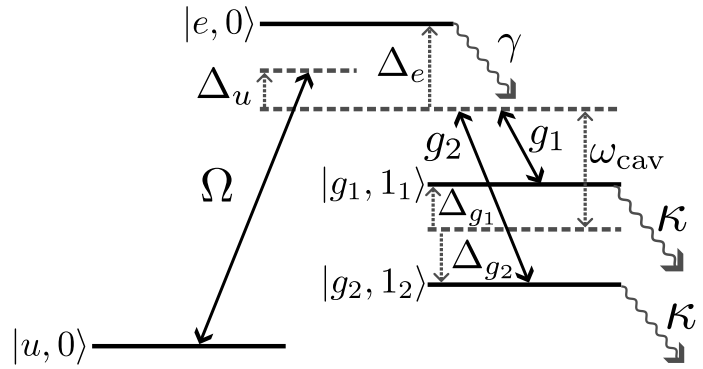


Figure 3.2: The level scheme for the multiple-level system, here depicted with two cavity decays. The Hilbert space comprises an initial ion state with no cavity photon $|u,0\rangle$, an excited ion state with no cavity photon $|e,0\rangle$, and an arbitrary number of stable ion states with a cavity photon $|g_j,1_j\rangle$, where the index j labels the cavity decay. The excited and the final states are coupled by the cavity couplings g_j. The cavity states are each detuned from a nominal central energy level by detuning Δ_{g_j} from which the excited state is itself detuned by Δ_e when including the cavity photon energy ω_{cav}. The initial and excited state are coupled by a laser field with time-dependent coupling Ω, again with nominal detuning Δ_u. Note that this setting of convention does not restrict this system to modelling monochromatic driving. The excited state amplitude decays at a rate γ due to spontaneous emission, and the cavity state amplitudes each decay at a rate κ due to field loss from the cavity.

$$
\begin{aligned}
|\Psi(t)\rangle &= \alpha_u |u,0\rangle + \alpha_e |e,0\rangle + \sum_{j=1}^{j_M} \alpha_{g_j} |g_j,1_j\rangle \,, \\
\dot{\alpha_u} &= -i\Delta_u \alpha_u - \Omega^* \alpha_e, \\
\dot{\alpha_e} &= -(\gamma + i\Delta_e)\alpha_e + \Omega\alpha_u + \sum_{j=1}^{j_M} g_j \alpha_{g_j}, \\
\dot{\alpha_{g_j}} &= -(\kappa + i\Delta_{g_j})\alpha_{g_j} - g_j \alpha_e,
\end{aligned}
\tag{3.16}
$$

where Δ_{g_j} is the detuning of state $|g_j,1_j\rangle$ from an arbitrary reference level, which, for numerical convenience, is best chosen near the centre of the manifold of states with occupied cavities.

The difficulty with optimising the outputs of this system is that specifying an output wavefunction (for example $\alpha_{g_1}(t)$) will intrinsically specify the others (for an appropriate initial condition) as all outputs couple only to one excited state $|e,0\rangle$. This interdependence is not straightforward to treat in the time domain as it involves both wavefunction terms and their time derivatives.

Instead, the wavefunctions $\alpha_{g_j}(t)$ can be specified as a sum

$$\begin{aligned} \alpha_{g_j}(t) &= \frac{1}{\sqrt{T_b}} \sum_n C_n^{(j)} e^{i\omega_n t}, \\ \omega_n &= \frac{2\pi}{T_b} n, \end{aligned} \tag{3.17}$$

of Fourier coefficients $C_n^{(j)}$ across a time domain of length T_b[1]. It will be helpful to write $\alpha_{g_j}(t)$ as a vector in Fourier space,

$$\boldsymbol{\alpha}_{\boldsymbol{g_j}}^{\boldsymbol{F}} = (C_0^{(j)}, C_1^{(j)}, C_2^{(j)}, ...). \tag{3.18}$$

The equations of motion (Eq. 3.16) relate the n$^{\text{th}}$ Fourier coefficient of the k$^{\text{th}}$ decay to the same coefficient of the j$^{\text{th}}$ decay through

$$C_n^{(k)} = \frac{g_k}{g_j} \frac{\kappa + i(\omega_n + \Delta_{g_j})}{\kappa + i(\omega_n + \Delta_{g_k})} C_n^{(j)} \equiv f_n^{(j\to k)} C_n^{(j)}. \tag{3.19}$$

The output probability can be maximised tracking only the $j = 1$ cavity wavefunction as the remainder of the level structure will be automatically encoded in these relations.

3.6.2 Calculating Probabilities

To maximise the output probability of the system, a selection of probabilities must be calculated in terms of the output vector $\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{F}}$. These probabilities are

- $P_{R_j}(t)$: The probability that a photon is emitted from level j before time t.
- $P_E(t)$: The probability of spontaneous emission before time t.
- $P_{C_j}(t)$: The probability that, at time t, there is a photon in the cavity corresponding to decay j.
- $P_e(t)$: The probability that the excited state is occupied at at time t.

These probabilities can all be written as expectation values of matrices $\boldsymbol{P}_{\boldsymbol{\zeta}}$ such that

$$P_\zeta(t) = \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{F}} \cdot \boldsymbol{P}_{\boldsymbol{\zeta}}(t) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{F}}, \tag{3.20}$$

for generic probability P_ζ. The explicit calculations for these matrices are presented in App. A.

[1] Because the photon wavefunctions start at zero amplitude, but are not necessarily zero at T, T_b must exceed T.

3.6.3 Enforcing the Initial Condition

The solution must satisfy the constraint that there is no cavity occupation at the beginning of the time window. Consider enforcing this condition on just the $j = 1$ cavity state.

$$0 = \alpha_{g_1}(0) = \frac{1}{\sqrt{T_b}} \sum_n C_n^{(1)}. \tag{3.21}$$

This condition can be guaranteed through

$$\begin{aligned} \boldsymbol{\phi_1^F} \cdot \boldsymbol{\alpha_{g_1}^F} &= 0, \\ \left(\boldsymbol{\phi_1^F}\right)_n &= 1. \end{aligned} \tag{3.22}$$

Equivalent conditions for the other cavity states can be expressed, leading a set of j_M conditions,

$$\begin{aligned} \boldsymbol{\phi_j^F} \cdot \boldsymbol{\alpha_{g_1}^F} &= 0 \ \forall j, \\ \left(\boldsymbol{\phi_j^F}\right)_n &= (f_n^{(1\to j)})^*. \end{aligned} \tag{3.23}$$

Assuming no cavity states are degenerate[2], the j_M vectors $\boldsymbol{\phi_j^F}$ define a j_M-dimensional subspace in which a valid solution vector $\boldsymbol{\alpha_{g_1}^F}$ should not lie. Using Gram Schmidt orthogonalisation, it is possible to produce a basis in which the final j_M states correspond to this subspace, along with a matrix $\boldsymbol{U}$ that can rotate a state in the Fourier basis to this basis. Additionally, define the projector $\boldsymbol{\Pi_{j_M}}$ which removes the last j_M components of a vector, and the reverse projector $\boldsymbol{\tilde{\Pi}_{j_M}}$ that takes a reduced vector and appends j_M coefficients with value zero. Any solution vector $\boldsymbol{\alpha_{g_1}^F}$ can now be projected to a valid solution

$$\boldsymbol{\alpha_{g_1}^P} = \boldsymbol{\Pi_{j_M}}\, \boldsymbol{U} \boldsymbol{\alpha_{g_1}^F}, \tag{3.24}$$

expressed as coefficients in the 'projected' basis, to which the superscript P refers. Probability matrices

$$\boldsymbol{P_\zeta^P} = \boldsymbol{\Pi_{j_M}}\, \boldsymbol{U} \boldsymbol{P_\zeta}\, \boldsymbol{U}^\dagger \boldsymbol{\tilde{\Pi}_{j_M}}, \tag{3.25}$$

[2]In the case of degeneracy, the number of conditions is equal to the number of nondegenerate states.

can then be defined over the new projected basis where ζ is a generic index that specifies the probability, with these projected matrices again labelled by subscript P. Within the projected basis, every state automatically satisfies the initial conditions, and therefore optimisation over this basis does not need to explicitly enforce the initial condition.

3.6.4 Normalising Probabilities

The (projected) matrix for total probability not in the initial state at time t

$$\boldsymbol{P}_{\boldsymbol{T}}^{\boldsymbol{P}}(t) = \sum_{j=1}^{j_M} \left(\boldsymbol{P}_{\boldsymbol{R}_j}^{\boldsymbol{P}}(t) + \boldsymbol{P}_{\boldsymbol{C}_j}^{\boldsymbol{P}}(t) \right) + \boldsymbol{P}_{\boldsymbol{E}}^{\boldsymbol{P}}(t) + \boldsymbol{P}_{\boldsymbol{e}}^{\boldsymbol{P}}(t), \tag{3.26}$$

can be calculated from other probability matrices. As for the three-level case, a photon shape is possible if the total probability not in the initial state remains below unity for all times. Therefore a particular probability $P_\zeta(t)$ can be expressed as

$$P_\zeta(t) = \frac{\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} \cdot \boldsymbol{P}_{\boldsymbol{\zeta}}^{\boldsymbol{P}}(t) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}}}{\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} \cdot \boldsymbol{P}_{\boldsymbol{T}}^{\boldsymbol{P}}(t_{\max}) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}}}, \tag{3.27}$$

where $0 \leq t_{\max} \leq T_p$ is the time for which $[\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} \cdot \boldsymbol{P}_{\boldsymbol{T}}^{\boldsymbol{P}}(t_{\max}) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}}]$ is maximised. This expression automatically normalises probabilities such that the solution can be physically realised.

3.6.5 Optimising Probabilities

To optimise probabilities, an iterative approach was used. Consider optimising a generic product of probabilities

$$V = \prod_{l=1}^{l_M} P_{\zeta_l}(t_l), \tag{3.28}$$

where l_M is the order of the probability product, ζ_l specifies the probability of the term in the product labelled by l, and t_l is the time at which the probability should be evaluated[3]. Each cycle of an iterative procedure begins with the current state $\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}}$. This vector is modified by adding a small correction

$$\delta\boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} = \epsilon \left[\left(\sum_{l=1}^{l_M} \boldsymbol{P}_{\boldsymbol{\zeta_l}}^{\boldsymbol{P}}(t_l) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} \right) - l_M \boldsymbol{P}_{\boldsymbol{T}}^{\boldsymbol{P}}(t_{\max}) \cdot \boldsymbol{\alpha}_{\boldsymbol{g_1}}^{\boldsymbol{P}} \right], \tag{3.29}$$

[3] in general, a sum of such products can be desired, but this is just a trivial extension.

where ϵ is set to a randomly-chosen positive small number each iteration to prevent the optimisation stalling and $t_{\max}$ is evaluated every iteration. This additional vector $\delta\boldsymbol{\alpha}^{\boldsymbol{P}}_{\boldsymbol{g_1}}$ is chosen because it matches the direction of steepest increase of V in the projected basis about the point representing the current trial solution on the assumption that $t_{\max}$ does not change. While the procedure technically allows the times t_l at which each probability is taken to be chosen, for maximising remote entanglement probability, the only interest is in outputs occurring by the full photon time T.

3.6.6 Calculating the Required Driving

Once a solution for $\alpha_{g_1}(t)$ has been obtained, the driving required to produce it can be calculated. In practice, the driving should be calculated for $\alpha_{g_1}(t) = (1-\epsilon)\alpha^S_{g_1}(t)$, where $\alpha^S_{g_1}(t)$ is the ideal solution and ϵ is a small number because the ideal solution is on the bounds of physical possibility, and therefore the required driving will tend to infinite strength.

Taking the equations of motion (Eq. 3.16), all variables except the driving strength Ω and the obtained cavity output functions $\left\{\alpha_{g_j}(t)\right\}$ can be eliminated to produce a differential equation

$$\begin{aligned}
\dot{\Omega}\left[\tilde{\kappa}\tilde{\gamma}\alpha_g + (\tilde{\kappa}+\tilde{\gamma})\dot{\alpha_g} + \ddot{\alpha_g} - g_1\sum_{j=1}^{j_M}(g_j\alpha_{g_j})\right] = \\
\Omega[i\Delta_u\tilde{\kappa}\tilde{\gamma}\alpha_g + (\tilde{\kappa}\tilde{\gamma} + i\Delta_u(\tilde{\kappa}+\tilde{\gamma}))\dot{\alpha_g} \\
+ (\tilde{\kappa}+\tilde{\gamma}+i\Delta_u)\ddot{\alpha_g} + \dddot{\alpha_g} - g_1\sum_{j=1}^{j_M} g_j(i\Delta_u\alpha_{g_j} + \dot{\alpha_{g_j}})] \\
+ |\Omega|^2\Omega(\kappa\alpha_g + \dot{\alpha_g})
\end{aligned} \tag{3.30}$$

where $\tilde{\gamma} \equiv \gamma + i\Delta_e$ and $\tilde{\kappa} \equiv \kappa + i\Delta_{g_1}$. These equations should be integrated from the initial conditions

$$\begin{aligned}
\alpha_e(0) &= -\frac{1}{g_1}\dot{\alpha_{g_1}}(0), \\
\alpha_u^{(\theta_0)}(0) &= \sqrt{1-|\alpha_e(0)|^2}e^{i\theta_0}, \\
\Omega^{(\theta_0)}(0,\theta_0) &= -\frac{(\tilde{\kappa}+\tilde{\gamma})\dot{\alpha_{g1}}(0) + \ddot{\alpha_{g1}}(0)}{g_1\alpha_u^{(\theta_0)}(0)},
\end{aligned} \tag{3.31}$$

where the superscript θ_0 indicates that there is freedom in the initial condition. The shape of the cavity output dictates that there might need to be population in $|e,0\rangle$ at time $t = 0$, meaning the initial state is a superposition of population in $|u,0\rangle$ and $|e,0\rangle$. The phase of this superposition is θ_0, and does not affect whether the photon can be produced, but does affect the drive Ω that produces it.

From a system initialised to $|u,0\rangle$, it is still possible to achieve the performance predicted, but the driving now consists of two parts. The first is an instantaneous (or in practice very fast) excitation that sets up the initial state of Eq. 3.31, and the second is specified by Eq. 3.30. The relative phase of the two parts of the drive determines θ_0, so must be precisely controlled.

It is important to note, however, that the conditions imposed have not restricted the driving in any way. As such, the drive may contain spikes of intensity where very fast changes in excited state amplitude are required to produce the optimum photon. These spikes persist even when the target amplitude is reduced from the theoretical maximum amplitude, and as such, finding the exact driving specified might not be the most useful step. It may be possible to implement conditions on the maximum intensity of the drive to obtain near-optimal photon shapes, but this has not been attempted.

3.6.7 Results

Comparison of numerical and analytical approach

For the case of the three-level system, an approach to derive upper and lower bounds for the maximum possible photonic output in time T was presented in Sec. 3.5. These bounds can be compared to the numerically optimised output following the approach of Sec. 3.6 for a system with $g = \kappa$ and $\gamma = 0.6\kappa$, with the results shown in Fig. 3.3. Examining first the comparison for different drive techniques over a range of times (panel a), it is seen that the numerical result always lies between the upper and lower bounds to photon extraction, which converge to the infinite time bound as time increases. The lower bound is not very tight, and generally (with the exception of very short times) the upper bound provides a much better estimate of optimised performance. An instant excitation approach gives near-optimal extraction

at short times, but the output saturates and cannot reach the maximum allowed by the system parameters. The linear drive curves are shown as an example of 'simple' driving procedures, which are sub-optimal for short-time extraction, but for extraction times beyond $10/\kappa$ perform very close to the optimum level. The population and output flux panels (b and c) show the the numerically optimised photon achieves performance much closer to the upper bound than the lower bound by deforming its shape slightly to have greater cavity population at earlier times. The results of this panel suggest that the upper bound is a reasonable tool to estimate the optimised output and photon shape, although this should be done with the understanding that the upper bound is not typically physically possible.

Optimisation of probability metrics

For remote entanglement, it is important to use an ion-cavity system with more than one cavity decay channel ($j_M > 1$) so that the final ion state can be entangled with the photonic state. For two identical nodes, the rate of remote entanglement is proportional to $P_{R_1}P_{R_2}$, where P_{R_1} is the probability the photon is emitted for one branch of the entangled ion-photon state, and P_{R_2} for the other branch. To this end, it is beneficial to optimise $P_{R_1}P_{R_2}$ for remote entanglement, but the numerical method can optimise more general metrics.

This capability is demonstrated in Fig. 3.4, which takes a system with $\gamma = 0.6\kappa$ and three cavity decays, such that $g_1 = \sqrt{1/3}\,\kappa$, $g_2 = -\sqrt{4/15}\,\kappa$, and $g_3 = \sqrt{1/30}\,\kappa$, chosen to reflect the $P_{\frac{3}{2}} \rightarrow D_{\frac{5}{2}}$ transition with the magnetic field orthogonal to the cavity axis in $^{88}Sr^{+}$. The magnetic field is set to give a Zeeman splitting of 5κ. The driving of this system was optimised for single photon emission (P_{R_1}), the two-component decay probability $P_{R_1}P_{R_2}$ pertinent for remote entanglement, and the three-component decay probability ($P_{R_1}P_{R_2}P_{R_3}$) for demonstrative purposes[4]. All optimisations were conducted for extraction time $T = 5/\kappa$. The results show that the same cavity system can be used to produce relatively pure single photon emission, or a binary ion-photon entangled state through changing the driving pulse, and can also produce a three-component decay superposition (although the cavity

[4]The two-component and three-component decays are not two-photon or three-photon events. In each case, the system produces at maximum one photon through the cavity.

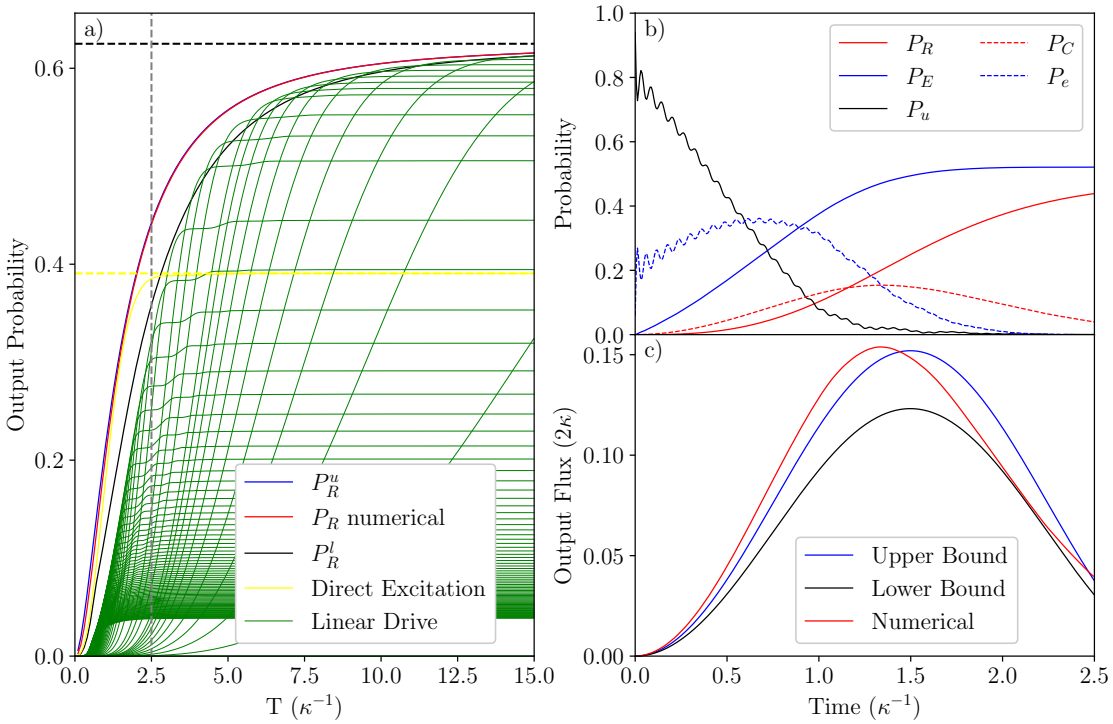


Figure 3.3: Comparison of the numerically-optimised performance and performance bounds for a three-level system with $g = \kappa$ and $\gamma = 0.6\kappa$. The numerical calculations used 101 Fourier basis states and $T_b = 1.3T$ for every output time. a) Comparison of the output probabilities achieved through a range of methods as a function of desired extraction time T. The blue and black lines mark the upper and lower bounds for the optimised extraction probability found in Sec. 3.5 respectively. The red line shows the numerically optimised output. In all cases, the numerically optimised output is between the upper and lower bounds of the maximised output, but for all but the shortest times, it essentially saturates the upper bound. The yellow line shows the extracted probability for direct excitation to the excited state, with the yellow dotted line marking the maximum extraction probability using this method, discussed in Sec. 3.2.2. The green lines show the output for a series of drives with linearly increasing amplitude at $\Delta_u = 0$, where the rate of increase of drive amplitude is varied between different lines. The infinite time extraction probability (discussed in Sec. 3.4) is marked by the dotted black line. The example time used for analysing the system populations and photon wavepackets in b) and c) is $T = 2.5/\kappa$, and is marked by the grey vertical line. b) System populations for the numerically optimised solution for output time $T = 2.5/\kappa$. The system is initialised in a superposition of initial and excited states before the population is either spontaneously emitted or collected through the cavity. c) The probability flux of the corresponding output wavepackets for the numerically optimised case (red), as well as the upper (black) and lower (blue) bounds for comparison. The numerically optimised output is slightly deformed from the shape recommended by the upper and lower bounds.

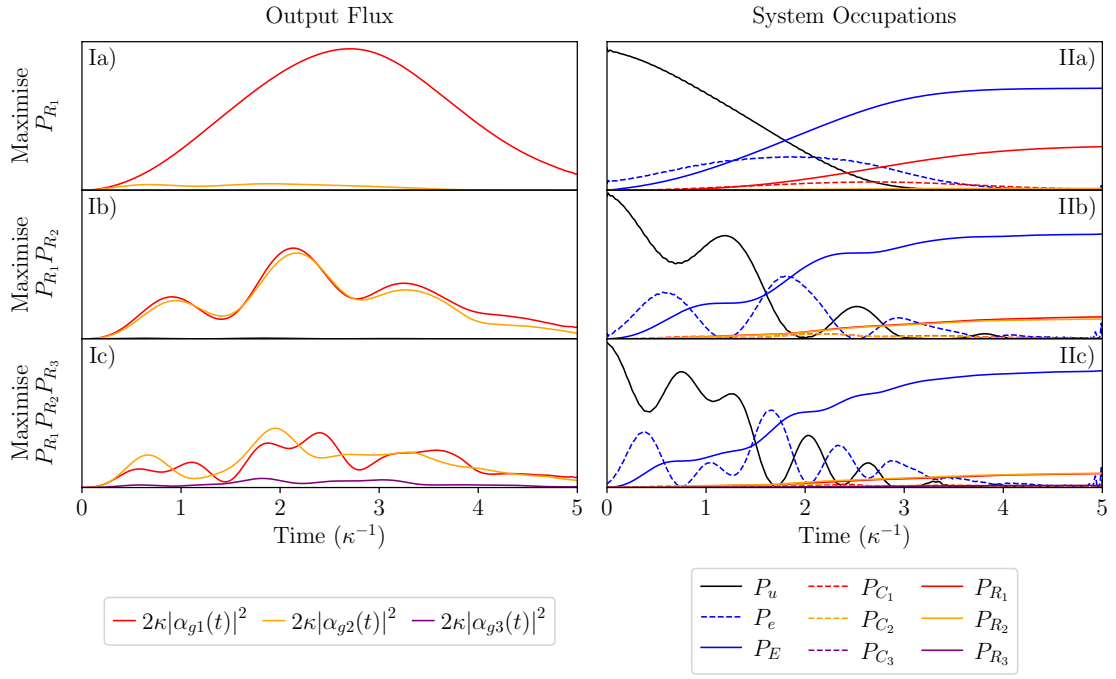


Figure 3.4: Example photon wavepackets (column I) and system populations (column II) for different maximisation targets of an emitter cavity system with $g_1 = \sqrt{1/3}\kappa$, $g_2 = -\sqrt{4/15}\kappa$, $g_3 = \sqrt{1/30}\kappa$, and $\gamma = 0.6\kappa$ over a time of $T = 5/\kappa$. The Zeeman splitting between $|g_1, 1_1\rangle$ and $|g_2, 1_2\rangle$, and $|g_2, 1_2\rangle$ and $|g_3, 1_3\rangle$, is $\Delta_Z = 5\kappa$. The numerical method used 201 Fourier basis states, with $T_b = 1.3T$. Row a) shows the case where the maximised quantity is the single photon emission probability P_{R_1}, row b) the two-component decay $P_{R_1}P_{R_2}$, and row c) $P_{R_1}P_{R_2}P_{R_3}$. As the number of different transitions that are targeted increases, the photon wavepackets gain more structure as different resonant frequencies must be matched simultaneously by the drive.

geometry is not optimal for this purpose). Finally, it should be noted that much more general products and sums of probabilities could be optimised, for example to discourage unwanted processes.

Effect of Zeeman splitting

The optimisation method permits the study of ion-cavity systems in a simplified setting because it removes the need to understand the driving that should be applied in a given situation. This capability was used to investigate how Zeeman splitting Δ_Z impacts the two-component decay probability $P_{R_1}P_{R_2}$ (proportional to the possible remote entanglement rate) for the system investigated in Sec. 3.6.7, but with the third decay channel removed for simplicity. This leaves a system with $g_1 = \sqrt{1/3}\kappa$, $g_2 = -\sqrt{4/15}\kappa$ and $\gamma = 0.6\kappa$. The behaviour of this system was investigated, using a sufficient Fourier basis for the problem according to the considerations of App. B, and

the results are shown in Fig. 3.5, which shows that an increasing Zeeman splitting reduces the possible remote entanglement rate from this system.

Simple pictures for the limiting cases can be used to understand the dependence of $P_{R_1}P_{R_2}$ on Δ_Z. Firstly, in the case of $\Delta_Z = 0$, the excited state actually couples directly to a single level

$$|g_{\text{eff}}, 1_{\text{eff}}\rangle = \frac{g_1 |g_1, 1_1\rangle + g_2 |g_2, 1_2\rangle}{\sqrt{g_1^2 + g_2^2}}, \tag{3.32}$$

with an effective coupling $g_{\text{eff}} = \sqrt{g_1^2 + g_2^2}$. Given that the ratio of the probability in this effective state that will be labelled P_{R_1} to the probability P_{R_2} is $g_1^2 : g_2^2$, one can calculate the success probability in infinite time

$$\begin{aligned} P_{\text{suc}}^{\max} &= \frac{g_1^2 g_2^2}{g_{\text{eff}}^4} P_{\text{emit}}^2, \\ C_{\text{eff}} &= \frac{g_{\text{eff}}^2}{2\kappa\gamma}, \\ P_{\text{emit}} &= \frac{2C_{\text{eff}}}{2C_{\text{eff}} + 1}, \end{aligned} \tag{3.33}$$

where P_{emit} is the probability that one system emits a single photon in infinite time.

Secondly, in the case that the Zeeman splitting becomes large, the production processes for the two components should be spectrally decoupled. The decay indexed by $j = \{1, 2\}$ then has a separate cooperativity $C_j = g_j^2/2\kappa\gamma$ and therefore an infinite-time ratio of cavity emission to spontaneous emission of $2C_j : 1$. This leads to an optimised success probability

$$P_{\text{suc}}^{\text{sep}} = \frac{f_{\max}(1 - f_{\max})2C_1 2C_2}{(2C_2 + 1 + 2f_{\max}(C_1 - C_2))^2} \tag{3.34}$$

where $0 < f_{\max} < 1$ is the proportion of total emission that decays through $|g_1, 1_1\rangle$ that maximises the expression.

The agreement of the data with these expressions is shown in panel b) of Fig. 3.5. The data does not extend to infinite extraction time, or to infinite Zeeman splitting Δ_Z, but the trend broadly matches the values expected from these simple models. In the context of remote entanglement generation, this means that simultaneous driving of decay channels separated by a small Zeeman splitting offers advantages in the total extraction probability, although, as discussed in Chapters 6 and 7, these conditions bring disadvantages that may outweigh the performance benefits.

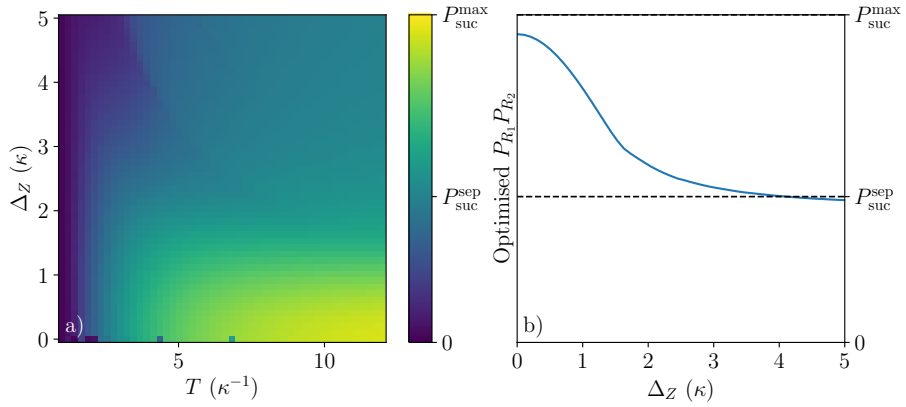


Figure 3.5: Investigation of the effect of Zeeman splitting on the two-component probability $P_{R_1}P_{R_2}$ optimised for time T for a system with $g_1 = \sqrt{1/3}\kappa$, $g_2 = -\sqrt{4/15}\kappa$ and $\gamma = 0.6\kappa$, where the two final states $|g_1, 1_1\rangle$ and $|g_2, 1_2\rangle$ are separated by Zeeman splitting Δ_Z. The numerical method used 161 Fourier basis states, with $T_b = 1.3T$ for all time values. a) The optimised $P_{R_1}P_{R_2}$ as a function of T and Δ_Z, showing that the optimised probability increases as the system is given longer to produce the output, but generally decreases as the splitting between the levels increases. b) The optimised output probability as a function of Δ_Z for just the longest time shown ($T = 12/\kappa$) compared to the expected infinite time output for zero-splitting case $P_{\text{suc}}^{\text{max}}$ and for the high-splitting case $P_{\text{suc}}^{\text{sep}}$. The optimised output approximately scans between these two limits, although does not match exactly because the time taken is not infinite, and the Zeeman splitting examined does not go to infinity.

3.7 Conclusions

This chapter has summarised different qualitative paradigms for extracting photons from emitter cavity systems, introduced a bounded approach for the extraction probability from a three-level system, and presented a numerical method applicable to more general systems and optimisation metrics.

The method of driving has a critical impact upon the performance and error vulnerabilities of photon production for remote entanglement. Direct excitation schemes are simple and fast, but surrender control and performance, whereas more gradual methods such as vSTIRAP or Raman driving can achieve higher output, but at the cost of a longer production time and sensitivity to temporal mixing errors. To maintain high performance and low temporal mixing errors over a range of cavity parameters, a driving scheme where the spontaneous decay from the excited state to the initial state has a low branching ratio should be combined with a gradual photon production method.

These non-instantaneous photon production processes can emit a single photon with a probability determined by the cooperativity over an infinite time period. In finite time, this emission probability reduces, and, as useful photons must be produced in finite time, it is important to know how the probability of emission depends upon the allowed emission time. A method to bound this probability for the three-level system was devised, and then extended to a numerical method for more general systems. This approach avoids the need to numerically optimise the driving pulse, while offering flexibility in the optimised metric. The numerical method was then used to examine aspects of photon production for remote entanglement, from which some conclusions can be drawn.

Firstly, for all but very short times, the maximum extraction probability from the three-level system is close to the upper bound, but for extraction times beyond $10/\kappa$, the maximum extraction probability becomes similar to the infinite time result, and simple drive pulses can achieve performance similar to optimised drives. Therefore, for extraction times beyond this threshold, it is reasonable to assume that the exact nature of the drive pulse will not have a critical impact on performance.

Secondly, for remote entanglement where two decay components must be collected, there is a performance advantage if the decays can be driven simultaneously with a small Zeeman splitting. The methods developed in this section allow this result to be obtained without assuming a particular driving scheme, and therefore the dependence of performance on the Zeeman splitting can be attributed only to the Zeeman splitting.

Though these conclusions are important, the limitations of the method presented should be understood. Firstly, physically realising the optimum output requires arbitrary control over the driving field, including effectively instantaneous transfers pulses at the start of, and potentially during, the process. Secondly, the usefulness of photon wavepackets for remote entanglement is not solely a question of rate; the wavepackets would ideally also offer low sensitivity to inevitable pathlength differences [226] through reasonably stable intensity and phase profiles- a criterion which is not trivial to satisfy [227]. Therefore, there would likely be a practical advantage to compromising the output probability somewhat to prioritise the simplicity of driving

and the robustness to temporal mismatch and other errors. These compromises would likely require involved simulation to optimise.

This investigation of cavity driving has provided important qualitative conclusions that will inform the remainder of the investigation, however, as will be discussed in the following chapter (Chapter 4), under reasonable assumptions for our proposed ion-cavity system, the exact characteristics of the driving can be largely ignored. Nevertheless, it is important that the origin of these conditions is understood.

4
Cavity Optimisation

4.1 Introduction

In designing a cavity for remote entanglement experiments, it is important to understand how to achieve the best performance under realistic constraints, particularly the residual transverse misalignment of the mirrors inherent to our proposed designs. The aim of this chapter is to develop and use a procedure for optimising the cavities in remote entanglement experiments that produces good results, but is intuitive enough to yield physical insight into the characteristics of optimal cavities and the most fruitful directions for engineering development.

Contribution and Acknowledgements

The content of this chapter is based upon the work published in [134], where more detail is presented. For that publication, Dr. Shaobo Gao and Dr. Jacob Blackmore wrote the modelling code, Dr. Jacob Blackmore generated the data with assistance from the author, Dr. Jacob Blackmore and produced the figures, and Dr. Shaobo Gao, Dr. Thomas Doherty, Dr. Jacob Blackmore, Dr Joseph Goodwin, and the author all contributed to developing the approach and the text. The author wrote this chapter's text and adjusted the figures to make the content of the publication suitable for presentation in the current work.

4.2 Optimisation Scenario

4.2.1 System

A fully comprehensive optimisation of the cavity for a remote entanglement experiment would be a complicated endeavour. Accounting for the full ion level structure, different possible level encodings and driving protocols, as well as common noise sources is a complex task and will not be attempted here owing to its difficulty and the likelihood of obfuscating the key physics. Instead, the model will be reduced to the simplest that can reasonably capture the trade-offs involved in cavity design, with the caveat that the optimisation is only intended to analyse these trade-offs and estimate performance, not to find a rigorous 'optimal' design. Indeed, the many possible metrics for an emitter-cavity system (photon extraction rate, ease of fabrication etc.) mean that an optimisation that tried to summarise performance in one complicated cost function would likely be less useful than a simple approach that clarifies the underlying physics.

The simplified model assumes the ion has the three-level structure discussed in Chapter 3, and is coupled to the mode of a cavity. The profiles are concave spherical caps with radius of curvature R and diameter D_M, and the centres of the mirrors are set at a nominal distance L from each other along the cavity axis. However, the cavity is assumed to be slightly misaligned, both transversely (with misalignment $M_\perp$) and longitudinally (with misalignment $M_\parallel$)[1]. The mode of the misaligned cavity tilts with angle ϕ and has half angle divergence θ

$$\begin{aligned} \phi &= \arctan\left(\frac{M_\perp}{2R - L - M_\parallel}\right) \approx \frac{M_\perp}{2R - L - M_\parallel}, \\ L_{\text{eff}} &= 2R - \frac{2R - L - M_\parallel}{\cos(\phi)} \approx L + M_\parallel, \\ \theta &= \sqrt{\frac{2\lambda}{\pi\sqrt{L_{\text{eff}}(2R - L_{\text{eff}})}}}, \end{aligned} \tag{4.1}$$

where L_{eff} is the effective length of the cavity mode, and the approximations hold in the case of small mode angle tilt. It is assumed that the ion is positioned at the

[1] the length must be tunable over $\lambda/2$ to lock to the desired resonance, but this does not necessarily rule out larger longitudinal misalignments.

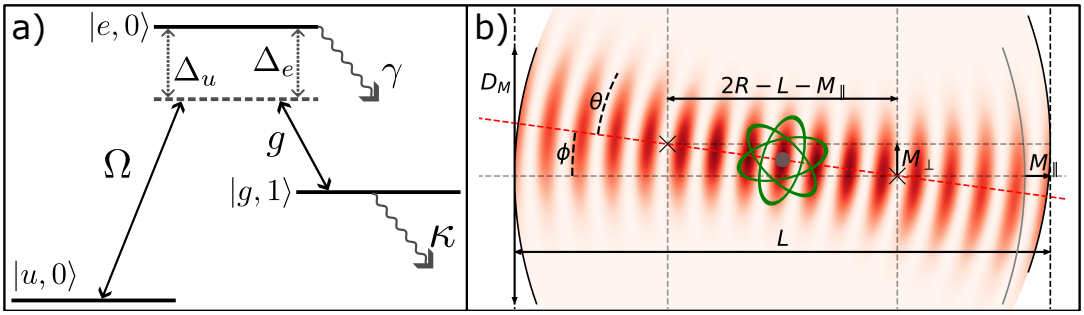


Figure 4.1: Schematic diagram of the system on which the optimisation of this chapter is based. a) A diagram of the level structure, identical to that shown in Chapter 3. The system consists of three levels that are the tensor product of an atomic state $\in \{|u\rangle, |e\rangle, |g\rangle\}$ and cavity photon occupation $\in \{|0\rangle, |1\rangle\}$. The system is initialised in state $|u,0\rangle$ and driven to $|g,1\rangle$ via a two-photon process through intermediate state $|e,0\rangle$. The transfer involves the absorption of photon from the laser drive that couples the initial and excited state with strength Ω and emission of a photon into the cavity mode that couples the excited state and final state with strength g. The detuning of transition between the excited and final states and the cavity resonance (Δ_e) and the detuning of the initial state from Raman resonance with the final state (Δ_u) are labelled for completeness, but play no role in the content of this chapter. The excited state amplitude decays with rate γ due to spontaneous emission, and the final state with rate κ due to decay of the cavity mode. b) The geometry of the misaligned cavity. Two spherical mirrors of radius of curvature R and diameter D_M are set a nominal length L apart. However, the mirrors are misaligned by $M_\parallel$ in the axial direction and $M_\perp$ in the transverse direction. The cavity mirrors are drawn in black, with the 'aligned' position of the right-hand mirror drawn in grey. A tilted mode forms in the misaligned cavity, with angle of propagation ϕ with respect to the nominal cavity axis, and mode divergence angle θ. It is assumed that the ion can be placed at the centre of the cavity mode that forms in the misaligned cavity.

centre of the mode that does form, for example using the techniques in [228]. The ion level structure and cavity geometry are depicted in Fig. 4.1. Generally, $M_\parallel$ and $M_\perp$ will be different from each other, but to simplify this analysis, the pessimistic case that both misalignments have the same value M is assumed.

4.2.2 Metric

Overview

Both the fidelity and rate of remote entanglement are of significance, but optimisation requires a single metric. The effects of cavity parameters on rate will be discussed later, but first the effects through which the cavity parameter choices impact mechanisms of fidelity loss will be analysed in turn.

Fidelity

Firstly, realistic ion level structure is more extensive than the three levels modelled here. Specifically, the initial state, excited state and final ground state will all typically be part of larger manifolds; indeed, the two photon remote entanglement protocols typically use more than one final ground state in order to produce ion-photon entanglement. This extra level structure means that the desired cavity process is not necessarily the only cavity process that is driven; typically, these extra unwanted processes will be disfavoured by geometry and mismatched resonance, but cannot all be eliminated completely. The particular cavity decays that cause fidelity loss and the mechanism of this fidelity loss depend upon the specifics of the chosen remote entanglement protocol and driving scheme, however, these unwanted cavity processes usually reduce fidelity to some degree. To resonantly select against these parasitic processes, they will typically have to be detuned from resonance by much more than a cavity half-linewidth κ (although the cavity coupling and drive do also play a role [192]). This means that the fidelity of the process is linked to the cavity parameters through κ. However, for the optimisation, it is assumed that the Zeeman splittings between different atomic states are sufficiently large that fidelity loss due to such 'spectator decays' is insignificant and the cavity parameters can be adjusted freely. For ions without nuclear spin, and therefore only fine structure, this is not an unreasonable parameter regime. It should, however, be noted that for other experimental platforms this is not the case. Particularly for schemes addressing the hyperfine structure, the Zeeman splittings might be limited by the breakdown of the linear Zeeman regime [229]. Such experiments would require κ lie below a threshold to guarantee the requisite resonant selectivity, and in such cases the optimisation procedures would differ (for example [230]).

The second important source of infidelity is temporal mixing due to specific spontaneous scattering effects during the photon generation process. As discussed in Chapter 3, this infidelity can be made negligible by choosing appropriate manifolds for the initial, excited and final states, and for the purposes of this optimisation we will assume that this has been done. It is important to note that if this assumption did not hold, the fidelity of the process would depend on exactly how the photons were

produced. In particular, fast driving schemes would be favoured, for which, as noted in Chapter 3, the photon extraction probability has a different dependence on cavity parameters compared to the adiabatic processes that will be assumed for this analysis.

Rate

Under the conditions above, it is assumed that the fidelity of the process is independent of the parameters of the cavity, and therefore the optimisation should maximise the entanglement rate. A simple extension of Eq. 2.26 in Chapter 2 gives the rate of remote entanglement as

$$R_{\text{ent}} = \frac{1}{2}\frac{\left(P_{\text{ext}}\eta_n\eta_d\right)^2}{\tau_{\text{att}}}, \tag{4.2}$$

where P_{ext} is the probability that a photon is extracted from the cavity, η_n the efficiency of routing this photon to a detector, η_d the efficiency of the detector and τ_{att} the attempt rate. This shows that entanglement rate depends upon the cavity through the probability of photon extraction and the rate at which entanglement attempts can be made. The attempt time is composed of several components

$$\tau_{\text{att}} = \tau_{\text{lat}} + \tau_{\text{prep}} + \tau_{\gamma} + \tau_{\text{prop}}, \tag{4.3}$$

where the times associated with latency, preparation to the initial state, photon production in the cavity and photon propagation to the detectors are given by τ_{lat}, τ_{prep}, τ_{γ}, and τ_{prop} respectively. In previous remote entanglement experiments, τ_{att} has been limited to 1 µs, of which 350 ns was preparation [96]. The requirement that for negligible temporal mixing the branching ratio to the initial state is small means that preparation times in the proposed experiment would likely be significantly longer. Therefore, the photon production time τ_{γ} will not significantly reduce the attempt rate while under ≈500 ns. As discussed in Chapter 3, the time required to extract a photon from a cavity with good efficiency is typically set by κ. As shown in Fig. 4.2, the extraction probability approximates the adiabatic limit for $\kappa\tau_{\gamma} > 10$ for a wide variety of system parameters, even when driven with a simple sine squared pulse. Thus for $\kappa > 10/500\,\text{ns} = 2\pi\times\,3.2\,\text{MHz}$, it is reasonable to assume that the adiabatic limit on photon production can be achieved without reducing the attempt rate. This

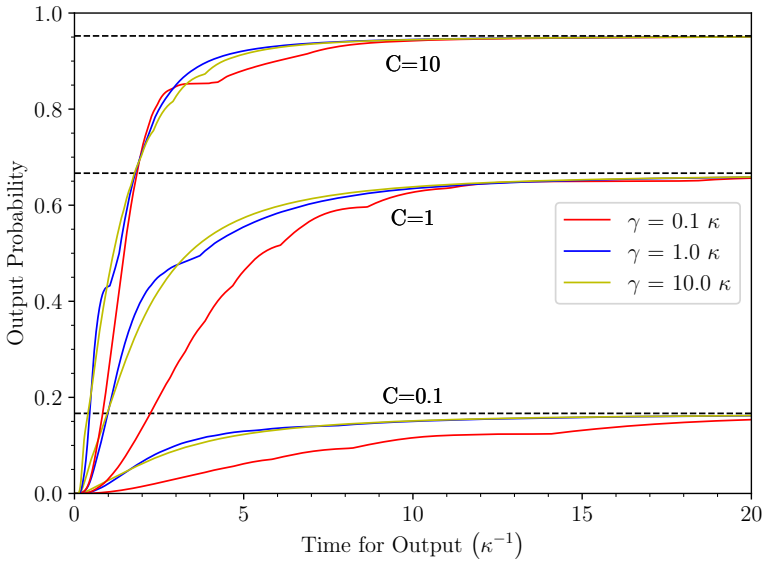


Figure 4.2: Finite time photon outputs for a range of ion-cavity coupling regimes using cavities with zero non-transmissive loss. The evolution of each system is simulated with a selection of pulses whose amplitude follows a sine-squared shape and is always resonant with the $|u,0\rangle \rightarrow |e,0\rangle$ transition. At each time point, the largest value that the output probability takes over all of the simulated pulses is recorded as a lower bound to the output achievable with a simple pulse. The systems simulated form three groups with cooperativities of: $C = 0.1$, 1.0 and 10.0, where the adiabatic limit $P_{\text{ext}} = 2C/(2C+1)$, valid for slow driving/infinite output time window and no non-transmissive cavity losses, is indicated with a dashed black line. Within each cooperativity set, the excited state amplitude decay rate γ takes three different values so that different coupling regimes may be compared. Generally, the photon output approaches the adiabatic limit before a time of $10/\kappa$. The cases where $\gamma = 0.1\kappa$ need more time (expressed in units of κ^{-1}) to approach the adiabatic output, but these cases are likely to have high κ anyway, and photon wavepackets longer than $10/\kappa$ are unlikely to present a bottleneck.

condition applies to all cavity decay rates considered in this section, however, for systems with different timescales, the length of the photon could have a significant impact on the entanglement rate, so this assumption is not universally applicable.

Given that the adiabatic limit to photon extraction probability can be attained without reducing the attempt rate significantly, there is no benefit to producing shorter photons, and the rate of the process can be maximised by simply maximising the extraction probability P_{ext}. This is therefore the metric that will be used to determine which cavity design is optimal.

4.3 Optimisation Overview

4.3.1 Separation of Geometric and Optical Properties

The parameters which describe the main properties of the cavity can be divided into two classes. The first contains fixed parameters that are generally desired to be zero, but due to fabrication tolerances cannot be made zero. This category contains the misalignment M and the rms roughness of the mirror surface σ, which determines the round-trip scattering loss $\mathcal{L}_{\text{scat}}$ through

$$\mathcal{L}_{\text{scat}} \approx 2\left(1 - e^{-\left(\frac{4\pi\sigma}{\lambda}\right)^2}\right), \tag{4.4}$$

where λ is the resonant wavelength [175].

The second category contains parameters which can be controlled, although generally only within certain limits. These are the mirror radius of curvature R, diameter D_M, transmission T and the cavity length L.

Taken together, these properties define a cavity system that has a maximum extraction probability in the adiabatic limit

$$\begin{aligned} P_{\text{ext}} &= \frac{2C}{2C+1}\frac{T}{\mathcal{L}} = \frac{2C}{2C+1}\frac{T}{T+\mathcal{L}_{\text{in}}}, \\ \mathcal{L}_{\text{in}} &= \mathcal{L}_{\text{scat}} + \mathcal{L}_{\text{clip}}, \end{aligned} \tag{4.5}$$

where $\mathcal{L}$ is the total round-trip loss of the system, comprising transmissive loss T and 'intrinsic' loss $\mathcal{L}_{\text{in}}$, which can itself be divided into $\mathcal{L}_{\text{scat}}$, the scattering and absorption loss due to mirror roughness and coating absorbance, and $\mathcal{L}_{\text{clip}}$ the loss due to incomplete confinement of the mode in the finite diameter of the mirror. The two factors in the product expression for P_{ext} encode a trade-off on the transmission T. A low transmission increases the cooperativity, and thus increases the probability that the photon is generated in the cavity, but a high transmission increases the second term, which represents the probability that a photon in the cavity is transmitted out rather than lost. Following [161], this trade-off can be optimised to give a maximum extraction probability

$$P_{\text{ext}} = 1 - \frac{2}{1+\sqrt{1+2C_{\text{in}}}}, \qquad C_{\text{in}} = \frac{g^2}{2\gamma\kappa_{\text{in}}}, \qquad \kappa_{\text{in}} = \frac{c\mathcal{L}_{\text{in}}}{4L}, \tag{4.6}$$
$$= 6\alpha_B \frac{\epsilon_q^V}{|\boldsymbol{\epsilon}^{\boldsymbol{V}}|}^2 \theta^2 \frac{1}{\mathcal{L}_{\text{in}}},$$

where C_{in} and κ_{in} are the intrinsic cooperativity and intrinsic decay rate respectively, α_B is the branching ratio of excited state decay to the chosen cavity transition, as defined in Chapter 2, and $\epsilon_q^V / \left|\boldsymbol{\epsilon}^{\boldsymbol{V}}\right|$ is the projection transverse to the cavity axis of the electric field with angular momentum q corresponding to the chosen transition. This maximum output occurs for a transmission of

$$T = \mathcal{L}_{\text{in}}\sqrt{1+2C_{\text{in}}}. \tag{4.7}$$

This means that the optimisation of P_{ext} is an optimisation of C_{in}, and thus of $\theta^2/\mathcal{L}_{\text{in}}$, which are defined purely by the geometry of the cavity. The transmission can then be set deterministically. Thus the geometric and optical properties of the cavity can be optimised sequentially.

The mode divergence angle θ can be increased by adjusting the geometry towards the concentric configuration $L_{\text{eff}} = 2R$ (see Eq. 4.1). This has two side effects. Firstly, it increases the size of the mode on the mirror, and secondly increases the tilt-angle ϕ of the mode away from the nominal axis for a given misalignment. Both effects will increase $\mathcal{L}_{\text{clip}}$. Therefore, both θ and $\mathcal{L}_{\text{in}}$ depend upon the geometry, and the ideal geometry optimises the compromise between the two.

4.3.2 Bounding the Geometrical Optimisation

The geometric optimisation can, however, be simplified with an approximation that becomes apparent when the clipping loss is evaluated. Assuming the x axis defines the direction of misalignment, the centre of the mode hits the mirror at a transverse coordinate $x_m = \phi L_{\text{eff}}/2$. As this approaches the edge of the mirror, the mode will begin to spill outside the spherical region of the mirror, leading to loss. This loss can be approximated by assuming the mode retains the same shape it would have in the absence of clipping and then calculating the fraction of mode power that falls

outside the confines of the mirror, as used in [231]. Extending the treatment of [185] to the off-axis case, the clipping loss is calculated through

$$\begin{aligned} u(z_M) &= u_0 \exp\left[-\left(\frac{(x-x_m)^2+y^2}{w^2}\right)\right] \\ &= u_0 \exp\left[-\left(\frac{\left(x-\phi\frac{L_{\text{eff}}}{2}\right)^2+y^2}{\left(\frac{\lambda}{\pi\theta}\right)^2+\left(\frac{L_{\text{eff}}\theta}{2}\right)^2}\right)^2\right], \\ \mathcal{L}_{\text{clip}} &= 1-\left(\int_{S_M}|u(z_M)|^2\,dA\right)^2, \end{aligned} \tag{4.8}$$

where $u(z_M)$ is the cavity mode function in the mirror plane, w is the waist of the mode at the mirror, S_M is the mirror surface and u_0 is a normalisation factor such that $\int_{S_\infty}|u|^2\,dA = 1$, where S_∞ is the surface of an infinite mirror. The squared integral in the expression for $\mathcal{L}_{\text{clip}}$ accounts for the two bounces per round trip.

This method of evaluating clipping loss is an approximation, and more accurate methods [232, 233] account for the change in mode shape as it begins to clip. Nevertheless, as discussed below, it is only necessary to evaluate clipping loss to the correct order of magnitude for the method presented here[2].

Evaluating clipping loss for a series of cavity geometries leads to Fig. 4.3, which shows the clipping loss as a function of cavity length and mirror radius of curvature for different mirror diameters. The vital conclusion is that the clipping loss changes by orders of magnitude for small changes in geometry, whereas θ changes much more gradually. Therefore, while $\mathcal{L}_{\text{clip}} \ll \mathcal{L}_{\text{scat}}$, the geometry should be optimised solely to maximise θ, and outside of this region, $\mathcal{L}_{\text{clip}}$ increases so rapidly that there is no significant improvement to be made in θ before the clipping loss dominates and the cavity does not confine light effectively. This means that the optimisation of C_{in} over cavity geometry can be approximated as an optimisation of θ over the geometries for which $\mathcal{L}_{\text{clip}} < \mathcal{L}_{\text{thresh}}$, where $\mathcal{L}_{\text{thresh}} \ll \mathcal{L}_{\text{scat}}$ is a threshold loss taken here to be 1 ppm. While $\mathcal{L}_{\text{thresh}}$ appears to have been set arbitrarily, $\mathcal{L}_{\text{clip}}$ is such a strong function of geometry that the exact value of the threshold loss has very little impact on the boundary region it defines, and therefore the optimisation produces near-identical results for any reasonable threshold loss.

[2]The validity of the clipping loss model will be discussed further in Chapter 5.

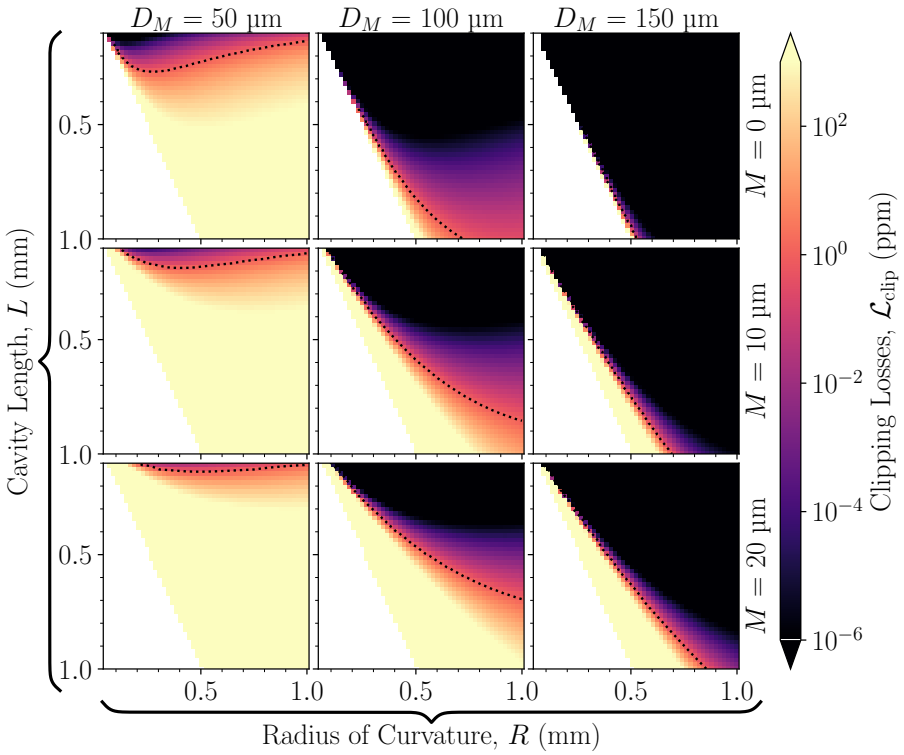


Figure 4.3: Clipping loss $\mathcal{L}_{\text{clip}}$ as a function of cavity length and mirror radius of curvature for a given mirror diameter D_M set by the column and misalignment M set by the row, where the resonant wavelength is 1033 nm. Geometries that do not fulfil the cavity stability criterion are left white as no classical mode is stably trapped. The black dotted line indicates the boundary of 1 ppm clipping loss, which is taken as the threshold clipping loss for this optimisation. In all cases (though particularly for the large diameter mirrors), the clipping loss changes by orders of magnitude within a small range of L and R around the boundary region. This means that the clipping loss rapidly transitions from negligible to unacceptable around the boundary, which makes it reasonable to treat the boundary contour as defining an allowed region within which to perform the optimisation. The region is larger if the mirror diameter is bigger, and smaller if the misalignment is greater.

The geometry must typically be optimised with respect to constraints associated with the manufacture. Here it will be assumed that two constraints apply. Firstly, the length must exceed a minimum length L_{min} that is set by the trap structure, optical access requirements, and the need to keep the ion away from the dielectric mirror surfaces to reduce heating and minimise disruption of the trap potential due to the cavity [118]. The second is that there is a maximum volume V_M

$$\begin{aligned} V_M &= \frac{\pi h}{6}\left(\frac{3D_M^2}{4} + h^2\right), \\ h &= R - \sqrt{R^2 - \left(\frac{D_M}{2}\right)^2}, \end{aligned} \tag{4.9}$$

of the mirror profile spherical cap, where h is the cap depth. This constraint

is motivated by the increase in milling time and difficulty in achieving smooth surfaces over larger scales. The constraint links the radius of curvature and diameter of the mirrors, limiting the possible geometries that can be realised. Note that specific methods of mirror construction may impose additional constraints or alter the constraint implemented. However, the method could be adapted to those scenarios simply.

Applying this constraint, the optimisation of θ takes the form given in Fig. 4.4. The milling volume limit is combined with the radius of curvature R to give the mirror diameter D_M, and from this, the bounded region defined by the clipping loss can be shown in (L, R) space. It is clear that the ideal point is at the minimum length, with the radius of curvature reduced towards a concentric configuration until clipping losses reach threshold.

With the best divergence angle obtained, the extraction efficiency and transmission can be calculated as a function of scattering loss using Eq. 4.6 and 4.7, as indicated in Fig. 4.5.

4.3.3 Summary of Optimisation Procedure

To summarise, the fixed parameters of the cavity are misalignment M and rms roughness of the mirror surface σ leading to scattering loss $\mathcal{L}_{\text{scat}}$. The controllable parameters are cavity length L, mirror radius of curvature R, diameter D_M and transmission T.

First, the geometry is optimised in the (L, R) space. The minimum length L_{min} is enforced and the maximum milling volume V_M defines a maximum diameter for a given radius of curvature. The limit $\mathcal{L}_{\text{clip}}$ <1 ppm, given misalignment M, defines a region in (L, R) space over which the divergence angle θ is optimised. This procedure pushes the length to the minimum, and reduces the radius of curvature towards the concentric configuration until the clipping loss boundary is reached.

With the geometry optimised, the transmission T is set through Eq. 4.7 in order to maximise the photon extraction efficiency. Thus, while the geometry is independent of $\mathcal{L}_{\text{scat}}$, its value does affect the optimum transmission and the optimised extraction efficiency.

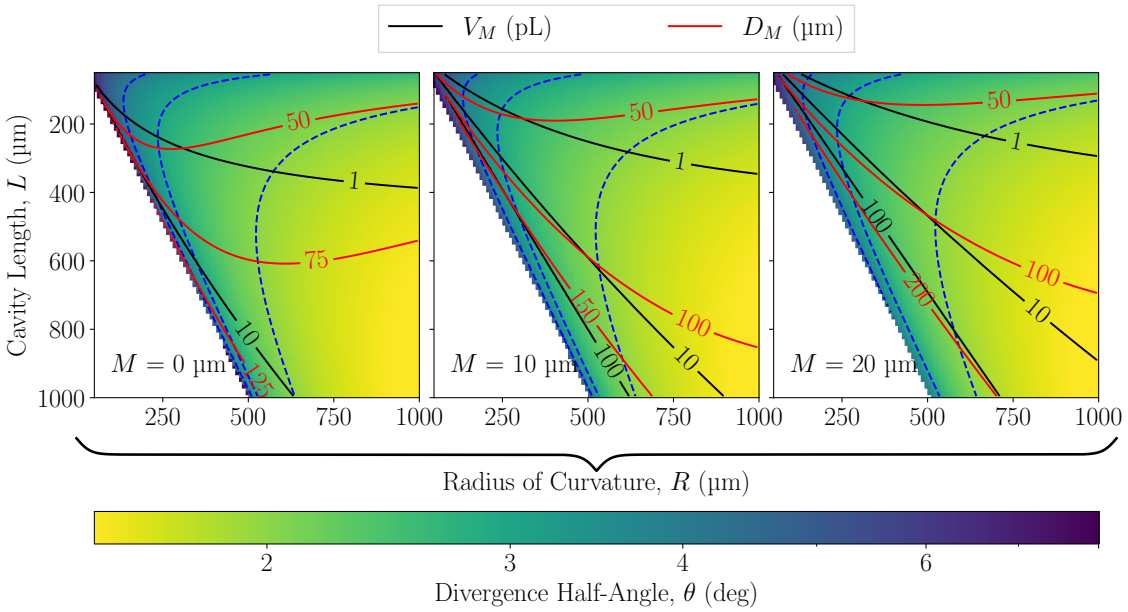


Figure 4.4: Diagrams of the clipping loss boundary approach for three different mirror misalignments. For each misalignment, the divergence angle θ can be calculated as a function of cavity length L and mirror radius of curvature R, where the white region indicates that no stable cavity mode can be formed. The clipping loss condition $\mathcal{L}_{\text{clip}} < \mathcal{L}_{\text{thresh}}$ defines allowed regions on these plots that depend upon the mirror construction. If the diameter D_M is fixed, the cavity geometry must remain within the relevant red contour to satisfy the clipping loss condition and, as used for the remainder of the section, if the milled volume V_M is constrained, the geometry must lie within the black contours. The optimum geometry is the highest divergence angle configuration within the relevant boundary The divergence angle is displayed by the colour map, with additional blue dashed contours to guide the eye. When the volume V_M is constrained, the minimum length is always preferred and then the smallest radius of curvature with sub-threshold clipping losses is chosen. For cavities with higher misalignment or more restrictive construction limits, the range of allowed geometries is smaller, and the performance is consequently reduced.

4.4 Expected Cavity Performance

The optimisation procedure presented in Sec. 4.3 can be used to predict the cavity performance for given M, $\mathcal{L}_{\text{scat}}$, L_{min}, and V_M. The results of this optimisation for a variety of these parameters are shown in Fig. 4.6, highlighting the compromises involved in cavity optimisation. In the ideal case that the minimum cavity length is small, the milled volume is large, and the misalignment is small, the cavity achieves a highly concentric configuration $R \approx 2L$ and the extracted probability is large. If the minimum cavity length or misalignment increases, the mirror diameter must increase to keep the mode confined. If the milled volume is reduced, the radius of curvature and diameter both increase. In all cases, the cavity moves away from

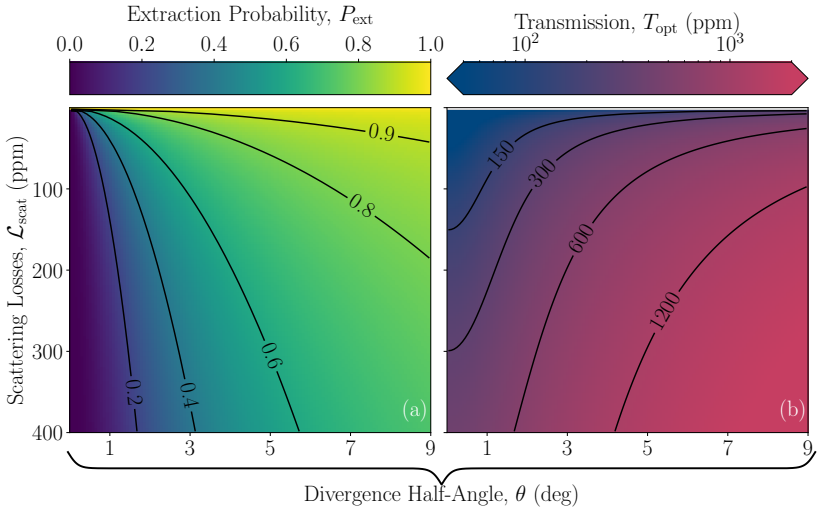


Figure 4.5: The a) extraction probability P_{ext} and b) optimum coating transmission T_{opt} that can be expected for a given divergence half angle and scattering loss under the optimisation procedure. The geometric and branching factor in the cooperativity $\alpha_B(\epsilon_q^V/\left|\boldsymbol{\epsilon}^{\boldsymbol{V}}\right|)^2$ is taken as 0.2, a value used for all subsequent figures. During the optimisation, θ is maximised under the condition that $\mathcal{L}_{\text{clip}} < \mathcal{L}_{\text{thresh}} = 1$ ppm, which, for the parameters used, means that the scattering loss $\mathcal{L}_{\text{scat}}$ is the dominant source of loss. From the maximised θ, the scattering loss can be used to calculate the T_{opt} that maximises P_{ext}, and the P_{ext} obtained in this case. Generally, for higher divergence angles, the cooperativity saturates and it becomes more beneficial to use higher transmission. Higher scattering loss also favours higher transmission so that a reasonable proportion of photons emitted into the cavity mode can be collected.

the concentric configuration until clipping losses can be tolerated. This results in a performance reduction, although the reduction is generally modest as the cavity parameters adjust to reduce clipping loss in the most efficient way to preserve P_{ext}.

Finally, it should be noted that a particular minimum length and misalignment demands a minimum milled volume from the mirror fabrication method. No mirrors with milled volume below this threshold will be able to confine the mode sufficiently to avoid crossing the clipping boundary. This three-sided relation can also be viewed from the two other directions: For a given milled volume and mirror misalignment, there is a maximum cavity length above which no cavity with acceptable clipping loss can be formed, and for a given milled volume and minimum cavity length, there is a maximum tolerable misalignment.

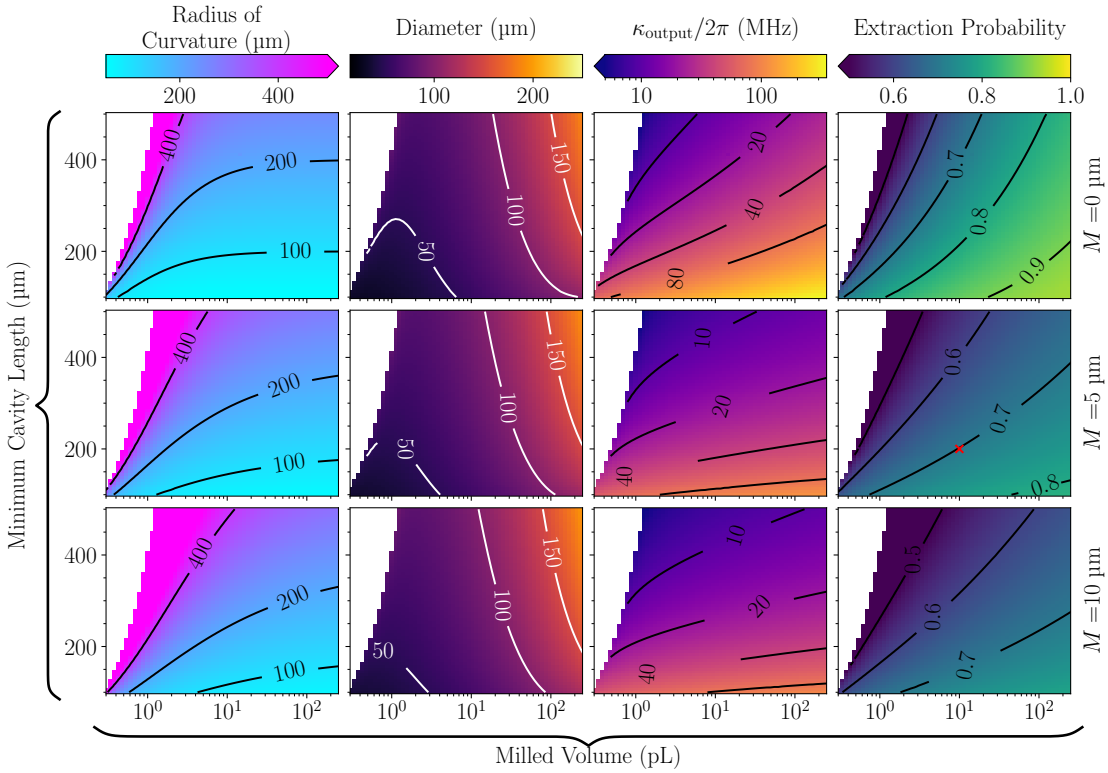


Figure 4.6: The cavity parameters and performance expected by the optimisation procedure as a function of milled volume V_M and minimum cavity length L_{min} for three different misalignment values. The resonant wavelength is always 1033 nm, and scattering loss is always 100 ppm, although changes to the scattering loss would not impact the optimised geometric parameters. The geometric and branching factor in the cooperativity $\alpha_B(\epsilon_q^V/\left|\boldsymbol{\epsilon}^{\boldsymbol{V}}\right|)^2$ is taken as 0.2. From left to right, the first three columns show the optimised mirror radius of curvature R, mirror diameter D_M and transmissive cavity decay rate $\kappa_{\text{output}} = Tc/4L$ respectively. Note the transmissive cavity decay rate provides a lower bound for the cavity loss rate and thus indicates that the assumption of Sec. 4.2.2 that photons could be extracted in the adiabatic regime without significantly extending the preparation time is reasonable. The right column shows the expected extraction probability P_{ext}. The white regions in each plot show combinations of cavity length, milled volume and misalignment for which there is no stable cavity with negligible clipping loss. The highest performance is found for the largest milling volumes and smallest cavity lengths at zero misalignment. Misalignment reduces performance, although these reductions are not necessarily very large as long as the geometry is adjusted accordingly. The optimised point taken to assess robustness in Fig. 4.7 is shown as a red cross on the furthest right graph of the middle row.

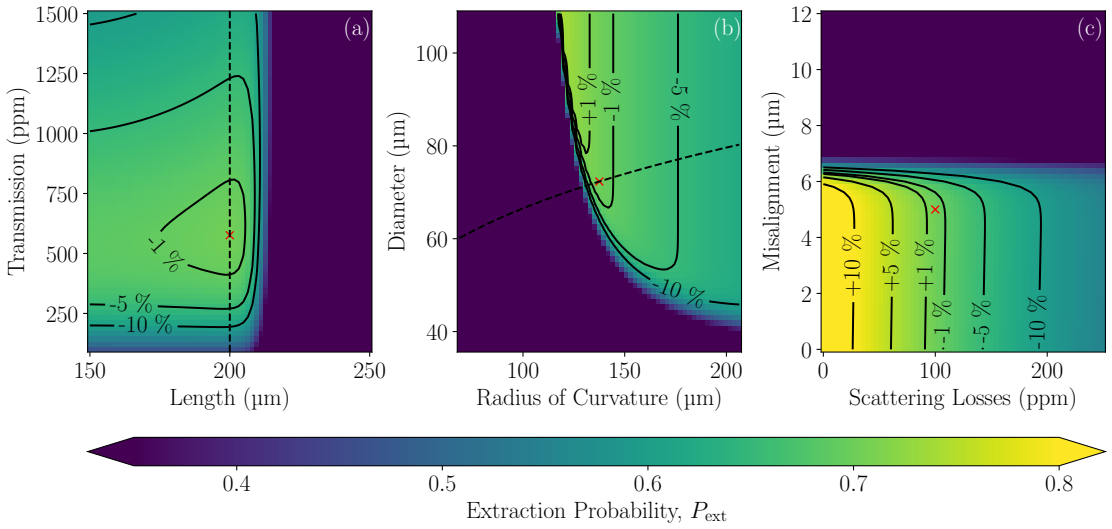


Figure 4.7: Change in performance expected from the solution marked with the red cross in Fig. 4.6 if the cavity parameters do not take their optimised value. a) percentage change in extraction probability given the transmission and length are not at their optimised value. Non-optimal choice of output transmission T reduces the performance from optimum, but gradually so. For the length, however, the situation is markedly asymmetrical. If the length is shorter than anticipated, clipping loss remains negligible, and the divergence angle reduces along with the extraction probability. If the length is longer, the cavity configuration is pushed beyond the clipping threshold and the performance rapidly deteriorates. b) differences in performance for changes in radius of curvature and diameter, where the dotted line shows the line of constant volume. Performance reduces rapidly if the radius of curvature or diameter are lower than intended. As the constant volume contour lies roughly orthogonal to the performance contours, the extraction probability must be compromised to move away from the sharp drop in performance while keeping a constant mirror size. c) extraction probability dependence on scattering loss and misalignment. If the misalignment is smaller than anticipated, performance is unaffected, but if larger, the extraction probability falls rapidly as clipping loss dominates. As scattering loss is purely parasitic, if the scattering loss is higher than anticipated, the extraction probability is worse, but if lower, the extraction probability is higher.

4.5 Robustness to Parameter Errors

As discussed, the optimisation procedure increases the cavity mode divergence angle as far as possible before clipping loss becomes significant, at which point clipping loss rapidly becomes unacceptable. This means that the 'optimised' design is very sensitive to errors in the controlled parameters that increase clipping loss. This is shown in Fig. 4.7, which calculates the change in cavity performance if either the controlled parameters were mis-set from the optimised values, or the fixed parameters took on different values to those assumed.

For the controlled parameters, mis-setting the parameters that define the geometry

(L, R and D_M) causes sharp decreases in performance in the directions that increase clipping loss (increase for L, and decrease for R or D_M) as the clipping threshold is breached and losses quickly dominate. A much gentler decrease is seen in the direction that reduces divergence angle for L and R, and no decrease at all for increasing the diameter. In contrast, the transmission has a gentle decrease in either direction from the optimum because mis-setting this value will not push the system beyond the clipping threshold, merely to a suboptimal compromise of photon generation and generated photon extraction.

In terms of the fixed parameters, if the misalignment is greater than the nominal value, the clipping loss will go over the threshold and the performance drops off rapidly. If the misalignment is smaller than nominal, no change in performance is seen, because the clipping loss is already negligible. Scattering loss was not involved in the optimisation of geometry, so any deviations of its value from the nominal does not have catastrophic impacts on performance, however the chosen transmission would be suboptimal.

Given these behaviours, it would be reasonable to choose a geometry that is inside the clipping loss boundary, rather than on the very edge as in the case of the optimum solution. This means that the cavity would still be able to perform well if L, D_M or R deviate from their nominal values due to fabrication defects, or the misalignment exceeded the nominal value, but would come at the cost of a slight reduction in performance. The degree to which performance should be compromised for the sake of robustness would depend upon the repeatability of the fabrication procedure.

4.6 Conclusion

This chapter has presented a simple procedure to optimise cavity design for remote entanglement using justifiable assumptions. The principal assumptions are that the fidelity is unaffected by the cavity decay rate, temporal mixing effects can be ignored, and the photons can be extracted with near-ideal efficiency without significantly increasing the attempt cycle time. Under these assumptions, the cavity has no impact on the process fidelity or the attempt rate of the protocol, and therefore the optimum cavity design is that which maximises the photon extraction probability.

For given losses and mode divergence angle, the optimum transmission of the cavity can be directly determined. Therefore, the geometry of the cavity is first optimised, and the transmission subsequently set. Optimising the geometry means balancing improvements in the divergences angle with increasing cavity losses as the configuration becomes more concentric. However, because the clipping loss is a very strong function of geometry, the cavity geometry space can be categorised into two parts, one for which the clipping loss is negligible and thus the cavity loss is fixed, and the other where the clipping loss is unacceptably high. The boundary region is extremely narrow for the microcavity geometries typical to ion-cavity network experiments, and thus no significant benefit could be gained from considering it. Using this insight, the geometry is optimised by maximising the cavity divergence angle over the region of negligible clipping loss. For fixed milling volume, this results in a cavity with the minimum permitted length, and mirrors pushed towards the concentric configuration as far as possible before clipping losses breach the threshold. In practice, it would be reasonable to choose a configuration slightly set back from the clipping loss boundary to make the performance less vulnerable to missetting of geometry.

This procedure is simple to apply and the approach can be modified to accommodate different fabrication constraints. The results suggest that it is possible to obtain good performance for misaligned cavities, as long as that misalignment is considered during the design process. However, underpinning this analysis is the evaluation of clipping loss as a function of misalignment, which utilised a simple model that did not account for the change in the cavity mode that the clipping loss induces. It is therefore important to understand if the clipping loss model is valid, and this will be discussed in the next chapter.

5 Mode Mixing

5.1 Introduction

The theory of the cavity-emitter interaction presented thus far has assumed that the mirrors have 'perfect' spherical shape. This is assumed because the confined modes of such a cavity are the standard propagating Gaussian beams introduced in Chapter 2. This theory works sufficiently well to understand the majority of cavity experiments.

However, strong emitter-cavity coupling, as desired for remote entanglement and other quantum information protocols, requires a small cavity mode volume. This is ensured by fabricating highly curved mirrors and forming a short cavity length. Unfortunately, traditional superpolishing, which can make very spherical cavity mirrors, is not able to produce mirrors of the tight radius of curvature required [184], and therefore other methods must be used. More detail on these approaches can be found in Sec. 2.5.2 of Chapter 2, but for the current discussion, only the shapes of the mirror profiles are relevant.

One of the leading approaches for fabricating micromirrors in recent quantum optics experiments is laser ablation. This method can produce very low roughness mirrors, and has been used to fabricate the mirrors in a series of leading cavity QED experiments [234–236]. However, the Gaussian beam of the ablation laser tends to form a Gaussian, rather than spherical, mirror surface [185]. There has been recent progress in modifying the laser ablation technique to make the mirrors

more spherical [237], but even so the spherical profile does not necessarily extend to the edges of the mirror. Focussed ion beam (FIB) milling is an alternative that offers finer control over the mirror profile [195], but thus far largely been interfaced with solid state qubits [238] using even smaller micromirrors than those suitable for integration with ion traps [196].

The envisaged quasimonolithic cavities with solely axial in-situ adjustment should provide stability and a small footprint at the cost of removing transverse alignment of the mirrors, either in use or at assembly. As discussed in Sec. 2.2.5 of Chapter 2, transverse mirror misalignment tilts the mode angle, causing the mode to intercept the mirror at away from its nominal centre. Of particular concern is whether non-spherical mirrors, which do not have constant curvature away from their optic axis, would be suitable for cavities with residual transverse misalignments as there is no guarantee that the 'clean' cavity physics that holds for aligned mirrors, or misaligned spherical mirrors, still applies. Thus, the need to tolerate mirror misalignment could rule out certain mirror shapes and thus fabrication methods, even if these methods have been used successfully in cavities with negligible transverse misalignments.

This chapter will first discuss existing methods to calculate the modes of optical cavities with non-spherical mirrors and residual transverse misalignment. This will start with a geometrical picture of the cavity mode and phase curvature to provide intuition about the expected solution. The method of mode mixing [233] will then be overviewed, and new techniques to facilitate its use to the present problem will be developed. The new techniques will subsequently be applied to understand the phenomena seen in cavities with Gaussian-shaped mirrors and transverse misalignment, and how this differs from the case where the mirrors are spherical.

Contribution and Acknowledgements

The content of this chapter is the subject of two publications in preparation. The work presented in this chapter is largely the author's own, with all results with the exception of Sec. 5.2 derived from the author's modelling code and mathematical approach. Dr. Peter Horak of the University of Southampton derived the formulae in Sec. 5.2 for the geometric approach to cavities with Gaussian mirrors, and wrote

a mode mixing code that used standard methods to provide a comparison during the development of the techniques discussed in this chapter.

5.2 Geometric Theory of Cavity Modes

The simplest understanding of the supported mode of a cavity can be obtained from the classical ray picture. In this picture, light is trapped in a cavity if a ray is perfectly retroreflected by both surfaces. In other words, the ray will be normal to both mirrors at the point of intersection.

Propagating light does not, however, behave as an infinitely thin ray, and the solutions to the paraxial approximation (Eq. 2.7 of Chapter 2) offer a more accurate description. This level of approximation, termed 'the geometric picture' is to assume that, as for a perfect spherical cavity, the cavity modes are Hermite-Gauss modes. Supported modes will have a phase curvature at the intersection of the classical ray with each mirror matches the local curvature of that mirror at intersection. Thus, a 'geometric' prediction of the cavity mode is made by first finding the central ray, and then calculating the radius of curvature of the each mirror where the mode intersects. The case of the spherical mirror, where the mode axis passes through the centres of the spheres that define the cavity mirrors, has already been detailed in Chapter 2 and used in Chapter 4. This case is simple because the curvature of a spherical mirror is constant across its surface.

However, as discussed, microfabricated cavity mirrors are not necessarily spherical; in particular, for certain fabrication methods, they may be better approximated by a Gaussian function. Gaussian mirrors can be parametrised by a central radius of curvature R and a $1/e$ width w_e, resulting in a mirror profile

$$\begin{aligned} z_G(x,y) &= D\left\{1-\exp\left[-\frac{x^2+y^2}{w_e^2}\right]\right\}, \\ D &= \frac{w_e^2}{2R}, \end{aligned} \tag{5.1}$$

where D is the depth of the mirror defined by the radius of curvature and the width.

While it has been implicitly assumed thus far that the ideal mirror profile is spherical, in the paraxial approximation, the phase front of the beam (and thus the

ideal mirror profile) is parabolic. A spherical profile matches the parabola in the small angle approximation, and it is often a better match for phase fronts than a parabola outside this approximation [239]. However, for mathematical simplicity in this analysis, spherical mirrors will be taken as parabolic, and Gaussian mirrors compared to the parabolic shape. This comparison can be made for a Gaussian mirror using a Taylor expansion

$$D\left\{1-\exp\left[-\frac{x^2+y^2}{w^2}\right]\right\}\approx\frac{x^2+y^2}{2R}\left(1-\frac{x^2+y^2}{2w_e^2}-...\right), \tag{5.2}$$

from which it is seen that the leading order correction to a parabolic curvature is a quartic reduction in curvature. This quartic correction relates to the parabolic term through scale $(r/w_e)^2$, where r is the transverse distance from the centre of the mirror. Therefore, the Gaussian mirror approximates a parabola well for $r \ll w_e$, but significantly deviates for $r \approx w_e$. Thus mirrors with bigger w_e remain better approximations to parabola over a larger region.

To calculate the mode expected for a misaligned Gaussian mirrors, take a coordinate system in which the two identical Gaussian mirrors are placed facing each other separated by L along the z axis with $z = 0$ at the centre. The two mirrors are then displaced by $\pm\Delta x/2$ for the mirror at positive and negative z respectively. The cavity axis is calculated to be normal to the mirror at both intersection points. The intersection point of the mode with the mirror at positive (negative) z is defined as $P_A = (x_a, y_a, z_a)$ ($P_B = (-x_a, -y_a, -z_a)$). If $x_m = x_a - \Delta x/2$ is defined as the distance of point P_A from the centre of the mirror, the solution can be simply calculated to satisfy

$$\frac{\Delta x}{2}=2D\frac{x_m}{w^2}e^{-x_m^2/w^2}\left[\frac{L}{2}+D\left(-1+e^{-x_m^2/w^2}\right)\right]-x_m, \tag{5.3}$$

which can be solved numerically for x_m and then used to calculate the coordinates of P_A. The radius of curvature of the mirror at P_A in x direction is given by

$$\begin{aligned} R_x &= -\frac{(1+z_G(x_m,0)'^2)^{3/2}}{z_G(x_m,0)''}, \\ z_G(x_m,0)' &= -2D\frac{x_m}{w^2}e^{-x_m^2/w^2}, \\ z_G(x_m,0)'' &= -D\frac{2}{w^2}e^{-x_m^2/w^2}+D\left(\frac{2x_m}{w^2}\right)^2 e^{-x_m^2/w^2}, \end{aligned} \tag{5.4}$$

where z'_G and z''_G are first and second derivatives of the mirror profile z_G (Eq. 5.1) with respect to x. The radius of curvature in the y direction is

$$\begin{aligned} R_y &= Re^{x_m^2/w^2}\cos\phi + x_m \sin\phi, \\ \sin\phi &= \frac{x_a}{\sqrt{x_a^2 + z_a^2}}, \end{aligned} \tag{5.5}$$

where ϕ is the angle of the cavity axis with respect to the z axis.

The results of Eqs. 5.3, 5.4 and 5.5 can be used to calculate the expected mode in cavities with Gaussian mirrors according to the geometric picture, which will differ from the mode expected from misaligned spherical mirrors. Firstly, the angle of the mode differs slightly because the normal to the surface is not always directed to the centre of the sphere with the on-axis radius of curvature. Secondly, the curvature of the mirrors varies across the mirror surface and takes different values in the two local principal axes (unlike for a spherical mirror, where the curvature is constant). When these radii of curvature differ, the mode calculated by geometric theory is elliptical; it has different phase curvatures, and thus different central waists in the two transverse directions. Finally, a Gaussian mirror is concave in the centre and is thus confining there. However, as the mirrors are misaligned beyond a critical misalignment Δx_c, the concavity becomes negative (at $x_m = \pm w_e/\sqrt{2}$). for intersection points here, or even further from the centre of the mirror, the mode will not be confined by the mirror, and no prediction for the mode properties can be made from geometric theory. The geometry and results of this section are summarised in Fig. 5.1.

Finally, it is important to note that while the methods discussed here take the example of spherical cap mirrors and Gaussian mirrors, the geometric theory presented can be applied to any smooth mirror shape to provide predictions.

5.3 Mode mixing by Mirror Transformation

In the geometric method (Sec. 5.2), the prediction of the fundamental mode does not account for the mirror shape beyond the local radius of curvature, essentially assuming that the radii of curvature at the intersection points continue to infinity. The deviation of the real mirrors from this assumed parabolic profile mean that true modes of these cavities will differ from those calculated by the geometric method.

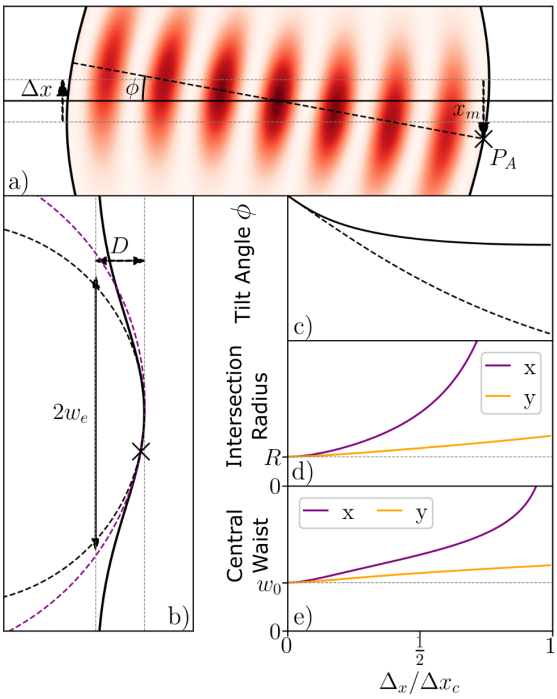


Figure 5.1: Schematic diagram of the Gaussian mirror geometry. a) An example mode in a cavity with Gaussian mirrors misaligned by Δx. The mode intersects the mirror with tilt angle ϕ at P_A , where the mirror has a different radius of curvature to the centre. b) A diagram of the dimensions of the mirror, showing the depth D and 1/e width w_e. Two dotted circles are marked; the one through the centre of the mirror (black dashed) shows the curvature at the centre, and the circle through the marked mode intersection point (purple dashed) shows the larger local radius of curvature at the intersection point. c-e) Mode parameters of a Gaussian cavity for Δx between 0 and Δx_c, the maximum misalignment for a stable mode. c) The mode angle (solid) compared to that of a spherical cavity with the same central curvature (dashed), showing the mode angle deviates less under misalignment for the Gaussian case. d) The curvatures at intersection in the x and y directions, showing both increase from their centre values but to different extents. e) The waists of the geometrically predicted mode, which is elliptical under misalignment.

This can be calculated through mode mixing [233], where the true fundamental mode is a superposition of different orders of transverse mode in the given basis. As a linear superposition of solutions to the paraxial equation, this mode itself satisfies the paraxial equation and is therefore a valid beam. This method has been applied to microcavities with non-spherical mirrors in non-misaligned configurations, where the mixing has been found to induce sporadic, severe drops in cavity finesse at particular resonant configurations [240], but mode mixing can also be harnessed to increase cavity coupling [241–243] or even tailor modes to have desired properties [244].

In this section, first the standard mode mixing theory will be introduced, and then methods to expedite the calculations, particularly for misaligned mirrors, will

be presented. These new methods will form the basis of a numerical investigation into the properties of misaligned cavities with non-spherical mirrors that will be presented later.

5.3.1 Mode Mixing Introduction

The standard mode-mixing formalism of Kleckner 2010 ([233]) calculates cavity modes by expressing the action of a round trip of the cavity as a matrix acting on a set of basis modes. It has been used successfuly to model the physics observed in different cavity experiments [240, 245], and will be outlined here before being developed to suit better the task of modelling misaligned cavities with non-spherical mirrors.

Generically, a paraxial light field may be expressed as a linear superposition of modes $u_s^{(\pm)}$, which are themselves solutions of the paraxial equation, where the label $\pm$ indicates whether the mode is moving towards positive or negative z and s is an index over all the modes in the basis. A mirror is encoded as a matrix whose elements detail how an ingoing mode in the ingoing basis is scattered to an outgoing mode in the outgoing basis. The action of the mirror is to imprint a phase front onto the ingoing mode because the incoming wave will travel further to meet the deeper parts of the mirror, which are generally in the centre, and thus acquires a differential phase across its wavefront. The mirror matrix components at positive z coordinate may be written

$$A_{s,t} = \int_{S_A} u_s^{(-)*} \exp\left(2ikz_A\right) u_t^{(+)} dS, \quad B_{s,t} = \int_{S_B} u_s^{(+)*} \exp\left(2ikz_B\right) u_t^{(-)} dS, \tag{5.6}$$

where k is the wavevector of the light, and S_A (S_B) and z_A (z_B) are the surface and profile of mirror A (B) (with the convention that a positive profile points towards the cavity centre for both mirrors). The integrals occur in a single transverse plane at the depth for which z_A is zero.

The calculations in this paper will use the Hermite-Gauss modes $u_{n_x,n_y}^{(\pm)}$ of classical optics, as introduced in Eq. 2.7 of Chapter 2. This basis is complete and orthonormal for each transverse plane. The basis states are also eigenmodes of propagation in the sense that a linear superposition of such modes propagating from one transverse plane at z_P to another at $z_P + d$ retains its coefficients (besides picking up a phasor

$e^{\mp ikd}$), as the evolution of the field upon propagation is encoded in the z-dependence of the basis functions themselves.

The round-trip matrix can be calculated from the two mirror matrices, taking advantage of the simplicity with which propagation can be expressed in the Hermite-Gauss basis:

$$M = BAe^{2ikL}, \tag{5.7}$$

with L as the cavity length. A generic mode supported by the cavity $|\Psi_i\rangle$ will be an eigenmode of the round-trip matrix M, with corresponding eigenvalue γ_i from the eigenmode equation

$$M\,|\Psi_i\rangle = \gamma_i\,|\Psi_i\rangle\,. \tag{5.8}$$

Both phase and amplitude information are contained within γ_i. The phase is the round-trip phase (modulo 2π) accrued by $|\Psi_i\rangle$. The phase is zero if the mode is on resonance. For the purposes of this chapter, the amplitude is more interesting as it gives the round-trip loss $1 - |\gamma_i|^2$ of the eigenmode within the cavity. A schematic diagram of the mode mixing method is shown in Fig. 5.2.

Using the mode mixing formalism, the eigenmodes of deformed cavities can be calculated. This process involves calculating many matrix elements through the integration of Eq. 5.6. A sensible approach to calculating the eigenmodes of cavities with non-spherical mirrors under misalignment would therefore appear to be to take the misaligned cavity, calculate the geometrically-expected mode using the theory of Sec. 5.2, and use this mode to define the basis in which to conduct the mode mixing calculation. This approach has the advantage that the basis of the mode-mixing calculation is always chosen 'sensibly' according to geometric considerations, and therefore it should be possible faithfully capture the resulting eigenmode with a relatively limited basis size.

However, performing calculations this way means that the basis is different for every misalignment and cavity length. Therefore, all of the matrix elements are calculated for each cavity configuration separately. Calculations performed this way are computationally arduous, making it hard to generate data with sufficient

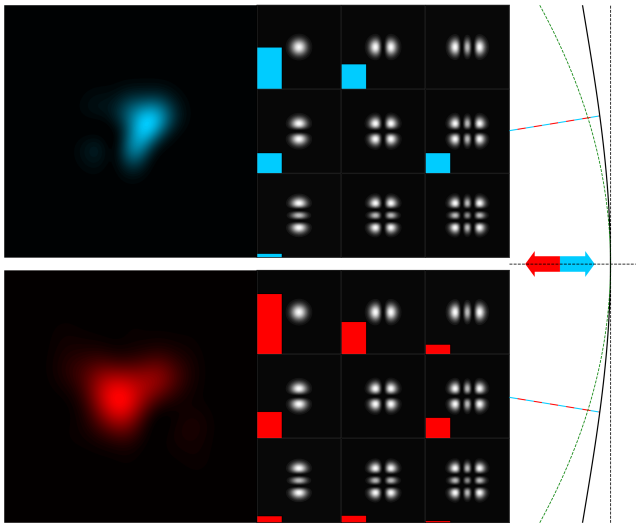


Figure 5.2: Diagram of mode mixing from a mirror. The incoming beam (blue, top) is formed from a linear superposition of basis states (here the basis states for n_x and n_y up to 2 are shown). The waist $w(z)$ of the basis as it approaches the mirror is marked with the blue-red dashed line. The amplitudes of the coefficients are depicted by the height of the bars adjacent to the corresponding basis states. The mirror imprints a phase front on the input beam reverses its direction. The mirror (black) is shown as Gaussian-shaped, with the parabolic component (green, dashed) marked. The output beam (red, bottom) is again expressed as a linear superposition of basis states (now in the outgoing basis). The intensity patterns of the modes are plotted in the transverse plane at which the basis states have their central waist. The action of the mirror can be summarised by the transfer amplitudes between the complete input basis and the complete output basis. Cavity eigenmodes are particular linear superpositions of states that are preserved under reflection from both mirrors.

resolution to understand the physics involved in these systems, particularly resonant interference effects.

The remaining parts of this section will detail numerical techniques that allow for the modelling of cavities with misaligned non-spherical mirrors to be performed more quickly.

5.3.2 Replacing Coordinates with Operators

The long-appreciated similarities between the Hermite-Gauss modes in any transverse plane and the simple harmonic oscillator wavefunctions [246, 247] inspire the writing of x or $\partial/\partial x$ in terms the ladder operators a_x and a_y, where a_x (a_y) reduces the n_x (n_y) index of the Hermite Gauss mode by 1. The operators for x and $\partial/\partial x$ are

$$
\begin{aligned}
x &= U_G^\dagger \frac{1}{2} w (a_x + a_x^\dagger) U_G, \\
\frac{\partial}{\partial x} &= \frac{1}{w_0}(a_x - a_x^\dagger), \\
(U_G)_{n'_x,n'_y,n_x,n_y} &= e^{i\Psi_G(n_x+n_y+1)} \delta_{n'_x,n_x} \delta_{n'_y,n_y},
\end{aligned}
\tag{5.9}
$$

where n_x and n_y (n'_x and n'_y) are the x and y indices of the input (output) modes of the matrix respectively. While the $\partial/\partial x$ operator is constant across all transverse planes, the matrix elements of x depend upon the z coordinate. The equivalent relations for y and $\partial/\partial y$ also hold, but with y-axis ladder operators replacing the x-axis operators. The derivations are given in App. C.

The action of the mirror is to reflect and imprint a phase front onto the ingoing mode (as expressed in Eq. 5.6). It is conceptually simpler to consider this sequentially: First imprinting the phase in the input basis, and then transforming the input basis to output basis with the operator

$$
U^{+\to-} = U_G^2 \exp\left(-2ik\frac{x^2+y^2}{2R}\right). \tag{5.10}
$$

The mirror matrix A can then be expressed

$$
A = U_G^2 \exp\left(-2ik\Delta_A\right), \tag{5.11}
$$

where $\Delta_A = z_A - (x^2+y^2)/2R$ is the deviation from the ideal parabolic surface of mirror A towards the centre of the cavity. If Δ_A (and the corresponding Δ_B) can expressed as a matrix without taking integrals, the mirror matrix A (B) can also be obtained without integrals, as discussed later in Sec. 5.3.5.

5.3.3 Calculating Polynomial Mirror Surface Profiles

The simplest Δ to evaluate is the case where Δ can be written as a power series in x and y, for which it is only necessary to calculate matrices of the arbitrary powers of x and y and sum each polynomial term with the appropriate coefficient. The main use of this method in the scope of this work was to test the numerics with a parabolic distortion (an adjusted radius of curvature) where the expected results are well-known.

5.3.4 Calculating the Gaussian Surface Profile

The Gaussian surface profile can also be expressed in the Hermite-Gauss basis without any integrals, but doing so requires a different approach, inspired by the appendix of [248] and detailed in App. D. The matrix elements of a unit Gaussian profile with e waist w_e in one dimension can then be written

$$\exp\left[-\frac{x^2}{w_e^2}\right]_{m',m} =$$

$$U_G^{\dagger}\left(1-\chi\right)^{-\left(\frac{m'+m+1}{2}\right)}\left(\frac{\chi}{2}\right)^{\frac{m'-m}{2}}\sqrt{m'!m!}\sum_{k=0}^{\left[\frac{m}{2}\right]}\frac{\left(\frac{\chi^2}{4}\right)^k}{\left(\frac{m'-m}{2}+k\right)!k!\left(m-2k\right)!}U_G,$$

$$\chi = -\frac{1}{2}\frac{w\left(z\right)^2}{w_e^2},$$

where m (m') is the index of the ingoing (outgoing) mode, $(m'-m)/2$ is an integer and $m' \geq m$. If $(m'-m)/2$ is not an integer, the matrix element is zero. If $m > m'$, the symmetry $\exp[-x^2/w_e^2]_{m',m} = \exp[-x^2/w_e^2]_{m,m'}$ may be used.

The true mirror profile is two dimensional, but as the Gaussian is an independent product of x and y, this can be found from the one-dimensional matrices by a simple tensor product.

After expressing the Gaussian itself, the deviation Δ from the parabolic surface can be written

$$\begin{aligned}\Delta &= D\left(1-\exp\left(\frac{x^2+y^2}{w_e^2}\right)\right)-\frac{x^2+y^2}{2R},\\ D &= \frac{w_e^2}{2R},\end{aligned} \tag{5.12}$$

where D is the depth of the Gaussian mirror at its centre and R is the radius of curvature at the centre of the Gaussian. Setting R in this manner effectively determines the basis of the calculation, which is chosen so that the lowest order mode of the basis is the eigenmode of a perfect spherical cavity with the same radius of curvature as the Gaussian mirror at the centre. This was done so that the mirror matrix tends to diagonal form in the limit that the Gaussian width grows large.

5.3.5 Taking the Exponent of the Surface Profile

With the surface profile deviation Δ expressed as a matrix, the surface profile phase matrix $\exp(-2ik\Delta)$ can be calculated. Initially, one might consider taking the matrix exponential of the surface profile deviation matrix to find the surface profile phase matrix, but this method cannot model losses. As Δ is a Hermitian matrix, the matrix exponent is unitary, and therefore every eigenvalue of the mirror matrix would have a modulus of precisely 1, meaning that the mirror is lossless. No matter how large a basis is chosen, Δ never models processes representing transfer from inside to outside the basis, and therefore there is no mechanism for power to leave the cavity.

To take the exponential in a way that can model losses, the following procedure is used. The matrix Δ is first evaluated in a larger basis than the intended simulation basis and then truncated to the size of the simulation according to specific rules. Each element of the Δ matrix represents a transfer from an input state to an output state. If the input state lies within the simulation basis, but the output state is outside, this matrix element encodes the loss that must be modelled. Therefore, for each input state, the sum over all the magnitudes of transfers to states outside the basis is calculated, evaluating the 'rate' at which field leaks from the input state to outside the simulation basis. This summed rate is then added as a negative imaginary number onto the diagonal element of the input state. This means that, when the matrix exponential is then taken in the simulation basis, the diagonal imaginary component becomes a loss from the input state, rather than a transfer. This procedure is conceptually very similar to simulating quantum systems with a non-Hermitian Hamiltonian, a common method in cavity QED [139], discussed in Chapter 2 and used in Chapter 3 and later in Chapter 6.

Mathematically, this procedure can be expressed in the following way; for a larger basis containing n_x and n_y up to $n_x^{\mathrm{max,H}}$ and $n_y^{\mathrm{max,H}}$ respectively, and the smaller simulation basis up to $n_x^{\mathrm{max,NH}}$ and $n_y^{\mathrm{max,NH}}$ respectively, the non-Hermitian Δ matrix

(that can then be used to take the matrix exponential) is written

$$\Delta_{n'_x,n'_y,n_x,n_y} = \Delta^{\mathrm{H}}_{n'_x,n'_y,n_x,n_y}, \quad n'_x, n_x \leq n_x^{\mathrm{max,H}},\ n'_y, n_y \leq n_y^{\mathrm{max,H}},\ \delta_{n'_x,n_x}\delta_{n'_y,n_y} = 0,$$

$$\Delta_{n_x,n_y,n_x,n_y} = \Delta^{\mathrm{H}}_{n_x,n_y,n_x,n_y} + i \sum_{n'_x=n_x^{\mathrm{max,NH}}+1}^{n'_x=n_x^{\mathrm{max,H}}} \sum_{n'_y=n_y^{\mathrm{max,NH}}+1}^{n'_y=n_y^{\mathrm{max,H}}} |\Delta^{\mathrm{H}}_{n'_x,n'_y,n_x,n_y}|, \tag{5.13}$$

where Δ^{H} is the Hermitian surface profile deviation matrix evaluated on the larger basis. The matrix exponential of the non-Hermitian Δ is then taken to find surface profile phase matrix $\exp(-2ik\Delta)$.

While this process is not mathematically identical to finding the true matrix $\exp(-2ik\Delta)$, it is valid in the limit that any amplitude that spills outside the simulation basis has negligible coupling back into the basis, which is expected for states outside the confining region of the mirror. The size of the larger basis must also be chosen so that all significant couplings from inside the simulation basis to outside are correctly modelled. In practice, this method of taking the exponential produces almost identical loss results to numerical integration for most cavity configurations, as shown later in Sec. 5.3.8.

5.3.6 Translating the Mirror

Now that the surface profile phase matrix is calculated, it is possible to write both mirror matrices and thus obtain the eigenmodes for a cavity. As discussed in Sec 5.3.1, it is possible to repeat calculations of mirror matrices for every misalignment, but this approach is computationally intensive. A better method, discussed in this section, is to translate the mirror operator after it has been evaluated in the aligned configuration to find the mirror matrix in the misaligned configuration without needed to calculate any of its elements again.

As depicted in Fig. 5.3 the action of a mirror translated by δ_x in the x direction on a given basis is equivalent to the action of the untranslated mirror on the basis states shifted in the opposite direction (i.e. by $-\delta_x$), as these are the same physical situation but for different choice of origin. This equivalence means that calculating the action of the translated mirror in the basis (i.e. representing the translated mirror as a matrix) is the same as translating the input and output bases in the compensating direction.

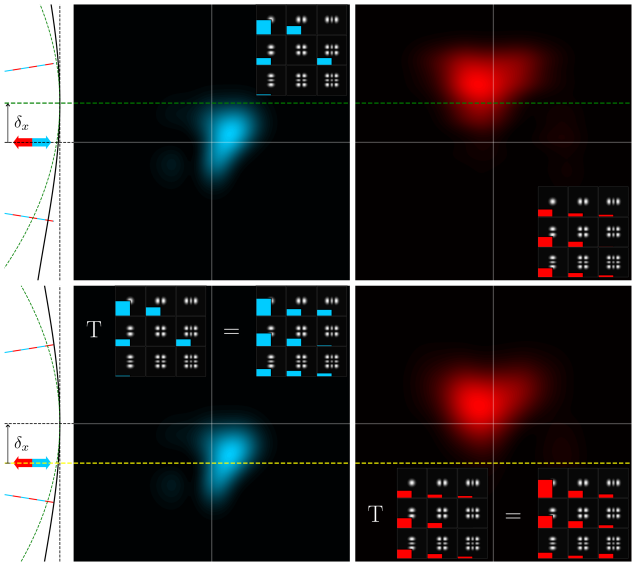


Figure 5.3: Diagram to explain the method for obtaining the mirror matrix of a mirror translated in the x direction by δ_x. Top: The input mode (blue, left) hits a translated mirror and scatters to the output mode (red, right). Both input and output modes are expressed as a linear superposition of basis states (here shown up to n_x and n_y of 2). The waist $w(z)$ of the basis as it approaches the mirror is marked with the blue-red dashed line. Bottom: The input mode and the mirror are now translated so that the mirror is in the nominal 'aligned' configuration. The input mode scatters to the same output mode as before, once the translation of the origin is accounted for. The equivalence of the physical process in the two scenarios means that the action of the translated mirror can be derived from the untranslated mirror by transforming the input and output bases to compensate for the mirror translation. In all cases, the intensity patterns are plotted in the transverse plane at which the basis states have their central waist.

The one-dimensional operator that translates the input and output bases is

$$
\begin{aligned}
T(\delta) &= \exp\left(\delta \frac{\partial}{\partial x}\right), \\
T_x(\delta)_{m',m} &= \sqrt{\frac{m!}{m'!}} \alpha^{m'-m} e^{-\frac{\alpha^2}{2}} L_n^{m'-m}\left(\alpha^2\right), \; m' \geq m, \\
T_x(\delta)_{m',m} &= \sqrt{\frac{m'!}{m!}} \left(-\alpha^{m-m'}\right) e^{-\frac{\alpha^2}{2}} L_m^{m-m'}\left(\alpha^2\right), \; m >' m, \\
\alpha &= \frac{\delta}{w_0},
\end{aligned}
\tag{5.14}
$$

where m (m') is the input (output) index in the one-dimensional basis and δ the translation in that dimension due to the operator. This is identical to the displacement operator of the simple harmonic oscillator [249], owing to the close similarity between the simple-harmonic and Hermite-Gauss bases. As this translation operator has

the same elements in the input and output bases, translating a mirror with matrix C by δ_x can be achieved through

$$C \rightarrow T_x^{\dagger}\left(-\delta_x\right) C T_x\left(-\delta_x\right). \tag{5.15}$$

Where T_x is formed from the tensor product of the one-dimensional translation in the x direction and the identity in the y direction. These translations can be especially convenient when scanning the misalignment of the mirrors, as the translation matrix can be calculated once for a single increment, and then successively applied to generate all of the mirror matrices. This method of translation radically speeds up the process of calculating eigenmodes of misaligned cavities, allowing much more data to be generated for the same computational exertion.

5.3.7 Spherical Cap Mirror

The finite diameter spherical mirrors modelled in Chapter 4, hereon termed 'spherical cap' mirrors, can also be studied in the mode mixing formalism. In these investigations, such a mirror is modelled as parabolic up to a certain transverse radius, and then completely non-confining beyond that. The mirror matrix would be diagonal were the mirror to extend to infinity, but the finite size introduces mode mixing. This 'hard edge' clipping is not calculated through operator methods, but instead the mirror matrices are formed by numerical integration on axis (where multiple symmetries allow for efficient evaluation), and are then translated as described in Sec. 5.3.6 to generate mirror matrices in a misaligned configuration.

5.3.8 Summary

The theory introduced in this section can be applied to a variety of mode-mixing scenarios to speed up the calculations. Firstly, for particular mirror shapes (where the deviation from the ideal parabolic profile can be expressed as a sum of polynomials in the transverse coordinates, or as a Gaussian function), the mode mixing matrix can be calculated through a combination of analytical results and a matrix exponential-removing the need for any integral to be taken. Secondly, once the mirror matrix has been obtained (whether through analytical methods or numerical integrals), the

mirror can be translated using operators (which also do not require integrals to calculate). This allows for investigation of how the specific mirror will affect cavity mode structure in the case of transverse misalignment. This is particularly pertinent for the cavities envisaged in this work, where the construction means that residual transverse misalignment cannot be eliminated. Finally, it should be noted that other transformations of mirrors may be possible, but here only transverse translation is performed because it produces relevant results for our application.

The more efficient computational procedure developed in this section can be compared to mode mixing calculations for Gaussian mirrors in the literature to check that it reproduces known results. Such a comparison is shown in Fig. 5.4, where the round-trip loss is calculated as a function of length for three different Gaussian-shaped mirrors with zero misalignment, and compared to published data. The results agree almost exactly, with the exception of the most highly concentric configurations. The methods presented for modelling translation allow for data to be generated as a function of misalignment as well, showing how the high loss peaks for aligned configurations evolve as the mirror is misaligned. Interestingly, there are some prominent features seen in the misalignment-length plots that do not figure for non-misaligned configurations. This implies that additional physics is involved for misaligned cavities that is not seen for perfectly aligned cavities, suggesting that it is important to model finite misalignment in the case that it is present, which applies to our proposed system and potentially to others owing to imperfect engineering tolerances. The origin of these extra features will be discussed in Sec. 5.4.

5.4 Results

5.4.1 Simulation Method

Procedure

The methods discussed in Sec. 5.2 and 5.3 can now be used to investigate the behaviour of cavities with both Gaussian and spherical cap mirrors under transverse misalignment. The comparison of the behaviour of these mirror shape types is made by taking two identical concave mirrors of nominal radius R (either Gaussian or spherical cap shaped), and calculating the cavity mode formed as the cavity length

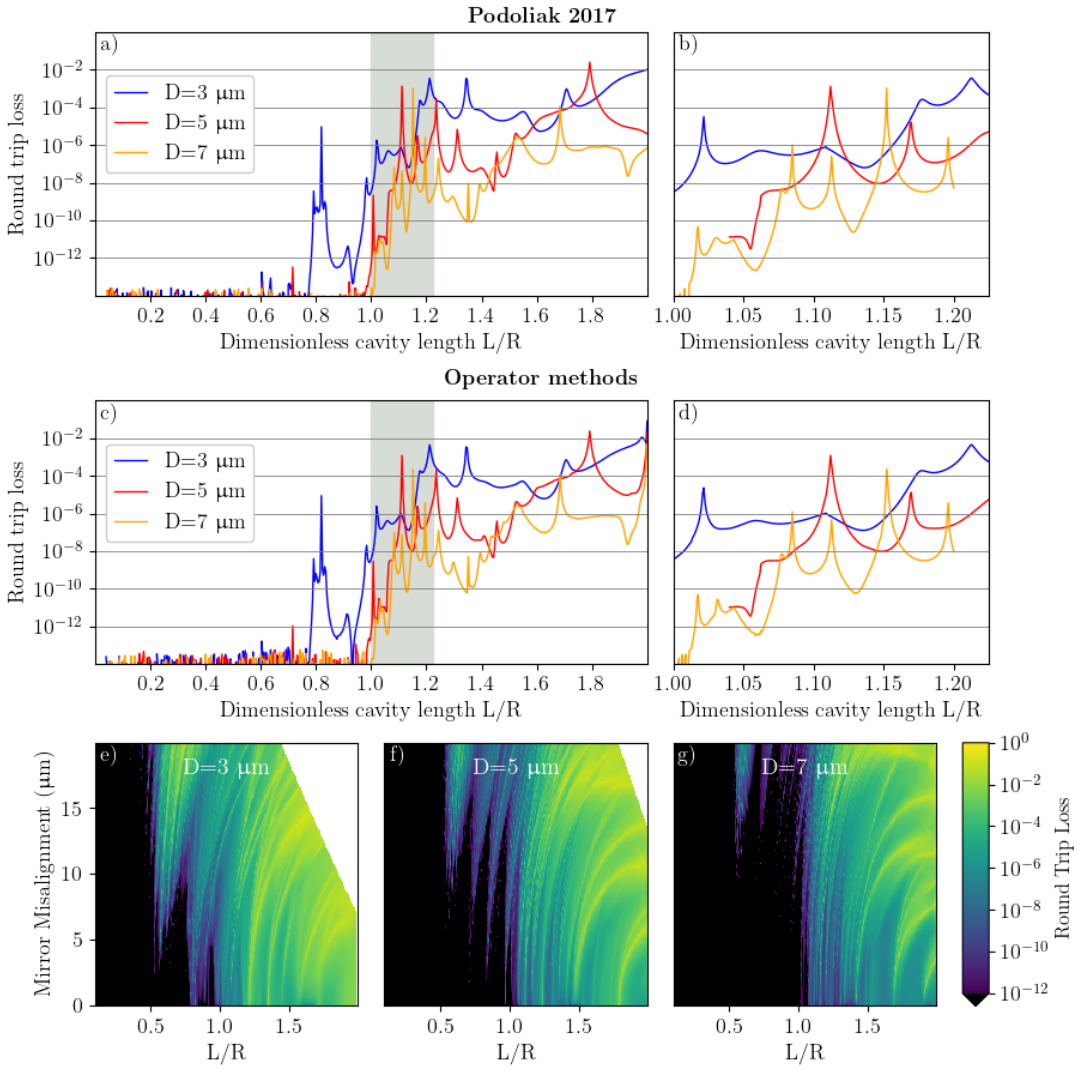


Figure 5.4: Comparison of round-trip loss data generated with standard methods and with the methods described in this section. All mirrors have a central radius of curvature of $R = 500$ μm, and the interrogation wavelength is 866 nm. a) Round-trip loss data for cavities with Gaussian mirrors of three depths as a function of cavity length over the whole stability region and b) in a small region of lengths to the concentric side of the confocal length, with data taken from [241]. The basis used was a Laguerre-Gauss basis up to $n = 30$. c) and d) Equivalent round-trip loss data as generated through the methods detailed in this section, using a basis of the first 30 even states in both Cartesian directions. As the basis types used are different, the methods operate over a different set of states, but they should agree up to convergence effects. The agreement is very good in the enlarged region (b and d), and good throughout the stability region with the exception of very concentric configurations where the methods developed here predict higher losses. e) f) and g) Round-trip loss as a function of length and misalignment for the three mirrors using the methods developed in this section. Configurations where the region inside one waist of the expected mode (as predicted by the theory of Sec. 5.2) would not be fully enclosed within the concave region of the Gaussian mirror are left white. Below round-trip losses of 10^{-12}, numerical noise becomes significant, and thus these results are left black. Certain features (such as the high loss bands appearing to originate from small misalignment at $L/R = 1.5$ μm for $D = 5$ μm) can be seen at finite misalignment, but not on a) or c).

and transverse misalignment of the mirrors is varied. At each length, the calculation begins by placing the mirrors on axis (in the basis selected by central radius R and the given cavity length). These mirrors are then successively misaligned, with the cavity mode calculated at each step. Once the misalignment scan is complete, the next length is chosen and the calculation repeated.

Choice of mode of interest

The mode mixing method produces a set of cavity eigenmodes $\{|\Psi_i\rangle\}$ and corresponding eigenvalues $\{\gamma_i\}$, and any of these modes could be chosen as the mode of interest. For the purposes of coupling to an emitter, it is ideal to have a mode that is both low loss to increase the cooperativity and of low transverse extent so that a relatively point-like emitter can couple to the mode strongly.

One option is that the mode of interest is chosen to be the mode with the lowest loss. Often, this mode resembles strongly the expected fundamental, but there are many configurations where the lowest loss mode has exotic spatial structure that would not be practical to couple to an emitter, or to a fibre to collect the generated photons. Generally, the higher the misalignment, the more likely the lowest loss mode will be very different from the expected mode, but this can also occur with no misalignment at all. An example of a mode that is low loss, but is not a good option for coupling to an emitter is shown in Fig. 5.5.

A more robust procedure, that will be used for this investigation, is to choose the mode of interest to be the one with the highest overlap with the fundamental mode expected from geometric theory. This procedure tends to choose modes that propagate along the expected axis and have simple transverse structure. The geometrically expected basis states will be denoted $\left\{\left|\Psi^G_{n_x,n_y}\right\rangle\right\}$, and therefore the geometrically expected mode $\left|\Psi^G_{0,0}\right\rangle$. The eigenmode chosen is the $|\Psi_i\rangle$ that maximises $|\left\langle\Psi_i\middle|\Psi^G_{0,0}\right\rangle|^2$. This chosen eigenmode is denoted $|\Psi_F\rangle$, with subscript F denoting that this mode is the eigenmode most similar to the geometrically-expected fundamental mode.

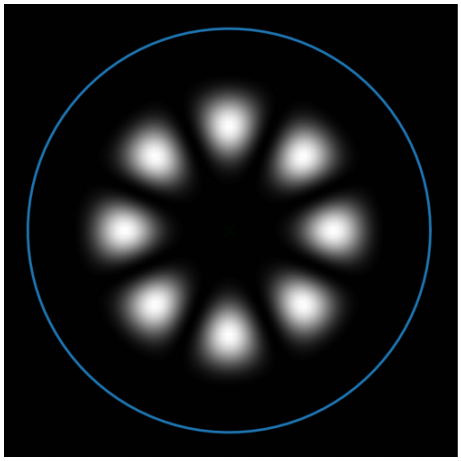


Figure 5.5: The intensity profile of the smallest loss cavity eigenmode in the mirror plane for the case of the $2w_e = 75$ μm Gaussian mirrors of central radius of curvature 200 μm and cavity length of 238 μm with no transverse mirror misalignment interrogated 1033 nm. The radius beyond which the mirror is not concave is marked by the circle. For this case, and for many other configurations explored, the lowest loss mode differs significantly in structure from a typical fundamental Gaussian beam.

Understanding the composition of the eigenmodes

The method presented calculates the cavity eigenmodes in a basis that lies along the nominal cavity axis. However, the mode of a misaligned cavity propagates at an angle with respect to the nominal cavity axis. In the calculation basis, this change of angle manifests as a complicated superposition of modes, which can obfuscate the physics as structurally simple but rotated modes (for example a Hermite Gauss (1,0) propagating along the expected axis) cannot be identified on inspection of the coefficients. This complicates efforts to distinguish between the mode rotation anticipated by geometric theory and the transverse mixing effects brought specifically by mode mixing. To overcome this, after calculation, the mode is transformed, using the methods of App. E, into a basis with the same waist as the nominal basis, but co-propagating with the mode, a basis that will henceforth be known as the 'co-propagating basis' and denoted $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$. This is done to remove the effect of mode pointing direction on the mode coefficients so that the physics of mode-mixing resonances can be seen more clearly.

Note that the co-propagating basis is not the basis of the expected mode, but instead the basis with the same waist as the initial calculation, rotated to match the axis of $|\Psi_F\rangle$. This is because transforming $|\Psi_F\rangle$ to the expected basis would leave a

residual angle if $|\Psi_F\rangle$ did not have quite the same angle as the expected mode $\left|\Psi^G_{0,0}\right\rangle$, which occurs in mode mixing because the mirror is treated as a phase surface rather than a geometric object (this point is discussed more in Sec. 5.4.2). Rotating the basis to match the angle of $|\Psi_F\rangle$ is a guaranteed way to completely remove effects of mode pointing from the mode coefficients. Note that when misalignment is zero, the basis of the expected mode $\left\{\left|\Psi^G_{n_x,n_y}\right\rangle\right\}$ does indeed match the co-propagating basis $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$.

Simulation options

To use the mode mixing method, the finite basis has to be selected. For the spherical cap simulations, the basis used was a tensor product of the first 50 n_x states and the first 25 even n_y states. For the Gaussian mirror, this was 100 n_x states and 50 even n_y states. To take the lossy exponential of Sec. 5.3.5, the Gaussian profile was initially expressed in a basis 15 states larger in each Cartesian direction, before being truncated to calculate the losses. These basis sizes were chosen as it was found that the calculations converged, with no significant visual difference to the data presented upon increasing the basis sizes. For simulations involving non-misaligned mirrors, it is possible to greatly reduce the size of the basis by using Laguerre-Gauss modes, realising that the mixed fundamental will be a superposition of only $m = 0$ states (where m is the azimuthal index of the Laguerre-Gauss basis), but once the rotational symmetry is broken by misalignment, a large state space is needed, which is where the translation methods developed in this chapter become very valuable.

5.4.2 Spherical cap Mirror

Overall summary

The behaviour of cavities with spherical cap mirrors is much simpler than those with Gaussian shaped mirrors because the radius of curvature of the mirror remains constant across its surface. For this reason spherical cap cavities will be discussed first and the concepts then applied to the Gaussian mirror case, but as many microcavity mirrors are indeed highly spherical, this is a useful case in its own right.

Given that spherical cap mirrors are perfectly spherical inside the clipping diameter of the mirror, a reasonable expectation is that the modes in cavities with spherical cap

mirrors will follow the results of the geometrical theory until the mirror is misaligned sufficiently for the geometrically-expected mode to begin clipping the edge of the cap.

To test this intuition, the mode of interest $|\Psi_F\rangle$ for cavities with mirrors of radius of curvature 200 μm and three different diameters was calculated for a scan of cavity length and mirror misalignment. For each cavity configuration, the round-trip loss was calculated, and is displayed in Fig. 5.6. Qualitatively, the loss behaves similarly to that expected from a simple clipping of the geometrically expected mode, with the loss increasing as the mode approaches the edge of the mirror.

The feature of the loss plots that cannot be captured in the simple geometric picture is the appearance of low loss bands at specific lengths. These bands arise due to transverse mode mixing, where $|\Psi_F\rangle$ becomes a weighted superposition of states of $\left\{\left|\Psi^G_{n_x,n_y}\right\rangle\right\}$, rather than just $\left|\Psi^G_{0,0}\right\rangle$ itself. These mode mixing effects tend to be strongest when different transverse modes have the same resonant frequency (i.e. acquire the same round-trip phase). For perfect spherical cavities, these coincidences occur at lengths where the transverse mode splitting of the cavity is a rational fraction of the free spectral range. As discussed in Sec. 2.2.4 of Chapter 2 (but repeated here for clarity), these lengths are

$$L_{p,q} = 2R\left[\frac{\tan^2\left(\frac{\pi q}{2p}\right)}{1+\tan^2\left(\frac{\pi q}{2p}\right)}\right], \tag{5.16}$$

for integer p and q, where q/p represents the ratio of the transverse mode splitting to the free spectral range. As seen in Fig. 5.6, the low-loss bands can all be attributed to a specific set of transverse resonances labelled by q/p, suggesting that transverse mode mixing occurs at these lengths and that it acts to reduce the losses. However, the impact of these resonances on the mode shape is minimal and the overlap of $|\Psi_F\rangle$ with the expected $\left|\Psi^G_{0,0}\right\rangle$ is close to unity everywhere within the bounds of the mirror. This means that the low loss bands would be advantageous for coupling to an emitter as the cavity loss is reduced (improving the cooperativity), but the mode shape is not disturbed significantly.

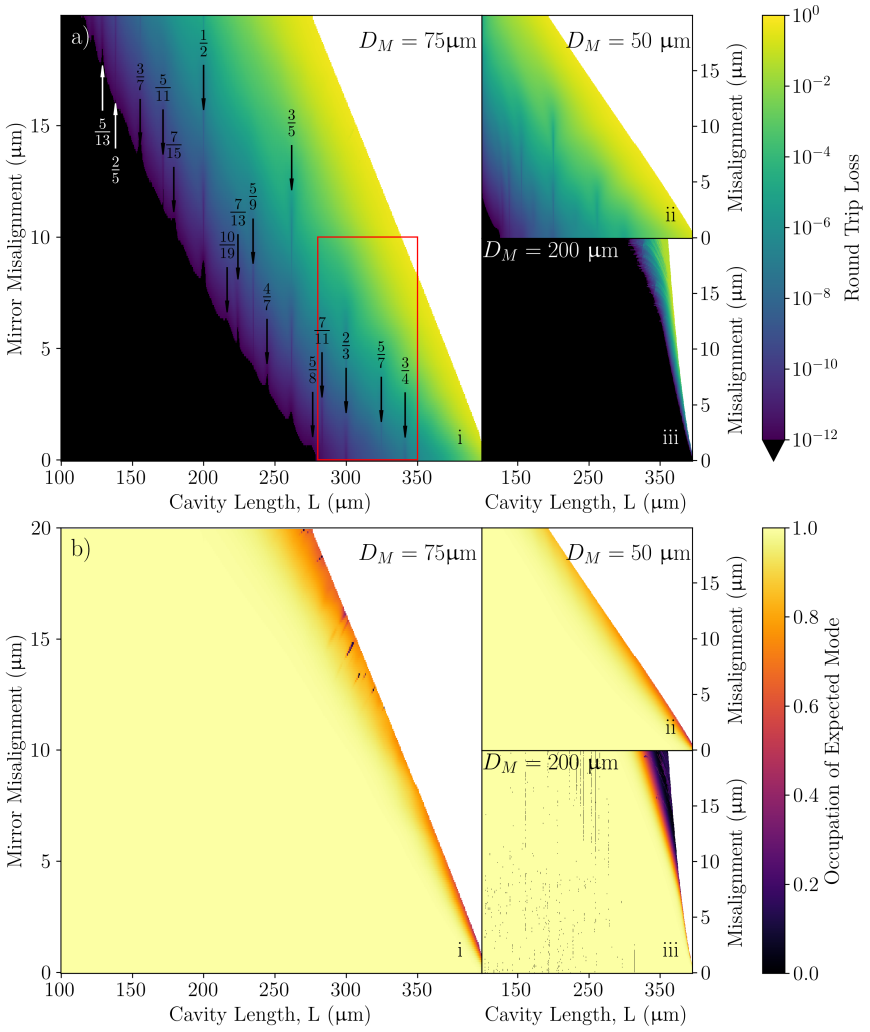


Figure 5.6: Round-trip loss and occupation of expected mode data for cavities with spherical cap mirrors as a function of cavity length and mirror misalignment. Each pair of mirrors forming the cavities has a central radius of curvature of 200 μm, with data shown for three different diameters D_M of the mirrors. The cavities are analysed with a wavelength of 1033 nm. Data is not calculated for translation-length combinations where the region inside one waist of the expected mode $\left|\Psi^G_{0,0}\right\rangle$ would not be fully enclosed within the curved region of the spherical cap, with this region left white. a) The loss dependence on cavity configuration. Losses below 10^{-12} are not shown, as below this level numerical noise begins to become significant. The losses generally obey a simple trend where the loss increases as the cavity mode impinges on the flats of the mirror (occurring as the data nears the white region). In addition to the broad trend, there are sharp dips in the loss at particular lengths; in ai), these features are labelled by the integer ratio q/p of transverse mode splitting to free spectral range to which the length corresponds. The red rectangle indicates the region taken for further analysis in Fig. 5.8. b) The occupation of the $|\Psi_F\rangle$ in $\left|\Psi^G_{0,0}\right\rangle$. Generally, while the mode is well confined within the mirror, the eigenmode is very similar to the expected mode, and the transverse resonances have no visible impact on the overlap.

Comparison with classical clipping

In Sec. 4.3.2 of Chapter 4 the classical clipping model for spherical cavities was presented. This method determines the round-trip loss on the assumption that the mode shape is unaffected as the mode clips. It is reasonable to expect that the losses predicted by the clipping model and by mode mixing should be of the same order of magnitude while the mode does not clip severely on the mirror as the mode shape should only suffer small perturbations as a result of clipping. The exact ratio of clipping losses to mode mixing losses is less clear because, outside of resonant effects, the losses and the disruption to the mode intensity distribution that could change those losses should be of the same order of magnitude as they both derive from the mode spilling outside the finite diameter of the mirror.

Losses calculated through the clipping and mode mixing methods are compared in Fig. 5.7. The two methods do indeed produce losses of the same order of magnitude with the exception of the low loss resonant bands. Generally, the clipping loss method underestimates the cavity loss, although the scale of the underestimate remains within an order of magnitude throughout. The biggest differences between the methods occur at the transverse resonances, where the clipping loss model overestimates the loss because the loss is reduced by transverse mode mixing, which the clipping model cannot invoke. For particular configurations the difference can surpass a factor of 10.

Specific example cases

More insight into the physics of cavities with spherical cap mirrors can be gained from studying the composition of $|\Psi_F\rangle$ at particular cavity configurations, as in Fig. 5.8. The first point point of interest studies $|\Psi_F\rangle$ as it begins to clip on the edge of the mirror. Here, the losses are high and the cavity mode is visibly non-Gaussian, with the mode occupations showing a broad cascade of population to higher order modes, which would eventually lead to loss as population leaves the finite basis. The other two configurations of interest study the mode coefficients around the low loss bands. At lengths not exactly on the transverse resonance, $|\Psi_F\rangle$ overlaps almost perfectly with $\left|\Psi^C_{0,0}\right\rangle$, confirming that for the spherical cap mirrors there is minimal transverse mode mixing away from length resonances and mode clipping. However,

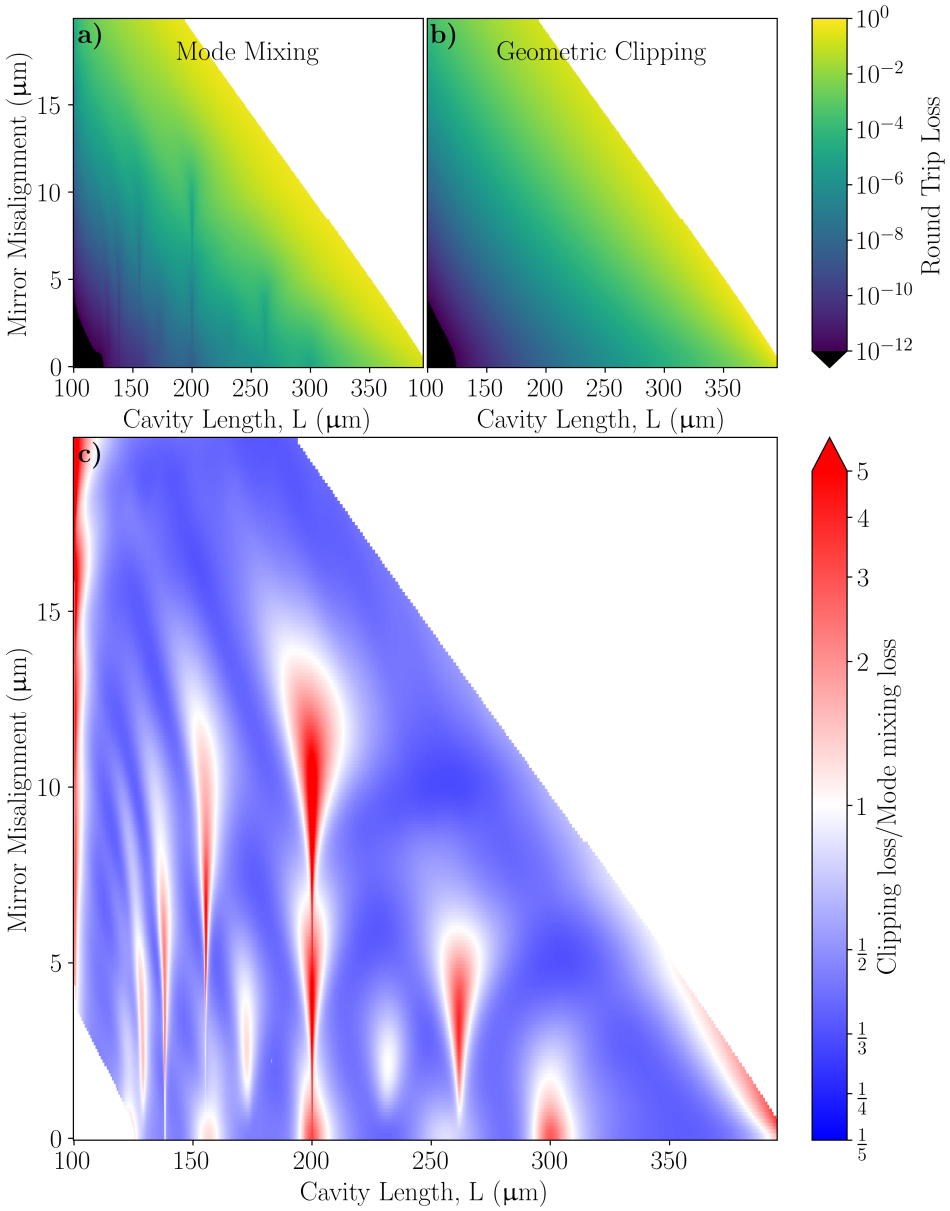


Figure 5.7: Comparison of the calculated clipping loss of cavities with spherical cap mirrors of diameter $D_M = 50$ μm using both mode mixing and the simpler classical clipping loss method. a) The round-trip loss as a function of cavity length and misalignment for mode mixing calculations and b) with classical clipping calculations. c) A comparison of the two methods, showing the ratio of the classical clipping loss to the mode mixing loss on a log scale. Red indicates that clipping loss exceeds that calculated by mode mixing, blue indicates the opposite, and white that the methods agree. Data is not shown for the case where the loss determined by either method is below 10^{-12}, as these results are vulnerable to numerical noise. For the majority of configurations, the two methods of evaluating the loss agree within an order of magnitude. The clipping method generally underestimates the loss, but can overestimate strongly at transverse resonances

at the resonant length, transverse mode mixing is seen, and the transverse modes that hybridise with the expected mode can be predicted from the resonance label q/p and symmetry considerations. First, the resonance label (in the case studied $q/p = 2/3$) sets the higher order modes that are resonant with the expected mode. These are the modes with n_x and n_y indices summing to multiples of p, which is 3 in the case examined. Secondly, symmetry constrains that, at zero misalignment (which is the case presented), only modes with even n_x and n_y indices have both the x and y parity required to overlap with $\left|\Psi^C_{0,0}\right\rangle$. Therefore, at zero misalignment and at the $q/p = 2/3$ resonant length, it can be expected that $\left|\Psi^C_{0,0}\right\rangle$ mixes with higher order modes where n_x and n_y indices sum to multiples of 6 with both n_x and n_y even, which is indeed borne out in the mode occupation patterns. The effect of this mode hybridisation is to make the mode physically more compact on the mirror (as can be confirmed numerically). Thus resonant features reduce clipping loss.

Shape of Resonant Features

The resonant features seen in the investigation of the spherical cap mirror exemplify a limitation to the mode mixing method, where the mirror is encoded as a 'phase surface' rather than a geometrical object. To understand the distinction, look back at Eq. 5.6, where the matrix of mirror A is calculated from its profile z_A. If the mirror profile were to change to $z'_A = z_A + v\frac{\lambda}{2}$, where $v \in \mathbb{Z}$, the mirror matrix would be unchanged because the propagation phase to and from the mirror (e^{2ikz_A}) is unaffected. However, this change of the mirror profile actually constitutes a change to the cavity length, and therefore should induce a change to the mirror matrix. The problem arises because the integral over the mirror surface only accounts for phases accrued between the nominal plane of the mirror and the surface, rather than the true changes that emerge upon propagation.

The consequences of this error can be seen in Fig. 5.6. Here the resonance bands are vertical, meaning that transverse resonances always occur at the same nominal length. But as the cavity is misaligned, the mode angle tilts and the true length of the mode differs from the nominal length (for all nominal lengths other than precisely

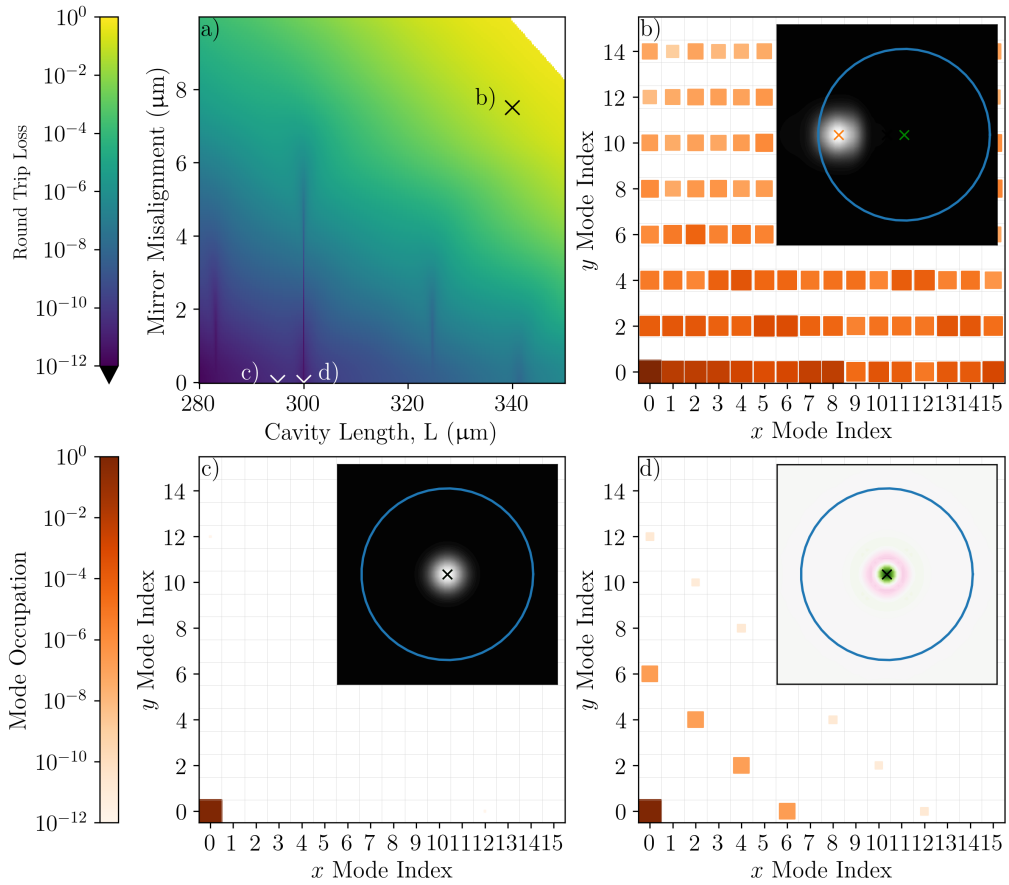


Figure 5.8: An investigation of specific length and misalignment configurations for cavities with spherical cap mirrors of diameter 75 μm. a): The round-trip loss as a function of length and mirror misalignment, on which are marked the 3 configurations of interest explored in the corresponding panels. b) c) and d): Breakdowns of the occupations of the cavity mode expressed in $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ for the corresponding configurations of interest. The insets depict the mode in the plane of the mirror at positive z. For a) and b), the mode intensity is plotted, but for c) the difference in intensity compared to the classical theory is shown to highlight the effect of transverse mode mixing. In each plot the circle represents the edge of the mirror. a): An example mode in the high loss region where light clips the edge of the mirror. Loss is seen as a cascade of occupation through higher transverse modes. b): The mode at zero misalignment away from the sharp dips in losses. The mode has near-complete overlap with $\left|\Psi^C_{0,0}\right\rangle$, which at zero misalignment matches $\left|\Psi^G_{0,0}\right\rangle$. c): The mode as the length is scanned onto a transverse resonance $q/p = 2/3$. Higher order transverse modes with n_x and n_y summing to multiples of 6 mix in with $\left|\Psi^C_{0,0}\right\rangle$. The reduction in round-trip loss at these lengths is explained by the residuals plot, which shows that the mode is more compact on the mirror than the expected mode.

confocal). This means that a model that properly accounts for the propagation to the mirror surface should not feature the vertical bands seen.

This deficiency in the mode mixing method is the reason why the method of removing rotation rotates the basis to the axis of the eigenmode $|\Psi_F\rangle$ rather than the axis of the expected mode $\left|\Psi^G_{0,0}\right\rangle$ even though transforming the basis to that of the expected mode would more clearly separate the geometric and mode mixing effects. The geometric model does properly account for the shape of the mirror, and therefore the mode angles predicted by the geometric model and mode mixing should subtly disagree.

In the context of the current investigation, this distinction between 'phase surface' and geometrical object is unlikely to change the general conclusions about the suitability of each type of mirror. However, the resonant features of misaligned cavities may not exactly follow the shapes depicted in the length-misalignment plots due to changes in the effective length of the cavity as the mode angle tilts, and this would be important to remember if comparing to experimental data.

Validity of Geometric Clipping Model

In Chapter 4, the geometric clipping model was used to optimise the parameters of cavities for remote entanglement. There is a difference between the losses predicted by the clipping model and the more sophisticated mode mixing calculation (shown in Fig. 5.7). However, as argued in Chapter 4, clipping losses transfer from tolerable to intolerable very sharply as the cavity is misaligned, and therefore clipping loss model is likely suitable for quickly determining which cavity configurations are viable, provided there is a safety factor in the chosen loss threshold $\mathcal{L}_{\text{thresh}}$. At specific lengths, it may be possible to tolerate higher misalignment than predicted by the clipping model by taking advantage of mode mixing, but that is not simple to arrange. The length would have to be tuned to simultaneously match the desired transition frequency and the appropriate transverse resonance, with the position of both dependent on the precise shape of the mirror.

5.4.3 Gaussian Mirror

Overview

As discussed in Sec.5.1, Gaussian mirrors are a significant case because mirrors formed from simple laser ablation are approximately Gaussian, and their investigation may also reveal general behaviour that can be seen in cavities with other non-spherical mirror shapes. The behaviour of cavities with Gaussian mirrors is expected to be more complex than for spherical cap mirrors because the curvature of a Gaussian mirror varies across its surface. This has two main effects. Firstly, at non-zero misalignment, the mode experiences two different principle radii of curvature, neither of which is the nominal radius of curvature. This results in an elliptical mode, a consequence predicted by the geometric treatment. Secondly, the mode predicted by the geometric picture will mix with others because its phase curvature will only match the mirror curvature precisely at the point of intersection, and therefore the geometrically predicted mode will not generally be a cavity eigenmode.

Example cavity loss structures are shown in Fig. 5.9, which can be contrasted with the spherical cap case. The low loss bands at constant length have been replaced with an array of high loss bands of varying width that spread across the length-misalignment landscape. The high loss bands do not exist at constant length, but instead curve toward longer length and split into multiplets as the misalignment increases. Some of these bands extend to zero misalignment, whereas others do not. The overlap of $|\Psi_F\rangle$ with the expected mode $\left|\Psi^G_{0,0}\right\rangle$ also contrasts with the spherical cap case; here, $|\Psi_F\rangle$ often differs significantly from $\left|\Psi^G_{0,0}\right\rangle$, even when expected to remain well within the confines of the mirror. While high mode distortion tends to be correlated with high loss, regions can be seen both where $|\Psi_F\rangle$ is not strongly distorted, but is highly lossy, and where $|\Psi_F\rangle$ is strongly distorted but not very lossy. High loss reduces the probability of photon extraction by reducing the cooperativity (see Chapters 2 and 4), and mode distortion reduces the coupling to fibre, and changes the coupling strength an emitter. Therefore, both of these effects would be problematic for remote entanglement. The features and behaviours seen will be discussed and analysed in the following paragraphs.

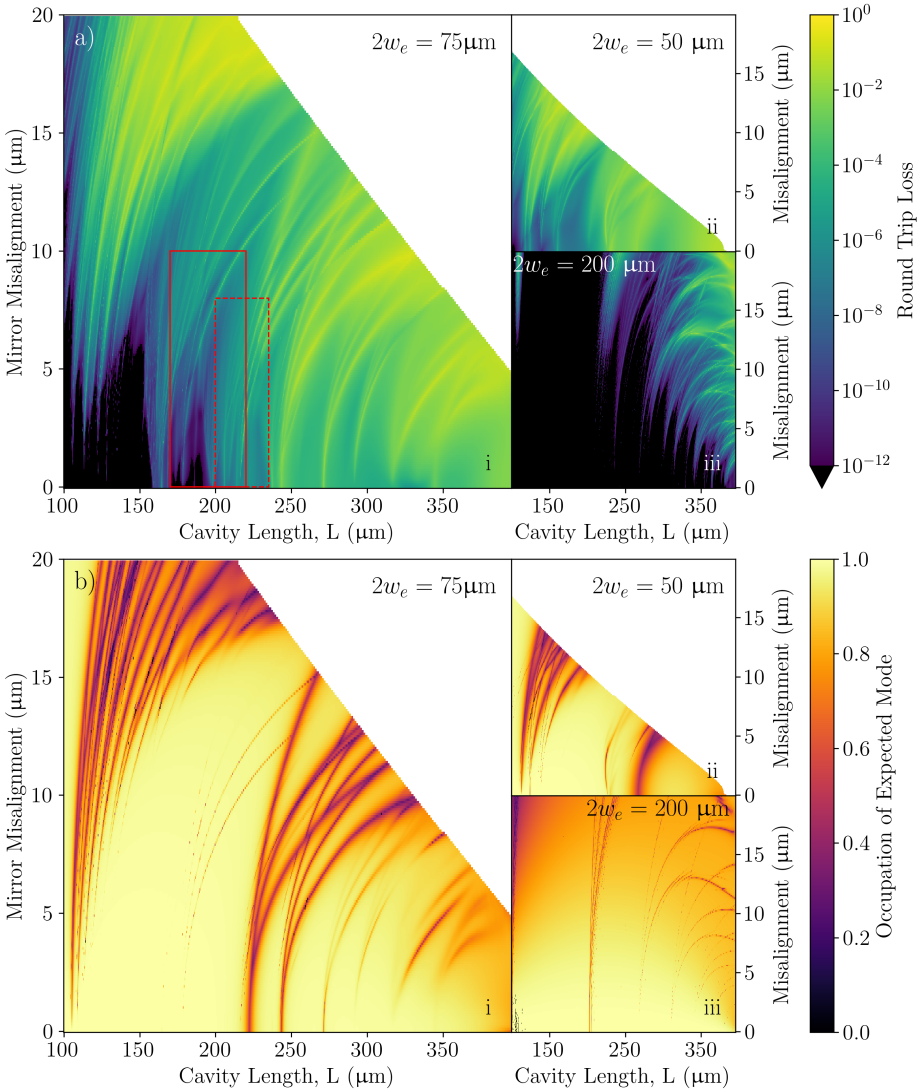


Figure 5.9: Round-trip loss (a) and overlap with the classically expected mode $\left|\Psi^G_{0,0}\right\rangle$ (b) as a function of cavity length and mirror misalignment for cavities with Gaussian mirrors. The mirrors of the cavity are always both identical and always have a central radius of curvature of 200 μm, but data is shown for three different Gaussian widths (as marked on the plots) and thus depths, examined at a wavelength of 1033 nm. Configurations where the region inside one waist of $\left|\Psi^G_{0,0}\right\rangle$ would not be fully enclosed within the concave region of the Gaussian mirror, are left white. Below round-trip losses of 10^{-12}, numerical noise becomes significant, and thus these results are left black. a) shows that Gaussian mirror cavities have high loss bands in length-misalignment space, with the sharpness and position of these bands depending upon mirror diameter. b) shows that these lossy bands often correlate with distortion of the cavity mode away from $\left|\Psi^G_{0,0}\right\rangle$, although there are regions where the cavity mode is distorted without high loss. The solid (dashed) red boxes in ai) indicate regions of interest that will be discussed in more detail in Figs. 5.12 (5.13).

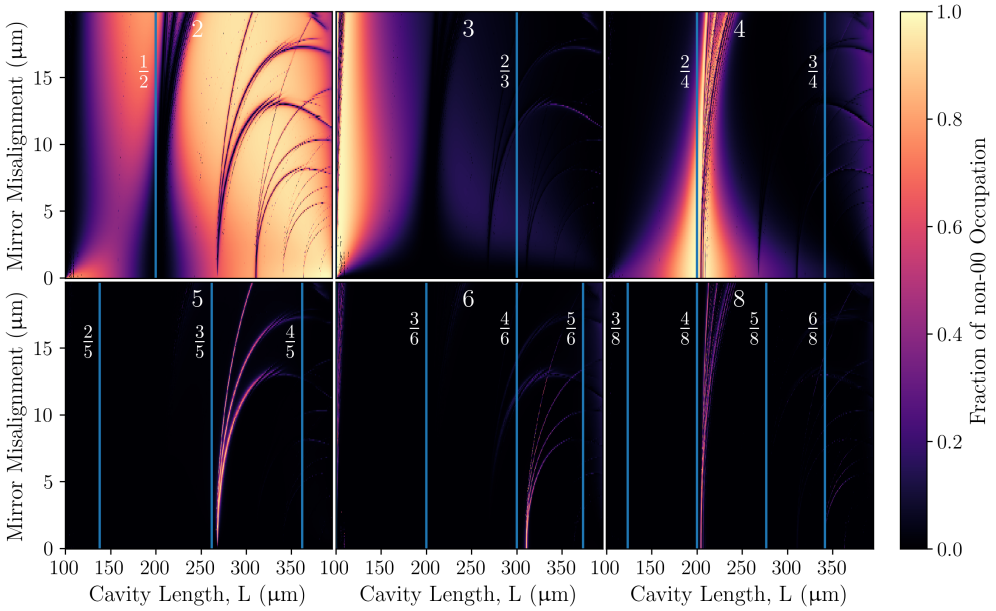


Figure 5.10: Occupation of different excitation indices $\mathcal{I}$ of the cavity mode in the co-propagating basis $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ over the length-misalignment map for Gaussian mirrors of $2w_e = 200$ μm and central radius of curvature 200 μm. Each panel corresponds to the labelled excitation index $\mathcal{I}$, with the lengths of resonance of each excitation index with the $\left|\Psi^C_{0,0}\right\rangle$ marked by vertical lines with the rational ratio q/p labelled. For each $\mathcal{I}$, the proportion of its occupation out of all not in the $\left|\Psi^C_{0,0}\right\rangle$ is plotted. Resonances are often dominated by a particular $\mathcal{I}$, for example for the 5 index near the q/p =3/5 line. As the mirror misalignment increases these resonances move to longer lengths and split into multiplets.

Understanding the loss features

Despite the more complicated structure, the features seen in the length-misalignment plots for Gaussian cavities can be attributed to resonances of higher-order transverse modes with $\left|\Psi^C_{0,0}\right\rangle$ as the low loss bands could in the spherical cap case. To understand the physics behind these features, it is easiest to begin with the $2w_e = 200$ μm case, depicted in Fig. 5.10, which is the mirror with the largest diameter presented and hence the best approximation to a spherical mirror. As transverse co-resonances (for perfect spherical cavities) apply to states with the same sum of n_x and n_y, this figure categorises the mode intensity in the co-propagating basis into different 'transverse excitations', labelled by $\mathcal{I} = n_x + n_y$. Different features in the length-misalignment configuration space can often be attributed to a specific $\mathcal{I}$, as discussed below.

Generally, most of the non-00 occupation is found in $\mathcal{I} = 2$ which represents a

low spatial frequency deformation of the mode shape. Around a ratio of $q/p = 1/3$, $\mathcal{I} = 3$ is dominant, and around $q/p = 1/2$, $\mathcal{I} = 4$ also becomes prominent; at this length, all modes of even $\mathcal{I}$ should be resonant with $\left|\Psi^{C}_{0,0}\right\rangle$ under classical theory.

For understanding the origin of the high loss bands, the most enlightening examples are found for $\mathcal{I} = 5$ and 6 around q/p of 3/5 and 4/6 respectively. Observing these features, there are four key points of interest, which will be discussed in turn

1. The features occur at a length slightly longer than the nominal resonant length at zero misalignment.

2. The features occur at longer length as the misalignment increases.

3. The features appear to split into multiplets as the mirror misalignment increases.

4. The feature attributed to $\mathcal{I} = 5$ does not extend to zero misalignment, but that of $\mathcal{I} = 6$ does.

The first point, that the features do not match exactly the length at which a transverse resonance would be expected, can be understood as a consequence of the mode exploring a non-zero extent of the mirror. This means that the 'effective curvature' seen by the mode will be some average of the curvature over the mode's extent on the mirror. As seen in Sec. 5.2, Gaussian mirrors have their highest curvature, the nominal curvature, in the centre, and therefore the effective curvature is always less than the nominal curvature, which pushes the resonant length longer in accordance with Eq. 5.16.

The second aspect, that the resonant features are observed at longer lengths as the misalignment increases, can also be understood in the same way. As the misalignment increases, the mode will intersect the Gaussian mirror further from its nominal centre, and therefore the curvature of the mirror will reduce (or equivalently, the radius of curvature will increase). Again, in accordance with Eq. 5.16, the increase in effective radius of curvature increases the resonant length.

The third point, the splitting of the resonances into multiplets as the misalignment increases, is a consequence of the elliptical curvature of a Gaussian surface away from the nominal centre. As discussed in Sec. 5.2, when the mode hits the Gaussian

mirror away from the nominal centre, the radius of curvature seen decreases, but also splits into two principle values: the lower curvature (higher radius) in the direction of the intersection point from the centre, and the higher curvature (the least reduced curvature) in the orthogonal direction. With two radii of curvature, the equation for the round-trip phase of different modes (Eq. 2.11 in Chapter 2) is adjusted to account for the different x and y curvatures

$$\phi_{\mathrm{RT}} = 4\pi\frac{L}{\lambda} - 4\left(n_x + \frac{1}{2}\right)\arctan\left(\sqrt{\frac{L}{2R_x - L}}\right) - 4\left(n_y + \frac{1}{2}\right)\arctan\left(\sqrt{\frac{L}{2R_y - L}}\right). \tag{5.17}$$

This equation does not permit a simple algebraic formula for the lengths at which transverse resonances occur, but it is clear that splitting of the radius of curvature into two values removes the degeneracy of transverse modes belonging to the same $\mathcal{I}$. Instead there are $(\mathcal{I}+1)$ resonances (the number of different n_x and n_y combinations in the excitation index). The mode with the most excitations in the direction corresponding to the smallest radius of curvature will be resonant with the fundamental at the shortest length. For Gaussian mirrors, because the radius of curvature increases from the centre value with misalignment, and increases more in the direction of misalignment (x) than the orthogonal direction (y), one would expect the resonances with higher n_y of a given multiplet to occur at shorter lengths. The multiplicity of the splittings can also be understood from this model. As mentioned before, there are $(\mathcal{I}+1)$ states of a given $\mathcal{I}$, and when the radii of curvature differ with misalignment, these states become non-degenerate. However, because the system always remains mirror symmetric about the y-axis (even under misalignment, as this only acts in the x direction), the cavity eigenmodes must have a definite parity in the y-direction. The eigenmode chosen is the one with greatest overlap with $\left|\Psi^G_{0,0}\right\rangle$, which is always even parity in the y-direction. Therefore the selected eigenmode $|\Psi_F\rangle$ must have even y-parity, and can only mix with states of the same even y-parity (i.e. even n_y). This means that the expected multiplicity of a peak is not $(\mathcal{I}+1)$, but is instead the number of even n_y states in the $\mathcal{I}$ manifold. For the $\mathcal{I}=5$ resonance, we would therefore expect the peak to split into 3 upon misalignment (as seen), with the coupling to higher order modes dominated by $\left|\Psi^C_{1,4}\right\rangle$, $\left|\Psi^C_{3,2}\right\rangle$ and then $\left|\Psi^C_{5,0}\right\rangle$ as length increases. By the same logic, for the $\mathcal{I}=6$ resonance, the multiplet should

be a quadruplet, corresponding to resonances with, $\left|\Psi^C_{0,6}\right\rangle$, $\left|\Psi^C_{2,4}\right\rangle$, $\left|\Psi^C_{4,2}\right\rangle$ and $\left|\Psi^C_{6,0}\right\rangle$, again ordered in increasing length. Finally, it is worth noting that the difference in radii of curvature in the x and y direction will also cause cavity birefringence (a splitting of $|\Psi_F\rangle$ into two orthogonal polarisations). Cavity birefringence cannot be accounted for in the mode mixing method, which is scalar, but could cause further problems for remote entanglement protocols, as will be discussed in Chapter 6.

The final point of note is that, while the feature of $\mathcal{I} = 6$ does extend to zero misalignment, the resonance of $\mathcal{I} = 5$ is only seen at non-zero misalignment. This observation can be understood by considering the symmetry of the states involved. Mode mixing occurs when the mirror mixes light from the $\left|\Psi^C_{0,0}\right\rangle$ (which has even parity in both Cartesian directions) to a higher order mode. At zero misalignment, the mirror has mirror symmetry in both directions about the point where the mode intercepts the mirror, and therefore mixing can only occur with modes of even parity in both the x and y directions. Such modes are present in $\mathcal{I} = 6$, but there are none for $\mathcal{I} = 5$. Therefore the $\mathcal{I} = 5$ feature cannot extend to zero misalignment even though there are modes resonant with $\left|\Psi^C_{0,0}\right\rangle$. At non-zero misalignment, the mode intersects the mirror away from its nominal centre. In the direction of the misalignment, the mirror is no longer symmetric about the intersection of the mode on the mirror due to the varying radius of curvature on either side of this intersection point[1]. This means that, at non-zero misalignment, $\left|\Psi^C_{0,0}\right\rangle$ can couple into modes with both odd and even parity in the direction of misalignment, allowing for coupling into states with $\mathcal{I} = 5$. In this way, the resonant features can be classed according to odd or even symmetry in the direction of misalignment. Calculations that only account for the perfectly aligned configurations (such as the previous calculations of [241]) therefore do not account for these odd-symmetry resonant features that are seen at finite misalignment. This was noted in the discussion of Fig. 5.4 of Sec. 5.3.8. It is important to consider odd-resonances, as any breaking of the mirror symmetry due to manufacturing imperfections leaves $\left|\Psi^C_{0,0}\right\rangle$ vulnerable to resonant mixing into odd modes.

The major unresolved question is why the resonant features are low loss for the spherical cap case but high loss for the Gaussian case. This is a difficult question to

[1] in the direction perpendicular to misalignment, the mirror is still mirror-symmetric about the intersection, as was important for understanding the multiplicity of the resonances.

answer because, within the mode mixing picture, the difference between increased loss and decreased loss is determined by how the different basis states interfere in the superposition. What can be said is that it is reasonable for mode mixing in the spherical cap case to produce low loss: Suppose a superposition out of the resonant states is more compact on the mirror than the geometric fundamental $\left|\Psi^G_{0,0}\right\rangle$. Because the transverse mode mixing of the spherical cap mirror originates entirely from the region outside the clipping radius, this more compact superposition will mix to other states less after one round trip than $\left|\Psi^G_{0,0}\right\rangle$ because less of the mode hits the mixing part of the mirror. While not a proof, this means that it is reasonable to suggest that cavities with spherical cap mirrors could have an eigenmode with lower expectation round-trip loss than $\left|\Psi^G_{0,0}\right\rangle$ at lengths where resonant mode mixing can occur. For the Gaussian mirror, it is much harder to make these arguments; the mode is generally always slightly mixed so that the equiphase surface of the mode at the mirror matches the slightly non-parabolic surface, and the mirror also cannot be neatly divided into a mixing and non-mixing part. This means that it is not necessarily expected for the same physics that produced low loss bands in for spherical cap mirrors to apply for Gaussian mirrors. This said, it would be interesting to find a more convincing explanation of why resonances correspond to low or high losses, and how that is affected by the mirror shape.

Further understanding of these effects can be found in Fig. 5.11 which displays the fraction of non-$\left|\Psi^C_{0,0}\right\rangle$ occupation for the first few $\mathcal{I}$ values at the three different widths shown in Fig. 5.9. This allows for the various effects discussed previously to be examined as a function of mirror width on a scale from a wide, highly spherical mirror ($2w_e = 200$ μm) to a smaller mirror that approximates a sphere only in a very small region about its nominal centre ($2w_e = 75$ μm).

The first aspect to study is the shift in position of resonance near zero misalignment. This was previously attributed to the mode sampling an average curvature over the mirror, which is always less than the nominal curvature. Thus, the resonances should shift further for the smaller waist Gaussian mirror whose curvature decreases more rapidly off-centre. This can be seen in the $q/p = 3/5$ feature of $\mathcal{I} = 5$ for the

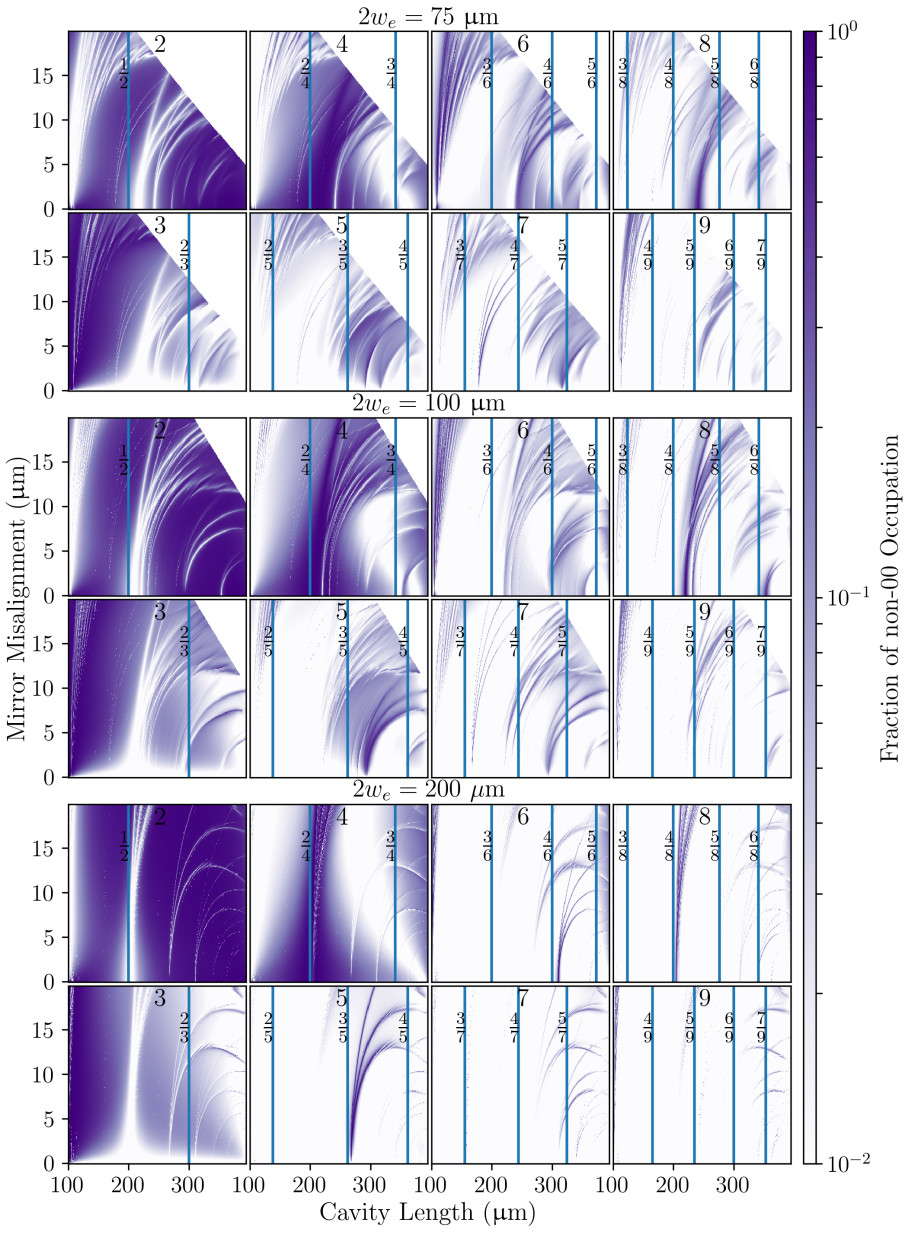


Figure 5.11: A summary of the non-$\left|\Psi^C_{0,0}\right\rangle$ occupation in the first few excitation indices $\mathcal{I}$ of mirrors with three different Gaussian widths but the same central radius of curvature of 200 μm. There are a set of 8 plots corresponding to each width, with the width labelled above each set. Each graph has its $\mathcal{I}$ labelled, and the labelled vertical lines mark the lengths at which transverse mode splitting to free spectral range ratio would give resonance of modes of that $\mathcal{I}$ with $\left|\Psi^C_{0,0}\right\rangle$ for a perfect spherical cavity. The label is the transverse mode splitting to free spectral range ratio of the expected resonance.

$2w_e = 200$ μm and $2w_e = 100$ μm datasets, where the length shift of the resonant feature is larger for the smaller mirror.

The second aspect of physics that this figure highlights is that this resonance shift is also dependent on the excitation index $\mathcal{I}$. In a perfect spherical cavity, transverse mode splittings are equally spaced, meaning that if a certain $\mathcal{I}$ is at resonance (for example $\mathcal{I} = 3$), all excitation indices with that index as a factor (such as $\mathcal{I} = 6$, $\mathcal{I} = 9$, etc.) are also resonant. However, for a Gaussian mirror with variable radius of curvature, the effective radius of curvature depends upon the extent of the mirror sampled. This means that larger modes (modes with higher $\mathcal{I}$) will experience more of the mirror and therefore have a higher effective radius of curvature. This lifts the degeneracy of resonances with different $\mathcal{I}$'s with common denominators. For the $2w_e = 200$ μm mirror, this shift is too small to observe, but in the case of $2w_e = 75$ μm, the region around a length of 100 μm clearly shows the step in length between the resonance of $\mathcal{I} = 6$ and $\mathcal{I} = 9$.

Specific examples

The physics of mode mixing in cavities with Gaussian mirrors can be understood further by analysing the specific example regions of $2w_e = 75$ μm indicated in Fig. 5.9, which are discussed in turn.

In Fig. 5.12 an $\mathcal{I} = 7$ multiplet is studied. As an odd-symmetry resonance, this feature only manifests at non-zero misalignment. As the misalignment increases, the feature splits into a quadruplet, which matches the number of even-n_y states within $\mathcal{I} = 7$. The mode composition of points in each of the four arms of the quadruplet was analysed. The dominant non-00 component in $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ for each point was, in order of increasing resonant length at a given misalignment, $\left|\Psi^C_{1,6}\right\rangle$ $\left|\Psi^C_{3,4}\right\rangle$, $\left|\Psi^C_{5,2}\right\rangle$ and $\left|\Psi^C_{7,0}\right\rangle$. The resonances with higher n_x occur at longer lengths, as these resonances are most affected by the strong reduction in curvature in the direction of misalignment, with the curvature reducing much less in the orthogonal y direction. Despite all features belonging to the $\mathcal{I} = 7$ resonance, each peak presents a dramatically different mode shape, as the higher order mode that mixes most strongly changes between the peaks of the multiplet.

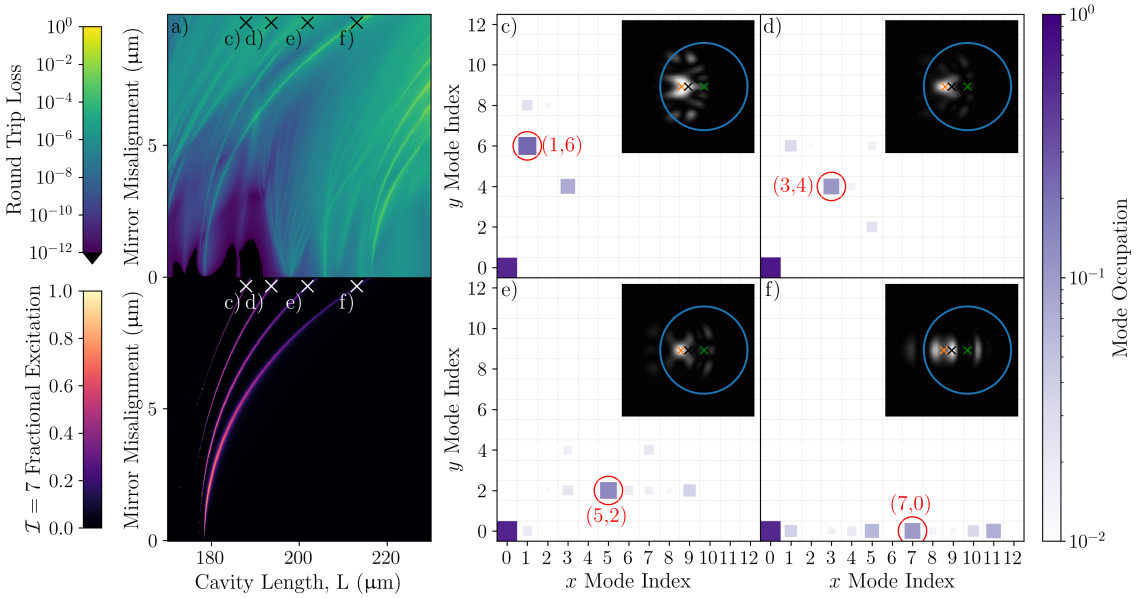


Figure 5.12: Study of a resonance feature with modes of $\mathcal{I} = 7$ for the $2w_e = 75$ μm mirrors. a) Cavity round-trip loss as a function of length and mirror misalignment and b) fraction of the non-$\left|\Psi^C_{0,0}\right\rangle$ occupation in $\mathcal{I} = 7$, showing that the multiplet has strong occupation in this excitation class. On a) and b), four example points, corresponding to (c-f) respectively, are marked. (c-f): Mode compositions and profiles of the configurations shown in panels (a) (b). The main axis shows the occupation of the $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ basis, with the inset showing the mode intensity on of one of the cavity mirrors. On each figure, the basis state with the largest non-$\left|\Psi^C_{0,0}\right\rangle$ occupation is ringed and labelled according to (n_x, n_y), with this n_x increasing across the multiplet in the direction of increasing length.

Fig. 5.13 studies the behaviour around the confocal configurations of the system, a region where previous studies have found many lossy resonances and areas of strong mode mixing (see, for example [241]). For this example, there is a very strong occupation of $\mathcal{I} = 4$ throughout, but with streaks of high loss scattered in. In contrast to the previous region, the interesting feature here is that it is hard to understand the high-loss streaks as any kind of particular resonance. This is because the mode mixing becomes so strong that the mode of interest $|\Psi_F\rangle$ is difficult to define. The mode of interest is selected as the cavity eigenmode with the greatest overlap with the expected mode $\left|\Psi^G_{0,0}\right\rangle$, a method chosen on the expectation that $|\Psi_F\rangle$ would usually be a perturbation of $\left|\Psi^G_{0,0}\right\rangle$. However, throughout this region, there are cases where the mixing is so strong that there is no 'good' choice for $|\Psi_F\rangle$. This leads to behaviour where $|\Psi_F\rangle$ can switch discontinuously between two orthogonal modes, as shown in the panel where there are two examples near-orthogonal modes of interest

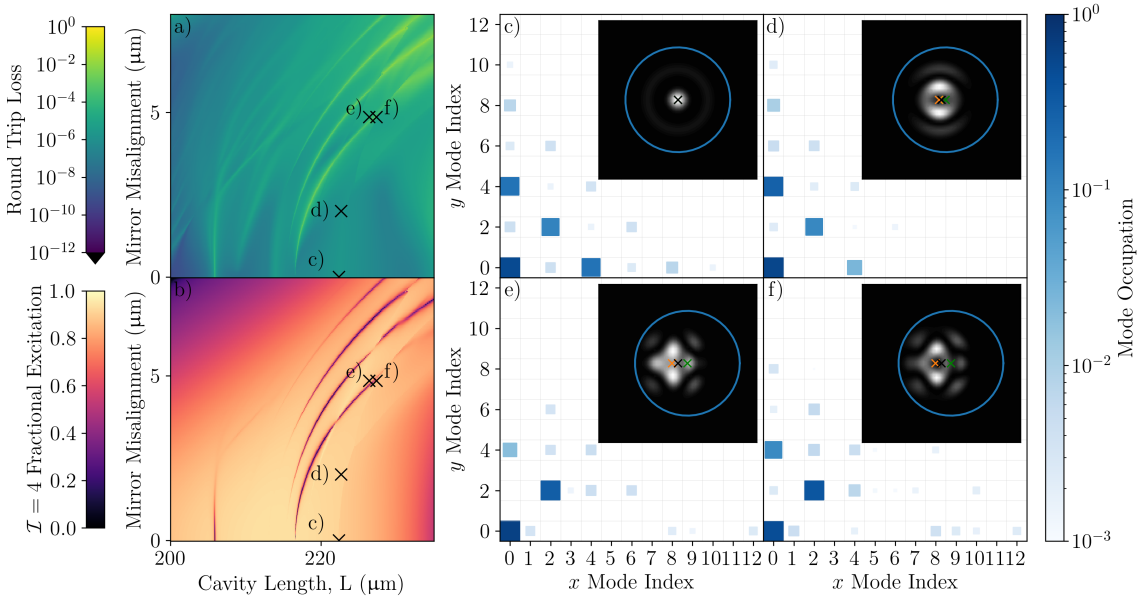


Figure 5.13: Study of the confocal configurations of cavities with $2w_e = 75$ μm mirrors. a) Round-trip loss as a function of length and translation. b) The fraction of the non-00 occupation of $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ with $\mathcal{I} = 4$; showing a very broad peak across the configurations shown. On a) and b), four example points corresponding to (c-f) respectively, are marked. (c-f): The mode compositions in $\left\{\left|\Psi^C_{n_x,n_y}\right\rangle\right\}$ for the configurations shown in panels (a) (b), with insets showing mode intensity. Points c) and d) show highly distorted modes, even though the round-trip loss is not noticeably high. These two modes have small overlap, with c) having an intensity peak in the centre, and d) a trough). e) and f) are two points across the same resonant feature that also show high distortion and almost orthogonal mode profiles. These figures demonstrate that the mode with greatest overlap to the expected mode $\left|\Psi^G_{0,0}\right\rangle$ can change discontinuously in regions of high mode distortion.

being selected from very close configurations within the parameter space.

Even though the round-trip loss of some of these deformed modes is not greatly increased, they would still present a big problem for cavity QED experiments. Firstly, the coupling of the emitter to the cavity mode would be affected by the deformation of the cavity mode, particularly for those modes with a reduced intensity in the centre. Secondly, to use the single photons extracted from a cavity, the photons must be collected, and this is done most conveniently into a single mode fibre. The distortion of the cavity mode will affect the maximum possible coupling into a single mode fibre, reducing the rate at which photons can be extracted from the cavity. This means that, particularly in this confocal region where large mode mixing tends to occur, the mode of interest $|\Psi_F\rangle$ should be calculated, not just to determine that the loss is acceptable, but also to check that the mode profile is suitable for the application.

Wavelength scanning

The calculations in this investigation were all conducted at a single wavelength (1033 nm). At this wavelength, the jth cavity eigenmode has an eigenvalue $\lambda_j = \gamma_j e^{i\phi_j}$. The round-trip phase ϕ_j is then ignored, and the loss calculated as $(L_{\mathrm{RT}})_j = 1 - |\gamma_j|^2$. This generates a loss value at every point on the plot that is consistent with the round-trip loss that would be seen where that wavelength is resonant. However, in a spectroscopy experiment, for a given cavity configuration the probe wavelength is scanned until the resonance is found, which raises the question of whether it is valid to calculate the results at a single wavelength. Changing the wavelength has two principal effects:

1. Altering the round-trip phase e^{-2ikL}
2. Changing the mirror matrix components

The first effect is usually the largest in size, and is the simplest to account for. The phase factor added to the round-trip matrix will not alter the eigenvectors, only the phases ϕ_j of the eigenvalues. Therefore, on the assumption that this effect is the only contribution, it is completely valid to calculate the losses without worrying exactly what the wavelength is, as the mode itself would not change when the wavelength is scanned to hit resonance. The second effect is the change in eigenvector composition due to wavelength, which occurs because, as wavelength is scanned, the basis modes intersecting the mirror change in waist, phase curvature, and phase relative to other modes.

Given the complexity of the calculations, it is difficult to argue analytically without assumptions that tweaking the wavelength to probe the cavity resonances does not affect the numerical conclusions significantly. Therefore, this assumption was tested numerically by conducting scans at three wavelengths. The first was a nominal central wavelength, and the other two had frequencies that differed from the central frequency on either side by half a free spectral range. This means that, at any point in the configuration space, the wavelength required to hit a resonance will lie within the bounds of the highest and lowest wavelength/frequency. The scans were taken over a sharp resonance feature, and are shown in Fig. 5.14. The wavelength

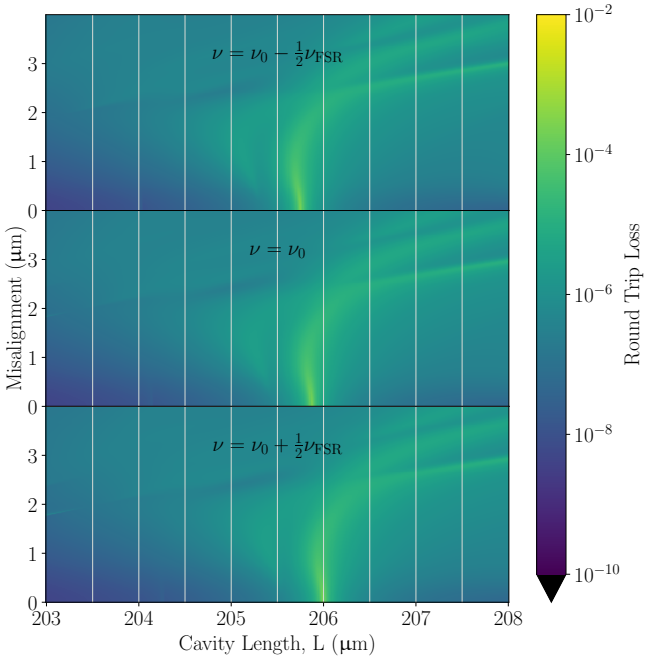


Figure 5.14: Round-trip loss in a small region of the length-misalignment configuration space for the $2w_e = 75$ μm Gaussian mirror cavity interrogated at three wavelengths/frequencies that cover a free spectral range of tuning. The resonance feature in these plots is seen to shift position with the change in probe frequency, but the structure does not distort significantly.

change is shown to affect the positions of the resonances, but keep the structure mostly intact. This suggests that calculating losses for a single wavelength is a valid approach to investigate the qualitative behaviours seen in the loss maps of these systems, but that comparisons of the exact positions of lossy resonances to experiment should account for the wavelength scanning. Additionally, for shorter cavities with larger free spectral ranges, the wavelength scans required to hit resonances will become more extensive, and below a certain length it will become important to consider these effects for even the qualitative conclusions.

5.5 Conclusion

The aim of this work was to understand the behaviour of cavities with non-spherical mirrors under misalignment. Because existing methods to perform these calculations are computationally intensive, new techniques were developed to expedite the calculations. These techniques exploit the similarity between the Hermite-Gauss functions that form a basis of cavity modes with the eigenfunctions of a simple

harmonic oscillator. This allows first for the mixing matrix of a mirror to be assembled directly from its profile rather than built piece-by-piece through integrals, and second for these mirrors to be translated efficiently by combining the relations between basis states with the existing mirror matrix. The combination of these techniques significantly reduces the time that would be required to calculate cavity modes as a function of both length and mirror misalignment, which allowed for data over this space to be calculated at high resolution, revealing rich physical behaviour.

These methods were then used to study the modes of cavities with spherical cap-shaped and Gaussian-shaped mirrors under transverse misalignment to understand how they would behave when coupled to an emitter. Generally, it was found that spherical cap mirrors behaved in a manner qualitatively similar to the simple geometric picture; mode loss increases as the cavity mode impinges on the edge of the mirror, and the mode shape is mostly very similar to the rotated fundamental mode expected from the geometric theory. Deviations from this classical theory are seen in sharp dips in the round-trip losses at lengths corresponding to the co-resonance of higher order transverse modes with the expected mode. The resonant higher order modes hybridise with the expected mode to reduce the intensity falling outside the finite diameter mirror and reduce the losses, although even in these cavity configurations, the overall shape of the mode remains mostly unaffected. Therefore, given a certain margin for safety, the properties of spherical cap cavities can be well understood from the simple geometric picture, and there are no unexpected physical effects that occur when the cavity is misaligned. This means that it is sufficient to choose a cavity geometry where, factoring in the expected misalignment, the geometrically expected mode is well confined within the mirror. If the length happens to correspond to a transverse resonance, there might even be a drop in clipping loss due to transverse mode mixing, but this does not distort the mode significantly, and would therefore be beneficial.

Changing to Gaussian shaped mirrors introduces a wealth of structure due to the variation in curvature across the surface. Instead of low loss bands at particular resonant lengths, the resonant features are high loss bands which split and curve over to longer length as the cavity mirrors are misaligned. Again, these bands can be

understood as resonances of the expected mode with higher order transverse modes, but due to the variable curvature across the mirror, the effective mirror curvature seen by the cavity mode tends to be lower for higher order basis states, and to reduce away from the central axis (most prominently in the direction of translation, but also in the orthogonal direction). These two effects alter the positions of resonance. Firstly, all resonances are shifted to longer lengths, with the higher order resonances shifted more. Second, as the misalignment increases, the resonances curve to longer length and split into multiplets. The multiplicity of the resonance can be predicted by counting the number of symmetry-allowed couplings of the resonant excitation index. The resonances can be categorised as odd or even, with even resonances allowed for zero misalignment, but odd resonances only observable when the cavity is misaligned.

To use Gaussian mirrors reliably for emitter cavity experiments, it is not merely necessary to set the configuration so that the geometrically expected mode remains in the confining region of the mirror; it is also important to avoid the high loss bands that can occur even when the mode should be well-confined by geometric theory. However, the positions of these bands depend strongly on the curvature of the surface, so may be hard to predict without a very repeatable fabrication process. Furthermore, the strong mixing that tends to occur at lengths slightly longer than confocal means that this region of the parameter space, promising for obtaining high coupling with low clipping loss, could be highly problematic. Misalignment, in addition to bringing the mode closer to the edge of the confining region of the mirror, also renders the mode vulnerable to odd-character transverse resonances, meaning that a design with residual transverse misalignment, such as ours, would have to 'avoid' even more high-loss resonances.

Finally, it should be noted that while this investigation has focussed cavities with Gaussian mirrors, many of the qualitative conclusions are likely to hold in the case of other non-spherical mirrors. Once the mirror matrix in the axial configuration is found for whatever mirror is of interest, the translation methods of Sec. 5.3.6 can be used to quickly understand the cavity's behaviour under mirror misalignment.

6

Polarisation Encoding and Polarisation Oscillations

6.1 Introduction

Two-photon heralded remote entanglement protocols function by producing entangled ion-photon pairs and then performing entanglement swapping through photonic measurement [168]. When using polarisation encoding (depicted Fig. 6.1), it is the polarisation degree of freedom that is entangled with the ion qubit state. Polarisation encoding is a common encoding for two-photon remote entanglement, used, for example, with trapped ions in free space [96, 250] and a recent cavity-enhanced entanglement demonstration [108]. Thus, polarisation encoding would appear a suitable choice for a cavity enhanced remote entanglement experiment.

However, polarisation is an easily manipulated degree of freedom, which leaves polarisation encoding susceptible to unwanted rotations. Specifically, this encoding is vulnerable to cavity birefringence, which will be discussed in the bulk of this chapter, as well as fibre noise and specific undesired transitions, which will be discussed in Sec. 6.6.

Contribution and Acknowledgements

The content of this chapter is based upon the research presented in [136], where further information can be found. For this publication, Dr. Ezra Kassa wrote the

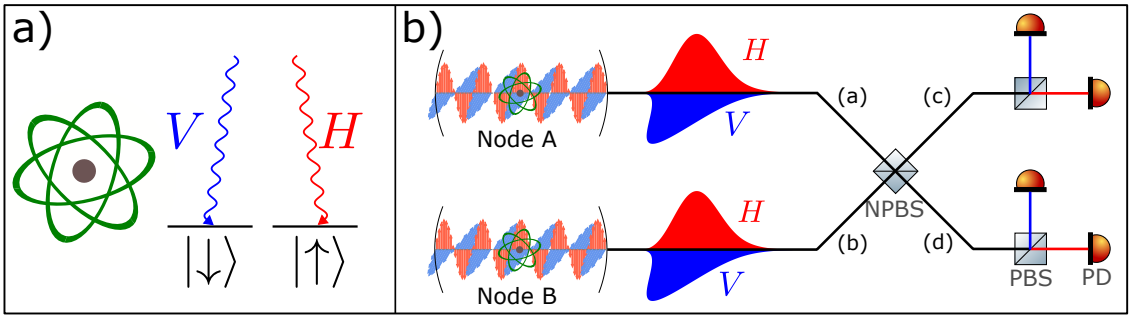


Figure 6.1: Diagram of polarisation-encoded remote entanglement. a) An ion in each node emits a photon whose polarisation is entangled with the final ion state. b) The photons from each node are interfered on a beamsplitter to remove information about which ion they came from, and then travel to a polarisation-resolving measurement. Upon a measurement of one horizontally-polarised (H) photon and one vertically-polarised (V) photon across the detectors in a single attempt, the ions are ideally projected into a fully entangled Bell state.

initial mathematics and discussion, with the author providing the modelling code and figures, with the exception of precursors to Figs. 6.1 and 6.4 (produced by Dr. Ezra Kassa), and the= initial version of Fig. 6.2 (Dr. Joseph Goodwin and Dr. Shaobo Gao). The results section of this publication was then substantially modified, with the text written by Dr. Joseph Goodwin and the author, and the author again managing the modelling and the figures.

6.2 Cavity Birefringence

Cavity birefringence shifts the two nominally-degenerate orthogonally-polarised fundamental modes of the cavity to different resonant frequencies. These eigen-polarisations are mutually orthogonal. For cavities with superpolished mirrors, birefringence has been attributed to stress encountered in coating or mounting [251], and with microcavities the birefringence is often higher [185, 192, 252–254] due to 'geometric birefringence': birefringence deriving from an asymmetric mirror shape [189]. Due to its detrimental impacts on experiments [102, 255], techniques to reduce the geometric birefringence in laser ablated setups have been developed [190, 191].

The origin of geometric birefringence cannot be understood from the paraxial approximation. However, as discussed in [256], the full Maxwell equations can be solved through series of corrections to the paraxial approximation in increasing orders of the small parameter $1/(kw_0)$. Each non-zero even order is a transverse

correction to approximate the second derivative that was ignored in the paraxial approximation. Each odd order is a longitudinal field to cancel the divergence of the previous even order (To satisfy Eq. Eq. 2.3 of Chapter 2). As discussed in [257], this longitudinal field changes the resonant frequencies of the cavity because it subtly alters the reflection condition that is used at the mirror. In the scalar picture, this was that the mirror formed an equiphase surface of the scalar field, but the true vector condition is that the incoming and reflected fields superpose to cancel the electric field component parallel to the mirror surface. If the mirror profile lacks symmetry, this reflection condition leads to geometric birefringence.

In the case of an elliptical mirror profile, taking just the first order longitudinal correction leads to a linear birefringence with splitting of

$$2|\boldsymbol{\delta_B}| = \frac{\nu_{\mathrm{FSR}}}{2\pi}\frac{R_1 - R_2}{kR_1R_2}, \tag{6.1}$$

where R_1 and R_2 are the principle radii of curvature of the elliptical mirror [189]. The difference in curvature could result from asymmetric milling or incidence off-axis on a non-spherical mirror (as discussed in Chapter 5). However, geometric birefringence can result from more general deformations than elliptical.

Whatever the origin, the central consequence of birefringence is that the degeneracy of the two fundamental cavity resonances is broken, and thus polarisations that are not 'eigenpolarisations' will evolve in time while in the cavity, a phenomenon termed 'polarisation oscillations' [258]. The potential for the photon polarisation to change between emission and detection presents problems for the polarisation-encoding scheme, as information about the ion state is transmitted in the photon polarisation. Photons leaving the cavity might not have the same polarisation as when they were first emitted, which can result in the wrong ion-ion state being produced after the photonic measurement, reducing the fidelity.

Finally, it should be noted that while cavity birefringence will be investigated as a detrimental effect for polarisation-encoded remote entanglement, there are some potentially useful applications of this phenomenon, such as cavity enhanced polarimetry [259], cavity stabilisation [260], increasing photon extraction [209] and

specific quantum protocols [261]. To this end, deliberately birefringent microcavities have been realised [237, 262].

6.3 Modelling Cavity Birefringence in Remote Entanglement

6.3.1 Cavity Geometry

In general, cavity birefringence will reduce the fidelity of polarisation encoded remote entanglement protocols, however, there are a many different factors to consider in order to evaluate the impact. The first is the relative orientation of the magnetic field and cavity axis, which affects atomic transitions that couple to the cavity modes (see Sec. 2.3.4 of Chapter 2). This in turn sets the resonant selectivity requirements, which the impacts the quantum dynamics of the coupled ion-birefringent cavity system. Therefore, to present a fair assessment of the impact of birefringence on polarisation encoded remote entanglement, it is important to compare the two principal options for ion-cavity coupling, the first with the magnetic field and cavity axis parallel, and the second with them orthogonal.

σ_-/σ_+ encoding

The first geometry considered has magnetic field colinear with the cavity axis, as shown in the upper row of Fig. 6.2. Such a geometry leads to a σ_-/σ_+ photon encoding scheme, as the π decay is unsupported by either cavity mode. While this geometry is compatible with a range of ions, it is most naturally used for ions with hyperfine structure, such as the case of $^{171}Yb^+$, shown here, where the symmetric Clebsch-Gordan coefficients give equal coupling on both decay channels. Because only two decay channels are geometrically allowed, there is no need to suppress additional channels with frequency selection, and thus this scheme can function close to Zeeman degeneracy. Upon emission from the cavity, a quarter-wave retarder transforms the polarisation $\sigma_-/\sigma_+ \rightarrow H/V$ for compatibility with the Bell state analyser.

Simulations of the σ_-/σ_+ encoding use $g_\uparrow = g_\downarrow = \kappa$ and $\gamma = 0.6\kappa$, where 2κ is the cavity linewidth and $g_\uparrow$ ($g_\downarrow$) is the cavity coupling strength between the excited state $|e\rangle$ and qubit state $|\uparrow\rangle$ ($|\downarrow\rangle$). The separation of atomic states is set to 0.5κ,

which is allowed to be small because the scheme does not rely on resonant selection, and therefore the field serves only to define a quantisation axis. The system is driven with a single laser frequency due to the small splitting of the qubit states. Note that the selection of $g_\uparrow, g_\downarrow \sim \kappa$ is near-ideal for optimal extraction efficiency [161], with the choice of $\gamma = 0.6g$ loosely based on currently achievable coupling strengths in ion-cavity systems [105].

H/V encoding

The second geometry sets the cavity axis orthogonal to the magnetic field, and the photon emission is directly encoded in an H/V basis via projection of σ_- and π decays onto the cavity mode, shown in the lower row of Fig. 6.2 with the example of $^{40}Ca^+$– or $^{88}Sr^+$–like Group-II ions. The Clebsch-Gordan coefficients for these channels are not equal, but, accounting for the angular dependence of the atom-field coupling, the coupling strengths can be made more equal using 'stretch' transitions at an extremum of a Zeeman manifold. Because σ_+ decays project to the same H-cavity photon polarisation as σ_-, the unwanted σ_+ decay channel must be suppressed through resonant selection, necessitating a large Zeeman shift. This in turn means that the two qubit channels should be driven by a bichromatic drive, leading to a more complex output photon wavepacket, as examined later.

For simulations of H/V encoding, accounting for the Clebsch-Gordan coefficients and orientation of cavity mode leads to $g_\uparrow = \sqrt{\frac{1}{3}}g_0$ and $g_\downarrow = -\sqrt{\frac{4}{15}}g_0$ with g_0 chosen such that $g_\uparrow^2 + g_\downarrow^2 = 2\kappa^2$. The two qubit states are separated in energy by 5κ so that the off-resonant decay channel ($P_{\frac{3}{2}}, m_J = 3/2 \rightarrow D_{\frac{5}{2}}, m_J = 1/2$) is strongly suppressed. This splitting between the qubit states means that a bichromatic drive must be used with the two separate frequencies driving two separate Raman process through the cavity mode. κ and γ remain unchanged from the σ_-/σ_+ case.

6.3.2 Simulating Production of Photon Wavepackets

The impact of cavity birefringence on the photon wavepackets produced by ion cavity systems was modelled for both photon production geometries. To focus as much as possible on the impact of birefringence, the simulations used a minimal model

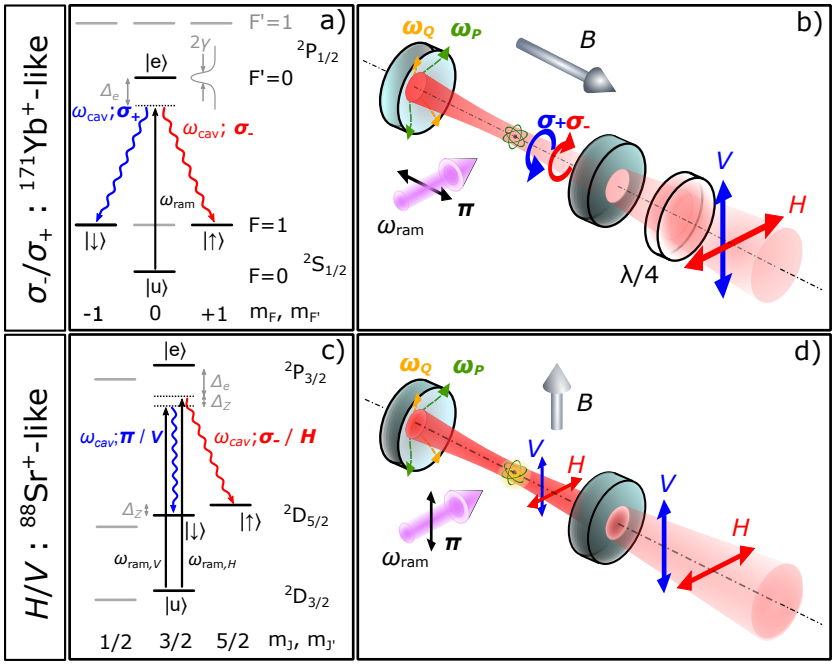


Figure 6.2: Schematic diagrams of the photon production geometries investigated. **(a)** Example level scheme for σ_-/σ_+ photon encoding, which may be implemented in isotopes with simple hyperfine structure e.g. $^{171}Yb^+$ and operated near Zeeman degeneracy, with a small field only to define the states. The ion is initialised in the ground state $|u\rangle$ and a single Raman laser at ω_{ram} drives simultaneous symmetric vSTIRAP transfers via $|e\rangle$ to the qubit states $|\uparrow\rangle$ / $|\downarrow\rangle$, accompanied by the emission of σ_-/σ_+ polarised photons. The cavity mode, in the absence of birefringence, has resonant frequency ω_{cav}, detuned by Δ_e from the qubit-excited state transition frequency. **(b)** The cavity lies parallel to the B-field to enhance the σ_-/σ_+ decays and forbid the π decay. If the cavity is birefringent, the cavity resonance ω_{cav} splits into two polarisation eigenmodes at frequencies ω_Q and ω_P. Upon exiting the cavity, the photon polarisations are transformed $\sigma_-/\sigma_+ \rightarrow H/V$ by a quarter wave plate. **(c)** Example scheme for H/V photon encoding, suitable for implementation in isotopes with a single valence electron and no nuclear spin e.g. $^{40}Ca^+$ or $^{88}Sr^+$. Lacking the level symmetry required for the σ_-/σ_+ scheme, the decay rates for this type of system are generally strongest and most balanced when transferring to the most extreme states of the lower manifold, here $|\uparrow\rangle = D_{5/2,+5/2}$ and $|\downarrow\rangle = D_{5/2,+3/2}$, accompanied by the emission of σ_- and π polarised photons respectively. Cavity decay may also occur via the σ_+ channel to the 'spectator' level $D_{5/2,+1/2}$, reducing the fidelity of the final state. To suppress this process, a substantial Zeeman splitting, Δ_Z, is applied. This requires the two vSTIRAP λ-systems to be addressed by dedicated frequency components $\omega_{\text{ram},H}/\omega_{\text{ram},V}$ of a bichromatic drive, with each λ system tuned to Raman resonance with the corresponding qubit state and cavity mode frequency ω_{cav}. **(d)** With the cavity axis orthogonal to the magnetic field, the σ_-/π decay channels project directly to H/V, while the difference in their angular emission patterns partially cancels the imbalance between of their Clebsch-Gordan coefficients.

to capture the ion cavity dynamics. Each node was modelled using just six states: Initial state, excited state and two qubit states accompanied with either cavity polarisation (H and V for the H/V scheme, and σ_- and σ_+ for the σ_-/σ_+ scheme) each. A non-Hermitian approach was used to integrate the equations of motion.

This simulation strategy ignored the effect of off-resonant excitation and spectator cavity decays, which would require a larger Hilbert space to model, and temporal mixing processes, as discussed in Chapter 3, which could be captured by a master equation approach. As concluded in Chapters 3 and 4, a suitable choice of driving levels renders the errors introduced by temporal mixing insignificant. Spectator decays can also be suppressed by using a large Zeeman splitting. Therefore the assumptions behind the model are valid in the case that these conditions are satisfied, but should also provide important insights even when not.

To capture the effects of cavity birefringence, both its magnitude (the splitting of the polarisation eigenmodes) and its 'orientation' (the polarisation of those eigenmodes), must be modelled. The orientation can be represented as a point on the Poincare sphere, which is shown in Fig. 6.3. This birefringence results in a Hamiltonian $H = \hbar\boldsymbol{\delta_B} \cdot \boldsymbol{\sigma}$ that acts on the polarisation subspace, where $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \sigma_3)$ is the standard vector of Pauli matrices, and $\boldsymbol{\delta_B}$ is the vector representing the birefringence on the Poincare sphere, pointing between the origin and the highest energy polarisation eigenmode. The length of this vector is $|\delta_B|$, meaning that the energy splitting between the two polarisation eigenmodes is $2\hbar|\delta_B|$ under this convention. Importantly, the basis in which $\boldsymbol{\sigma}$ is expressed impacts how a given physical birefringence $\boldsymbol{\delta_B}$ is translated into a matrix in the Hamiltonian, a flexibility that is used in these simulations. It should be noted that the physical birefringence, as expressed in terms of ψ and χ with respect to cardinal polarisations, is unaffected by the choice of basis.

To present a consistent Hamiltonian for the two schemes (σ_-/σ_+ and H/V), the flexibility in the choice of basis for $\boldsymbol{\sigma}$ is exploited to ensure that photons are always emitted into single basis states. For the H/V scheme, the directions $(\boldsymbol{e_1}, \boldsymbol{e_2}, \boldsymbol{e_3})$ that specify the qubit basis are $(\boldsymbol{e_+}, \boldsymbol{e}_{\mathrm{RCP}}, \boldsymbol{e_H})$ on the Poincare sphere, and for the σ_-/σ_+ scheme $(\boldsymbol{e_1}, \boldsymbol{e_2}, \boldsymbol{e_3}) = (\boldsymbol{e_H}, \boldsymbol{e_+}, \boldsymbol{e}_{\mathrm{RCP}})$, where $\boldsymbol{e_i}$ represents a unit vector in the given

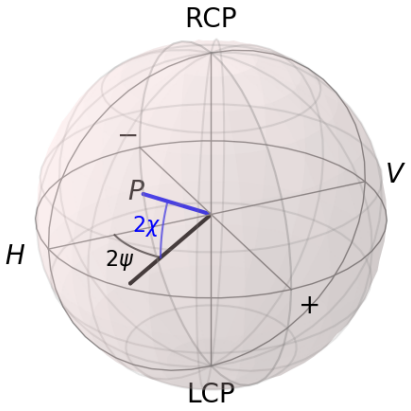


Figure 6.3: The Poincare sphere, including the standard angles ψ and χ. A general polarisation P is represented by a point on the surface of the sphere, with ψ and χ specifying the orientation of P with respect to the cardinal polarisations, where RCP and LCP are right and left circularly polarised, respectively. The propagation direction is taken so that that an atomic emission σ_- (σ_+) corresponds to RCP (LCP) in the σ_-/σ_+ scheme. Birefringence can be represented as a vector in this coordinate system, pointing towards the highest energy polarisation eigenmode, with a magnitude $|\delta_B|$ that is half the birefringence splitting. Photons in the cavity then evolve under the Hamiltonian $H = \hbar\boldsymbol{\delta_B}\cdot\boldsymbol{\sigma}$.

basis. In this basis $\boldsymbol{\delta_B}$ has components $(\delta_{B,1}, \delta_{B,2}, \delta_{B,3})$. The atomic states used were $|u\rangle$, a stable initial atomic state, $|e\rangle$, an excited atomic state, and the qubit states $|\downarrow\rangle$ and $|\uparrow\rangle$. The basis for each scheme is summarised in Tab. 6.1, where $|1_\uparrow\rangle$ ($|1_\downarrow\rangle$) is the scheme-dependent cavity state corresponding to a single occupation of the photon polarisation expected to correlate with the spin state $|\uparrow\rangle(|\downarrow\rangle)$

Significance of state	$\boldsymbol{\sigma_-/\sigma_+}$ **scheme state**	$\boldsymbol{H/V}$ **scheme state**
Initial state, $\|u\rangle\|0\rangle$	$\|u\rangle\|0\rangle$	$\|u\rangle\|0\rangle$
Excited state, $\|e\rangle\|0\rangle$	$\|e\rangle\|0\rangle$	$\|e\rangle\|0\rangle$
Qubit $\|\uparrow\rangle$, Exp. pol., $\|\uparrow\rangle\|1_\uparrow\rangle$	$a^\dagger_{\sigma_-,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$	$a^\dagger_{H,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$
Qubit $\|\uparrow\rangle$, Rot. pol., $\|\uparrow\rangle\|1_\downarrow\rangle$	$a^\dagger_{\sigma_+,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$	$a^\dagger_{V,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$
Qubit $\|\downarrow\rangle$, Exp. pol., $\|\downarrow\rangle\|1_\downarrow\rangle$	$a^\dagger_{\sigma_-,\mathrm{cav}}\|\downarrow\rangle\|0\rangle$	$a^\dagger_{H,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$
Qubit $\|\downarrow\rangle$, Rot. pol., $\|\downarrow\rangle\|1_\uparrow\rangle$	$a^\dagger_{\sigma_+,\mathrm{cav}}\|\downarrow\rangle\|0\rangle$	$a^\dagger_{V,\mathrm{cav}}\|\uparrow\rangle\|0\rangle$

Table 6.1: The basis used to simulate birefringent cavity dynamics and photon production in the $\boldsymbol{\sigma_-/\sigma_+}$ and $\boldsymbol{H/V}$ schemes, where 'Exp. pol.' stands for 'Expected polarisation', and 'Rot. pol.' for 'Rotated polarisation'. The basis used for the simulations is different for the two schemes so that the Hamiltonian that determines the dynamics is similar.

The Hamiltonian (written in the bases detailed in Tab. 6.1) is

$$\begin{aligned} H &= H_{\mathrm{I}} \otimes I_{\mathrm{P}} + I_{\mathrm{I}} \otimes H_{\mathrm{P}} + H_{\mathrm{C}}, \\ H_{\mathrm{I}} &= (\Delta_e - i\gamma)\,|e\rangle\langle e| + \frac{1}{2}\Delta_Z(|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|) + (-i\Omega(t)\,|u\rangle\langle e| + \text{h.c.}), \\ H_{\mathrm{P}} &= \delta_{B,3}(|1_\uparrow\rangle\langle 1_\uparrow| - |1_\downarrow\rangle\langle 1_\downarrow|) + ((\delta_{B,1} - i\delta_{B,2})\,|1_\uparrow\rangle\langle 1_\downarrow| + \text{h.c.}), \\ H_{\mathrm{C}} &= (ig_\uparrow\,|e\rangle\,|0\rangle\,\langle 1_\uparrow|\,\langle\uparrow| + ig_\downarrow\,|e\rangle\,|0\rangle\,\langle 1_\downarrow|\,\langle\downarrow| + \text{h.c.}), \end{aligned} \tag{6.2}$$

where H_{I} and I_{I} (H_{P} and I_{P}) are the subsystem Hamiltonian and identity for the ion (cavity photonic) Hilbert space, and H_{C} is the Hamiltonian for the ion-cavity coupling. This Hamiltonian is expressed in a frame rotating with the average energy of the cavity resonances. The initial state has zero detuning to fix the system at Raman resonance. The frequency splitting of the qubit states is Δ_Z and Δ_e is the frequency detuning of the Raman drive from the initial state-excited state transition frequency. $\Omega(t)$ is the time dependent coupling of the laser drive. In the σ_-/σ_+ scheme, where only one tone is applied due to the low qubit splitting, $\Omega(t)$ is linearly increasing and has a central frequency of zero to address qubit states equally. In the H/V scheme, the driving contains bichromatic components, with $\Omega(t) = (\Omega_\uparrow + \Omega_\downarrow)$ where $\Omega_{\uparrow,\downarrow} = |\Omega_{\uparrow,\downarrow}(t)|\exp(\pm i\Delta_Z t/2)$ are the time-dependent coupling strengths of the bichromatic components of the laser drive in cavity-Raman resonance with the $|\uparrow\rangle, |\downarrow\rangle$ states, and $g_{\uparrow,\downarrow}$ the respective cavity couplings. Following the convention of this document, $\gamma = \Gamma/2$ is the amplitude decay rate of the excited state due to spontaneous emission and κ is the cavity field amplitude decay rate.

As birefringence is taken here as an unexpected defect most likely with randomly varying magnitude and orientation across nodes, the drive frequencies are not updated in the presence of birefringence, and in this case both cavity eigenmodes will be symmetrically detuned from Raman resonance. Note also that if the two cavity couplings $g_\uparrow$ and $g_\downarrow$ are not equal (e.g. due to different Clebsch-Gordan coefficients), the laser drive strengths of the two frequency components are adjusted to equalise the probabilities of decay to either qubit state.

The Hamiltonian was integrated numerically using the Schrödinger equation $\partial\,|\Psi_N(t)\rangle\,/\partial t = -iH\,|\Psi_N(t)\rangle$ to calculate the quantum state $|\Psi_N(t)\rangle$ of each network node, where $N \in \{A, B\}$ labels the node. This state evolution was then converted into

photon wavepacket functions through the procedure detailed in Tab. 6.2, where $f^N_{s,p}(t)$ is the photonic wavepacket associated with ion state $s \in \{\uparrow,\downarrow\}$ and polarisation $p \in \{H,V\}$ emitted from node $N \in \{A,B\}$ at time t. These functions are not normalised, so $\sum_{N\in\{A,B\}}\sum_{s\in\{\uparrow,\downarrow\}}\sum_{p\in\{H,V\}}\int_0^\infty |f^N_{s,p}(t)|^2\,dt = P_R$, where P_R is the probability of any photon emission occurring during the photon generation process.

Output wavepacket	$\boldsymbol{\sigma_-/\sigma_+}$ **scheme**	$\boldsymbol{H/V}$ **scheme**
$f^A_{\uparrow,H}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{\sigma_-,cav}\,|\Psi_A(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{H,cav}\,|\Psi_A(t)\rangle$
$f^A_{\downarrow,V}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{\sigma_+,cav}\,|\Psi_A(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{V,cav}\,|\Psi_A(t)\rangle$
$f^A_{\uparrow,V}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{\sigma_+,cav}\,|\Psi_A(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{V,cav}\,|\Psi_A(t)\rangle$
$f^A_{\downarrow,H}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{\sigma_-,cav}\,|\Psi_A(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{H,cav}\,|\Psi_A(t)\rangle$
$f^B_{\uparrow,H}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{\sigma_-,cav}\,|\Psi_B(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{H,cav}\,|\Psi_B(t)\rangle$
$f^B_{\downarrow,V}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{\sigma_+,cav}\,|\Psi_B(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{V,cav}\,|\Psi_B(t)\rangle$
$f^B_{\uparrow,V}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{\sigma_+,cav}\,|\Psi_B(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\uparrow|\,a_{V,cav}\,|\Psi_B(t)\rangle$
$f^B_{\downarrow,H}(t)$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{\sigma_-,cav}\,|\Psi_B(t)\rangle$	$\sqrt{2\kappa}\,\langle 0|\langle\downarrow|\,a_{H,cav}\,|\Psi_B(t)\rangle$

Table 6.2: Conversion between node states $|\Psi_A(t)\rangle$ and $|\Psi_B(t)\rangle$ and the photon wavepackets $f^N_{s,p}$ analysed by the photonic measurement.

6.3.3 Photonic Measurement

The wavepackets $f^N_{s,p}$ yield a state before photonic detection of

$$|\chi\rangle = \left[\left(\sum_{s\in\{\uparrow,\downarrow\}}\sum_{p\in\{H,V\}}\left(\int f^A_{s,p}(t)a^\dagger_p(t)\,dt\right)|s\rangle\right) \otimes \left(\sum_{s'\in\{\uparrow,\downarrow\}}\sum_{p'\in\{H,V\}}\left(\int f^B_{s',p'}(t')b^\dagger_{p'}(t')\,dt'\right)|s'\rangle\right)\right]|0\rangle\,, \tag{6.3}$$

where $a^\dagger_p(t)$ ($b^\dagger_p(t)$) is the creation operator for a photon with polarisation p in port a (b) of the Bell state analyser at time t (featured in Fig. 6.1). The Hilbert space used is the tensor product of the qubit in node A, the qubit in node B and the photonic modes.

The beam splitter transforms the input creation operators $a^\dagger_p(t)$ and $b^\dagger_p(t)$ to the output port creation operators $c^\dagger_p(t)$ and $d^\dagger_p(t)$ according to the beamsplitter transformation

$$\begin{pmatrix} a^\dagger_p(t) \\ b^\dagger_p(t)\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & -1\\ 1 & 1\end{pmatrix}\begin{pmatrix} c^\dagger_p(t) \\ d^\dagger_p(t)\end{pmatrix}. \tag{6.4}$$

Success of the remote entanglement procedure is signalled by a two-photon detection 'herald' with H-polarised photon detected at time t_H and V-polarised photon detected at time t_V, If, for concreteness, both photon detections occur at output node c of the beamsplitter, the resulting ion-ion state

$$\begin{aligned}|\psi_{cc}(t_H,t_V)\rangle &= \langle 0|\, c_H(t_H)c_V(t_V)\,|\chi\rangle\,,\\ &= \frac{1}{2}\sum_{s\in\{\uparrow,\downarrow\}}\sum_{s'\in\{\uparrow,\downarrow\}}\left\{f^A_{s,H}(t_H)f^B_{s',V}(t_V)+f^A_{s,V}(t_V)f^B_{s',H}(t_H)\right\}|s\rangle\,|s'\rangle\,,\end{aligned} \tag{6.5}$$

is produced. The fidelity of this state with the desired Bell state $\left|\Psi^+_{\rm Bell}\right\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)$ is given by

$$\mathcal{F}_{cc}(t_H,t_V) = \frac{\left|\left\langle\Psi^+_{\rm Bell}\middle|\psi_{cc}(t_H,t_V)\right\rangle\right|^2}{|\langle\psi_{cc}(t_H,t_V)|\psi_{cc}(t_H,t_V)\rangle|}. \tag{6.6}$$

Corresponding relations holds for fidelity $\mathcal{F}_{dd}$ with $\left|\Psi^+_{\rm Bell}\right\rangle$, while $\mathcal{F}_{cd}$ and $\mathcal{F}_{dc}$ express fidelity with $\left|\Psi^-_{\rm Bell}\right\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$.

First, consider the fidelity produced in a system with no birefringence. In these systems, there are no polarisation oscillations, so $f^N_{\uparrow,V}(t) = f^N_{\downarrow,H}(t) = 0$ for $N \in \{A,B\}$ and the polarisation is always correlated with the ion state as expected. In this case, $\mathcal{F}_{cc}(t_H,t_V)$ takes on a simplified form

$$\mathcal{F}_{cc}(t_H,t_V)|_{\boldsymbol{\delta}^A_B=\boldsymbol{\delta}^B_B=\mathbf{0}} = \frac{1}{2}\left(1+\frac{f^A_{\uparrow,H}(t_H)f^B_{\downarrow,V}(t_V)f^{A*}_{\downarrow,V}(t_V)f^{B*}_{\uparrow,H}(t_H)+\text{c.c.}}{\left|f^A_{\uparrow,H}(t_H)f^B_{\downarrow,V}(t_V)\right|^2+\left|f^A_{\downarrow,V}(t_V)f^B_{\uparrow,H}(t_H)\right|^2}\right), \tag{6.7}$$

where $\boldsymbol{\delta}^A_B$ ($\boldsymbol{\delta}^B_B$) is the birefringence vector of node A (B). The fidelity is unity if $f^A_{\uparrow,H}(t_H) = f^B_{\uparrow,H}(t_H)$ and $f^A_{\downarrow,V}(t_V) = f^B_{\downarrow,V}(t_V)$, but reduces if the wavepackets from the nodes are distinguishable.

Now introduce birefringence, but just to node B, as depicted in Fig. 6.4. In this case $f^A_{\uparrow,V}(t) = f^A_{\downarrow,H}(t) = 0$, but no corresponding relation holds for node B. The ion-ion state produced upon the detection of two photons at node C is

$$\begin{aligned}|\psi_{cc}(t_H,t_V)\rangle\,|_{\boldsymbol{\delta}^A_B=\mathbf{0}} =& \frac{1}{2}\left[f^A_{\uparrow,H}(t_H)\,|\uparrow\rangle\,(f^B_{\downarrow,V}(t_V)\,|\downarrow\rangle + f^B_{\uparrow,V}(t_V)\,|\uparrow\rangle)+\right.\\ &\left. f^A_{\downarrow,V}(t_V)\,|\downarrow\rangle\,(f^B_{\uparrow,H}(t_H)\,|\uparrow\rangle + f^B_{\downarrow,H}(t_H)\,|\downarrow\rangle)\right],\end{aligned} \tag{6.8}$$

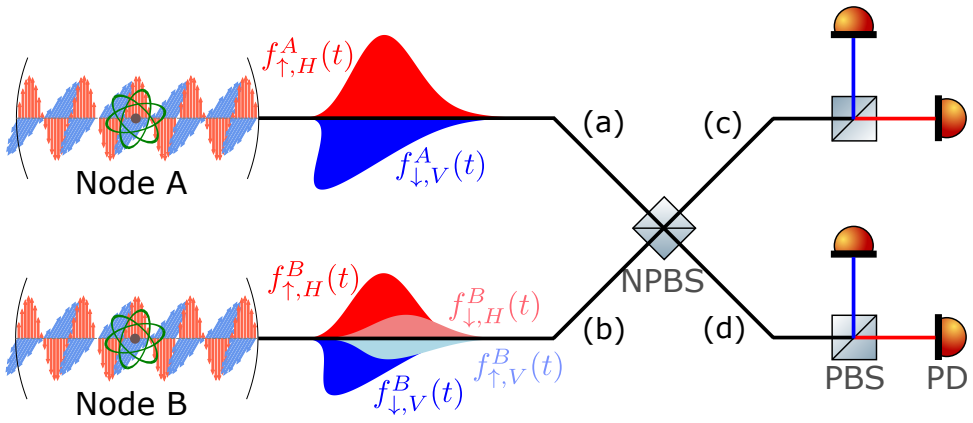


Figure 6.4: Diagram of polarisation-encoded remote entanglement in the case where node B is birefringent. Cavity birefringence distorts the wavepackets from node B (as compared to the non-birefringent node A), and introduces components where the ion spin and photon polarisations are not correlated as intended. The H-polarised wavepackets (vertical, above the line, in red and pink) follow the red path to the detectors, and the V-polarised wavepackets (diagonal, below the line, in blue and light blue) follow the blue paths.

which shows how the resultant state is distinct from the desired Bell state $\left|\Psi^+_{\text{Bell}}\right\rangle$ due to the contributions of $f^B_{\uparrow,V}(t_V)$ and $f^B_{\downarrow,H}(t_V)$ that arise from polarisation rotation within the cavity. It appears that the post-herald ion state in node B has been transformed according to

$$\begin{aligned}
|\downarrow\rangle &\rightarrow \frac{\left(f^B_{\downarrow,V}(t_V)\,|\downarrow\rangle + f^B_{\uparrow,V}(t_V)\,|\uparrow\rangle\right)}{\sqrt{\left|f^B_{\downarrow,V}(t_V)\right|^2 + \left|f^B_{\uparrow,V}(t_V)\right|^2}} \equiv |\downarrow'\rangle \\
|\uparrow\rangle &\rightarrow \frac{\left(f^B_{\uparrow,H}(t_H)\,|\uparrow\rangle + f^B_{\downarrow,H}(t_H)\,|\downarrow\rangle\right)}{\sqrt{\left|f^B_{\uparrow,H}(t_H)\right|^2 + \left|f^B_{\downarrow,H}(t_H)\right|^2}} \equiv |\uparrow'\rangle\,,
\end{aligned} \tag{6.9}$$

but it is important to recognise that this transformation is only unitary if $\langle\uparrow'|\downarrow'\rangle = 0$, which is not guaranteed. This effective non-unitary transformation will be seen to complicate efforts to mitigate the impact of birefringence in Sec. 6.5.

The fidelity with the desired Bell state is given by

$$\begin{aligned}
\mathcal{F}_{cc}(t_H,t_V)|_{\boldsymbol{\delta}^A_B=\mathbf{0}} = \frac{1}{2}\Bigg(1 + &\frac{f^A_{\uparrow,H}(t_H)f^B_{\downarrow,V}(t_V)f^{A*}_{\downarrow,V}(t_V)f^{B*}_{\uparrow,H}(t_H) + \text{c.c.}}{|\langle\psi_{cc}(t_H,t_V)|\psi_{cc}(t_H,t_V)\rangle|} \\
&- \frac{\left|f^A_{\uparrow,H}(t_H)f^B_{\uparrow,V}(t_V)\right|^2 + \left|f^A_{\downarrow,V}(t_V)f^B_{\downarrow,H}(t_V)\right|^2}{|\langle\psi_{cc}(t_H,t_V)|\psi_{cc}(t_H,t_V)\rangle|}\Bigg)\,.
\end{aligned} \tag{6.10}$$

The expression is very similar to the case without birefringence (Eq. 6.7), but there is an extra term with an explicit reduction in fidelity due to the rotated components of the polarisation from the birefringent node. If both nodes are birefringent, even more complicated expressions result, but their analytical forms are not highly revealing.

6.4 Correcting for Birefringent Effects

As seen in Sec. 6.3, cavity birefringence in a network node can lead to a temporal variation of the output photon polarisation, which reduces the fidelity of the remote entanglement protocol. Even if the birefringence cannot be made negligible at source, there is potential for this fidelity loss to be reduced, and three methods for doing so are considered. Firstly by attempting to 'undo' the polarisation oscillation directly via an active correction of the photonic state, secondly through post-selection of the the subset of heralds with photon detection times expected to yield the highest fidelity entangled pairs, and thirdly though post-detection correction of the ion-ion state, via a local spin-qubit unitary determined by the photon arrival times. The basic principle of these schemes is depicted in Fig. 6.5 and each will be discussed in turn for the remainder of this section.

6.4.1 Pre-herald Photonic Qubit Correction

The most physically direct approach to reducing the fidelity loss caused by the distortion of the emitted wavepackets is to restore those wavepackets before the interference-based measurement with a time-dependent linear optical retarder (such as a Pockels cell) between the node and the non-polarising beamsplitter. The distortions to the photonic waveforms introduced by birefringence can be categorised in two parts, which will be discussed in turn.

The first class of photonic distortion is the introduction of time-dependence to the photon polarisation correlated with each atomic state. In the case of zero birefringence, the qubit spin state $|\uparrow\rangle$ ($|\downarrow\rangle$) is correlated with a single node output polarisation H (V) at all times. A birefringent node produces a much more complicated state with spin-polarisation correlations that evolve in time, as expressed in Tab. 6.2. To recover the ideal ion spin-photon correlation, a transformation that rotates the polarisation

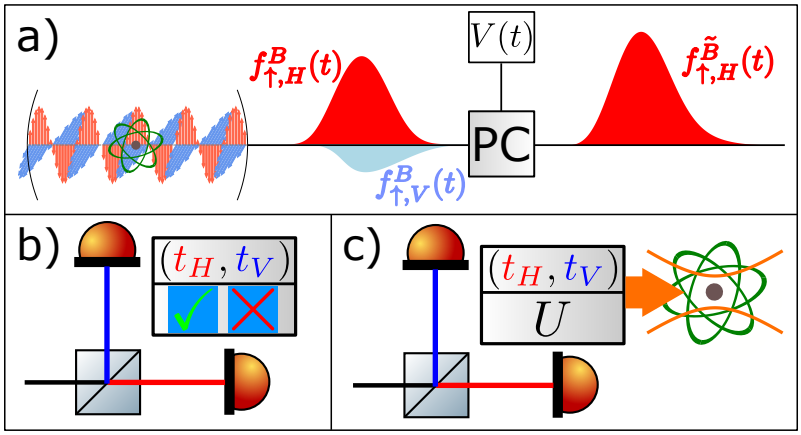


Figure 6.5: Schematic of the birefringence correction schemes discussed. a) Photonic corrections: A Pockels cell with time dependent voltage is used to restore the spin photon correlations (here shown for the case of spin $|\uparrow\rangle$). For the corrections modelled in this section, the Pockels cell has a fixed orientation and the voltage $V(t)$ is linear. b) and c) Post-detection corrections after a two photon detection event of H-polarised photon at t_H and V-polarised photon at t_V. b) Windowing: The detection times t_H and t_V are used to decide if the trial will be considered a success. If the expected fidelity of the event is low, the event will not be considered a success, which raises average fidelity at the expense of success probability. c) Spin correction: Based upon detection times t_H and t_V, an arbitrary unitary is applied to the ion in the birefringent node to restore the fidelity as well as possible.

component correlated with $|\uparrow\rangle$ to H, and that correlated with $|\downarrow\rangle$ to V, at all times t is required. For a birefringent node B this may be written

$$\begin{aligned} f^{B}_{\uparrow,H}(t)b^{\dagger}_{H}(t) + f^{B}_{\uparrow,V}(t)b^{\dagger}_{V}(t) &\rightarrow f^{\tilde{B}}_{\uparrow,H}(t)b^{\dagger}_{H}(t), \\ f^{B}_{\downarrow,H}(t)b^{\dagger}_{H}(t) + f^{B}_{\downarrow,V}(t)b^{\dagger}_{V}(t) &\rightarrow f^{\tilde{B}}_{\downarrow,V}(t)b^{\dagger}_{V}(t), \end{aligned} \tag{6.11}$$

where $f^{\tilde{B}}_{\uparrow,H}(t)$ and $f^{\tilde{B}}_{\downarrow,V}(t)$ are arbitrary functions of time[1]. Such a polarisation transformation is only possible via a unitary operation if the Jones vectors for the photonic components correlated with $|\uparrow\rangle$ and $|\downarrow\rangle$ remain mutually orthogonal at all times; explicitly

$$\begin{pmatrix} f^{B}_{\uparrow,H}(t) & f^{B}_{\uparrow,V}(t) \end{pmatrix} \cdot \begin{pmatrix} f^{B}_{\downarrow,H}(t) \\ f^{B}_{\downarrow,V}(t) \end{pmatrix} = 0; \qquad \forall t > 0, \tag{6.12}$$

a condition which only holds for certain special cases, such as the complete symmetry of the emitter-cavity system with respect to the permutation of $\{|H\uparrow\rangle, |V\downarrow\rangle\}$. For cases and times when this orthogonality does not hold, no unitary operation on

[1]Note that in [136] (the publication on which this chapter is based), similar notation is used for an unrelated purpose.

the node output can simultaneously undo the polarisation transformation of both $|\uparrow\rangle$-correlated and $|\downarrow\rangle$-correlated components. Nonetheless, in practice (for example the cases discussed in Sec. 6.5.3) photonic corrections can restore the correlation of the photon polarisation and qubit state effectively, even if they cannot do so perfectly.

The second photonic impact of birefringence is to modify the time dependence of the overall emitted wavepacket amplitude. The shape of this wavepacket is governed by the coupled interactions of driving laser, ion, and cavity field, and when the polarisation of the intra-cavity field rotates due to birefringence, the cavity-field/ion coupling is modified, altering the system dynamics [209]. The result is that even if the spin-polarisation entanglement could be fully restored, the wavepackets emitted from nodes with distinct degrees of birefringence will still differ, because a unitary photonic correction cannot change the overall amplitude of the wavepacket at any given time. This inherent distinguishability between the photonic wavepackets from the two nodes is, for the cases we consider, the principal limit to the fidelity that may be recovered via photonic correction.

For the corrections presented in Sec. 6.5, photonic corrections are implemented with a device, for example a Pockels cell, driven such that the retardance varies linearly in time. The cell orientation, initial retardance and retardance ramp rate are optimised such that the integrated probabilities of the 'anticorrelated' photon wavepackets $f^N_{\uparrow,V}(t)$ and $f^N_{\downarrow,H}(t)$ over all time for node N are minimised. This is not the highest performance photonic correction scheme possible, but has the considerable advantage that the photonic correction could be calibrated through the emission of the corrected node alone, and thus the simple three-parameter optimisation need only be done once per network node, and requires only single photon detection events.

Higher fidelities could be achieved using more complex retardance profiles or multiple Pockels cells at different orientations. One could also optimise on the two-node fidelity for higher performance, but the need to calibrate on two-photon coincidences for every pairwise combination of nodes in the network would drastically increase the calibration time required. In simulations, the improvement seen from using these more complicated photonic corrections was generally negligible compared to the simpler approach selected.

6.4.2 Windowed Post-Selection

In the two-photon polarisation-encoded remote entanglement protocol, repeated entanglement attempts are made until a suitable 'H–V' two-photon herald, after which the ions are deemed entangled. The fidelity of a herald is a function of detection times t_H and t_V (as seen in Eq. 6.10). By deeming a trial successful only if the fidelity at those detection times exceeds a certain threshold (Fig. 6.5(b)), the average fidelity of successful trials will improve at the expense of rate. The efficacy of this approach depends upon the accuracy with which the detection times can be measured. Windowing may be used in conjunction with the other correction methods in this section; in these cases the fidelities *after* application of those corrections are used to determine the windowing regions.

6.4.3 Spin-Qubit Rotation

The reduction in remote entanglement fidelity arises because the produced ion-ion state does not match the target state. The fidelity of the process might therefore be improved post-detection through local unitaries applied to the ion-spin qubits (Fig. 6.5(c)). As shown in Eq. 6.9, cavity birefringence causes more than a trivial unitary rotation of a qubit in the birefringent node. This means that the two-qubit state $|\psi_{cc}(t_H, t_V)\rangle$ is not generally fully entangled for all t_H and t_V, and no local operations on the qubits post-herald could completely restore the fidelity, however, they can significantly improve it. Because the target state is fully entangled, local corrections need only be applied to one qubit, and the optimal unitary is uniquely determinable (see discussion in App. F). However, the optimum local correction unitary depends upon the detection times t_H and t_V, and, to implement, these corrections would have to be calibrated for all pairs of nodes, detector combinations and photon arrival times.

6.5 Effect of Cavity Birefringence on Remote Entanglement Fidelity

6.5.1 Analytical suggestions

The equations of motion of the general system are far too complicated to offer analytical solutions in all but the most simplified cases. Nevertheless, even highly simplified models may be able to reveal general trends in the expected fidelity of remote entanglement with birefringent systems.

For these models, it will be assumed that $g_1 = g_2 = g$ to increase the symmetry of the system, $\Delta_Z = 0$ to remove off-resonant effects, $\boldsymbol{\delta_B} = (\delta_{B,1}, 0, 0)$ for node B so that birefringence causes only polarisation oscillations, and that node A has zero birefringence. For concreteness, the states considered will be the ones simulated in the H/V scheme.

The first simplified model assumes that both nodes are initialised in the desired superposition of $(1/\sqrt{2})(a^\dagger_{H,\text{cav}}\left|\uparrow\right\rangle + a^\dagger_{V,\text{cav}}\left|\downarrow\right\rangle)\left|0\right\rangle$, and that the ion-cavity coupling immediately reduces to zero after this initialisation (for example by moving the ion to a node of the standing wave). While crude, these assumptions allow for the output wavepackets to be expressed analytically

$$f^B_{\uparrow,H}(t) = f^B_{\downarrow,V}(t) = \sqrt{\kappa}\cos(\delta_{B,1}t)e^{-\kappa t}, \quad f^B_{\downarrow,H}(t) = f^B_{\uparrow,V}(t) = -i\sqrt{\kappa}\sin(\delta_{B,1}t)e^{-\kappa t}. \tag{6.13}$$

This leads to an average fidelity after, for example, a $\{(\text{c}),(\text{c})\}$ herald

$$\langle \mathcal{F}_{cc}(t_H, t_V)\rangle_{\text{heralds}} = \left(\frac{1}{1+\frac{\delta^2_{B,1}}{4\kappa^2}}\right)^2 \approx 1 - \frac{1}{2}\left(\frac{\delta_{B,1}}{\kappa}\right)^2. \tag{6.14}$$

The second simplified model corresponds to a slow driving of the system. Two of the coupled Hamiltonian equations of motion can be expressed

$$\begin{aligned} \frac{d\beta_{H,\uparrow}}{dt} &= -g\beta_e - \kappa\beta_{H,\uparrow} - i\delta_{B,1}\beta_{V,\uparrow}, \\ \frac{d\beta_{V,\uparrow}}{dt} &= -\kappa\beta_{V,\uparrow} - i\delta_{B,1}\beta_{H,\uparrow}, \end{aligned} \tag{6.15}$$

where $\beta_{s,p}(t) = \langle 0|\langle s| a_{p,cav}\left|\Psi_B(t)\right\rangle$ for $p \in \{H,V\}$, $s \in \{\uparrow,\downarrow\}$, and $\beta_e(t) = \langle 0|\langle e|\Psi_B(t)\rangle$. Equivalent equations of motion hold for $\beta_{V,\downarrow}$ and $\beta_{H,\downarrow}$ respectively.

If the system is driven slower than all rates in the Hamiltonian, the wavefunction coefficients can be assumed independent of time. Doing this leaves the following relation,

$$\beta_{V,\uparrow} = -i\frac{\delta_{B,1}}{\kappa}\beta_{H,\uparrow}. \tag{6.16}$$

This results a constant remote entanglement fidelity

$$\mathcal{F}_{cc}(t_H, t_V) = \frac{1}{1+\frac{\delta_{B,1}^2}{\kappa^2}} \approx 1 - \left(\frac{\delta_{B,1}}{\kappa}\right)^2. \tag{6.17}$$

The main purpose of these models is not to provide an accurate description of the physics, but to show that the ratio of birefringence to cavity decay rate appears crucial for determining the fidelity. This ratio sets the scale of the infidelity for both simplified models. The first model was designed to approximate a fast photon generation process that placed the photon into the cavity immediately, whereas the second model is appropriate for adiabatic photon generation. Despite the radically different timescales of photon emission, the ratio of birefringence to cavity decay rate determines the average fidelity loss, which is subtly different from previous suggestions that it is ratio of birefringence to inverse temporal length of the photon that determines the infidelity [258] (although, of course, the inverse temporal length of the photon will tend to correlate with the cavity decay rate). Anecdotally, it was found that this ratio produced good estimates for the scale of the average infidelity even for the more complicated situations simulated with asymmetric couplings, non-zero Zeeman splittings and variable orientation. For this reason the birefringence will always be quoted in units of κ for the numerical investigation.

6.5.2 Structure of Numerical Investigations

To calculate the impact of birefringence numerically, the remote entanglement protocol was simulated as discussed in Sec. 6.3 to calculate the photon temporal wavepackets emerging from each node, and thus the state $|\chi\rangle$ of the joint ion-photon system before measurement. The post measurement two-qubit states produced as a result of detection events at times t_H and t_V were calculated, and the fidelity of these states with the target Bell state determined, as given in Eq. 6.6. Finally, each of the mitigation strategies described in Sec. 6.4 was assessed.

First a selection of concrete examples illustrating the effect of cavity birefringence on the photonic wavepackets and the resultant entanglement fidelity as a function of the detection times were investigated to understand the general phenomena introduced by birefringence. The corrections that should be applied to restore the fidelity were then calculated and assessed. Subsequently, the dependence of achievable fidelity on birefringence orientation and magnitude was studied systematically. It was found that the choice of encoding, birefringence magnitude, and birefringence orientation are all important determinants of the impact of birefringence upon the achievable fidelity pre- and post-correction, as well as the difficulty of applying corrections.

6.5.3 Illustrative Examples

The first illustrative examples study the effect of birefringence on the output wavepackets of nodes and the corrections that should be applied for a range of specific cases. Considering the σ_-/σ_+ and H/V schemes in turn, in each case one node has zero birefringence while the other experiences a birefringence of magnitude $|\delta_B|$ whose principal axes lie at an angle of $\pi/4$ to the photon emission axes on the Poincare sphere. This intermediate angle is chosen for the presented case as it results in general dynamics that combine the effects of birefringence as diagonal and off-diagonal components in the simulated Hamiltonian.

The results of this initial analysis are shown in Figs. 6.6 and 6.7, which will be discussed in turn. Fig.6.6 shows the ion-photon wavepackets produced by the node before **(a)** and after **(b)** photonic corrections, and indicates the optimal photonic **(c)** and spin **(d, e)** corrections. As expected the wavepackets from a birefringent node exhibit a time-dependent polarisation, with the 'rotated' component of the polarisation larger for higher birefringence. The wavepacket shape in the H/V scheme is more complicated than for the σ_-/σ_+ scheme because of the bichromatic drive necessitated by the Zeeman splitting.

Upon application of photonic corrections, though the correlation between photon polarisation and ion spin state is almost perfectly restored, the shape of the wavepacket from the birefringent node remains distorted due to the effect of birefringence on the photon generation process, as discussed in Sec. 6.4.1. The

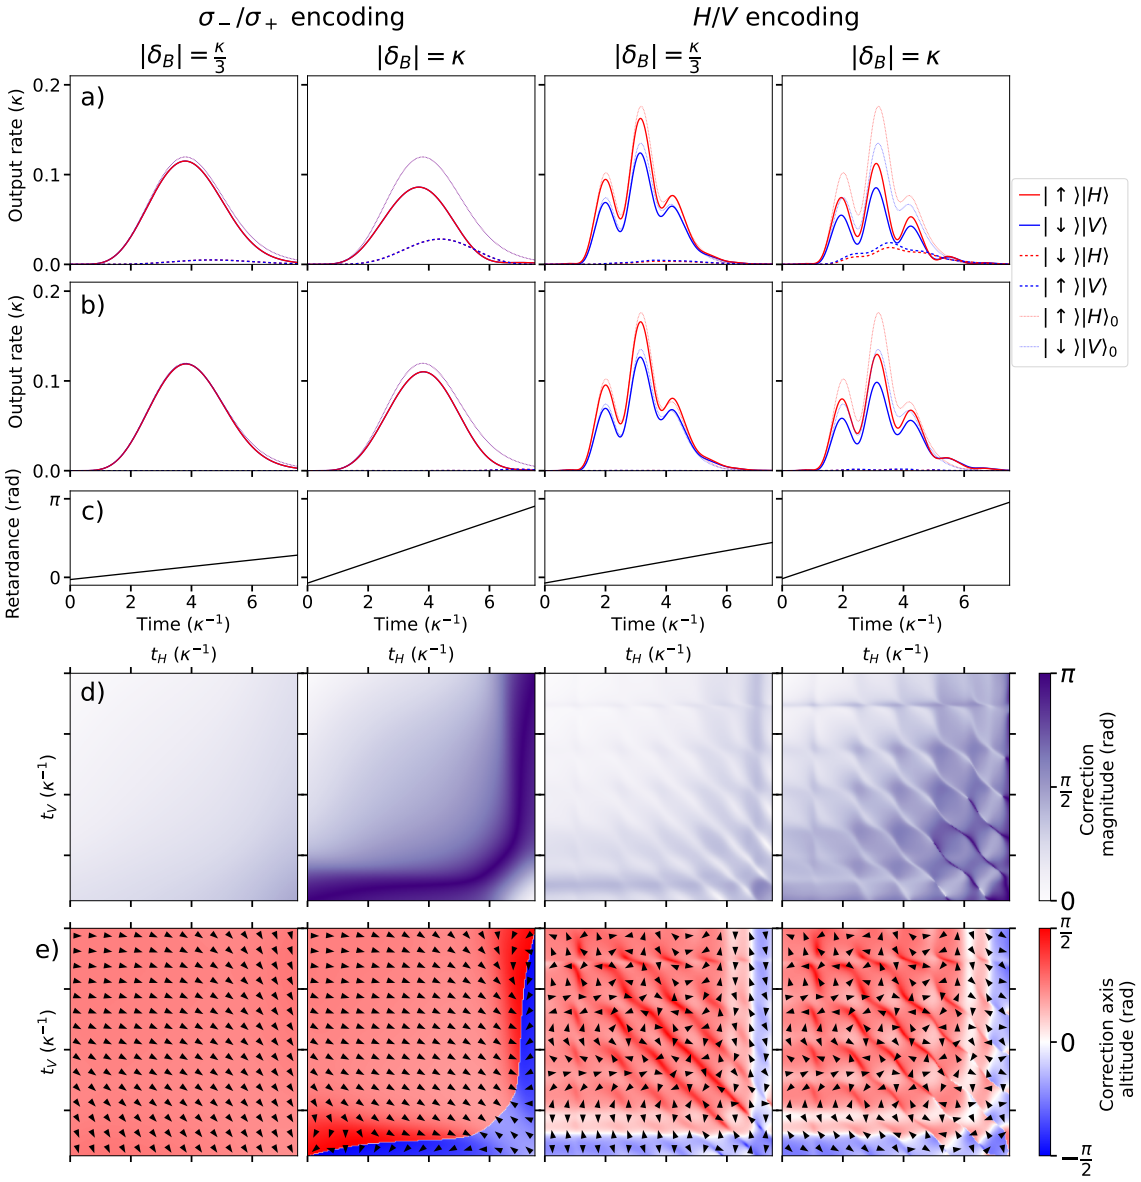


Figure 6.6: Examples of node output wavepackets and pre-herald photonic or post-herald spin corrections for the σ_-/σ_+ (Columns 1 & 2) and H/V (Columns 3 & 4) systems with weaker or stronger birefringence (left and right of each set of two columns respectively). In each case, one node has zero birefringence while the other has birefringence of labelled magnitude $|\delta_B|$, applied at an angle of $\pi/4$ to the atomic emission axes on the Poincare sphere. Rows **(a)-(e)** compare the wavepackets and the possible correction operations. **(a)** Pre-correction output of the birefringent node. The solid red and blue lines show the output rates of the atom-photon states correlated as expected in the absence of birefringence; the heavy dashed lines show the output rates of the opposite correlations. The photon polarisation is seen to vary during the production process, with the 'rotated' components more prevalent later in the production, and stronger for higher birefringence. The faint dotted lines indicate the output of the non-birefringent node. *Caption continued next page.*

Continued from previous page. **(b)** Output of the birefringent node upon photonic correction; the spin-polarisation correlation is near-perfectly restored but the wavepacket profile deviates from the non-birefringent case. **(c)** Linearly varying retardance as a function of time for the Pockels cell implementing the corrections, with retardance ramp-rate and fast-axis orientation chosen to optimally restore the expected photon correlations (Sec. 6.4.1). **(d)** Magnitude of optimal single spin-qubit unitary correction (see Sec. 6.4.3), applied to the ion in the birefringent node after a $\{(c),(c)\}$ herald leaves the system in the imperfect state given by Eq. 6.10. **(e)** Corresponding orientation of the unitary rotation axis on the Bloch sphere, with the arrows and heatmap indicating the azimuthal and altitudinal angles respectively.

retardance slew rates required for the presented photonic corrections are readily achievable for κ typical of ion-cavity experiments, and enacting the simple three-parameter node calibration necessary for these corrections may be worthwhile given the fidelity gains this confers.

For post-herald spin-corrections, the optimal unitary parameter heatmaps show that these corrections vary with both herald detection times in magnitude and orientation on the Bloch sphere. The correction magnitudes generally grow with increasing $|\delta_B|$ and the magnitude and correction axis orientation vary more severely in the H/V scheme. These post-herald spin corrections would be much costlier to implement in a practical quantum network than photonic corrections would be, requiring separate calibration of the unitary for each pair of detectors and pair of nodes over all possible (t_H, t_V) pairs. For the σ_-/σ_+ scheme, the correction angle and magnitude typically vary slowly with the two detection times, and therefore the correcting gates would not need to be calibrated over a fine mesh of herald times. For the H/V case however, the rapid variation in magnitude and orientation of the correcting unitary with detection times (t_H, t_V) means a very fine-grained calibration would be necessary, rendering this correction method all but impractical for systems using a bichromatic drive.

Fig. 6.7 shows how remote entanglement fidelity is impacted by birefringence before and after photonic correction, spin-qubit correction, and herald windowing. Both with or without corrections, the fidelity is a strong function of the herald times with the highest fidelities generally found at low (t_H, t_V), where the impact of birefringence remains small. Photonic or spin corrections yield higher fidelities over a larger range of detection times, and are generally more effective for the σ_-/σ_+

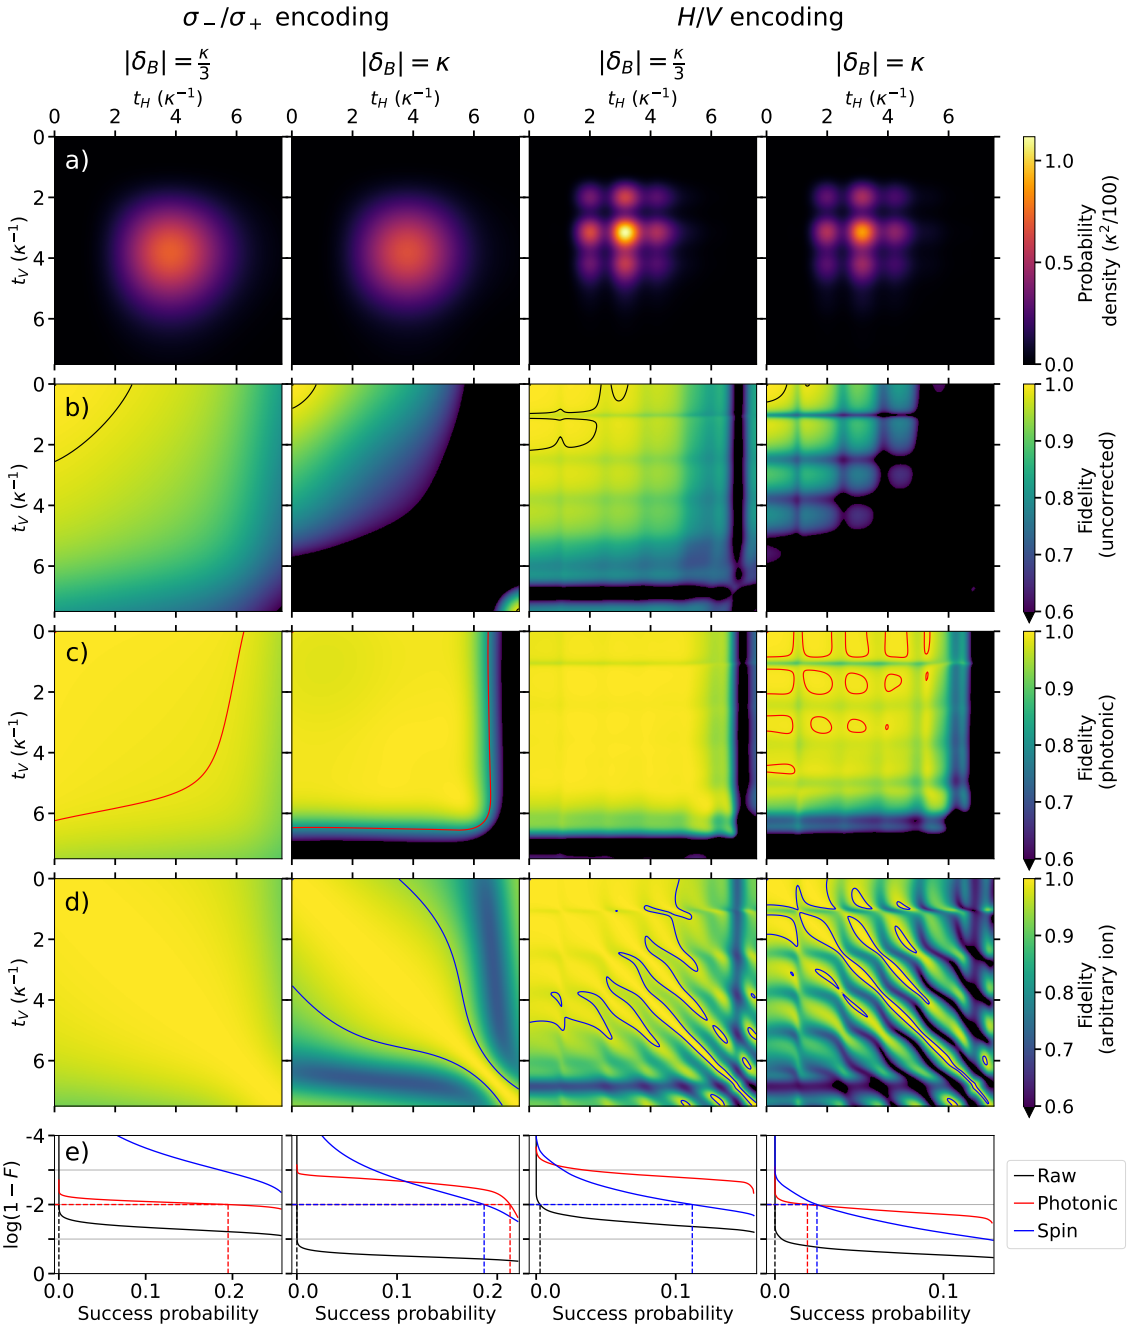


Figure 6.7: Comparison of fidelities achievable for the correction schemes shown in Fig. 6.6, where the significance of the columns is equivalent. All heatmaps display data for {(c), (c)} heralds only, but Row **(e)** summarises performance across all possible heralds. **(a)** Herald probability density against t_H, t_V, indicating the combinations of detection times with the most significant contribution to the entanglement success rate. Note that this probability density will be slightly altered by the photonic corrections. *Caption continued next page.*

Caption continued from previous page. Rows **(b)-(d)** show fidelity as a function of t_H and t_V for different correction schemes. The contours show the window regions that yield an average 99% entanglement fidelity with optimum success probability; if no windowing is required to achieve this average fidelity, no contour is drawn. **(b)** Raw, uncorrected fidelity. **(c)** Fidelity after photonic corrections. **(d)** Fidelity after arbitrary post-measurement single-spin correction. Row **(e)** details the fidelity/rate trade-off provided by post-herald detection-time windowing alone or in combination with pre- or post-herald corrections. Maximum success probability and minimum fidelity is achieved when no windowing is applied. Increasingly stringent selectivity of high-fidelity regions increases fidelity at the expense of success probability. The dotted lines, where visible, represent the 99% fidelity intercept, showing the severity of windowing required to reach this fidelity.

scheme, although no class of correction fully restores fidelities at all times. For photonic corrections, fidelity is primarily limited not by the residual rotation of polarisation components (which is strongly suppressed, as seen in Fig. 6.6), but by distortion in the wavepacket shape; the resultant distinguishability leads to a post-herald ion-ion state that is not maximally entangled. For the spin qubit corrections, unit fidelity is not achieved because the post-herald ion-ion state is not in general maximally entangled as a consequence of the birefringent dynamics. The local ion spin corrections cannot increase the degree of entanglement of the ion-ion state, but can nevertheless often notably increase the fidelity. It should be noted that the data presented in the heat maps is taken solely from $\{(\mathrm{c}),(\mathrm{c})\}$ heralds. In general, the herald probability and fidelities as a function of (t_H, t_V) do depend subtly on the heralding detectors, although this is not usually a strong effect.

To be practically useful, the fidelity of remote entanglement may need to exceed a certain threshold (e.g. for error correction or efficient distillation [94, 263]). If the requisite fidelity cannot be achieved over the average of all detection times, 'windowing' (Sec. 6.4.2) can serve to trade entanglement rate for fidelity by rejecting two-photon coincidences that would produce low-fidelity entanglement. High fidelities (contours of $\mathcal{F} = 0.99$ are indicated) can be achieved with this method even in the presence of strong birefringence, although the entanglement rate is drastically reduced for high birefringence. Active correction schemes increase the acceptable window region and lessen the reduction in rate. Nevertheless, if the birefringence is too high, no windowing strategy can produce entanglement with useful fidelity and rate, even when utilising correction methods.

6.5.4 Dependence on Birefringence Orientation

The fidelity obtained with and without correction depends not only on the magnitude of the birefringence but also its orientation with respect to the atomic emission polarisations. This is shown in Fig. 6.8, which gives the average fidelity obtained without windowing for different correction schemes for a birefringence of constant magnitude $|\delta_B| = \kappa$, where the orientation of the birefringence is varied to explore the full range of projections of the birefringence onto the atomic emission axes for a given birefringence magnitude. This exhibits the full range of possible fidelities to make the results complete. Because of this, the birefringence orientations examined differ for the two schemes. In the case of the σ_-/σ_+ scheme, the ellipticity of the birefringence is varied, with $\chi = \pm\pi/4$ indicating circular birefringence and $\chi = 0$ linear birefringence; the tilt angle of the ellipticity is uninteresting in the case of circularly polarised atomic decay, and is thus fixed at $\psi = 0$. In the case of the H/V scheme, the range of behaviour is observed through applying a linear birefringence ($\chi = 0$) and varying the angle of the highest energy cavity polarisation eigenmodes from H polarised to + polarised to V polarised. Additional data is also shown for the case of ion spin unitary operations restricted to rotations about the z-axis of the Bloch sphere, or about an optimally selected axis in the equatorial plane of the Bloch sphere, in order to understand how birefringence angle affects the type of spin correction that is useful.

The first observation is that the fidelity varies strongly with the angle of the uncontrolled birefringence. Secondly, the type of post-measurement spin correction that is useful also varies with angle. When the birefringence is aligned to the atomic emission axes, the effect is purely an energy shift for the polarisations correlated with the qubit states, with no polarisation oscillations, and the correction required is thus a Pauli$-Z$ rotation to counteract a phase shift in the resultant ion-ion state. When the birefringence is orthogonal on the Poincare sphere to the atomic emission axes, qubit rotations about an equatorial axis on the Bloch sphere are required to counteract the impact of polarisation oscillations. To correct for intermediate orientations of the uncontrolled birefringence requires spin corrections around arbitrary (and varying) axes, making calibration challenging. Finally, it should be noted that if

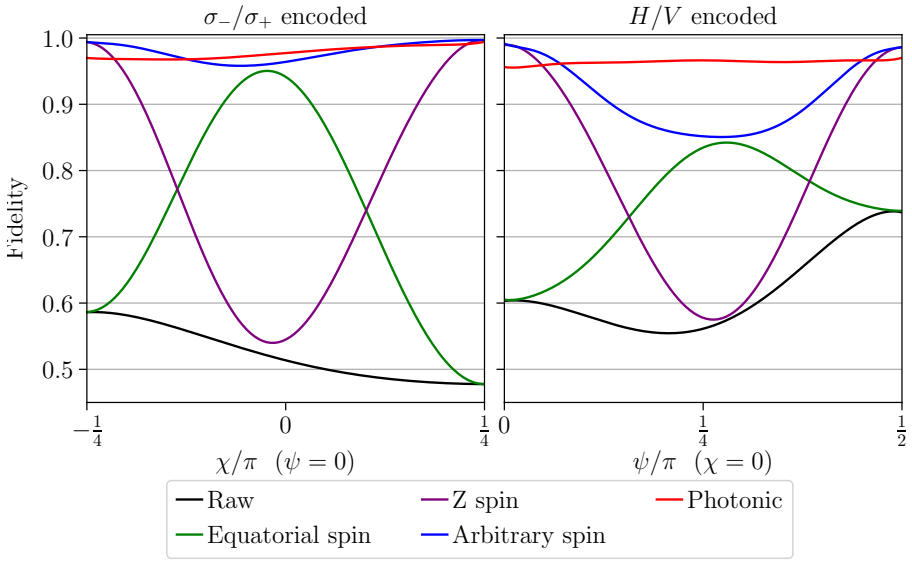


Figure 6.8: Dependence of fidelity on the orientation of the atomic decay polarisations to the birefringent cavity eigenmode polarisations for σ_-/σ_+ and H/V schemes with a variety of correction strategies, all without detection-time windowing. The orientation of the birefringence is aligned to the higher-energy polarisation eigenmode, parameterised by its azimuthal (ψ) and altitudinal (χ) angles on the Poincare sphere. In all cases, one node has zero-birefringence, while the second experiences a birefringence of magnitude $|\delta_B| = \kappa$, the orientation of which is varied. As well as considering fidelity without corrections ('Raw') and after photonic and arbitrary spin corrections, two additional classes of spin correction are shown. The 'Z spin' correction stipulates that the qubit unitary must be a rotation about the z axis on the Bloch sphere, and the 'Equatorial spin' that the rotation is about an (optimal) equatorial axis. For the σ_-/σ_+ scheme, the fidelity with photonic and arbitrary spin corrections is generally higher and varies less with birefringence orientation than the H/V case, though interestingly the raw fidelity is generally lower. For both schemes, when the polarisation eigenmodes of the birefringent cavity align with the atomic emission (the left and right limits of the range), the optimal arbitrary spin correction is a z-type unitary as the birefringence induces a phase error with no polarisation rotation. When the birefringence is orthogonal to the atomic emission directions (centre of both plots), equatorial spin corrections become much more effective than z-spin corrections, indicating that the spin qubit populations must be partially swapped post-herald to counteract the polarisation rotation. Neither specific class of spin correction is suitable for an arbitrary birefringence.

the birefringence orientation approaches alignment to the atomic emission axes, the vanishing polarisation rotation makes it increasingly difficult to optimise a photonic correction scheme by inspection of the output wavepacket alone, as these corrections are found by nulling these rotated polarisation components.

6.5.5 Dependence on Magnitude of Birefringence

The previous section discussed the impact of birefringence orientation on entanglement fidelity. However, if originating from manufacturing imperfections, the angle of the birefringence may not be predictable. It is therefore important to consider the range of fidelities achievable for a given magnitude of birefringence, as shown in Fig. 6.9, to provide an estimate of the maximum tolerable $|\delta_B|$ for a given remote entanglement application. Here, again one node has zero birefringence and the other an uncontrolled birefringence. The angle of the uncontrolled birefringence is selected from across the Poincare sphere to indicate the range of fidelities with angle, and the fidelity is averaged over linear birefringence angles to estimate the performance in the case that the birefringence is uncontrolled and linear, for example the archetypal case where an elliptical deformation of constant magnitude but random orientation is present.

The data highlights that without correction strategies the fidelity drops rapidly as birefringence magnitude increases and therefore aggressive windowing and a correspondingly large reduction in rate would be necessary to attain useful fidelities for large birefringences. The spin corrections yield superior fidelity at large birefringences, but for modest birefringence the photonic corrections out-perform spin corrections at restoring fidelity, with the additional advantage of substantially simpler calibration. Both correction types are generally more effective for the σ_-/σ_+ encoding.

6.5.6 Summary of Results

The numerical investigation has shown how cavity birefringence affects the wavepackets emitted from the network nodes. Photonic corrections act to repair these wavepackets before the interference-based measurement, and qubit spin corrections repair the ion-ion state after. Neither restores fidelity fully while both add complexity. For photonic corrections this may be manageable, but for the spin corrections the complexity of the calibration procedures required makes the approach impractical. Windowing improves fidelity, but the reductions in rate can be dramatic if the unwindowed fidelity is low. Overall, the interplay of birefringence orientation, magnitude, correction capability and windowing prohibits sweeping and rigorous

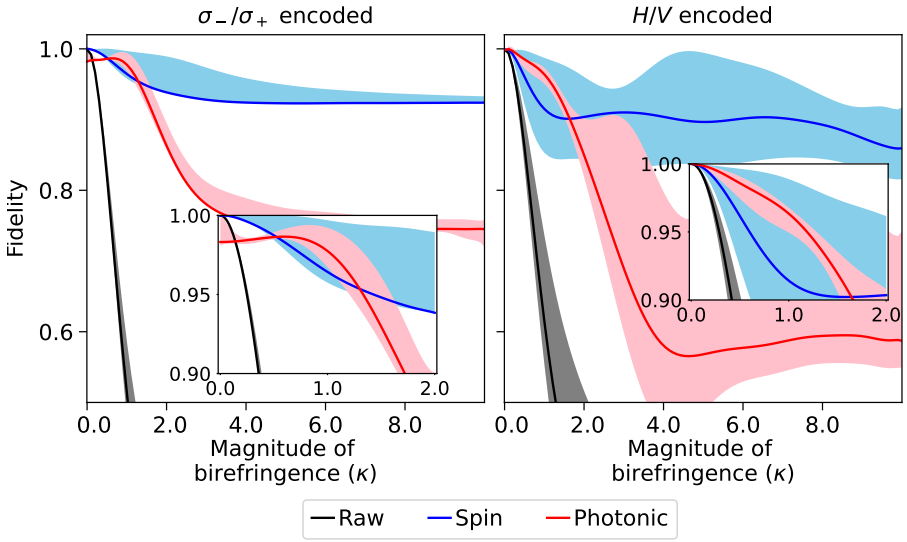


Figure 6.9: Fidelity versus birefringence magnitude $|\delta_B|$ for σ_-/σ_+ and H/V schemes without detection time windowing. In each case one node is non-birefringent and the other experiences a birefringence of variable magnitude and angle sampled from the corresponding parameter range in Fig. 6.8 (this covers the full range of fidelities). The fidelity for raw uncorrected heralds (black), arbitrary post-measurement spin corrections on the birefringent node (blue) and photonic corrections (red) are shown. The shaded range around each line covers the range of fidelities over all of the angles sampled. The solid line indicates the average fidelity for all *linear* birefringences, to estimate of the performance for the case where linear birefringence dominates. The inset on each plot details the region of birefringence $|\delta_B| < 2\kappa$, where fidelities $\mathcal{F} \gtrsim 0.9$, necessary for efficient entanglement distillation [94, 263], can be maintained. For both schemes, fidelity reduces very strongly with increasing birefringence if there are no corrections, and at very high birefringence, arbitrary spin corrections maintain fidelity much better than photonic corrections. For birefringence values up to the order of κ, photonic corrections typically perform well, with the exception of the lowest birefringence values for the σ_-/σ_+ scheme. This anomaly is due to the simple Pockels cell optimisation procedure picking unhelpful orientation angles that restore that restore spin-polarisation correlations correctly but introduce phase errors in the final state; In a real system, this could be avoided by only using photonic corrections if the raw spin-polarisation correlation is below a certain threshold.

conclusions, suggesting simulation is necessary to predict the specific properties of a given system.

The central practical conclusion of the numerical investigation is that any significant birefringence ($|\delta_B| \gtrsim 0.5\kappa$) without mitigating strategies is likely to reduce the remote entanglement fidelity in polarisation-encoded schemes below 90%. The most obvious solution is to reduce the birefringence of the cavity: with $|\delta_B| \ll 0.1\kappa$, the associated loss of ion-ion fidelity is typically $< 1\%$, which would likely not be the dominant source of remote entanglement infidelity. Where this is not possible, the

reduction in fidelity can be mitigated via correction strategies, windowing or both.

While the magnitude of birefringence remains modest ($|\delta_B| < 2\kappa$) the fidelity may be partially restored using local post-measurement operations or photonic corrections at the cost of increased experimental complexity and, for spin-qubit rotations, a formidable calibration overhead. The effectiveness of the corrections varies, but generally corrections for the σ_-/σ_+ scheme are simpler and better-performing than those for the H/V scheme. As a last resort for attaining an acceptable entanglement fidelity, windowing on the detection times can strongly improve the fidelity by post-selecting only heralds whose fidelity is not reduced strongly by birefringence. However, the windowing region can be complicated and the entanglement rate is often severely reduced. For very high birefringence, none of the strategies considered are likely to yield good fidelity and rate.

6.5.7 Deliberately Applied Birefringence

For completeness, it is important to note that while, hitherto, cavity birefringence has been treated as an unwanted, uncontrolled fabrication error, there is scope to fabricate deliberately birefringent cavities. As discussed, this has been done for other quantum information applications [237, 262], but could also be applied usefully in the context of remote entanglement.

The most obvious use is to apply a deliberate birefringence aligned with the atomic emission polarisations (this affects the diagonal component $\delta_{B,3}$ in Eq. 6.2). The advantage of this is that the cavity photon is much less vulnerable to polarisation oscillations caused by stray birefringence in an analogous manner to the operation of polarisation-maintaining fibre. The node would still be vulnerable to uncontrolled birefringence aligned with the applied birefringence, but the resulting phase shift would be much simpler to correct than the errors exhibited in the numerical investigation.

It is also possible (as noted in [258]) to achieve pure and stable output polarisations by applying a strong birefringence to *deliberately induce* polarisation oscillations. In this case the polarisation output is exactly opposite to the case of zero birefringence. However, the output efficiency of the photon production is strongly reduced, so this

phenomenon appears to be more a curiosity of use only if the polarisation could not be protected by a deliberate diagonal birefringence.

6.6 Further Issues with Polarisation Encoding

This chapter has focussed on the problems caused by cavity birefringence in polarisation encoded remote entanglement schemes. However, polarisation encoding also suffers from other disadvantages that would discourage its use in remote entanglement experiments.

The first major disadvantage is that polarisation-encoded photonic qubits are vulnerable to fibre birefringence. This means that long distance entanglement would require careful calibration and cancellation of fibre polarisation noise. This problem has led experiments to encode photonic information in time-bin qubits for long-distance entanglement experiments, even when the atom naturally emits a polarisation bit [264, 265].

The other disadvantage appears in the context of the extended level structure that was ignored for the simulations in this section. If the H/V encoding scheme is chosen, then both H and V polarised photons must be collected from the cavity, which would ordinarily correspond (in the system of Fig. 6.2) to σ_- and π-polarised decay respectively. However, there is also a σ_+ decay channel that is not forbidden by angular momentum selection rules or geometric arguments. This is usually suppressed by being an inherently weaker transition and being at a substantial detuning from Raman resonance, but any decays to this state that result in cavity emission of a H-polarised photon will result in fidelity loss. In encodings where only one polarisation need be used, this can be chosen to be V, corresponding to the π decay, and the $\sigma_\pm$ decays can both be rejected, rendering such encodings invulnerable to this form of spectator decay. For the σ_-/σ_+ scheme, only two cavity decay channels are geometrically allowed, so polarisation encoding does not suffer from spectator decays in this case.

6.7 Conclusions

Polarisation encoded remote entanglement schemes suffer from a number of practical complications that must be understood before designing an experiment. Mirror birefringence induces time dependence to the polarisations of the network photons, which can damage the fidelity of the remote entanglement process even at modest birefringence. Correction strategies can reduce fidelity losses, but do so imperfectly and at the expense of increased complexity. Considering this problem and the additional vulnerabilities of polarisation encoding to fibre noise and certain spectator decays, it is reasonable to conclude that encodings other than polarisation may be preferable, such as time-bin encoding discussed in Chapter 7. It should be noted that other encodings are not entirely immune from the impact of cavity birefringence as the distortion of wavepacket shape still harms the fidelity. However, they should be less sensitive.

7 Time-Bin Encoding

7.1 Introduction

Two-photon remote entanglement schemes entangle stationary ion qubits with travelling photonic qubits, followed by a projective measurement of the photonic state to leave an entangled ion-ion state. Such schemes can encode the photonic qubit in various degrees of freedom. In Chapter 6, polarisation encoding was discussed. This encoding has been used for remote entanglement experiments with refractive optics [96, 250] and cavity-enhanced collection with superpolished mirrors [108], but might be vulnerable to the impacts of cavity birefringence, which is particularly prevalent in microcavities with highly curved mirrors. An alternative encoding that is less affected by cavity birefringence is time-bin encoding. This chapter will summarise this encoding scheme and discuss its advantages and disadvantages, before presenting a detailed scheme for time-bin encoding with alkaline earth ions such as $^{88}Sr^{+}$.

Contribution and Acknowledgements

Secs. 7.2, 7.2, and 7.4 summarise discussion contained within the literature. Sec. 7.5 then presents schemes that the author devised and are, to the author's knowledge, not found in the literature.

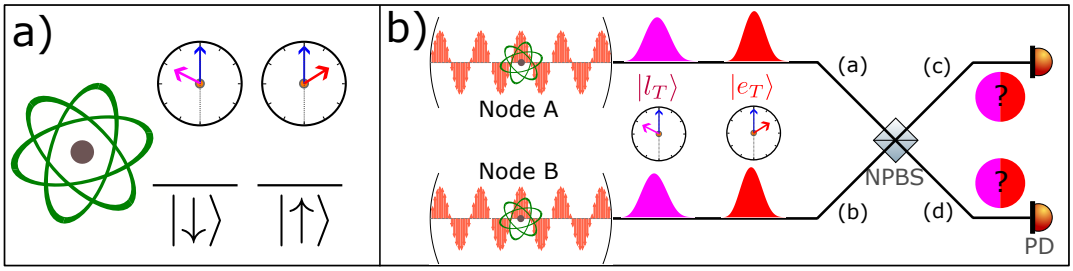


Figure 7.1: Schematic diagram of the time-bin encoding scheme. a) Each trapped ion is driven such that its qubit state is entangled with the time-bin of a single photon. b) The single photons obtained from each node are routed to a non-polarising beamsplitter to remove which-path information and then detected at time-resolving detectors. Upon measurement of one photon in the early time-bin and one photon in the late time-bin, the ions are projected into an entangled state.

7.2 Principle of Time-Bin Encoding

The time-bin encoding scheme follows the same basic principles as the polarisation encoding scheme of Chapter 6, but with the photonic information encoded in a time-bin qubit rather than a polarisation qubit. Thus, the two stationary ion qubits are manipulated such that they lie in a superposition of one qubit state (for convention, $|\uparrow\rangle$) after emission of a photon in the first, 'early' time-bin (denoted $|e_T\rangle$) and the opposite qubit state $|\downarrow\rangle$ after emission of a photon in the 'late' time-bin $|l_T\rangle$. Such entanglement of a matter qubit with a photonic time-bin qubit has been demonstrated in a variety of systems, such as NV centres [266] and atomic spin waves [267], and entanglement with time-bin qubits is also common in purely photonic settings [268, 269]. As with the polarisation-encoded scheme, the photons emitted from the two stationary qubits are routed to a non-polarising beam splitter to erase which-path information. Upon measurement of one photon in the early time-bin, and one photon in the late time-bin, the ions are projected to an entangled state. The scheme is depicted in Fig. 7.1, and described in greater detail in [168, 270].

7.3 Advantages of Time-Bin Encoding Over Polarisation Encoding

The first major advantage of time-bin encoding is that it should be less vulnerable to cavity-birefringence-induced polarisation oscillations. This is because the information is not encoded in the polarisation degree of freedom, and therefore polarisation

rotation does not directly lead to errors. However, time-bin encoding is not entirely invulnerable to birefringence for two reasons. Firstly, the birefringence may cause the energy levels of joint ion-photon states to differ between nodes, which, if not properly cancelled, could cause unexpected phase errors in the final state. Secondly, the distortion of photonic wavepackets caused by the effect of birefringence on the photon production process (discussed in Chapter 6) will reduce the indistinguishability of the photons from different nodes, thus reducing the entanglement fidelity. Therefore, time-bin encoded schemes should be much less vulnerable to birefringence than polarisation encoded schemes, but could still be impacted by highly birefringent cavities.

For similar reasons, time-bin encoded photons are also less vulnerable than polarisation encoded photons to fibre birefringence. In polarisation-encoded schemes, the superposition of polarisation must be preserved, and thus (perhaps counterintuitively) non-polarisation maintaining fibre must be used to avoid the rapid accumulation of relative phase between the polarisation modes that is inherent to polarisation-maintaining fibre. However, care must then be taken to ensure the polarisations are maintained during transmission or any rotation corrected via appropriately controlled retarders, which would require active stabilisation [271–273]. In contrast, time-bin encoded photons can share a single polarisation mode [265, 267], and polarisation maintaining fibre can be used, making the encoding robust against fibre noise and suitable for much longer-range transmission.

A further advantage is that only one cavity polarisation need be used to generate the photons, whereas for polarisation encoding both orthogonal cavity modes were required. This allows the suppression of errors from off-resonant processes, particularly in the case where the magnetic field is orthogonal to the cavity axis. In this geometry, all three angular momentum components of the electronic decay are supported by the cavity, with the π decay projecting to V polarisation and $\sigma_{\pm}$ decays projecting both to H. When using polarisation encoding, typically the weaker σ decay is not wanted, and should be separated in frequency space from the desired transitions to suppress decays on this channel. For time-bin encoding in this geometry, it is possible to only use π decays for both qubit states, and therefore

to filter out all H-polarised photons emerging from the cavity, removing unwanted off-resonant decays from causing a detection herald.

A technical advantage of time-bin encoding is that it only requires two photon detectors, as opposed to the four used in polarisation encoding. In practical experiments with finite budgets, this could be important as high-quantum efficiency detectors can be expensive.

Finally, a more speculative advantage is that time-bin encoded photons are easier to frequency-convert than polarisation encoded photons [264] as the frequency conversion process typically functions best for a particular polarisation orientation [274]. This would be advantageous for longer distance applications where conversion of photons to the telecom band may be desired.

These advantages have led to the use of time-bin encoded photons, even in experiments that initially produce a polarisation encoded photonic qubit [264].

7.4 Disadvantages of Time-Bin Encoding Compared to Polarisation Encoding

Though the advantages of time-bin encoding are compelling, its use would bring its own set of difficulties.

The most obvious disadvantage of time-bin encoding is that, as the two photon wavepackets must leave the cavity sequentially, the time taken for each attempt cycle is longer than if the photon wavepackets were emitted simultaneously. The magnitude of this disadvantage depends on the percentage of the attempt cycle that is used to generate the photons. As discussed in Chapter 4, for the case of remote entanglement with trapped ions using microcavities, the short cavity length and long preparation time mean that the photon generation process does not generally constitute a large fraction of the attempt time, and therefore this disadvantage is not serious. However, if the photon emission timescale were the dominant component of the attempt time, the slowdown mandated by time-bin encoding would be more concerning.

A second disadvantage is that, compared with polarisation encoding, the photon generation process can be less efficient due to the sequential nature of driving. A performance advantage for systems with simultaneous driving of multiple cavity

decays, such as polarisation encoding, was seen in Chapter 3. However, it is important to note that this required the two final states to be close in energy. To avoid significant spectator decays in a polarisation encoded scheme, a large Zeeman splitting may be required, and in this case, the output probability tends to the limit that the photons are generated in separate processes- exactly the same limit as applies to time-bin encoded schemes.

Another disadvantage is that, while polarisation qubits are simple to manipulate (requiring just waveplates and beamsplitters), time-bin qubits are more complicated to control [275]. Though the time-bin encoded remote entanglement protocol itself requires only time-stamped measurement, demonstrating entanglement between an ion and a time-bin photonic qubit would require coherent manipulation of the time-bin superposition. This can be achieved using a Franson-like interferometer [266, 276]. Here, beam splitters and fibre delay lines are used to split and recombine the photon wavepackets across different time-bins. Such manipulation is perfectly possible, but the extra complexity and expense is a disadvantage.

Time-bin encoding is also very vulnerable to temporal mixing of the photon wavepacket. As described in Chapter 3, temporal mixing is caused by spontaneous emission and subsequent cavity emission, which can lead to distinguishability of the interfering photons. This distinguishability usually takes the form of a delayed photon (although the photon wavepacket will also be deformed). For time-bin encoding, such a distortion is especially problematic because the information is directly encoded in the delay of the photon. As an example, this means that *any* spontaneous emission during the generation of the late photon followed by emission from the cavity would lead to a completely separable ion-photon, and therefore ion-ion, state. As discussed in Chapter 3, the level structure utilised can be chosen to make temporal mixing negligible, but this can be complicated for the sequential driving of a time-bin encoded process, as detailed in Sec. 7.5.5.

A final more general point is that time-bin encoded entanglement generation is sensitive to fluctuations in the system parameters between the early and late photon bins, such as photonic pathlength variations of magnetic field fluctuations. Single-bin encodings are also sensitive to system fluctuations between the earliest and latest

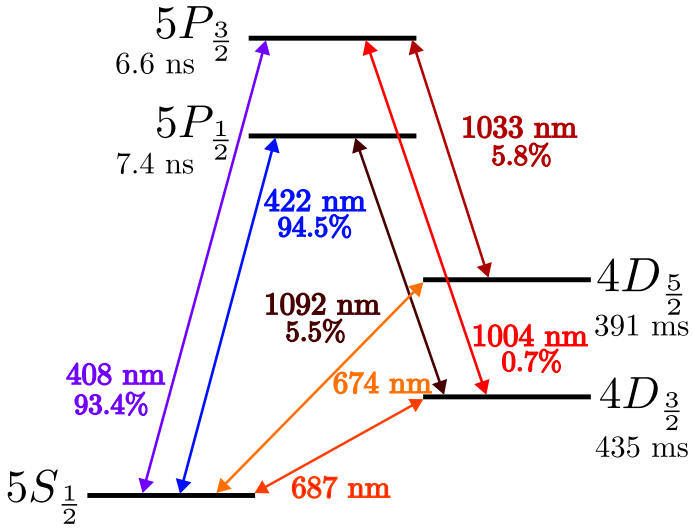


Figure 7.2: Selected manifolds of the fine level structure of the $^{88}Sr^{+}$ ion. Manifolds are labelled by their angular momentum and total electron spin quantum number J, and, where the state decays, the expected lifetime of the states in that manifold. Transitions are indicated on the coloured arrows with their corresponding wavelength. Dipole transitions also have the branching ratio from the excited manifold to the lower manifold indicated as a percentage.

possible detection times, but the larger range of detection times in the time-bin case will increase this sensitivity by a factor of two or more.

7.5 Detailed Protocol

7.5.1 Introduction

Compared to polarisation encoding, time-bin encoding has advantages and disadvantages. On balance, the reduced sensitivity to birefringence and the suppression of most off-resonant errors are persuasive enough to suggest this encoding be investigated further by considering specific schemes to implement the protocol. The remainder of this section will therefore discuss how best to perform time-bin encoded remote entanglement using $^{88}Sr^{+}$ as an example case, the level structure of which is shown in Fig. 7.2.

7.5.2 Manifolds

As discussed in Chapter 2, the cavity should couple to a dipole transition for fast photon extraction, and the advantages of an infrared photon (i.e. higher

mirror finesse and lower fibre losses) would favour one of the dipole transitions of $P_{\frac{1}{2}} \rightarrow D_{\frac{3}{2}}$, $P_{\frac{3}{2}} \rightarrow D_{\frac{3}{2}}$ or $P_{\frac{3}{2}} \rightarrow D_{\frac{5}{2}}$.

To reduce the errors associated with temporal mixing (discussed in Chapter 3), the decay from the excited state to the initial state should have a low branching ratio. This means that the initial state should also lie within the D manifolds (rather than the $S_{\frac{1}{2}}$ manifold). However, complications can arise if the initial and final manifolds are the same. Firstly, it becomes possible to have cavity decay followed by re-excitation and subsequent cavity decay, and secondly, it is difficult to distinguish through measurements of the atomic state whether a cavity decay has occurred or not. For these reasons, it would be best if the initial and final manifolds were different.

This narrows down the possibilities of the three-level transfer $|u\rangle \rightarrow |e\rangle \rightarrow |g\rangle$ to $D_{\frac{5}{2}} \rightarrow P_{\frac{3}{2}} \rightarrow D_{\frac{3}{2}}$ or $D_{\frac{3}{2}} \rightarrow P_{\frac{3}{2}} \rightarrow D_{\frac{5}{2}}$. Of these, the latter is preferable because the cavity decay has a higher branching ratio, and the decay to the initial state that leads to temporal mixing is notably weak. Thus, the schemes considered for the remainder of this section will all have initial state $|u\rangle$ in $D_{\frac{3}{2}}$, excited state $|e\rangle$ in $P_{\frac{3}{2}}$, and final state $|g\rangle$ in $D_{\frac{5}{2}}$.

7.5.3 Pulse Shaping

As discussed in Chapter 4, for the parameter regime considered it should be possible to obtain output probabilities near the adiabatic limit without compromising significantly on rate. Therefore, the question of the best way to drive the systems to maximise entanglement rate will be ignored on the grounds that any improvements over simpler driving protocols are likely to be marginal, and the adiabatic limit provides a good approximation to the expected performance. This does not mean that careful control of the driving procedure would not be important; Shaping the driving pulse can offer control over the output wavepacket phase and amplitude to make remote entanglement more resistant to errors, and it is important also not to drive too fast to avoid introducing off-resonant processes. Nevertheless, for this section these considerations will not be discussed.

7.5.4 Geometry

As discussed in Chapter 2, the relative orientation of the magnetic field and cavity axis controls how strongly atomic decays (labelled by the angular momentum of the emitted photon π or $\sigma_{\pm}$) couple to the two orthogonal cavity modes. Typically, special configurations are chosen to enhance particular transitions and rule out others. These configurations either have the cavity axis parallel to the magnetic field, or the cavity axis perpendicular to the magnetic field, and these will be discussed in turn.

If the cavity axis is parallel to the magnetic field, only $\sigma_{\pm}$ decays are supported by the cavity, and they are supported with the strongest possible overlap. If only one angular momentum decay was desired, a polarising beamsplitter could be placed after the quarter-wave plate on the cavity output to reject the other polarisation, reducing off-resonant errors. If the birefringence is linear (as would be caused by an elliptical manufacturing defect, or mirror misalignment for Gaussian-shaped mirrors), then birefringence always introduces polarisation oscillations. Even though time-bin encoding is not sensitive errors in the polarisation itself, fidelity would be lost to the distortion of the wavepacket resulting from the intra-cavity polarisation rotation.

For the cavity axis perpendicular to the magnetic field, π-transitions couple to the (conventionally) V polarised cavity mode, and $\sigma_{\pm}$ transitions to the H polarised cavity mode, with an overlap diluted by a factor a $\sqrt{2}$. Given this, it would be natural to choose the scheme to produce only π decays, and to reject the $\sigma_{\pm}$ decays with a polarising beamsplitter, again reducing off-resonant errors. A linear birefringence would cause a mix of polarisation oscillation and a resonant mode shift depending on its orientation with respect to H and V. However, the linear cavity modes and the fact only one of the two cavity polarisations need be used opens up options to deliberately apply birefringence along the atomic emission axes by fabricating intentionally elliptical mirrors. In this case, polarisation oscillations can be rendered negligible by applying a much stronger deliberate birefringence than the random birefringence arising from fabrication defects and aligning the axis of this birefringence to the emission polarisation direction. Because only one mode is used, the cavity length could then be tuned to shift the cavity frequency to the desired value, meaning that it wouldn't be necessary to control the deliberate

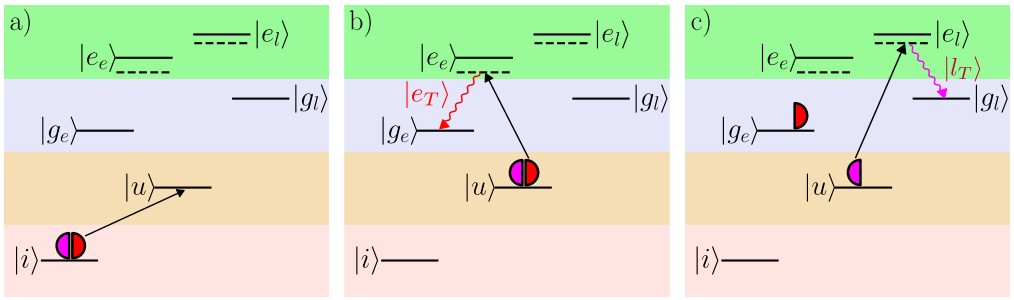


Figure 7.3: Schematic diagram of a time-bin encoding scheme using a single initial state. a) The prepared state $|i\rangle$ is transferred to the initial state $|u\rangle$. b), the early photon $|e_T\rangle$ is produced by a two-photon process, leaving the ion in state $|g_e\rangle$. The probability transferred through this process should be carefully controlled. c) The late photon $|l_T\rangle$ is produced through a two-photon cavity-enhanced process using the remaining population.

birefringence to a high degree of precision. Depending upon the expected magnitude of the random, uncontrolled birefringence, this potential to reduce errors associated with birefringence could be a highly advantageous.

7.5.5 Driving Protocol and Temporal Mixing

As discussed previously, choosing an appropriate initial state can reduce the errors associated with temporal mixing to negligible levels. However, the sequential nature of time-bin encoding means that care must be taken to avoid inadvertently introducing the error.

To see why, consider precisely how a time-bin encoding protocol could be enacted within the level structure of an ion[1]. Firstly, the ion would be initialised to a state $|i\rangle$ in the $S_{\frac{1}{2}}$ manifold through optical pumping. This population could then be transferred through a quadrupole pulse, Raman transfer or STIRAP process to the initial state $|u\rangle$ of the three-level system $|u\rangle \rightarrow |e\rangle \rightarrow |g\rangle$ with states in the $D_{\frac{3}{2}}$, $P_{\frac{3}{2}}$, and $D_{\frac{5}{2}}$ manifolds respectively. In general, the three-level systems could be distinct for the early and late photons, so the levels comprising the 'early' three-level system will be denoted $\{|u_e\rangle, |e_e\rangle, |g_e\rangle\}$, and the 'late' $\{|u_l\rangle, |e_l\rangle, |g_l\rangle\}$.

The simplest idea would be to have a single initial state ($|u_e\rangle \equiv |u_l\rangle$), and then drive half the population for the early photon, and the remaining half for the late photon as depicted in Fig. 7.3. However, there are two key limitations.

[1]Note that there are many different sequences that could be used, but this is one of the more straightforward and natural examples for this choice of initial and final states.

Firstly, the partial driving of the first process loses a nice property of full transfer by vSTIRAP. In a full transfer, the initial state is the eigenstate for zero drive strength, the final state is the eigenstate when the drive Rabi frequency far exceeds the cavity coupling. As long as the transfer is adiabatic throughout, the exact way in which the drive Rabi frequency evolves, or even the precise final value, is not important. This means that fine calibration is not necessary to enact the transfer: small fluctuations in e.g. the cavity coupling g will not change the initial and final state. In contrast, the partial transfer to produce the early photon would need to be careful calibrated, and would be vulnerable to fluctuations in the cavity coupling and drive strength. Because of this, it is tempting to split the population into two halves before the cavity is involved, so that the finely calibrated process that produces the superposition need only involve the ion and the control lasers.

To solve the first problem, the initial states of the two three-level systems ($|u_l\rangle$ and $|u_e\rangle$) should be distinct. However, if population is immediately moved to a superposition of both initial states, then $|u_l\rangle$ is occupied while the early photon is being driven and the laser that drives the early process can inadvertently couple to the population earmarked for the late photon. This can lead to photons emitted in the early time-bin that are not correlated with $|g_e\rangle$, reducing the fidelity of the process. To avoid this off-resonant driving, it is tempting to leave the half of the population for the late photon in the prepared state $|i\rangle$ while the early process is being driven, before then transferring this population to $|u_l\rangle$ and driving the late process. This would mean that the 'late population' cannot couple to the early drive, so avoids the miscorrelation that could result from the previous suggestion. However, any spontaneous decay from $|e_e\rangle$ to $|i\rangle$ (occurring with rate $\Gamma_{e_e,i}$) during the early process will result in a temporal mixing error when the population is then transferred from $|i\rangle$ to $|u_l\rangle$ and the late photon driven. This error process is depicted in Fig. 7.4. As $|i\rangle$ lies in the $S_{\frac{1}{2}}$ manifold and $|e_e\rangle$ in the $P_{\frac{3}{2}}$ manifold, the branching ratio of this decay is generally strong, and therefore the temporal mixing would again be a source of significant error.

This problem can be avoided by careful use of level structure. If the angular momentum projection m_J of the prepared state $|i\rangle$ and early excited state $|e_e\rangle$ differ

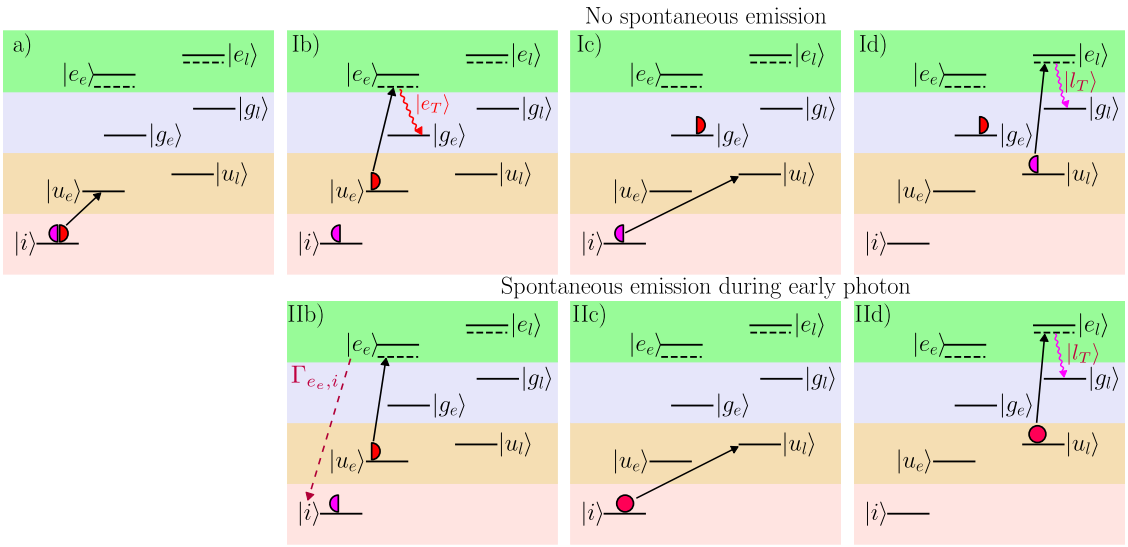


Figure 7.4: Schematic diagram of a time-bin encoding scheme where the population for the late process is kept in the preparation manifold while the early process is driven. a) Half the population is transferred from the prepared state $|i\rangle$ to the initial state $|u_e\rangle$. The early photon process is then attempted. If there is no spontaneous emission, path I is followed, leading to b) generation of the early photon $|e_T\rangle$, c) transfer of the remaining half of the population to $|u_l\rangle$, and d) production of the late photon $|l_T\rangle$. If this process proceeds without mishap, the qubit encoded in the ion is entangled with the time-bin of the single photon emitted through the cavity. If there is spontaneous emission to the prepared state during the driving of the early photon, path II is followed. IIb) Spontaneous emission transfers the half of the population intended for the early photon to the prepared state, which is where the half of the population for the late photon has been stored. IIc) The whole population is now transferred to $|u_l\rangle$. IId) The late photon process is driven. If a photon is emitted into the cavity via path II, the final ion state is $|g_l\rangle$, and photon state is $|l_T\rangle$. This means that the ion qubit and photon time-bin are not entangled, and the entanglement swapping protocol will not function.

by more than 1, then the spontaneous decay from $|e_e\rangle$ to $|i\rangle$ that causes temporal mixing is not dipole allowed and $\Gamma_{e_e,i} = 0$ (all the spontaneous decay from $|e_e\rangle$ is directed to the other state in the $S_{\frac{1}{2}}$ manifold, which does not take part in the process). Such a scenario can be arranged, and avoids temporal mixing errors from spontaneous decay during the early process. This requirement significantly narrows down the possible schemes that can be enacted.

7.5.6 Specific Schemes

Coupling strengths

Having discussed features of a time-bin encoded driving protocol that would avoid certain error classes, it is now possible to combine the criteria imposed by these

features with the level structure of $^{88}Sr^{+}$ in order to present possible schemes. To decide which fine structure levels to use, it is important to consider the cavity coupling strengths afforded by different transitions and cavity geometries. As discussed in Chapter 2, the cavity coupling g contains two factors related to the geometry of the experiment and the specific fine levels involved:

- $\left\langle J_e, m_{J_e} \middle| J_g, 1, m_{J_g}, q \right\rangle$: The Clebsch-Gordan coefficient of the specific transition coupled to the cavity mode.
- $\epsilon_q^V / \left|\boldsymbol{\epsilon}^{\boldsymbol{V}}\right|$: The alignment between the dipole of the transition and the cavity mode polarisation.

These can be combined into a factor

$$f_g = \left\langle J_e, m_{J_e} \middle| J_g, 1, m_{J_g}, q \right\rangle \frac{\epsilon_q^V}{\left|\boldsymbol{\epsilon}^{\boldsymbol{V}}\right|}, \tag{7.1}$$

such that the cavity coupling g varies according to geometry only through factor f_g. The cooperativity, which controls the limit to photon generation efficiency (as seen in Chapters 2 and 4), is proportional to $|f_g|^2$. These relative cooperativities are shown in Fig. 7.5 for both the case where the cavity axis is aligned to the magnetic field and the case where they are perpendicular.

Cavity axis parallel to magnetic field

If the cavity axis is parallel to the magnetic field, the most eye-catching scheme to use is an 'end to end' protocol utilising both of the strongest transitions available in the manifold. In this scheme, depicted in Fig. 7.6, the early and late photons couple to oppositely-circular cavity modes.

The main advantage of such a scheme is that the transitions are inherently strong. As well as increasing the probability of photon emission, this strength also reduces any residual errors from temporal mixing and allows for the cavity transmission to be elevated, which reduces the errors introduced by birefringence. On the negative side, as with all schemes that use $\sigma_{\pm}$ photons, linear cavity birefringence always causes polarisation rotation, so an uncontrolled birefringence would disrupt the wavepacket shape and reduce fidelity, and a deliberate linear birefringence could not be used

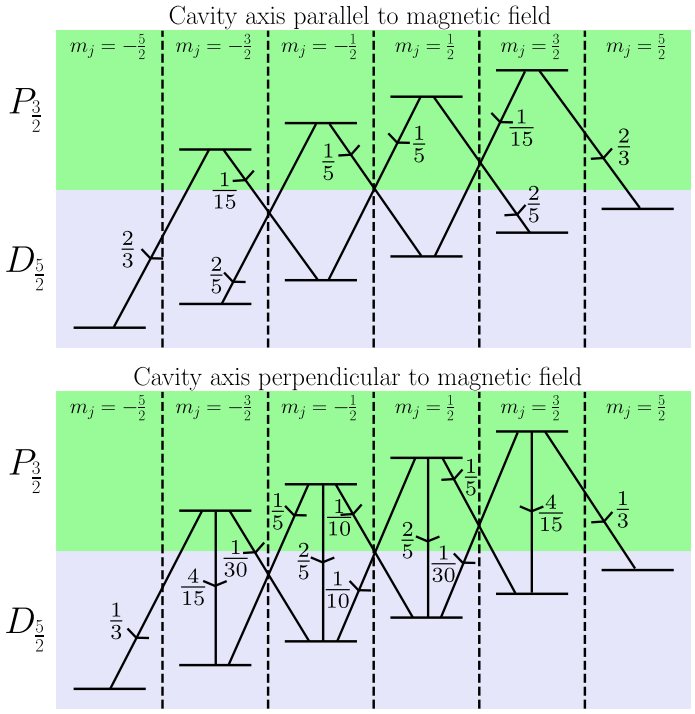


Figure 7.5: The relative cooperativities $|f_g|^2$ of the cavity transition with (top row) cavity axis parallel to the magnetic field, and (bottom row) cavity axis perpendicular to the magnetic field. When the cavity axis and magnetic field are parallel, π polarised decays are not allowed, and therefore each excited state can only decay with a change in angular momentum $\Delta m_j = \pm 1$. When the cavity axis and magnetic field are perpendicular, π and $\sigma_\pm$ decays are allowed, although the $\sigma_\pm$ decays are diminished in strength due to their projection onto the cavity modes. For this geometry, π decays tend to be relatively strong.

to keep the intra-cavity polarisations constant. There are two other downsides specific to this scheme. Firstly, the early and late photons are oppositely-polarised. Therefore any filtering to remove the off-resonant cavity decays would have to switch the permitted and rejected polarisations between the two time-bins. This can be done with a Pockels cell, but does increase experimental complexity. Secondly, the angular momentum difference of the final superposition is very large, meaning the state is highly vulnerable to magnetic field fluctuations. This superposition would typically be mapped back to a ground state qubit post-entanglement, but even so, magnetic field fluctuations before the transfer could introduce serious errors.

Cavity axis perpendicular to magnetic field

If the cavity mode is perpendicular to the field, a promising method to drive the ion-photon entanglement is to use π polarised atomic decays (coupling to the V

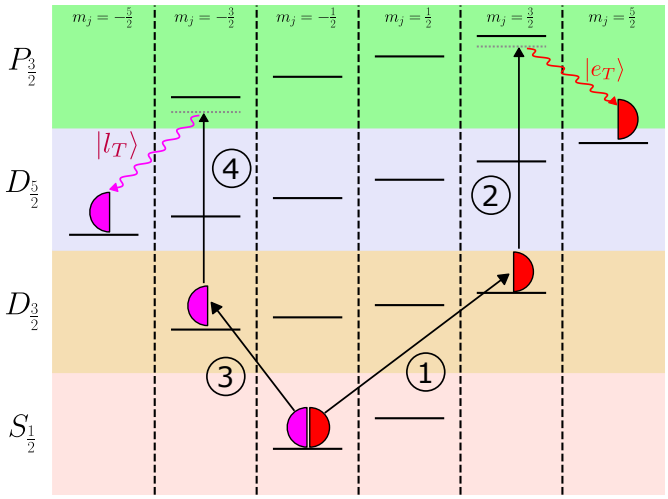


Figure 7.6: Diagram of a possible time-bin encoded scheme for the case where the cavity axis is parallel to the magnetic field. The driven transitions are depicted by straight arrows, cavity photon transition by waved arrows, and the population by fractions of a circle that represents the whole population. In this scheme, the early photon population is first prepared and then the corresponding two-photon process driven. Crucially, there is no allowed dipole decay between the excited level for the early photon and the prepared state, preventing temporal mixing. Once the early photon has been driven, the late photon process is executed. The early and late processes utilise the two strongest decays in the manifold, but the final superposition features a large separation of angular momentum in the ion electronic state, which would make this superposition very vulnerable to magnetic field noise at the ion.

polarised cavity mode), which allows for the rejection of all σ polarised atomic decays through a polarising beamsplitter. These schemes are also compatible with an applied deliberate birefringence along the H or V direction to reduce errors associated with polarisation oscillation.

The first suggested method of driving the process is the 'symmetric decay' scheme depicted in Fig. 7.7 (panel a). Here, the early and late photons are both produced by π decays and the final states have angular momentum $\pm 3/2$. The final superposition has lower field sensitivity than the 'end to end' scheme, but also inherently weaker cavity coupling.

A less symmetric procedure (Fig. 7.7 panel b) could confer advantages. Here, the early and late photons transfer to final states with angular momentum $+3/2$ and $+1/2$ respectively, significantly reducing the field sensitivity of the final superposition. This scheme does lack the symmetry of the 'symmetric decay', but it should be possible to balance the cavity drive strengths and the population split between early

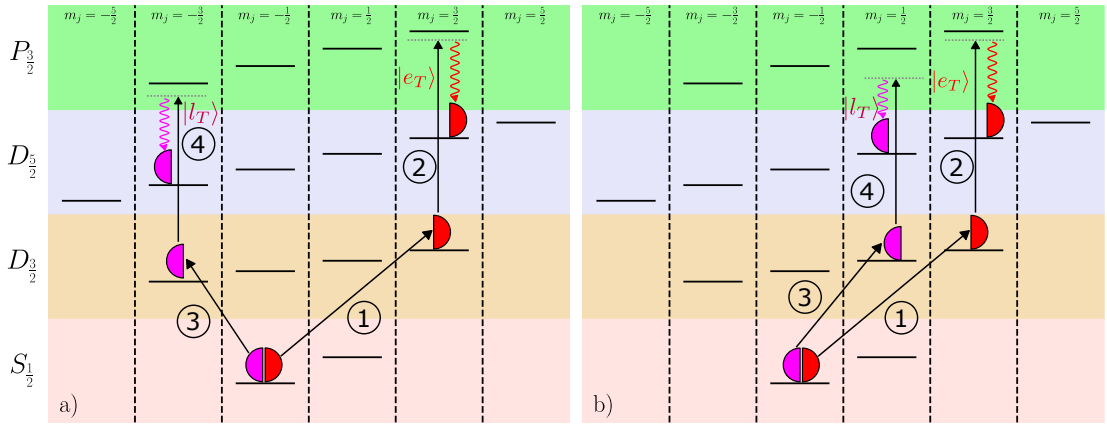


Figure 7.7: Diagram of a possible time-bin encoded schemes for the case where the magnetic field is perpendicular to the axis. Both schemes share a common structure. First, half the population is transferred from the prepared state to the initial state of the early process. The early process is then driven via an excited state that cannot decay to the prepared state. After the early process, the late photon generation process is driven. In both cases, the cavity photon has an angular momentum projection of zero onto the quantisation axis. a) A 'symmetric' configuration for this scheme in which the early and late three-level systems have opposite angular momenta. b) An 'asymmetric' configuration in which the early and late three-level systems have different angular momenta. The 'symmetric' scheme has greater symmetry, but the 'asymmetric' scheme benefits from a stronger cavity coupling on the late transition, and a lower magnetic field sensitivity in the final superposition.

and late portions to counteract the difference in cavity coupling.

7.6 Conclusion

This chapter has introduced time-bin encoding as an alternative to polarisation encoding for producing remote entanglement. This encoding has notable advantages: It is less vulnerable to cavity and fibre birefringence, undesired off-resonant processes can be filtered off by polarisation and, depending on the cavity geometry, birefringence may be deliberately applied to reduce the negative effects of uncontrolled birefringence. Time-bin encoding is, however, not without its disadvantages. The protocol tends to take longer than polarisation encoding and is more vulnerable to temporal mixing, and time-bin qubits are much more difficult to manipulate coherently than photonic qubits. However, the advantages conferred by this encoding mean that it is worthwhile attempting to understand how a reasonable scheme would work.

The sequential population transfers of time-bin encoding introduce complications

to the physics. Firstly, it is tempting to avoid splitting the early and late populations with cavity processes, as any fluctuations in ion-cavity coupling would lead to errors. However, it is not ideal to leave the late population in the manifold of the initial state while driving the early photon due to the potential for driving unwanted cavity processes. This means it is an attractive proposition to leave the late population in the $S_{\frac{1}{2}}$ manifold while the early photon is driven, but the levels have to be chosen carefully to avoid introducing temporal mixing errors.

Combining these criteria leaves a few options for the time-bin encoding protocol. Some options use the cavity parallel to the magnetic field, and some perpendicular. Evaluating which scheme is ideal is a complicated endeavour.

This conclusion leads to the most obvious area of further work. The three schemes presented in Sec. 7.5.6 have different advantages and disadvantages. Some use strong cavity decays, others weak, some are vulnerable to birefringence, others to magnetic field noise. It is impossible to say which one would be preferable in all situations. The best way of making this comparison is probably through detailed numerical simulation of the entire level structure, modelling cavity birefringence, as well as magnetic field noise and other imperfections.

However, once numerical modelling has been performed on the suggested schemes, it is worth looking back at some of the earlier assumptions in the argument. It was suggested that splitting the early and late population during the cavity processes would be vulnerable to fluctuations in the cavity coupling, and therefore the early and late three-level systems should have distinct initial states. However, splitting to distinct initial states increases the number of mapping pulses required, which will come with its own fidelity and attempt rate costs. Additionally, it was also argued that it would be dangerous to have $|u_l\rangle$ populated while the early process is driven in case the late process is excited during the early time-bin. However, the extent to which the late process would be driven inadvertently depends upon the relevant Zeeman splittings and the polarisation purity of the driving laser, so it is not necessarily impossible for a scheme to operate at high fidelity even if $|u_l\rangle$ is occupied during the early driving. Also, the extra procedures required to avoid this particular issue increase the length of the remote entanglement scheme, require a superposition be maintained across manifolds

for longer (placing stringent requirements on laser coherence), and necessitate a mapping pulse that itself must avoid off-resonant excitation to keep temporal mixing errors at bay. For the qualitative discussion of this chapter, the final schemes have been proposed because they avoid particular classes of error that are assumed to be problematic, and indeed narrowing the options in this manner is necessary to provide any coherent narrative. However, the assumptions made in the qualitative analysis about what constitutes an acceptable error and what is unacceptable are not necessarily well-founded. Therefore, it would be worthwhile to additionally model schemes that were initially ruled out on qualitative grounds, in order to understand under what circumstances these qualitative restrictions are justified and evaluate if schemes initially deemed unpromising could offer useful performance.

8 Atomic Sources

8.1 Background and Theory of Atomic Sources

8.1.1 Introduction

The aim of the work undertaken for this thesis was to lay some of the foundations required for a cavity-enhanced quantum network of trapped ion nodes. The work presented thus far has focussed on the physics of the ion-cavity interaction, but as important is the development of technologies that will make the cavity-integrated nodes scalable. This chapter presents work undertaken to this end, and concerns the development of a prototype miniature atomic source.

Ion trap experiments require a source of neutral atoms which can then be ionised (for example by two-stage photoionisation) and then trapped [277]. There are a handful of approaches to producing these neutral atoms. This chapter will first introduce the different types of source that are commonly used in ion traps, before discussing in more detail the theory of thermal atomic sources. The advantages of moving towards optically-heated miniature atomic sources in general will then be presented, followed by a discussion of the use of neutral fluorescence as a sensitive probe of the performance of these sources that will be used in this chapter. Finally, investigations into two atomic sources will be detailed, along with an overall conclusion and outlook.

Contribution and Acknowledgements

The results of this chapter are partially contained in [135], where more detail can be found on the first of the atomic sources discussed in this chapter. The work detailed herein was a collaborative effort. The MITAS experiment was a joint collaboration with ColdQuanta UK. The experimental investigation was conducted by the author, Dr. Shaobo Gao, Dr. Joseph Goodwin, and Dr. Timothy Ballance. The original text of the publication ([135]) was written mainly by Dr. Joseph Goodwin and Dr. Shaobo Gao, with Dr. Shaobo Gao producing the diagrammatic figures. The author analysed the data for the original submission, and then substantially modified the text with Dr. Joseph Goodwin following peer-review. The 'Pico-Oven' source was conceived by Dr. Joseph Goodwin, with further prototyping from Dr. Jacob Blackmore and Dr. Thomas Doherty. The sources were prepared and assembled by the author, Dr. Jacob Blackmore, and Dr. Thomas Doherty. The author designed and assembled the experiment chamber and the excitation-collection system. Dr Jacob Blackmore built the initial heating system, and Dr. Thomas Doherty the final heating system. The author conducted the experiments with assistance from Dr. Jacob Blackmore and Dr. Thomas Doherty, and analysed the data.

8.1.2 Approaches to Sources

The most common type of source used in ion trap experiments is the thermal source [278], in which the target metal is heated to sufficient temperature that neutral atoms are liberated from the surface. Historically these sources have been resistively heated, but this section will motivate the development of an optically heated source, with additional in depth discussion in later sections.

An alternative approach is to use laser ablation to generate the necessary neutral atoms. In this process a pulsed 'ablation laser' is directed at a target of the desired metal, and a high-energy pulse used to liberate atoms from the surface [279]. The main advantages of this approach are that a flux of neutral atoms can be generated near-instantaneously (as there is no need for the source to warm up), only small quantities of the source metal are needed as only small quantities of atoms are produced at a time [280], and the overall energy delivered to the system for loading is low, which

is particularly important for cryogenic systems [281, 282]. On the other hand, the technique requires a pulsed laser and for high-intensity pulses to be delivered to the target. In addition, the plume of atoms produced tends to exhibit a very high 'stream' velocity meaning only a minority of atoms can be trapped [283], and the produced plume can contain charged impurities which could affect trapping potentials.

A final option is to use a magneto-optical trap (MOT) as a source of atoms that can then be transferred to the ion trap [284]. While the MOT provides slow-moving atoms without impurities, it would itself require a source of neutral atoms. The additional footprint and experimental complexity involved with operating such a loading scheme means that this method is not a scalable option for ion trap quantum computing experiments, particularly for networked architectures where each node should be as simple as possible.

8.1.3 Miniature Optically-Heated Atomic Sources

Conventional resistively-heated thermal atomic sources present a number of disadvantages which the miniature, optically-heated sources detailed in this chapter are designed to overcome. Most fundamentally, the Joule heating is inherently inefficient as the low-resistance current supply wiring simultaneously strongly conducts heat away from the source, raising the power required to maintain a certain temperature. In the context of an experiment, the high powers required cause unwanted thermal and compensation drifts in room-temperature systems and more serious problems in cryogenic systems, to the extent that ovens have had to be placed outside the low-temperature stage [71, 285, 286] to avoid issues such as large temperature drifts after loading [287]. In addition, the high-power feedthroughs for the heating wires provide a natural failure point in the vacuum system, with ceramic elements that are vulnerable to damage. Diagnosing and solving such in-vacuum problems can cost a lot of time, particularly as experiments grow in complexity. Recent work has focused on making these resistive sources more compact [288], but the nature of the heating method means that the power requirements are still relatively high.

In contrast, optically heated sources offer compelling advantages. The heat is provided via a laser beam, removing the need for the source to be connected

to the rest of the vacuum system by a conductive wire, enabling the design of sources with greatly reduced heating power requirements. Secondly, the heating power can be provided through one of the existing viewports on the chamber rather than a dedicated feed-through, reducing the number of failure points on the system. Finally, as the power requirements are modest, the heating light can be provided by a simple continuous-wave diode laser of near-arbitrary wavelength, with minimal requirements on coherence.

Further improvements can be yielded by miniaturising the source and collimating the output atomic beam. A smaller source reduces the surface area through which heat can be lost through radiation, and may be placed closer to the trapping zone. In addition, a collimator can very effectively direct the beam to the desired location, with a minimal quantity of atoms delivered outside of the trapping zone. The use of collimated sources in atomic physics experiments is not new [289, 290], but use in ion trap experiments is rare because they are not critical to the functioning of experiments and add complexity during fabrication. The improvements in thermal and atomic delivery efficiency from using miniature collimated sources can reduce the operating power requirements by orders of magnitude, and even more drastically reduce the total outflux of atoms into the vacuum system. The promise of a reduced physical footprint, greater mechanical reliability, greatly reduced thermal drifts and negligible impact on the background pressure during operation represent an enticing combination, particularly as ion traps move towards scalable and reliable systems.

8.1.4 Theory of Thermal Sources

Thermal sources function by heating a sample of the target metal to a sufficient temperature that atomic vapour will be liberated from the surface. However, many ion trap experiments utilise singly ionised group-II metals, which react strongly with atmospheric oxygen and water vapour to form an oxide layer. This oxide layer forms a barrier that prevents atomic vapour from evolving from the surface of the metal, and this layer must therefore be removed in vacuum before the source can be used. The breaking of this layer is achieved by heating the source up to a high temperature (much higher than the standard operational temperature), a procedure

known as 'cracking' [291]. Once the oxide layer has been broken, the source can be used normally, while UHV conditions are maintained.

Thermal sources typically enclose the target metal in a housing with an outlet hole through which an effusive flux of atomic vapour can be emitted. The oven enclosure, the solid target and the gas of target atoms within the enclosure are all assumed to be in thermal equilibrium, meaning that at an oven temperature T, the gas of atoms within the enclosure is at vapour pressure $p_V(T)$ (this can be obtained from standard tables for the element of interest [292]), and thus, following the ideal gas law, has a corresponding number density $n_o(T) = p_V(T)/(k_B T)$. These atoms can then effuse through the aperture in the casing of the source.

The first source (described later in Sec. 8.2) had a simple aperture (a hole in the casing with effectively zero depth), for which the atomic flux density through the aperture is $\Phi = n_o(T)\bar{v}(T)/4$, where $\bar{v} = \sqrt{8k_B T/m\pi}$ is the mean speed of atoms in the oven. The atomic flux density on the half-sphere r away from the source is $\phi(r, \beta) = \phi_0(r)\cos\beta$, where β is the angle to the normal to the aperture (which becomes the atomic beam axis). Flux conservation of the half sphere with radius r determines axial flux density $\phi_0(r)$:

$$\Phi A_o = \int_0^{\pi/2} \phi(r, \beta)(2\pi r \sin\beta) r d\beta = \phi_0(r)\pi r^2, \tag{8.1}$$

where A_o is the area of the aperture. The flux density along the axis is also related to the target region number density n by $\phi_0(r_o) = \bar{v}n$, where r_o is the distance between the oven aperture and the target region. The target region number density is then

$$n(T) = \frac{\phi_0(r_o)}{\bar{v}(T)} = \frac{n_o(T)A_o}{4\pi r_o^2} = \frac{p(T)A_o}{4\pi k_B T r_o^2}. \tag{8.2}$$

The second source (discussed later in Sec. 8.3) used a cylindrical collimator in place of a simple aperture in order to produce a much narrower beam. Following the results derived in [293], cylindrical collimators are characterised by the ratio G of length to diameter. The flux density out of the collimator Φ_C is reduced from the flux density into the collimator Φ by a factor W, the transmission probability of the collimator whose values are tabulated [294], but tends towards $W = 4/(3G)$ for large G [293]. Theoretical derivations for the angular distribution of the collimated beam

depend upon models for roughness of the collimator surface, but several models suggest an angular half-width H of approximately $1.70/G$ in radians [293]. For our investigation, the transverse beam profile was measured to calculate an effective area of the beam A_b, defined as the ratio of the integrated number density over the transverse beam profile to the number density in the centre of the beam. Using this assignment, the atomic flux from a uniform beam of the central density and transverse cross-section A_b would have the same flux as the measured beam. Assuming the angular distribution does not change significantly with temperature, the number density at the centre of the beam can be calculated

$$n(T) = \frac{Wp(T)A_o}{4k_BTA_b}. \tag{8.3}$$

A typical target number density for ion trap experiments is $100\,\mathrm{cm}^{-3}$, which should be sufficient to load an ion in a few seconds with $\sim 1\,\mathrm{mW}$ of power in the second stage ionisation laser [277]. Our experiments were only able to measure a significant fluorescence signal from beams with much higher number densities, but by determining an effective thermal model for the source from these higher flux beams, it is possible to extrapolate to lower temperatures and thus determine the heating power required to achieve the target densities.

8.1.5 Theory of Measuring Fluorescence

The atomic density n that the ovens provide can be inferred from the rate of fluorescence observed from the atoms under excitation. In our experiments, neutral calcium atoms are excited on the $4s^2\,{}^1S_0 \rightarrow 4s4p\,{}^1P_1$ transition, a strong dipole transition with negligible (in the context of this experiment) branching ratio to states outside the ground manifold. The magnetic field is sufficiently weak that the different Zeeman sublevels are not spectrally resolvable. Under these conditions, a beam of intensity I and detuning Δ from a resonant transition with FWHM Γ produces a steady state scattering rate from a single atom of

$$r(\Delta) = \frac{I/I_{\mathrm{sat}}}{1 + I/I_{\mathrm{sat}} + 4\frac{\Delta^2}{\Gamma^2}}\frac{\Gamma}{2}, \tag{8.4}$$

where $I_{\text{sat}} = \hbar(2\pi f)^3\Gamma/12\pi c^2$ is the saturation intensity, with f the transition frequency.

Because the magnetic field is not strong enough to spectrally resolve the Zeeman sublevels, the atoms scatter photons in the dipole pattern determined by the polarisation of the input radiation. The proportion of photons scattered that are collected by the detector within the collection region is then made up of three efficiencies: η_d, the efficiency with which the detector converts input photons to counts, η_g, the proportion of the solid angle from which the detector collects photons and η_p, the intensity of photons scattered in the direction of the collection with respect to isotropic emission.

Photons are collected over a collection volume V_C over which the laser intensity can vary (the atomic beam density could also vary, but is approximated as constant over the measured region in our analysis). In addition, the detuning Δ of the laser frequency from the atomic resonance is affected by the Doppler shift due to atomic motion in the direction of the laser beam. To minimise Doppler shifts, the excitation beam is oriented orthogonal to the atomic beam, but residual Doppler shifts still result from the divergence of the atomic beam about the central axis. This means that the atomic density n will be split into classes $n_\Delta(\Delta, T) = np_\Delta(\Delta, T)$, where $n_\Delta(\Delta, T)d\Delta$ is the number density of atoms for which the laser is detuned from the atomic resonance between Δ and $\Delta + d\Delta$, and therefore $p_\Delta(\Delta, T)$ is the corresponding probability distribution.

Integrating over the collection volume, the expected count rate at a temperature T is given by

$$C(T) = \eta_d\eta_g\eta_p \int_{V_C} \int r(\Delta) n_\Delta(\Delta, T) dV d\Delta. \tag{8.5}$$

Rearranging, the following expression converts the measured count rate into the number density

$$n(T) = \frac{C(T)}{\eta_d\eta_g\eta_p \int_{V_C} \int r(\Delta) p_\Delta(\Delta, T) dV d\Delta}. \tag{8.6}$$

Therefore, by measuring the count rate and comparing to the thermal model that predicts number density as a function of temperature, both the number density in

the collection region and temperature of the source can be inferred as a function of heating laser power.

8.1.6 Thermal Models for Atomic Sources

The temperature behaviour of the thermal sources was fitted with a simple thermal model. This model treats the source's thermal losses as resulting from conduction with conductivity K and radiation with emissivity ϵ_{rad} from a single body with one temperature T. This results in a power loss

$$P_{\text{loss}} = K(T - T_0) + \epsilon_{\text{rad}}\sigma S(T^4 - T_0^4), \tag{8.7}$$

where σ is the Stefan-Boltzmann constant, S is the surface area of the oven and T_0 the ambient temperature of the surroundings (assumed to be room temperature). Equating of the power absorbed to the power lost gives an expression linking the input laser power P_{laser} to the temperature of the source that is inferred from measurements,

$$P_{\text{laser}} = \frac{P_{\text{loss}}}{\epsilon_{\text{abs}}} = \frac{K}{\epsilon_{\text{abs}}}(T - T_0) + \frac{\epsilon_{\text{rad}}}{\epsilon_{\text{abs}}}\sigma S(T^4 - T_0^4), \tag{8.8}$$

where ϵ_{abs} is the fraction of heating laser power that is absorbed. ϵ_{abs} may differ from the emissivity coefficient ϵ_{rad} because the heating beam and radiated light will have different wavelengths, and because radiation will occur over the whole surface, whereas heating is only directed at one part. Using the model, measurements of temperature as a function of laser power can be used to fit coefficients K/ϵ_{abs} and $\epsilon_{\text{rad}}/\epsilon_{\text{abs}}$. These indicate the mechanisms by which the atomic source loses heat, allowing for the behaviour of the source to be extrapolated to temperatures at which the neutral fluorescence cannot be directly measured, and indicated where improvements can be made to improve thermal efficiency.

8.2 MITAS

8.2.1 Introduction and Construction

The first source tested was named 'MITAS' (Miniature Ion Trap Atomic Source), and was tested as part of a project with ColdQuanta. The source was tested as part of a compact vacuum system with integrated UHV pumps; the manufacture of the

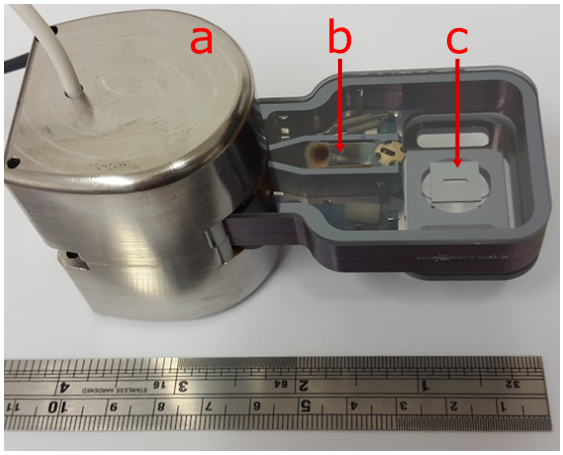


Figure 8.1: The ultrahigh vacuum system containing the MITAS source, showing (a) Penning cell ion pump; (b) optically-heated calcium oven; (c) the dummy trap target region.

system was performed by ColdQuanta and will not be discussed in detail. The oven itself was a stainless steel cylinder, 10 mm long, 2 mm in diameter and with a tube thickness of 0.1 mm, which was filled with approximately 20 mg of granular calcium. Both ends were closed with stainless steel plugs, one of which had a 0.75 mm aperture through its centre. The oven was mounted on a framework of glass rods and folded steel supports (of proprietary design) to reduce thermal conductivity between the oven and the rest of the vacuum system. The atomic beam from the oven passed through a second 0.75 mm-aperture on its way to the target region, which was a slot in a silicon wafer designed to mimic an ion trap.

The heating laser had a wavelength of approximately 780 nm, and was capable of supplying up to 7 W. The heating laser was focused onto a small piece of the oven tube that was intentionally darkened to increase absorption.

8.2.2 Collection of Fluorescence

Upon heating, the oven produced a beam of calcium atoms, which was detected through resonance fluorescence spectroscopy on the $4s^2\,{}^1S_0 \rightarrow 4s4p\,{}^1P_1$ transition at a wavelength of 423 nm. To generate atomic fluorescence, a resonant beam of $1/e$ amplitude waist $w_0 = 40\,\mu$m from an external-cavity diode laser (Toptica DLPro HP) was passed through the target region orthogonal to the atomic beam. The resonance fluorescence emitted by the atoms interacting with the excitation beam was collected

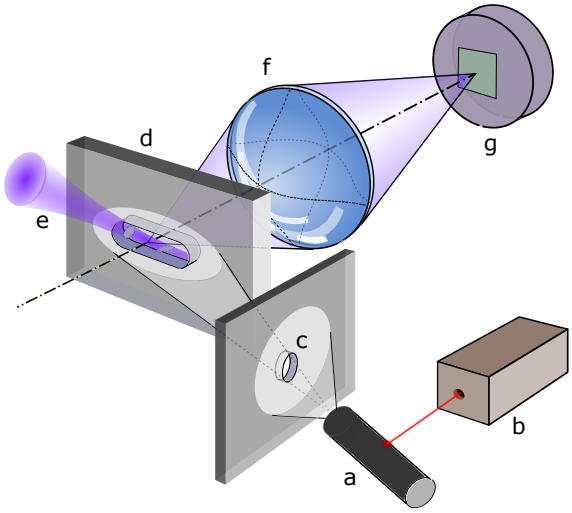


Figure 8.2: Schematic diagram of apparatus for optical heating of calcium oven and resonance fluorescence spectroscopy of emitted atomic beam (figure reproduced from [135]). (a) Stainless steel oven tube filled with calcium; (b) 780 nm optical heating laser; (c) atomic beam passing through second collimation aperture; (d) dummy ion trap denoting the target region; (e) 423 nm spectroscopy laser beam orthogonal to the atomic beam; (f) NA ~ 0.3 collection lens that images with magnification $m = 10$ to (g) photon-counting PMT detector. Figure produced by Dr. Shaobo Gao and reused with permission.

at 45° to the excitation beam with an imaging system of NA 0.299(3) and focused onto a photomultiplier tube (PMT) at a magnification of 10. A 500 µm aperture was placed immediately before the PMT so that the photon collection volume had a diameter of 50 µm. A schematic diagram of the oven and optical systems is given in Fig. 8.2.

The heating laser power applied to the oven was varied and, from the count rate measured on the PMT, the number density in the target region was determined through Eq. 8.6. The imaging system had a detection efficiency $\eta_d = 25(3)\%$ (including losses in the optics), with a geometric efficiency $\eta_g = 2.28(5)\%$ determined by the numerical aperture. The uncalibrated but constant polarisation of the excitation beam means the parameter encoding the strength of emission in the direction of the detector $\eta_p = 1.1(4)$ has a large uncertainty. Despite this systematic uncertainty affecting the number density, the temperature can be determined relatively precisely because the number density is such a strong function of temperature.

The excitation field is orthogonal to the atomic beam axis, so the only Doppler shifts Δ are due to the divergence of the atomic beam across the interaction

region, with

$$\Delta(\theta, T) \approx 2\pi\theta \left(\frac{\bar{v}(T)}{\lambda} \right), \tag{8.9}$$

where θ is the small angle between the atom velocity and the axis of the atomic beam in the plane containing the axis of the excitation beam, and λ is the wavelength of the excitation field.

To estimate the Doppler shift probability distribution $p_\Delta(\Delta, T)$ the collection region is assumed to be point-like given its size and distance from the source. An atom traversing the collection region therefore has a direction determined solely by the point in the oven aperture from which it originated. The collimation aperture is circular, and assuming a uniform emission probability over this circle, the Doppler-shift probability distribution

$$p_\Delta(\Delta, T) = \frac{2}{\pi\Delta_{max}} \sqrt{1 - \left(\frac{\Delta}{\Delta_{max}} \right)^2}, \tag{8.10}$$

can be calculated, where $\Delta_{max}(T) = \Delta(\theta_{max}, T)$ is the maximum Doppler shift, achieved for atoms with the greatest angle $\theta_{max} = 0.044$ between the atomic velocity and beam axis, derived from the size of the first aperture (0.75 mm) and the distance to the collection region (17 mm).

The collection region V_C, and the intensity of the excitation beam over that region is determined by the beam geometry. The excitation beam had a waist of 40 µm, and thus a Rayleigh range of $z_0 = 12$ mm, which is much longer than the length of the excitation beam in the collection region. Therefore, the excitation beam is considered to have constant waist for all points where fluorescence is collected.

The imaging system has a magnification of $M \approx 10$ and passes through a 500 µm diameter aperture immediately before the PMT. Therefore counts can be collected from up to 25 µm away from the imaging axis in the imaging plane. The depth of field is assumed to be infinite because it is much larger than the length along the imaging axis over which atoms are excited by the laser.

With these assumptions, the collected volume and the intensity within that collection volume can be defined. To do this, two coordinate systems are set up according to Fig. 8.3. The coordinates x, y and z are belong to the imaging system,

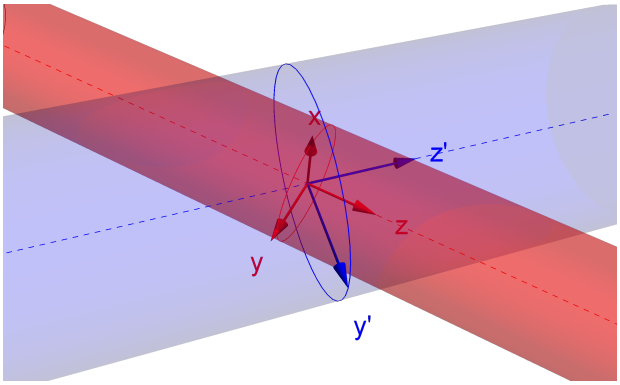


Figure 8.3: Diagram of the coordinate system geometry used to calculate the excitation intensity over the collection volume. The collection region is shown in red and the excitation beam (shown here as a cylindrical contour) is blue. The atoms are excited in the overlap of the excitation beam and the collection.

with z along the imaging axis and the origin in the focal plane. The coordinates x', y' and z' belong to the excitation beam, with z' along the beam axis and x' and y' as transverse coordinates. The two coordinate systems share origin and x' is chosen to coincide with x without loss of generality.

The intensity of the excitation field in its coordinate system can be written

$$I(x', y', z') = \frac{2P}{\pi w_0^2} \exp\left[-\frac{2\left(x'^2 + y'^2\right)}{w_0^2}\right], \tag{8.11}$$

where P is the laser power and there is no z' dependence because the Rayleigh range is large. The transformation to the collection coordinate system is achieved through $x' = x$ and $y' = y\cos\alpha + z\sin\alpha$ with $\alpha = 45°$ the angle between the excitation beam axis and imaging axis. Finally, the imaging Cartesian coordinates x, y are converted to cylindrical polar coordinates ρ, ϕ using $x = \rho\cos\phi$ and $y = \rho\sin\phi$ to more simply encode the radial clipping of the PMT aperture in the imaging plane. This leaves the intensity as a function of imaging coordinates

$$I(\rho, \phi, z) = \frac{2P}{\pi w_0^2} \exp\left[-\frac{2}{w_0^2}\left((\rho\cos\phi)^2 + (\rho\cos\alpha\sin\phi + z\sin\alpha)^2\right)\right]. \tag{8.12}$$

This intensity can be used in Eq. 8.6, with the limits on the collection volume $0 < \rho < 25\,\mu\text{m}$, $0 < \phi < 2\pi$, and $-\infty < z < +\infty$.

8.2.3 Thermal Performance

The oven was heated with varying the heating laser power P, and the fluorescence at the centre of the beam was measured when excited directly on resonance. By

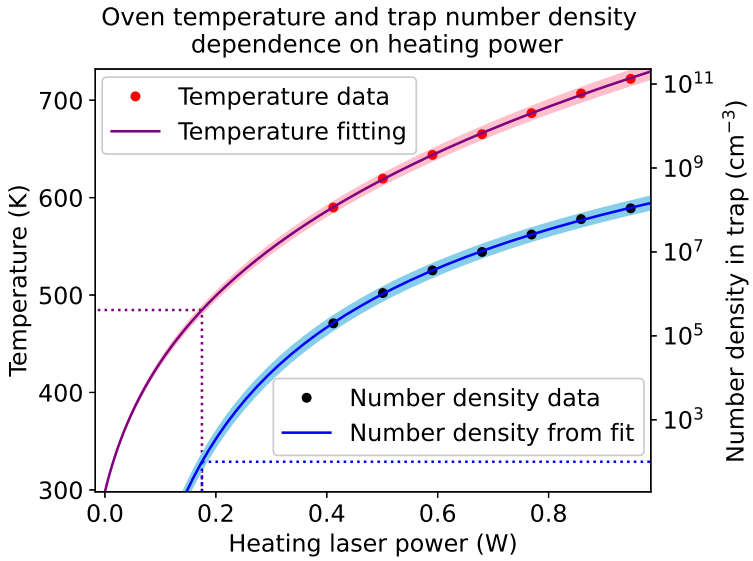


Figure 8.4: Dependence of oven temperature and atomic number density in the collection region on optical heating power. A simple thermal model is fit to this data with shaded regions around the fit line indicating the $\pm 1\sigma$ bounds of the model fit, with the uncertainty dominated by the unknown but constant polarisation of the excitation beam. This fit predicts that 175 mW of heating power is required to achieve the target atomic density ($100\,\mathrm{cm}^{-3}$) in the trap region, whose fluorescence is too low to measure directly with the apparatus.

inferring number density from the measured fluorescence with Eq. 8.6, and using the predicted dependence of number density on temperature (Eq. 8.3), the number density in the collection region and temperature of the source was calculated as a function of laser power, with results shown in Fig. 8.4.

A thermal model of the form of Eq. 8.8 is fit to the dependence of source temperature on laser power, yielding the parameters

$$\begin{aligned} \frac{K}{\epsilon_{\mathrm{abs}}} &= (5.7 \pm 6.2) \times 10^{-5}\,\mathrm{WK}^{-1}, \\ \frac{\epsilon_{\mathrm{rad}}}{\epsilon_{\mathrm{abs}}} &= 0.89 \pm 0.07. \end{aligned} \tag{8.13}$$

The fitted thermal model allows for the measured data to be extrapolated to number densities which could not be detected directly using the apparatus. According to the fitted thermal model, a target density of $100\,\mathrm{cm}^{-3}$ can be achieved with a source temperature of 485(1) K, requiring a heating power of 175(10) mW. Despite large uncertainties in the measured atomic density, these parameters can be predicted relatively precisely due to the strong dependence of atomic density on temperature

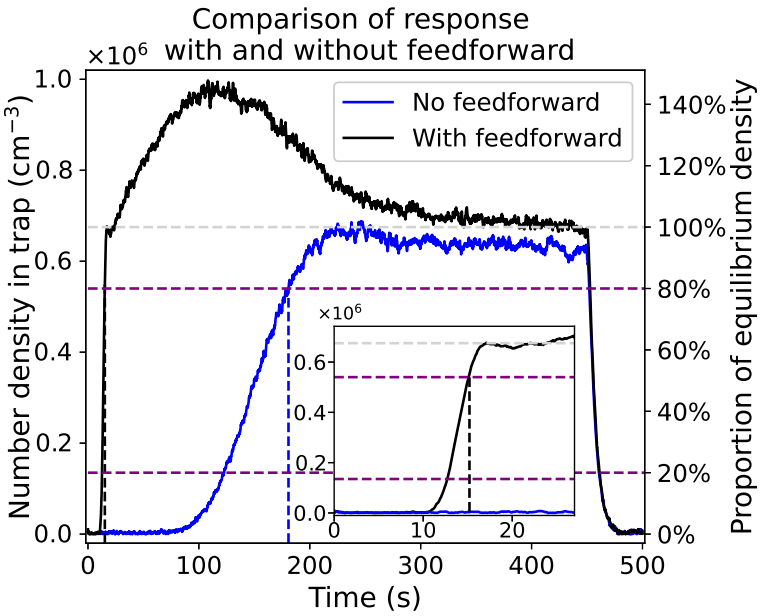


Figure 8.5: Temporal response of number density in collection region after heating laser turn-on, with and without feedforward to optical power. At a constant heating power of 500 mW, the oven reaches 80% of the equilibrium output in ∼3 min. With a simple two-level turn-on profile (3 W for 11.5 s, 0 W for 6.5 s and 500 mW until turned-off) this response time can be reduced to ∼15 s (as shown more closely in the inset). Upon turning off, the number density reaches 20% of its steady state value in 11 s.

and hence power. Radiative losses are predicted to dominate at typical operating temperatures, with conductive losses consistent with zero and comprising roughly 10% of losses at 485 K.

8.2.4 Temporal Response

In addition to the steady state behaviour of the oven, the temporal response of a source is important as it limits how quickly an ion can be loaded when starting from a cold source. To understand the temporal response, the fluorescence counts (and thus number density) were measured as a function of time for a constant 500 mW of heating power applied from $t = 0$s onwards, and found the number density reached 80% of its steady-state value ($\sim 6.8 \times 10^5\,\mathrm{cm}^{-3}$) after 181 s, as shown in Fig. 8.5. The turn-on time reduces if more power is used while the oven is cold to heat it rapidly. Specifically, a simple feedforward sequence consisting of 3 W of laser power applied for 11.5 s, 0 W for the next 6.5 s and the equilibrium 500 mW until turn-off was tested. This procedure takes only 15 s to reach 80% of the steady state value, and is shown in black in Fig. 8.5. It is likely that an optimised sequence could

turn on the oven even quicker and with less overshoot, but the existing feedforward scheme demonstrates the principle. In both cases, the number density reduces quickly after the power is turned off, reaching 20% of the equilibrium level in 11 s. The temperature decreases at a much lower rate, but a reduction in temperature of just a few Kelvin is enough to drastically reduce the number density.

8.2.5 Predicted Lifetime

Using the relation between number density in the collection region and flux density out of the oven (Eq. 8.1), the total flux

$$J_{\text{atom}} = \Phi A_o = n(T)\bar{v}(T)\pi r_0^2, \tag{8.14}$$

can be calculated from the number density in the collection region. To achieve the target density, the oven emits $J_{\text{atom}} = 46 \times 10^6 \text{s}^{-1}$, at which flux the 20 mg of calcium in the oven should last for 200,000 years, far longer than the expected lifecycle of most experiments.

8.3 'Pico Oven'

After the promising results of the previous atomic source, development of an even smaller atomic source was undertaken. The aim of this investigation was to develop sources with even lower power requirements that will be useful for future ion-trap experiments within the group. This investigation took place in separate apparatus, but the tests conducted were broadly similar.

8.3.1 Source Design

The prototype source was fabricated using femtosecond laser writing followed by etching, according to the design of Dr Joseph Goodwin shown in Fig. 8.6. This produced a two-part design consisting of (i) a central cylindrical crucible (of inner diameter 0.75 mm and depth 1 mm) separated from outer housing by high thermal impedance spokes, and (ii) a crucible cap with similar spokes and housing. After loading the crucible with the source metal, the cap could be placed on top and held with light pressure, which would in turn seal the oven via a compliant mechanism

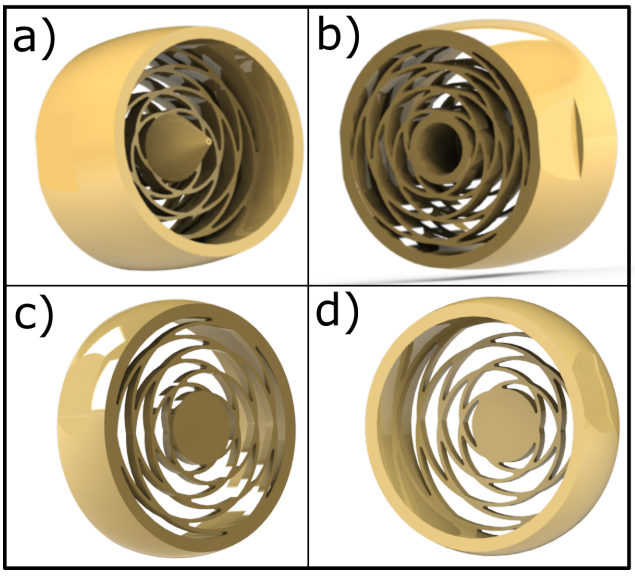


Figure 8.6: Images of the oven designs. The front 'collimator' piece of the design viewed from in front a) and behind b). This piece contains a central volume for source metal and a thin collimator through which the atomic beam is to be emitted. The back 'cap' piece, again from in front c) and behind d). A raised platform of the cap seals the crucible of the collimator under tension when the pieces are pushed together. Both pieces are attached to their outer housing through long, thin glass spokes in order to reduce the impedance of the thermally conductive paths.

in the cap piece. The sealed crucible has one exit route, which is through a long collimator on the front of the crucible that acts to produce a narrow atomic beam. The collimator is 1 mm long and 30 µm in diameter, producing a collimator transmission probability $W = 0.0376$.

Compared to the MITAS source, this atomic source presents several compelling advantages. Firstly, it is smaller, and therefore may be placed closer to the ion trap. This reduces the size of the atomic beam as it passes through the loading zone, meaning a lower flux from the oven is required for a given loading rate. Secondly, despite the complex physical form of the parts, assembly is very simple; the source consists of just two pieces that require only small pressure applied over a wide area to seal. Thirdly, the thermal isolation is intrinsic to each part of the source by virtue of the spokes, rather than with a fiddly and fragile heterogeneous assembly as for MITAS. Finally, the integrated collimator massively reduces the outflow of atoms required to sustain a given number density at the target, offering very long lifetimes despite the low volume of source metal enclosed, and an extremely low impact on the background pressure during operation.

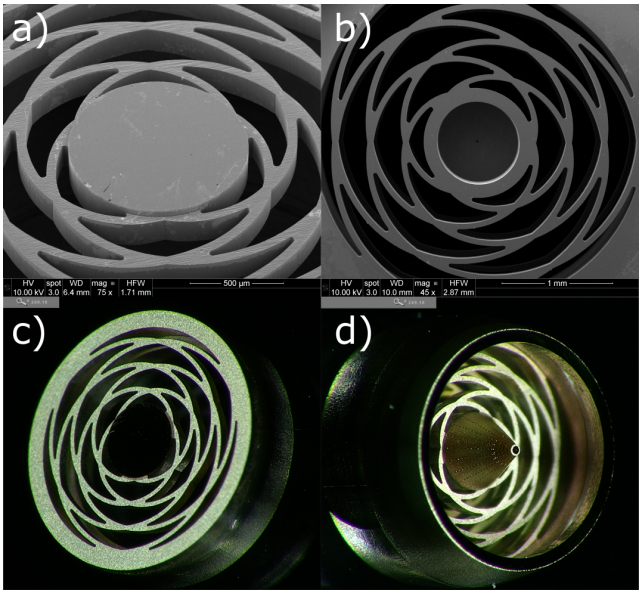


Figure 8.7: SEM and optical images of the cap and collimator pieces. a) SEM image of the cap showing the raised platform that seals the crucible. b) SEM image of the crucible on the collimator part. c) and d) optical images of the cap and collimator parts respectively after coating. SEM images taken by Dr Jason Brown.

8.3.2 Source Preparation

To prepare the samples for use, both of the pieces were hand agitated for 7 minutes in acetone, isopropanol and deionised water sequentially before undergoing 2 minutes of plasma etching to remove any remaining contaminants. The pieces were then sputter-coated with titanium (10 nm) and gold (100 nm) on both sides. Images of the samples are shown in Fig. 8.7.

To load the oven, a small bead of mesh -16 granular calcium was selected and placed with tweezers into the crucible. After inspection to ensure that the metal was not protruding from the crucible and that there was no more space to insert another granule, the cap piece was placed on top of the crucible. Fig. 8.8 shows calcium metal in the crucible of our prototypes.

8.3.3 Oven mounting and Beam Geometry

The collimator and cap parts are located in a threaded holder screwed into a mount with a central threaded hole. To seal the crucible, pressure is applied on the cap piece using a spring (Springmasters D20800) sandwiched between two PTFE cylinders to make the force uniform. The spring is kept under tension with a custom grub

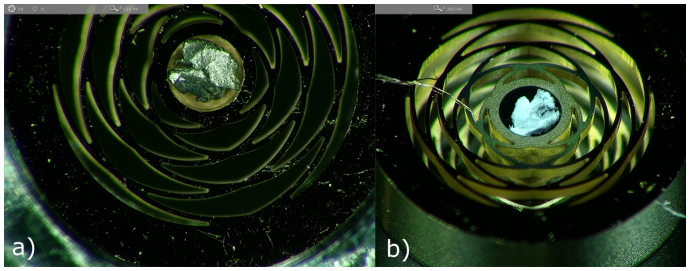


Figure 8.8: Images of the crucible containing fresh calcium (left) and calcium left to oxidise over several days (right). Even upon brief exposure to atmosphere, an oxide layer forms that must be cracked before the oven is operational.

screw threaded into the back of the holder and tightened to a calculated position predicted to seal the oven. The holder is mounted to a modified blanking flange at the base of a vacuum chamber.

This mounting structure leaves access for the heating to be applied directly down the collimator, with the atomic beam propagating back along the same direction. The excitation beam runs across the diameter of the chamber orthogonal to the atomic beam to reduce Doppler shifts. The atomic fluorescence is collected vertically through a large viewport in the top of the chamber. The oven mounting mechanism and beam geometry are shown in Fig. 8.9.

Before loading, the oven holder, vacuum chamber and internal screws and dowel pins were UHV-cleaned in detergent, acetone, isopropanol and deionised water in an ultrasonic bath for 7 minutes in each liquid, while the uncoated viewports were hand-agitated. The coated viewports were not cleaned as they are supplied clean. The system was then attached to its ion pump and baked at 200° for 2 days. In operation, the pressure of the system ranged from 10^{-7}mbar to 10^{-9}mbar, having been opened regularly to switch oven prototypes. This level of background gas is far below the level of gas required to affect the atomic beam dynamics (pressures of 10^{-5}mbar are sufficient to give mean free paths in excess of 1 m).

8.3.4 Heating and Oven Imaging System

Heating is provided by a 785nm diode laser (Thorlabs FPL785S-250), capable of supplying 250mW of optical power into a single mode fibre. This fibre is attached to the heating system, which is described in Fig. 8.10.

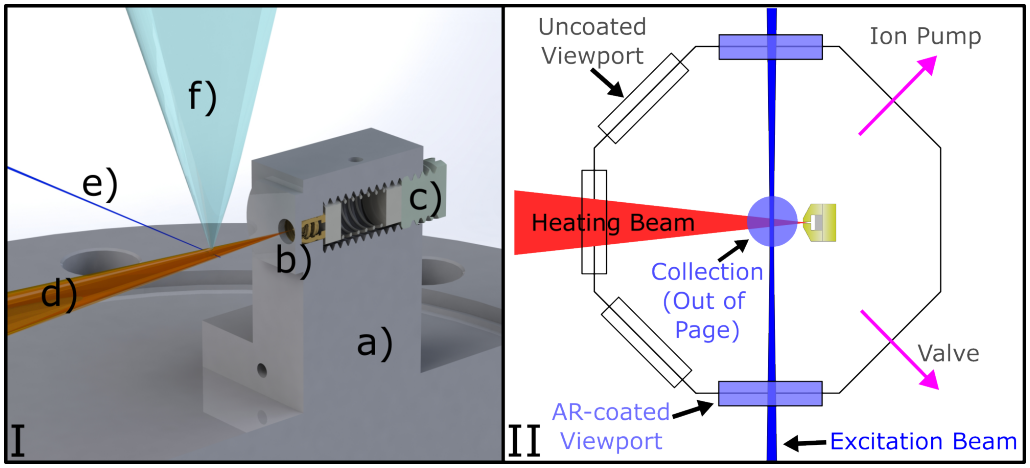


Figure 8.9: Diagrams showing the orientation of the beams and the wider vacuum chamber. I: Diagram of the source mounting and beam geometry inside the chamber. Labelled are a) the source mount b) the source in its threaded holder c) the sealing mechanism consisting of two PTFE cylinders, a spring and a grub screw. The beam geometry consists of d) the heating and atomic beams which are coaxial, e) the excitation beam that causes the atoms to fluoresce and f) The cone of the photons collected. II: Diagram of the vacuum chamber, with the heating, excitation and collection beams marked. The centre of the chamber is the octagon (Kimball Physics MCF450M-SphOct-Ec2A8). At its base is a blanking flange which was modified to mount the source holder. The heating beam viewport and the two either side are uncoated viewports (Scanwel CFVPZ034), whereas the excitation beam enters and exits via anti-reflection coated viewports. The fluorescence is collected through the top viewport (Scanwel CFVPZ114). The remaining ports of the octagon are used to attach the valve for pumping down (Lewvac AVM16CF), the ion pump (Agilent X3602-64001) and one blanking flange.

The primary function of the heating setup is to focus the optical power into the atomic source collimator, but the design has two key features offering significant practical advantages.

The first is the photodiode pickoff that that monitors the optical power through the system. By calibrating measurements of beam power against photodiode voltage and accounting for transmission losses into the chamber, the pickoff allows for the heating power to be monitored continuously without interrupting experiments.

The second feature is the imaging arm, confocal to the heating beam, that uses light reflected from the oven to produce an image on a camera. This enables direct imaging of the heating beam at its focus, making it easy to adjust the heating beam as required. To align the heating beam, first, the illumination source was turned on and the longitudinal adjuster of the focusing lens used to bring the image of the front of the collimator into focus on the camera, which would in turn move the focus

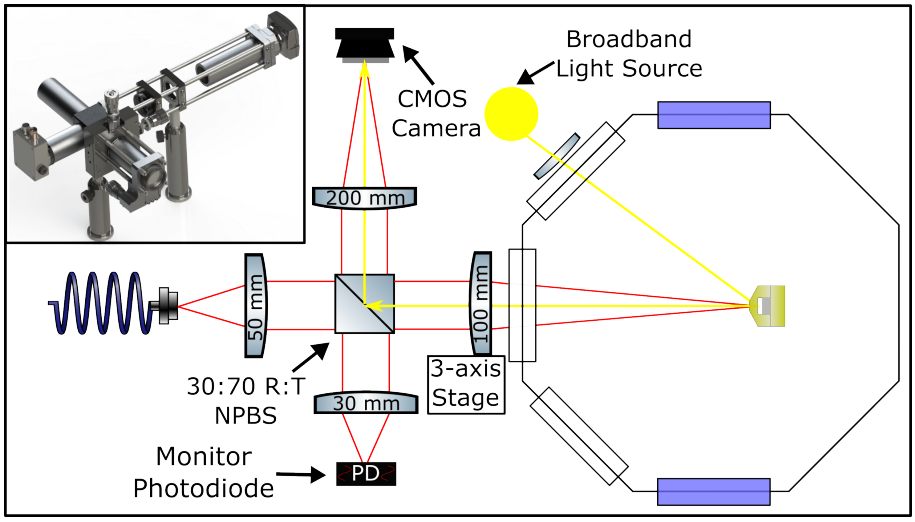


Figure 8.10: Diagram of the heating beam setup used to apply optical power to the atomic source, with the inset showing a render of the system. The optical power is collimated out of the fibre (Thorlabs AC254-050-B) and sent onto a 30:70 R:T non-polarising beamsplitter (Thorlabs BS020). The reflected light is attentuated and focused (Thorlabs AC254-030-B) onto a photodiode (Thorlabs PDA36A-EC) to monitor the power. The transmitted light is focused onto the oven collimator with a lens (Thorlabs AC254-100-B-ML) that can be adjusted transverse (Throlabs ST1XY-D/M) and parallel to the beam (Thorlabs SM1ZA) for fine control over the focus position. Any light reflected from the chamber is directed through the beamsplitter to an imaging system, where the light is focused (Thorlabs LA1708-B-ML) onto a CMOS camera (Thorlabs DCC1545M). The oven surface may additionally be illuminated through a viewport at 45° to the heating beam using a lamp (Thorlabs QTH10/M). Scatter from this illumination is imaged on the camera.

of the heating beam to this plane. Then the focusing lens was adjusted transverse to the heating beam to steer it onto the centre of the collimator. In addition, magnetic mounts allow for the simple switching of the camera and a power meter to measure the heating power reflecting from the oven surface. By comparing the reflected heating power when the heating beam was directed down the collimator to the when heating beam reflected off the surface immediately surrounding the collimator output aperture, it was found that 97% of the light was passing into the collimator.

8.3.5 Excitation-Collection System

To collect fluorescence from the atomic beam, the atoms need to be excited, and the resulting fluorescence detected. The excitation-collection system performs these two functions, with the excitation beam orthogonal to the atomic beam to reduce Doppler shifts, and the collection direction orthogonal to both. Because the fluorescence need only be collected from the atoms that are excited, the excitation and collection

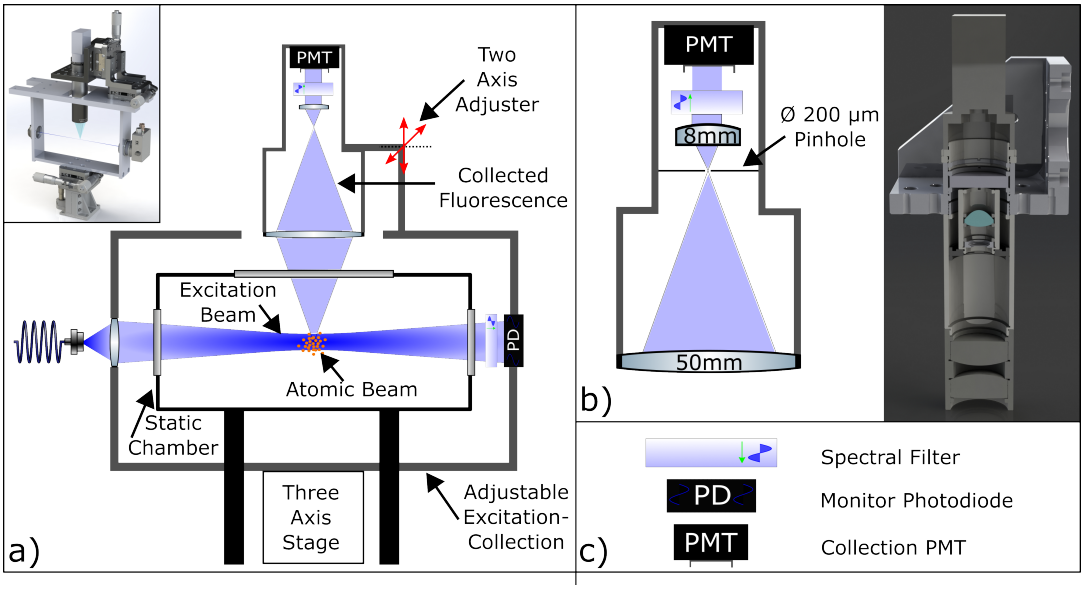


Figure 8.11: Diagrams of the excitation-collection setup. a) Diagram of the entire excitation-collection system, with render inset. b) Diagram and render of just the collection part of the apparatus. c) Key for symbols used in the diagrams. The excitation light is focused (Thorlabs PAF2-A4A) to a 40 µm waist in the centre of the chamber where it intersects with the atomic beam, producing fluorescence. The excitation beam then passes out of the chamber through an AR coated viewport, followed by a spectral filter (Thorlabs FES0650) to remove scattered heating light and onto a calibrated photodiode (Thorlabs PDA36AEC) to monitor the power of the excitation beam in real time. The atomic fluorescence is collected through the top of the chamber with a NA 0.301 lens (Thorlabs MAP105050-A, f =50 mm). This projects the light to a secondary focus with magnification 1, where a 200 µm diameter pinhole (Thorlabs P200HD) is used to restrict the field of view of the imaging system. The light is then collimated (Thorlabs ACL12708A, f =8 mm) through a narrow-band spectral filter (Semrock FF01-433/24) to remove stray heating or ambient light, and counted on a PMT (Hammamatsu H10682-210). The focus of the collection system can be aligned in the two transverse directions of the excitation beam (Throlabs LX10/M) to make sure the system is collecting light from the most intense part of the excitation beam. Once this is calibrated, the entire excitation-collection system is then scanned as a single unit both vertically (Thorlabs MVS005/M) and in the two horizontal directions (Thorlabs LX20/M) to allow for the atomic beam to be profiled at different points.

systems are mounted to the same rigid frame so that they can be scanned in unison across the atomic beam to study the beam shape and extent. The configuration of the combined excitation-collection system is described in Fig. 8.11.

To set up the system, the fibre collimator that focused the light to a waist of 40 µm at the centre of the chamber was adjusted before mounting to focus at the desired distance from the fibre facet in free space (with the desired distance mapping to the centre of the chamber, accounting for focal length shift through the viewport).

A second fibre port was adjusted this way and temporarily mounted in place of the excitation-monitoring photodiode, where good coupling from the input fibre port to the second fibre port showed that the input light was well-directed across the chamber, and produced approximately the correct waist diameter.

The excitation beam polarisation is orthogonal to the collection axis to maximise the collected counts. If light were collected only directly along the collection axis, η_p (the probability of emission in the direction of collection compared to an average over all angles) would be 3/2. However, the collection over a small range of angles slightly reduces this probability to 1.47 (see App. G).

The geometric efficiency $\eta_g = 2.35\%$ of the imaging system is calculated from the clear aperture 22.9 mm and working distance 37.8 mm of the collection lens.

The detection efficiency η_d was predicted to be 0.23 on the basis of the PMT datasheet and expected transmission losses in the path. This efficiency was calibrated by placing an uncollimated excitation beam fibre at the focus of the collection system and collecting the photons from a beam of known power onto the PMT (through a calibrated neutral density filter). The detection efficiency was indistinguishable from that predicted within the 5% error on the beam power (the beam power used to take data was much more stable, but the calibration procedure mechanically disturbed the fibre from the excitation laser, causing the uncertainty on this calibration).

8.3.6 Initial Alignment

The experimental protocol requires that the excitation beam passes through the atomic beam, and the collection system is aligned to collect photons emitted from those excited atoms. Once the system is collecting photons from the atomic beam, the alignment of the system is simple, iteratively scanning (i) the position of the complete excitation-collection rig, and (ii) the relative alignment of the collection optics with respect to the excitation beam, to maximise the signal. However, before the initial observation of the signal, the three regions (atomic beam, excitation beam and collection) must be co-located with sufficient accuracy to detect fluorescence in the presence of atoms. To do this, the heating beam was first pointed so that it was focused onto the oven collimator. Then the vacuum chamber was temporarily

removed and the excitation turned on. A piece of paper was used to image both excitation and heating beams, and the excitation-collection rig adjusted vertically until they overlapped. Finally, a hair was placed at the intersection of both heating and excitation beams, and the collection adjusted relative to the excitation to maximise the collected signal. The chamber was then returned and the heating beam tweaked to focus again on the oven collimator. In practice, this procedure was able to achieve a level of alignment that gave sufficient signal to proceed with optimisation immediately after the oven had cracked.

8.3.7 Experimental Control

The experiment is controlled using the artiq control system [295] and a simple set of Sinara hardware. The control hardware contained a 'Kasli' board to schedule experiments, a 'DIO SMA' board to collect counts and trigger the laser scan, a 'Sampler' to read the monitoring photodiodes and the laser scan voltage and a function generator (Agilent 33220A) to produce the voltage for the laser scan.

8.3.8 Cracking

The oxide layer that forms on the loaded calcium before the chamber is evacuated must be cracked before the oven can operate. This was achieved by heating the source with high power, and waiting for a spike in fluorescence which indicates that the oxide layer has been fractured and the bare metal is emitting atoms. It is vital that this sudden emission of atomic flux produces detectable fluorescence as the source metal depletion is many orders of magnitude faster at cracking temperatures than at operation temperatures. Thus, an excitation beam waist of $\sim 1\,\mathrm{mm}$ was used during cracking to make sure that the excitation beam was intersecting the atomic beam. It is also possible to modify the collection to place the PMT directly at the focus of the relay lens, but for the final prototype this was unnecessary. For the final prototype, the spike in fluorescence occurred at a power of $325\,\mathrm{mW}$. Upon observing this spike the heating power was quickly reduced and the excitation and collection alignment optimised to maximise counts.

8.3.9 Spectroscopy

For the MITAS experiment, where theory predicted the angular distribution of the atoms exiting the aperture, the sole measurement of interest was the fluorescence collected from the centre of the beam for a laser tuned onto resonance. For this source, however, a key result, both in its own right and to calculate the thermal behaviour of the source, is the transverse shape of the beam; rough theoretical predictions (such as the angular halfwidth $H \approx 1.70/G$) do exist, but it was decided not to assume that they held in this system during the analysis. Thus, the transverse profile of the beam had to be measured, and there was no prior-known beam shape from which to predict the Doppler broadening. For this reason, each measurement taken was a spectral measurement in which the count rate was recorded as the laser frequency scanned over the resonance. For each run, the Doppler frequency distribution ($p_\Delta(\Delta, T)$) was assumed to be a Gaussian distribution with a width that was fitted independently for each trial, and the number density was then inferred from the fit.

To run the spectroscopy sequence the excitation laser frequency was scanned by applying a linear ramp voltage to spectroscopy laser's piezo adjuster. The voltage of this scan and the fluorescence rate were recorded in real time. In addition, the wavemeter (High Finesse WS6) monitored the change in laser frequency. While this monitoring was not fast enough to assign a frequency to each measurement as it was made, a quadratic fit of frequency to scan voltage was later applied to convert the scan voltage to a frequency. This procedure gives a good estimate of the frequency, but there is still some distortion. This can be seen most directly in Fig. 8.12, where the predicted frequency of the isotope shift between calcium 40 and 44 is shifted by roughly 13%. For the purposes of this experiment, such a correction would not have important quantitative implications because the number density changes by orders of magnitude in the data. However, more precise measurements could be obtained by locking the laser to a cavity and scanning with, for example, an AOM.

For every measurement, the open viewports of the chamber are shielded with black card, and the collection system protected from stray light with further card. In practice this is vital for improving the signal to noise ratio of the system.

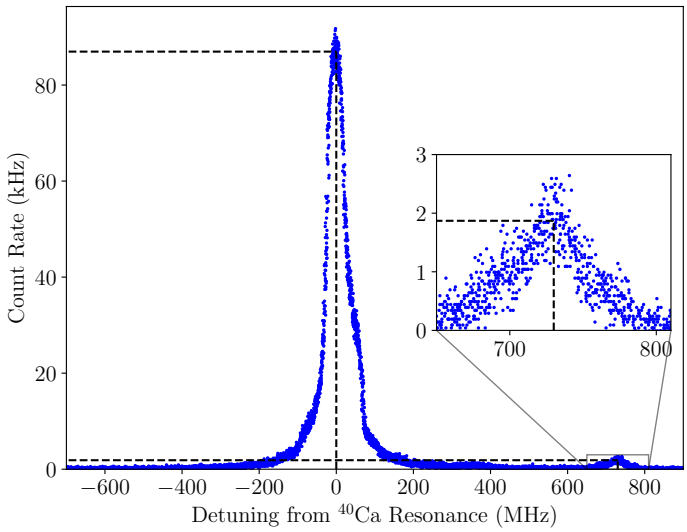


Figure 8.12: Wide spectrum of the atomic beam, showing a central peak from calcium 40 and a second peak from calcium 44. The predicted number fraction from the calcium 44 peak (2.1 %) is consistent with the expected value (2.09%), but the isotope shift from the graph (730 MHz) is significantly different from expected (842 MHz) due to the non-linearity of the frequency scan. Expected values from [277].

8.3.10 Summary of Characterisation

The performance of this atomic source was evaluated by converting the measured count rate to an atomic number density in the same manner as the MITAS source (Eq. 8.6). The parameters used in this equation have been discussed earlier in this section, but are summarised in Tab. 8.1.

Parameter	**Symbol**	**Value**
Geometric Efficiency	η_g	0.0235
'Emission direction factor'	η_p	1.47
Detection Efficiency	η_d	0.23
Collimator Transmission Probability	W	0.0376

Table 8.1: Properties of the pico-oven test system for use with Eq. 8.6.

Using the above details, the number density in the collection region was determined from the measured spectra, which was in turn used to assess the atomic beam profile and the thermal performance of the source.

8.3.11 Results

Structural evaluation

A series of prototypes were evaluated for this investigation. Though the results presented pertain to just the final prototype, the series of tests allowed the mechanical suitability of the source to be tested.

The sources were found to be relatively robust; of the five prototypes tested during the investigation, there were no breakages during coating, cleaning or loading. This may seem surprising for parts formed from thin pieces of glass, but the laser etching process is known to remove surface fractures from the substrate, increasing the failure stress [296]. Manual loading is fiddly, however with standard laboratory equipment e.g. tweezers and microscopes, the chamber can be opened, refilled and under vacuum in less than an hour.

The 'cracking' of the oxide layer is probably the part of the procedure that represents the greatest risk to the source as it is heated to a considerably higher temperature than at operation. For this process, a larger excitation beam waist is useful to make sure that, once the source is cracked and atoms are being emitted, the fluorescence is measured and the heating power reduced quickly. When testing the first prototype, the evidence of 'cracking' was missed and the oven was heated until the structure fractured, which later inspection confirmed had occurred after the calcium inside had melted.

After operation at high temperatures for extended periods, the collimator and cap pieces can fuse together. In operation, this is probably advantageous to prevent the leaking of source metal through any orifice other than the collimator, but it does mean that care must be taken if trying to reuse the sources. Anecdotally, exposing the source to atmosphere for a few days seemed to weaken the bond and thus allow the pieces to be separated and the source reused. This suggests the bond is formed from molten calcium wicking between the pieces and solidifying, rather than a thermocompression bond from the gold surfaces, because the calcium oxidises when exposed to atmosphere.

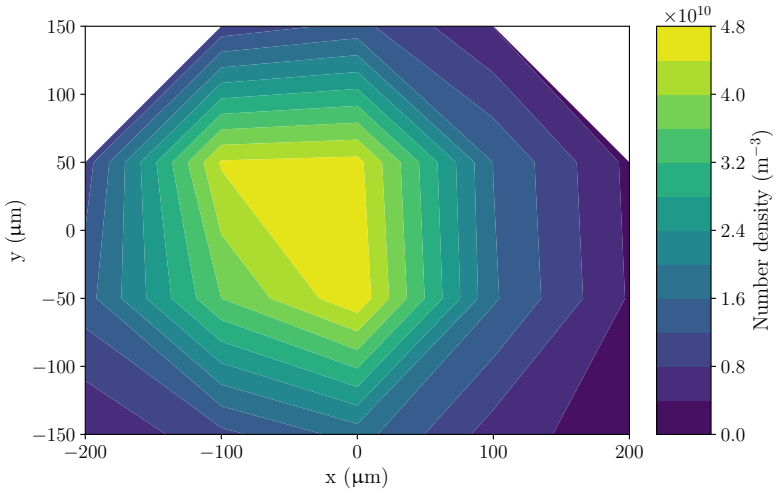


Figure 8.13: The transverse profile of the beam at a distance of 10 mm from the collimator front. The beam has an approximate area of $6.4 \times 10^{-8}\,\mathrm{m}^2$.

Beam profile

The transverse beam profile was measured at a distance of 10 mm away from the collimator end. This measurement was conducted by applying a constant 80 mW of heating power and raster-scanning the excitation-collection system through a transverse plane of the beam, taking a fluorescence spectrum at each point from which the number density was inferred. The beam profile is shown in Fig. 8.13. By performing a discrete integration over the beam profile, an area of $6.4 \times 10^{-8}\,\mathrm{m}^2$ (250 µm$\times$ 250 µm) of the beam was determined. Assuming a Gaussian beam profile, this would correspond to a $1/e$ beam half-angle of 1.7°. This profile can be compared to the area that would be obtained from the cosine distribution of a simple aperture the same distance away of $3.1 \times 10^{-4}\,\mathrm{m}^2$ (18 mm$\times$ 18 mm). Thus, the integrated collimator improves the atomic efficiency of the source by almost four orders of magnitude by emitting a much more directed atomic beam. Note that the length scale of the beam ($\sim$250 µm) does not far exceed the scale of the fluorescence collection volume (200 µm$\times\sim$80 µm$\times\sim$80 µm), and therefore the collection volume would have to be reduced to characterise beams much more compact than this.

Thermal performance

With the beam area determined, the excitation-collection system was then centred on the intensity peak of the atomic beam, again 10 mm from the collimator facet, and

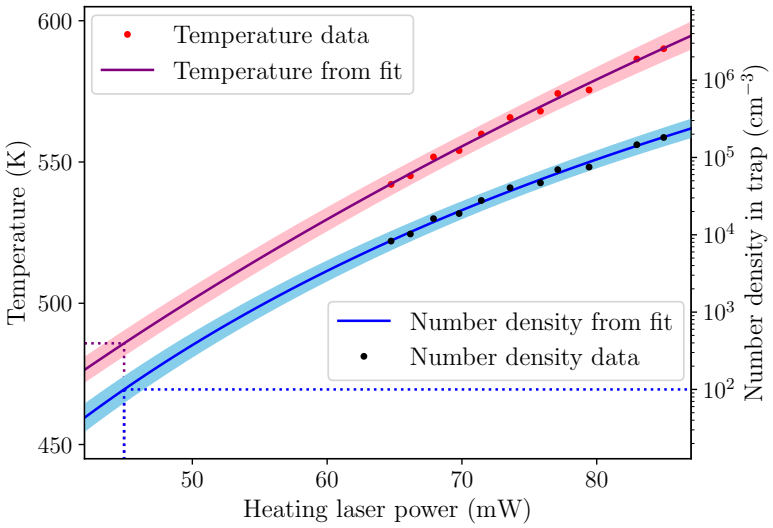


Figure 8.14: The atomic number density 10 mm from the source collimator and the inferred source temperature as a function of applied heating power. A simple thermal model is fit to the temperature data with shaded regions around the fit line indicating the $\pm 1\sigma$ bounds of the model fit. As in the case of MITAS, this uncertainty is dominated by systematic errors, mainly from the area of atomic beam and the waist of the excitation beam. The extrapolated thermal model predicts that 45 mW of heating power would achieve the target atomic density (100 cm^{-3}) in the trap region.

the heating power was varied, allowing the source to come into thermal equilibrium at each point. For each heating power, a spectrum was taken, and the number density inferred as before. These number densities were used to infer the temperature of the source through Eq. 8.3, which was then fit to a thermal model with Eq. 8.8, yielding parameters

$$\begin{aligned} \frac{K}{\epsilon_{\text{abs}}} &= (1.4 \pm 0.1) \times 10^{-4}\,\text{WK}^{-1}, \\ \frac{\epsilon_{\text{rad}}}{\epsilon_{\text{abs}}} &= 0.45 \pm 0.05. \end{aligned} \tag{8.15}$$

This fit predicts a heating power of 45(1) mW is required to reach a source temperature of 486 K and provide the target number density. The conductive and radiative losses comprise 60(5)% and 40(5)% of the total loss at this temperature respectively.

When discussing the results of the thermal model, it is important to note that the metallic coating on the parts was not uniform. It should be close to the nominal thickness of 10 nm titanium followed by 100 nm gold on the faces of the parts (which were normal to the sputtering direction), but the sides of the parts had a notably



Figure 8.15: An optical image of the side of the central crucible and cap, which fused on use and were then detached from the surrounding spokes. The coating is noticeably thinner on the sides of the crucible.

thinner coating (see Fig. 8.15). This thinner coating would have three effects. Firstly, as the thinner coating makes the sides of the pieces semi-transparent, it is possible some heating power applied to the source would not be reflected enough times before escaping to be absorbed by the source. This would reduce the efficiency of heating (i.e. reduce ϵ_{abs}). The second is that the emissivity would be affected; the emissivity depends significantly on factors such as surface roughness, but titanium and glass both have generally higher emissivity than gold, so the radiative losses would likely increase as a result of incomplete coating. The final effect is to reduce the conductive losses as there is less high-conductivity metal deposited on the spokes.

With this in mind, the conductive losses predicted from the structure range from $1\times10^{-4}\,\text{WK}^{-1}$ in the case of uniform spoke coating to $0.5\times10^{-4}\,\text{WK}^{-1}$ in the case that there is no side coating on the spokes. Given that the side coating thickness should lie between these two limits, this suggests that $0.36 < \epsilon_{\text{abs}} < 0.71$; an estimate of 50% will be used for speculation about possible performance improvements in Sec. 8.4.1.

Finally, it should be noted that the behaviour of the source could not always be attributed to a simple thermal model. For example, hysteresis was observed where, after heating to a higher temperature, the oven could not produce the same number density at lower powers as before. The level of hysteresis was fairly insignificant over the orders of magnitude change in number density observed in the experiment, but nevertheless indicates that the true behaviour of the source is more complex than modelled. These hysteresis effects may be due to changes internal to the crucible

(redistribution of the calcium or even local melting) upon applying high heating powers, but it is not possible to test these hypotheses in the current apparatus.

Lifetime

At the target temperature, the oven is predicted to emit 3300 atoms per second. If the entire internal volume of the source was filled with calcium metal, this metal would take 100 million years to deplete under continuous operation, which is not a realistic concern. Despite the smaller size of the source compared to MITAS, the lifetime is predicted to be much longer owing to the greatly improved efficiency in direction of atoms towards the trapping region.

8.4 Discussion

8.4.1 Evaluation of Sources

This chapter has detailed the development of two miniature optically-heated atomic sources. The first source demonstrated that such miniature sources could achieve suitable number densities with much lower powers than typical resistively-heated ovens (for example 1.5 W in [71]) while removing the need for a high-current feed through. The second source built upon this by making the source yet smaller, making the mechanical structure simpler with in-built thermal isolation, and drastically improving the directionality of the source. If there were no further development, the tested prototype would be expected to reduce heating power requirements to 45 mW, while dramatically reducing the total number of atoms emitted into the vacuum system.

Looking ahead, there are a few simple design changes that could improve the performance of the source further. The main improvements would be to gold-coat the central crucible much more thoroughly while leaving the spokes masked, which would simultaneously improve the power absorption, conductive losses and radiative losses. The source could also be moved much closer to the ion trap and even integrated into the structure. The ultimate aim would be to use a collimator with an output aperture of the same diameter as the trapping region and to collimate the thermal beam so that it effectively did not expand before reaching the trapping region (extrapolating

our beam size, this would give a distance between collimator facet and trapping region of <1 mm). In this way, every atom that is emitted has the possibility of being trapped. Given the much-enhanced atomic beam geometric efficiency, the source could potentially be run continuously without poisoning the vacuum. In this scenario, thermal atomic sources would have moved from bulky, complex components that take time to produce atoms and introduce thermal and compensation drifts on each use to small, simple, low power devices, offering instantaneous loading with no intermittent drifts because they are run continuously. With low power requirements and no warm-up time, many of the key disadvantages of thermal sources compared to ablation would be mitigated, while the advantages in avoiding the use of a pulsed laser or encountering charged fragments in the atomic plume are retained.

Assuming that the suggested fabrication improvements are made, and that the oven is integrated sufficiently close for the beam area A_b to match the collimator area A_0, we can imagine a future 'Pico-Oven 2.0' and calculate its characteristics, which is compared to tested sources in Tab. 8.2. The table shows that the source performance has scope to improve further; the fabrication changes could allow for the operating optical power to be reduced to just 11 mW, and the atomic efficiency gained from moving the source closer to the trap lends a further two orders of magnitude to the lifetime. Given the unnecessarily long predicted lifetime, the source could be miniaturised further if desired.

Field	**MITAS**	**Pico Oven**	**'Pico Oven 2.0'**
ϵ_{abs}	N/A	N/A (est. 0.5)	1.0
κ/ϵ_{abs} (WK^{-1})	5.7±6.2 $\times 10^{-5}$	1.4±0.1 $\times 10^{-4}$	5.0 $\times 10^{-5}$
$\epsilon_{rad}/\epsilon_{abs}$	0.89±0.07	0.45±0.05	0.15
Target temperature (K)	485	486	440
Target power (mW)	175±10	45±1	11
Lifetime (years)	200,000	100,000,000	10,000,000,000

Table 8.2: Summary of properties of the two sources discussed in this chapter, and a third source, 'Pico Oven 2.0', with projected performance based upon realistic improvements to the Pico Oven. With simple improvements, the power requirements are predicted to reduce even further, and the atomic efficiency could improve a further two orders of magnitude.

8.4.2 Improvements to Experiment

The experiments detailed in this section were conducted for the purposes of ascertaining whether a miniature optically-heated atomic source would be viable for the group's future ion traps. As such, the main goal of the experiment was to assess the usability of the source and get a rough estimate of the power required before deployment into an ion trap. Due to the large dependence of number density on temperature, this meant that an accurate measurement of the number density was not required, and thus the steps necessary to achieve an accurate measurement of number density were not taken. In the future, it may be worthwhile to optimise the design of these sources, for example to reduce further the power requirements. If characterising the sources for this purpose, improvements to the experiment such as those detailed below may yield more accurate measurements.

The first step would be to improve the frequency scanning of the laser for spectroscopy. In the pico-oven tests, the frequency was scanned by applying a voltage sweep to the piezoelectric actuator of the spectroscopy ECDL and monitoring the frequency on a wavemeter, which was then fitted to synchronise it with the experiment. This approach was good enough to detect frequency differences of roughly 10%, which is sufficient for order-of-magnitude determinations of number density, but not to quantitatively analyse the lineshape. A better solution would be to lock the laser to a stable cavity and scan the frequency with an AOM. With this setup, the frequency accuracy might be sufficient to use the spectral lineshape as a measure of the velocity distribution within the beam, and the central frequency as a measure of the average velocity. In this way, spectral measurements (which essentially view the velocity profile in the direction of the excitation laser) and beam profile measurements could be combined to form a consistent picture of the atomic beam shape and its expansion.

The second key advance would be to mount the excitation-collection system on motorised stages to allow for automated scanning, and thus allow a much higher quantity of beam profile data to be taken. This could be combined with a tighter focus of excitation beam and a smaller diameter collection pinhole to more finely probe of the beam profile.

Thirdly, the powers of the excitation and heating lasers could be stabilised by feeding back on the pick-off photodiodes. This would be particularly important for the heating laser as the number density varies so strongly with heating power.

The laser waist size was the key contributor to the nominal systematic uncertainty in the data. This was because, after setting up the collimators outside of the chamber so that they would produce the desired waist at the desired location (adjusting for the focal shift of the viewport), there was no method of measuring the laser waist in the chamber. This could potentially be solved by mounting a precision pinhole along the beam axis (but further from the source than measurements could be taken), which would allow for the waist of the beam to be determined in-situ through transmitted power, and also calibrate the imaging-collection position to the front of the collimator.

After these steps, the next obvious improvements would be to reduce the signal to noise ratio during measurements. After cardboard shielding of the chamber's open ports, the dominant contribution to count rate noise was caused by the excitation beam. If a vacuum-compatible absorptive coating could be applied to the components on the inside of the chamber, this might reduce the intensity of any residual scatter reflecting around the chamber. To improve the signal, a large numerical aperture collection system could be utilised, either through bigger optics, or through a smaller chamber where the collection objective can operate at a smaller working distance. By combining the above steps and averaging out noise over longer collection times, it may be possible to measure number densities towards the target number density directly, rather than having to extrapolate over two orders of magnitude.

Finally, there are two more outlandish changes that could be made to the setup for a more serious investigation of the behaviour of atomic sources. The first would be to hold multiple sources with the chamber, so that different designs could be tested without the need to open the chamber between each one. Secondly, it may be possible to measure the temperature of the source directly using a thermal imaging camera, which would additionally give information of the temperature distribution around the source. This would however, require special viewports and optics designed for infrared thermal radiation, and a change to the holder design to give a line of sight to the source from the thermal imaging camera.

9 Conclusion and Outlook

9.1 Criteria for Successful Cavity-Enhanced Remote Entanglement

To integrate an optical cavity into an ion trap and perform cavity-enhanced remote entanglement is a challenging goal; while simplified models appear straightforward, there are many complications that could hamper experimental performance and non-trivial decisions to be made. The research presented on these subjects clarifies the set of criteria that should be fulfilled for a viable remote entanglement demonstration.

9.1.1 Criterion 1: Negligible Temporal Mixing

The first key criterion is that the remote entanglement protocol should not suffer significant fidelity loss due to the temporal mixing that results from spontaneous emission followed by cavity emission. There are three main ways to achieve this:

1. Use a cavity with sufficiently high cooperativity that the probability of spontaneous emission is negligible.

2. Drive the photon production so quickly that the probability of spontaneous emission during driving is negligible.

3. Choose an initial state with a low branching ratio from the excited state to reduce the probability that spontaneous emission can be followed by cavity

emission.

Of these, the first option is ideal but if using a Lambda-type system with strong branching on the drive arm (such as the $S_{\frac{1}{2}} \rightarrow P_{\frac{1}{2}} \rightarrow D_{\frac{3}{2}}$ three-level system), this demands a cooperativity beyond that hitherto achieved with trapped ion-cavity systems. The second option is achievable, but requires the use of fast driving pulses, sacrifices a potentially significant proportion of performance, and surrenders the fine control over the photon wavepacket envelope afforded by slower driving. The third option, assumed in this work, uses the level structure to render temporal mixing negligible, but comes at the cost of additional mapping pulses during each entanglement attempt. Finally, even if the three-level systems involved are well-chosen to reduce temporal mixing, it is important to check that other mapping pulses do not re-introduce these errors, as was critical for certain time-bin encoding schemes discussed in Chapter 7.

9.1.2 Criterion 2: Cavity Design Tolerates Residual Misalignment

It is possible to actively align cavities at the point of assembly, or to continuously adjust the transverse alignment of cavity mirrors in-situ, so that the cavity mirrors have small transverse misalignment. However, there are mechanical advantages in stability and simplicity to designs that forego active transverse alignment and integrate cavity mirrors directly into a monolithic structure with only length tuning available. With such structures, the transverse misalignment is not necessarily negligible, and it is important that the cavity is designed so that the intended protocol can cope with the residual misalignment.

Consequence 1: Extraction of Light

If the cavity mirrors are transversely misaligned, the cavity mode will form at an angle to the nominal cavity axis. For cavity constructions where the light is extracted through a fixed fibre (for example, fibre cavities), the extraction efficiency could be greatly reduced if the mode propagates at a significant angle to the one for which the collection is optimal. This means that, for cavity constructions that might involve

considerable mode tilt, it is advisable to separate the cavity and the collection, for example by collection of the photon mode leaving the cavity through a separate fibre whose angle and position can be adjusted.

Consequence 2: Mirror Shape

Cavity QED experiments have traditionally used superpolished mirrors with highly spherical profiles. However, this technique is incompatible with the small sizes and tight radii of curvature crucial for microcavity experiments. Several different methods are now used to fabricate microcavity mirrors, and each has their own characteristics. Perhaps the most commonly-used of these methods is laser ablation, which at its crudest produces an approximately Gaussian mirror profile, and even in more refined forms (e.g. multiple pulse 'dot milling') may yield significantly aspherical profiles. Though both spherical and Gaussian mirrors have the same parabolic profile at the centre of the mirror, the shape of the mirror does impact the mode that forms, particularly when the mode is tilted.

The structure of the cavity mode can be calculated through geometric methods, which predicts the cavity mode accounting for only the surface gradient (i.e. angle) and second derivative (i.e. curvature). This method predicts modes that intersect both mirrors normal to their surface and whose waists are determined by the local curvature of the mirror where the mode intersects. For a spherical mirror, the surface curvature is always constant and so, while misalignment tilts the mode, the mode profile remains isotropic. In contrast, for Gaussian mirrors the curvature is not constant, and misalignment introduces mode ellipticity and birefringence due to anisotropic curvature where the mode intersects the mirror.

The geometric model is helpful for providing intuition about the phenomena observed with misaligned cavities, but the cavity mode is more properly calculated through a mode mixing method, which does not restrict the solution to fundamental Gaussian beams. A method has is presented in Chapter 5 for fast calculation of both mirror matrices associated with Gaussian mirrors and translation matrices for arbitrary mirrors. These techniques facilitate efficient calculation of cavity properties as a function of mirror misalignment.

In the spherical mirror case, the mode mixing results produce mode profiles that are generally very similar to those derived with a geometric model, with associated losses within an order of magnitude. The effect the geometric model cannot capture is the hybridisation of transverse modes at particular, predictable lengths, in the case of the spherical mirror acts to *reduce* loss. However, this hybridisation has little impact on the mode shape, and reduces the losses, so is arguably beneficial. This means both that cavities with spherical mirrors do not suffer from unexpected physical effects as the cavity mirrors are misaligned, and that the geometric model is able to capture, to a reasonable degree of accuracy, the behaviour that is seen. This is advantageous as the geometric model is considerably simpler and more intuitive.

These conclusions are in contrast to those for the case of Gaussian-shaped mirrors, where the variation of curvature across the mirror surface leads to much stronger mode distortion than observed with spherical mirrors. Mode hybridisation is also seen, but it now leads to *increased* loss and generally deforms the mode much more severely. The cavity configurations for which hybridisation occurs can be understood through an 'effective curvature' dependent upon the mode order, position on mirror surface and direction on mirror surface, leading to a much more complicated resonant structure than for cavities with spherical mirrors. Resonances can be characterised by parity, and thus mirror misalignment renders the cavity vulnerable to odd parity resonances (in addition to the even resonances that are present even for aligned mirrors), meaning that misaligned cavities with Gaussian mirrors are especially disfavoured.

The conclusions of this investigation are that spherical mirrors are preferable, given their more easily predictable and simpler physics, high performance throughout the region where the cavity mode is well-confined on the mirror and lack of misalignment-induced birefringence. Gaussian mirrors *could* be used, but care would be required to make sure that configurations of significant hybridisation are avoided and that any birefringence is tolerable.

Consequence 3: Optimisation of Geometry and Transmission

Given that a spherical mirror shape has been chosen, the geometrical and optical parameters of the cavity still need to be set, and ideally optimised. However, the

question of what constitutes 'optimised' does not have a clear answer as performance is truly a balance of photon extraction probability, entanglement attempt rate and entanglement fidelity. Alone, these factors are complicated enough, but taken together it becomes difficult to understand how to optimise cavity parameters in a specific circumstance, and even harder to attain the intuition of the trade-offs involved in the optimisation that would be needed to identify the most promising areas for technical development.

With a set of specific assumptions, that are justifiable in the context of this work, the optimisation of cavity geometry can be made intuitive by yielding a single metric with a simple mathematical form. These assumptions are the following:

- The fidelity of remote entanglement does not depend on the cavity geometry. This requires that fidelity loss from temporal mixing and off-resonant processes in the cavity driving are negligible.
- The entanglement attempt rate is not significantly impacted by the cavity parameters. This requires that photon generation time constitutes a small part of the overall attempt cycle time.

Under these assumptions, the remote entanglement fidelity and the attempt rate are not dependent on the cavity geometry and the cavity parameters should be optimised to maximise the metric of photon extraction probability.

Using this metric means that further simplifications can be performed. Firstly, the cavity transmission and the cavity geometry can be separated, with the geometry determining the maximum photon extraction probability and the transmission set to achieve this. The optimisation of geometry is, strictly, a compromise; pushing towards a concentric configuration increases mode divergence, which strengthens the cavity coupling, but also increases losses as the mode is more divergent and more angled. However, the trade-off between these two factors can be simplified to a bounded optimisation because the geometry-induced losses rapidly transition from negligible to unacceptable across a narrow 'threshold' region of the geometry parameter space. Therefore, optimisation of the cavity geometry is reduced to optimisation of the mode divergence angle within the region where the clipping loss is negligible.

For the case where cavity construction is limited by a minimum cavity length (to permit sufficient optical access and reduce heating from the mirror dielectric) and a maximum milled volume (set by the constraints of the mirror fabrication process), the optimised geometry has the minimum allowed length and the minimum possible radius of curvature possible without inducing intolerable clipping loss. When the mirrors are misaligned, the optimised geometry adjusts to reduce clipping loss at the expense of performance, but provided this misalignment is accounted for, the reduction in performance is not severe.

Practical considerations may mean that the optimised configuration is not ideal in reality. Firstly, the constraints imposed are somewhat simplified; it may be that smaller diameter mirrors can be placed closer to the ion, resulting in a minimum length dependent on the mirror diameter, or it could be that the difficulty of mirror manufacture is not purely a function of the volume, but another function of the mirror curvature and diameter, or even that the eventual mirror roughness depends upon the geometry milled. Secondly, as the optimised solution is found by pushing the geometry to the boundary where clipping loss becomes important, the optimised solution is very vulnerable to a slight mis-setting of parameters that would increase clipping loss. Therefore a more detailed optimisation process could account for the fabrication process, understanding how mirror geometry affects the difficulty of manufacture and surface roughness, as well as determining variances in the cavity parameters the set how close to the clipping loss boundary to set the geometry. This stage of optimisation would require detailed fabrication characterisation that is unlikely to be generally applicable.

9.1.3 Criterion 3: Protocol Tolerates Residual Cavity Birefringence

Cavities can be birefringent, which introduces a complicating aspect of physics that is not found in free-space implementations of remote entanglement. Birefringence can have several root causes, such as anisotropic stresses from mirror mounting or coating, or a mirror surface that lacks mirror symmetry about the centre of the incident mode. Whatever its cause, in Chapter 6 it was shown that cavity birefringence can

modify the polarisation of a photon as it is being generated, meaning the emitted photon polarisation does not necessarily correlate with the atomic process that occurred, and also impacts the intensity profile of the emitted wavepackets. These effects reduce the fidelity of remote entanglement, and particularly so for the case of polarisation-encoded remote entanglement, where information about the ion qubit is directly written into a polarisation qubit.

Correction strategies can be used to reduce the impact of birefringence. This work studied three methods. The first used a time-dependent optical retarder such as a Pockels cell to restore the polarisation correlations at the output of each node. This method generally recovers the polarisation correlations effectively and for a reasonable level of added complication, but cannot restore the fidelity completely. The second method was a detection-time dependent local operation on the ion qubits post measurement. This can also restore the fidelity effectively, but would be much more difficult to calibrate. Finally, it is possible to post-select on the detection times to proceed only in the case of high-fidelity entanglement, but at the cost of reduced rate.

These correction strategies make tolerable a greater range of birefringence values, but they cannot maintain high fidelity and high rate for all birefringence values. In the case that the cavity birefringence is too great for fidelity loss to be mitigated adequately without excessive additional complexity or reduction in rate, more drastic action is required. Using time-bin encoding as opposed to polarisation encoding should substantially reduce the vulnerability of the entanglement scheme to polarisation oscillations. Additionally, using time-bin encoding with a cavity whose axis is aligned perpendicular to the magnetic field allows for only one of the cavity polarisations to be used, and a strong linear birefringence to be deliberately applied to reduce polarisation oscillations resulting from random fabrication errors.

9.1.4 Criterion 4: Suitable Driving Scheme

The previous three criteria were related to avoiding specific problems in cavity-enhanced remote entanglement. Provided these criteria are satisfied, what remains is to design a viable scheme for remote entanglement within the level structure of the chosen ion. In this process, the following questions should be considered:

Will the photonic qubit be polarisation encoded or time-bin encoded? This affects the errors to which the scheme is vulnerable and the detailed driving protocols that can be used. Polarisation encoding should be faster (both increasing rate and decreasing errors from drifts) and the photonic qubits are more easily manipulable. However, time-bin encoding should be less sensitive to cavity or fibre birefringence and offers more flexibility to filter out unwanted processes photonically or to utilise deliberately applied birefringence.

Should the cavity axis be parallel or orthogonal to the magnetic field? This choice impacts the transition strengths between the fine levels and whether the cavity modes to which the decay couples are linearly or circularly polarised. A cavity axis along the field generally gives access to the strongest transitions, and reduces the number of possible cavity-enhanced transitions, whereas a cavity axis orthogonal to the field yields linearly polarised cavity modes, presenting distinct advantages for applying deliberate birefringence.

Which transitions will be used for the cavity decays? Generally, these would be the strongest available, to not only increase the rate of the desired process, but decrease the probability of most unwanted processes. However, other factors to consider are the sensitivity to magnetic field fluctuations (which generally favours using states with similar m_J values) and the likelihood of off-resonant processes.

How can the impact of off-resonant processes be minimised? Ideally, the protocol would be designed to remove the possibility of undesired processes as far as possible. The final scheme for time-bin encoding (Chapter 7) used two such methods, first by 'hiding' population in another manifold when it does not need to be driven, and secondly by filtering out unwanted cavity decays based on their polarisation so that erroneous processes do not lead to heralds. However, it is practically impossible to remove the possibility of all undesired processes in this way. In this case, increasing the Zeeman splitting, improving the polarisation purity, and driving transitions sufficiently slowly can all help. This aspect should be carefully evaluated before designing an experiment, as such changes are not necessarily possible once the experiment has been commissioned; the Zeeman splitting could

be dictated by the clock field of another species, and the laser systems would not necessarily be simple (or cheap) to replace.

9.2 Atomic Sources

The work finished with details of the development of a miniature thermal source for ion trap experiments. While the work appears to stand separate from the rest of the thesis, the goal of developing scalable quantum networks is a common one. The development of these atomic sources is an example of the engineering development necessary to make a simple, scalable and robust node. This source uses the capabilities and flexibility afforded by 3D microfabrication to yield significant advantages over existing models. These advantages stem from the following principal factors:

- **Optical heating:** An inherently more efficient way of heating the oven than e.g. resistance heating, where the method of energy delivery itself provides a high thermal conductivity route for energy escape. Using optical heating also removes two or more high-current feedthroughs, improving vacuum system reliability.

- **Thermally isolated crucible:** The thermal impedance from the crucible to the rest of the vacuum chamber is high as heat can only conduct away through long, thin support spokes.

- **Low emissivity coating:** Radiative losses are reduced by the gold coating.

- **Collimating aperture:** The collimator greatly improves the efficiency with which atoms are directed to the desired location, reducing the number of atoms delivered into the vacuum system for each ion loaded.

These factors lend the source a series of practical advantages: It is small, it can be heated with low power to a high temperature quickly, and it reduces the overall number of atoms emitted into the vacuum system. The small size means that one could consider integrating the source directly into the trap structure, with the low power requirement reducing the cost and complexity of the heating apparatus, as well as thermal drifts inside the vacuum chamber. The increased atomic efficiency

means that the source lifetime is, for all practical considerations, unlimited, and the total number of atoms emitted into the vacuum system is low, raising the possibility of continuous generation of a weak neutral beam, which would further reduce loading times and thermal drifts.

9.3 Areas for Further Investigation

While the above results and conclusions should prove useful for the realisation of a cavity-enhanced ion trap quantum network, further effort, both experimental and theoretical, is required to develop this technology to the level of practical utility. Some important areas for future investigation are identified and discussed in this section.

9.3.1 Simulation of Remote Entanglement Scheme

Perhaps the most obvious immediate area of future work is to produce a detailed simulation of the remote entanglement scheme that accounts for the physical effects discussed (temporal mixing, birefringence, off-resonant interactions, time-bin and polarisation encoding) in a single unified framework. In this work, each effect has been studied individually in order to understand them in the simplest context, which is useful for extracting qualitative principles. However, such an atomised procedure is limited in two principal ways.

Firstly, identifying the optimum protocol requires evaluating a balance of competing factors as no single scheme outperforms the others in every aspect, as exemplified by the trade-offs discussed in Sec. 9.1.4. Only a single model with all the relevant physics could determine the best compromises in practice- compromises whose impact on the eventual fidelity and rate is not necessarily minimal.

Secondly, extracting principles from simplified models is potentially subject to biases if conclusions are drawn too hastily. While it makes qualitative sense that time-bin encoding is less vulnerable to cavity birefringence than polarisation encoding, it would be nice to see this emerge from simulations broad enough to cover both schemes. Additionally, in Chapter 7, it was qualitatively argued that certain schemes were preferable as they removed certain errors, for example excitation of the late

photon generation process during the early time bin. Again, this argument would be more conclusive if backed up by disinterested simulation.

To account for temporal mixing and off-resonant interactions, these simulations would have to capture the ion-cavity dynamics using the full relevant level structure with a master equation approach as opposed to the non-Hermitian Hamiltonian used in Chapters 3 and 6. The wavepackets extracted from these systems would then have to be propagated through the non-polarising beamsplitter and then subsequently measured, as done in Chapter 6, but also applied to time-bin encoding. After building the capabilities necessary to simulate a single instance of a remote entanglement experiment, different parameters could be varied to give insight into the trade-offs involved, requiring the development of analysis tools to extract the salient information in a transparent manner. Though the development of such simulations would be involved, it does not involve physics that is fundamentally different from that encountered in this thesis. Additionally, the cost of developing these simulations is likely far less than the effort that will be expended developing a real-world remote entanglement experiment. While it is impossible to simulate everything prior to developing the experiment, the case for a unified simulation of the physics discussed in this thesis is persuasive.

9.3.2 Mode Mixing

Results pertaining to the mode mixing investigation have formed a major part of the arguments presented in this thesis. The central conclusion– that Gaussian mirrors introduce complications that would best not be encountered in a remote entanglement experiment– is likely to be well-founded, however the faithfulness of the technique as a simulation method could be improved.

Firstly, the method could be extended to the basis of vector cavity modes and the use of the vector reflection condition so that geometric or misalignment-induced birefringence could emerge directly from the mode mixing calculations. Secondly, the full effects of the change in propagation distance due to mirror topography (rather than just the wavefront phase) could be modelled. Finally, the method could be

extended to many orders of the paraxial approximation, increasing accuracy for cavities approaching the concentric configuration.

With these improvements, the mode mixing method could be applied more widely than has been done here. It could be used to investigate mode structure and loss in cavities with mirror profiles beyond spherical and Gaussian as a function of misalignment, or to design mirrors for cavities that are less vulnerable to misalignment. The extension to vector modes could also open up investigations into unusual physical phenomena, such as a possible coupling between transverse mode profile and polarisation, even if applications are not immediately obvious.

9.3.3 Experimental Development

In the longer term, however, the most obvious further work (indeed, the entire motivation for the current work) is to develop a network node for modular quantum computation, with the first major goal being the demonstration of cavity-enhanced remote entanglement. The first step in this process would be to develop the fabrication of the spherical micromirrors and characterise their loss and birefringence. From this data, the cavity geometry will be designed on the basis of the optimisation procedure presented in Chapter 4, and the detailed protocol formulated on the basis of the unified simulation discussed in Sec. 9.3.1. Once the apparatus required to interface trapped ions with an optical cavity is fashioned, it will be time to demonstrate the technology, with ion-photon entanglement and finally cavity-enhanced ion-ion remote entanglement the most significant goals for single and two-node systems respectively.

While achieving these would represent a major milestone, they are not the end of the process. Once remote entanglement is demonstrated in simple, single ion experiments, it would be time to take the remaining steps towards a modular quantum computer. These steps include the inclusion of more than two nodes to the network, the use of entanglement purification, and the trapping of multiple species within a single node. As these feats are achieved, the potential applications of cavity-enhanced trapped ion quantum networks in quantum information applications will multiply. Eventually these systems might reach the refinement and scale required

to perform useful calculations as a networked quantum computer, exploiting the advantages that the modular architecture bestows.

The task ahead is, without question, challenging. It is likely to be possible only with a combination of theoretical understanding of the physics encountered and concerted effort to simplify and integrate the trap design and fabrication at each stage. However, nothing found so far suggests that it is impossible, and the possibilities afforded by the control light and matter at this fundamental level surely merit the effort expended in trying.

References

[1] J. S. Bell. “On the einstein podolsky rosen paradox”. In: *Physics Physique Fizika* 1.3 (1964), p. 195.

[2] A. Steane. “Quantum computing”. In: *Rep Prog Phys* 61.2 (1998), p. 117.

[3] R. P. Feynman. “Simulating physics with computers”. In: *Feynman and computation*. CRC Press, 2018, pp. 133–153.

[4] S. Lloyd. “Universal Quantum Simulators”. In: *Science* 273.5278 (1996), pp. 1073–1078.

[5] T. D. Ladd et al. “Quantum computers”. In: *Nature* 464.7285 (2010), pp. 45–53.

[6] D. Deutsch. “Quantum theory, the Church–Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400.1818 (1985), pp. 97–117.

[7] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys Rev A* 52.4 (1995), R2493–R2496.

[8] A. M. Steane. “Error correcting codes in quantum theory”. In: *Phys Rev Lett* 77.5 (1996), p. 793.

[9] *SFCS ’94: Proceedings of the 35th Annual Symposium on Foundations of Computer Science*. USA: IEEE Computer Society, 1994.

[10] L. K. Grover. “Quantum Mechanics Helps in Searching for a Needle in a Haystack”. In: *Phys Rev Lett* 79.2 (1997), pp. 325–328.

[11] A. W. Harrow, A. Hassidim, and S. Lloyd. “Quantum Algorithm for Linear Systems of Equations”. In: *Phys Rev Lett* 103.15 (2009), p. 150502.

[12] C. D. Bruzewicz et al. “Trapped-ion quantum computing: Progress and challenges”. In: *Appl Phys Rev* 6.2, 021314 (2019), p. 021314.

[13] G. Wendin. “Quantum information processing with superconducting circuits: a review”. In: *Rep Prog Phys* 80.10 (2017), p. 106001.

[14] L. Childress and R. Hanson. “Diamond NV centers for quantum computing and quantum networks”. In: *MRS Bull* 38.2 (2013), pp. 134–138.

[15] P. Kok et al. “Linear optical quantum computing with photonic qubits”. In: *Rev Mod Phys* 79.1 (2007), pp. 135–174.

[16] J. J. Morton et al. “Embracing the quantum limit in silicon computing”. In: *Nature* 479.7373 (2011), pp. 345–353.

[17] N. Ohlsson, R. Krishna Mohan, and S. Kröll. “Quantum computer hardware based on rare-earth-ion-doped inorganic crystals”. In: *Opt Commun* 201.1 (2002), pp. 71–77.

[18] L. Henriet et al. “Quantum computing with neutral atoms”. In: *Quantum* 4 (2020), p. 327.

[19] F. M. Penning. “Die Glimmentladung bei niedrigem Druck zwischen koaxialen Zylindern in einem axialen Magnetfeld”. In: *physica* 3.9 (1936), pp. 873–894.

[20] L. S. Brown and G. Gabrielse. “Geonium theory: Physics of a single electron or ion in a Penning trap”. In: *Rev Mod Phys* 58.1 (1986), pp. 233–311.

[21] W. Paul. “Electromagnetic traps for charged and neutral particles”. In: *Rev Mod Phys* 62.3 (1990), pp. 531–540.

[22] P. K. Ghosh. *Ion traps*. 1995.

[23] D. Leibfried et al. “Quantum dynamics of single trapped ions”. In: *Rev Mod Phys* 75.1 (2003), pp. 281–324.

[24] D. J. Berkeland et al. “Minimization of ion micromotion in a Paul trap”. In: *J Appl Phys* 83.10 (1998), pp. 5025–5033.

[25] D. J. Wineland, R. E. Drullinger, and F. L. Walls. “Radiation-pressure cooling of bound resonant absorbers”. In: *Phys Rev Lett* 40.25 (1978), p. 1639.

[26] J. Cirac et al. “Laser cooling of trapped ions: The influence of micromotion”. In: *Physical review A* 49.1 (1994), p. 421.

[27] F. Diedrich et al. “Laser Cooling to the Zero-Point Energy of Motion”. In: *Phys Rev Lett* 62.4 (1989), pp. 403–406.

[28] W. Neuhauser et al. “Localized visible Ba^+ mono-ion oscillator”. In: *Phys Rev A* 22.3 (1980), pp. 1137–1140.

[29] J. C. Bergquist et al. “Energy and Radiative Lifetime of the $5d^9 6s^2\,{}^2D_{\frac{5}{2}}$ State in Hg II by Doppler-Free Two-Photon Laser Spectroscopy”. In: *Phys Rev Lett* 55.15 (1985), pp. 1567–1570.

[30] J. C. Bergquist et al. “Observation of Quantum Jumps in a Single Atom”. In: *Phys Rev Lett* 57.14 (1986), pp. 1699–1702.

[31] J. I. Cirac and P. Zoller. “Quantum computations with cold trapped ions”. In: *Physical review letters* 74.20 (1995), p. 4091.

[32] A. Steane. “The ion trap quantum information processor”. In: *Applied Physics B: Lasers and Optics* 64.6 (1996), pp. 623–643.

[33] D. T. C. Allcock et al. “omg blueprint for trapped ion quantum computing with metastable states”. In: *Appl Phys Lett* 119.21, 214002 (2021), p. 214002.

[34] J. E. Christensen et al. “High-fidelity manipulation of a qubit enabled by a manufactured nucleus”. In: *npj Quantum Inf* 6, 35 (2020), p. 35.

[35] N. Akerman et al. “Universal gate-set for trapped-ion qubits using a narrow linewidth diode laser”. In: *New J Phys* 17.11 (2015), p. 113060.

[36] T. Ruster et al. “A long-lived Zeeman trapped-ion qubit”. In: *Appl Phys B: Lasers Opt* 122.10, 254 (2016), p. 254.

[37] A. Sørensen and K. Mølmer. “Quantum computation with ions in thermal motion”. In: *Physical review letters* 82.9 (1999), p. 1971.

[38] D. Leibfried et al. “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”. In: *Nature* 422.6930 (2003), pp. 412–415.

[39] C. Zhang et al. “Submicrosecond entangling gate between trapped ions via Rydberg interaction”. In: *Nature* 580.7803 (2020), pp. 345–349.

[40] P. Schindler et al. “A quantum information processor with trapped ions”. In: *New J Phys* 15.12 (2013), p. 123012.

[41] F. Mintert and C. Wunderlich. “Ion-Trap Quantum Logic Using Long-Wavelength Radiation”. In: *Phys Rev Lett* 87.25 (2001), p. 257904.

[42] T. P. Harty et al. “High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit”. In: *Phys Rev Lett* 113.22 (2014), p. 220501.

[43] C. J. Ballance et al. “High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits”. In: *Phys Rev Lett* 117.6 (2016), p. 060504.

[44] J. P. Gaebler et al. “High-Fidelity Universal Gate Set for $^{9}Be^{+}$ Ion Qubits”. In: *Phys Rev Lett* 117.6 (2016), p. 060505.

[45] T. P. Harty et al. “High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves”. In: *Phys Rev Lett* 117.14 (2016), p. 140501.

[46] C. R. Clark et al. “High-Fidelity Bell-State Preparation with $^{40}Ca^{+}$ Optical Qubits”. In: *Phys Rev Lett* 127.13 (2021), p. 130505.

[47] R. Srinivas et al. “High-fidelity laser-free universal control of trapped ion qubits”. In: *Nature* 597.7875 (2021), pp. 209–213.

[48] C. H. Baldwin et al. “High-fidelity light-shift gate for clock-state qubits”. In: *Phys Rev A* 103.1 (2021), p. 012603.

[49] T. R. Tan et al. “Multi-element logic gates for trapped-ion qubits”. In: *Nature* 528.7582 (2015), pp. 380–383.

[50] A. C. Hughes et al. “Benchmarking a High-Fidelity Mixed-Species Entangling Gate”. In: *Phys Rev Lett* 125.8 (2020), p. 080504.

[51] M. A. Sepiol et al. “Probing Qubit Memory Errors at the Part-per-Million Level”. In: *Phys Rev Lett* 123.11 (2019), p. 110503.

[52] P. Wang et al. “Single ion qubit with estimated coherence time exceeding one hour”. In: *Nature communications* 12.1 (2021), pp. 1–8.

[53] V. M. Schäfer et al. “Fast quantum logic gates with trapped-ion qubits”. In: *Nature* 555.7694 (2018), pp. 75–78.

[54] C. Cohen-Tannoudji and A. Kastler. “I Optical Pumping”. In: ed. by E. Wolf. Vol. 5. Progress in Optics. Elsevier, 1966, pp. 1–81.

[55] D. J. Wineland et al. “Double-resonance and optical-pumping experiments on electromagnetically confined, laser-cooled ions”. In: *Opt Lett* 5.6 (1980), pp. 245–247.

[56] W. Nagourney, J. Sandberg, and H. Dehmelt. “Shelved optical electron amplifier: Observation of quantum jumps”. In: *Phys Rev Lett* 56.26 (1986), pp. 2797–2799.

[57] T. Sauter et al. “‘Quantum jumps’ observed in the fluorescence of a single ion”. In: *Optics communications* 60.5 (1986), pp. 287–292.

[58] A. Keselman et al. “High-fidelity state detection and tomography of a single-ion Zeeman qubit”. In: *New J Phys* 13.7 (2011), p. 073027.

[59] C. Ryan-Anderson et al. “Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code”. In: *arXiv e-prints*, arXiv:2208.01863 (2022), arXiv:2208.01863.

[60] L. Egan et al. “Fault-tolerant control of an error-corrected qubit”. In: *Nature* 598.7880 (2021), pp. 281–286.

[61] L. Postler et al. “Demonstration of fault-tolerant universal quantum gate operations”. In: *Nature* 605.7911 (2022), pp. 675–680.

[62] H. Häffner, C. Roos, and R. Blatt. “Quantum computing with trapped ions”. In: *Phys Rep* 469.4 (2008), pp. 155–203.

[63] R. Ozeri. “The trapped-ion qubit tool box”. In: *Contemp Phys* 52.6 (2011), pp. 531–550.

[64] C. Monroe and J. Kim. “Scaling the ion trap quantum processor”. In: *Science* 339.6124 (2013), pp. 1164–1169.

[65] K. A. Landsman et al. “Two-qubit entangling gates within arbitrarily long chains of trapped ions”. In: *Phys Rev A* 100.2 (2019), p. 022332.

[66] P. H. Leung and K. R. Brown. “Entangling an arbitrary pair of qubits in a long ion crystal”. In: *Phys Rev A* 98.3 (2018), p. 032318.

[67] T. Choi et al. “Optimal Quantum Control of Multimode Couplings between Trapped Ion Qubits for Scalable Entanglement”. In: *Phys Rev Lett* 112.19 (2014), p. 190502.

[68] P. H. Leung et al. “Robust 2-Qubit Gates in a Linear Ion Crystal Using a Frequency-Modulated Driving Force”. In: *Phys Rev Lett* 120.2 (2018), p. 020501.

[69] A. R. Milne et al. “Phase-Modulated Entangling Gates Robust to Static and Time-Varying Errors”. In: *Phys Rev Applied* 13.2 (2020), p. 024022.

[70] D. F. James. *Quantum dynamics of cold trapped ions with application to quantum computation*. Tech. rep. 1997.

[71] G. Pagano et al. “Cryogenic trapped-ion system for large scale quantum simulation”. In: *Quantum Sci Technol* 4.1 (2018), p. 014004.

[72] G.-D. Lin et al. “Large-scale quantum computation in an anharmonic linear ion trap”. In: *Europhys Lett* 86.6 (2009), p. 60004.

[73] J. Zhang et al. “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator”. In: *Nature* 551.7682 (2017), pp. 601–604.

[74] R. Yao et al. “Experimental realization of a multiqubit quantum memory in a 218-ion chain”. In: *Phys Rev A* 106.6 (2022), p. 062617.

[75] C.-Y. Shih et al. “Reprogrammable and high-precision holographic optical addressing of trapped ions for scalable quantum control”. In: *npj Quantum Inf* 7.1 (2021), pp. 1–8.

[76] D. Kielpinski, C. Monroe, and D. J. Wineland. “Architecture for a large-scale ion-trap quantum computer”. In: *Nature* 417.6890 (2002), pp. 709–711.

[77] B. Lekitsch et al. “Blueprint for a microwave trapped ion quantum computer”. In: *Sci Adv* 3.2 (2017), e1601540.

[78] H. Kaufmann et al. “Scalable Creation of Long-Lived Multipartite Entanglement”. In: *Phys Rev Lett* 119.15 (2017), p. 150503.

[79] Y. Wan et al. “Quantum gate teleportation between separated qubits in a trapped-ion processor”. In: *Science* 364.6443 (2019), pp. 875–878.

[80] J. M. Pino et al. “Demonstration of the trapped-ion quantum CCD computer architecture”. In: *Nature* 592.7853 (2021), pp. 209–213.

[81] P. T. Dumitrescu et al. “Dynamical topological phase realized in a trapped-ion quantum simulator”. In: *Nature* 607.7919 (2022), pp. 463–467.

[82] E. Chertkov et al. “Holographic dynamics simulations with a trapped-ion quantum computer”. In: *Nat Phys* 18.9 (2022), pp. 1074–1079.

[83] H. A. Fürst et al. “Controlling the transport of an ion: classical and quantum mechanical solutions”. In: *New J Phys* 16.7 (2014), p. 075007.

[84] R. Bowler et al. “Coherent Diabatic Ion Transport and Separation in a Multizone Trap Array”. In: *Phys Rev Lett* 109.8 (2012), p. 080502.

[85] R. T. Sutherland et al. “Motional Squeezing for Trapped Ion Transport and Separation”. In: *Phys Rev Lett* 127.8 (2021), p. 083201.

[86] C. Simon and W. T. M. Irvine. “Robust Long-Distance Entanglement and a Loophole-Free Bell Test with Ions and Photons”. In: *Phys Rev Lett* 91.11 (2003), p. 110405.

[87] C. Monroe et al. “Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects”. In: *Phys Rev A* 89.2 (2014), p. 022317.

[88] C. H. Bennett et al. “Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels”. In: *Phys Rev Lett* 76.5 (1996), pp. 722–725.

[89] W. Dür and H.-J. Briegel. “Entanglement Purification for Quantum Computation”. In: *Phys Rev Lett* 90.6 (2003), p. 067901.

[90] E. T. Campbell. “Distributed quantum-information processing with minimal local resources”. In: *Phys Rev A* 76.4 (2007), p. 040302.

[91] L. Jiang et al. “Distributed quantum computation based on small quantum registers”. In: *Phys Rev A* 76.6 (2007), p. 062323.

[92] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin. “Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links”. In: *Phys Rev X* 4.4 (2014), p. 041041.

[93] Y. Li and S. C. Benjamin. “Hierarchical surface code for network quantum computing with modules of arbitrary size”. In: *Phys Rev A* 94.4 (2016), p. 042303.

[94] R. Nigmatullin et al. “Minimally complex ion traps as modules for quantum communication and computing”. In: *New J Phys* 18.10 (2016), p. 103028.

[95] G.-D. Lin and L.-M. Duan. “Equilibration and temperature distribution in a driven ion chain”. In: *New J Phys* 13.7 (2011), p. 075015.

[96] L. J. Stephenson et al. “High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network”. In: *Phys Rev Lett* 124.11 (2020), p. 110501.

[97] P. Kurpiers et al. “Quantum Communication with Time-Bin Encoded Microwave Photons”. In: *Phys Rev Appl* 12.4 (2019), p. 044067.

[98] C. J. Axline et al. “On-demand quantum state transfer and entanglement between remote microwave cavity memories”. In: *Nat Phys* 14.7 (2018), pp. 705–710.

[99] E. T. Campbell, B. M. Terhal, and C. Vuillot. “Roads towards fault-tolerant universal quantum computation”. In: *Nature* 549.7671 (2017), pp. 172–179.

[100] H. J. Kimble. “Strong interactions of single atoms and photons in cavity QED”. In: *Phys Scr* 1998.T76 (1998), p. 127.

[101] S. Ritter et al. “An elementary quantum network of single atoms in optical cavities”. In: *Nature* 484.7393 (2012), pp. 195–200.

[102] D. Niemietz et al. “Nondestructive detection of photonic qubits”. In: *Nature* 591.7851 (2021), pp. 570–574.

[103] T. Stolz et al. “Quantum-Logic Gate between Two Optical Photons with an Average Efficiency above 40%”. In: *Phys Rev X* 12.2 (2022), p. 021035.

[104] G. R. Guthöhrlein et al. “A single ion as a nanoscopic probe of an optical field”. In: *Nature* 414.6859 (2001), pp. 49–51.

[105] H. Takahashi et al. “Strong coupling of a single ion to an optical cavity”. In: *Physical review letters* 124.1 (2020), p. 013602.

[106] A. Stute et al. “Tunable ion–photon entanglement in an optical cavity”. In: *Nature* 485.7399 (2012), pp. 482–485.

[107] P. Kobel, M. Breyer, and M. Köhl. “Deterministic spin-photon entanglement from a trapped ion in a fiber Fabry–Perot cavity”. In: *npj Quantum Inf* 7.1 (2021), pp. 1–7.

[108] V. Krutyanskiy et al. “Entanglement of trapped-ion qubits separated by 230 meters”. In: *arXiv e-prints*, arXiv:2208.14907 (2022), arXiv:2208.14907.

[109] M. J. Madsen et al. “Planar ion trap geometry for microfabrication”. In: *Appl Phys B: Lasers Opt* 78.5 (2004), pp. 639–651.

[110] J. Chiaverini et al. “Surface-Electrode Architecture for Ion-Trap Quantum Information Processing”. In: *arXiv e-prints*, quant-ph/0501147 (2005), quant–ph/0501147.

[111] S. Seidelin et al. “Microfabricated Surface-Electrode Ion Trap for Scalable Quantum Information Processing”. In: *Phys Rev Lett* 96.25 (2006), p. 253003.

[112] M. G. House. “Analytic model for electrostatic fields in surface-electrode ion traps”. In: *Phys Rev A* 78.3 (2008), p. 033402.

[113] J. H. Wesenberg. “Electrostatics of surface-electrode ion traps”. In: *Phys Rev A* 78.6 (2008), p. 063410.

[114] M. D. Hughes et al. “Microfabricated ion traps”. In: *Contemp Phys* 52.6 (2011), pp. 505–529.

[115] J. M. Amini et al. “Toward scalable ion traps for quantum information processing”. In: *New J Phys* 12.3 (2010), p. 033031.

[116] G. Wilpers et al. “A monolithic array of three-dimensional ion traps fabricated with conventional semiconductor technology”. In: *Nature nanotechnology* 7.9 (2012), pp. 572–576.

[117] D.-I. ’. Cho et al. “A review of silicon microfabricated ion traps for quantum information processing”. In: *Micro and Nano Systems Letters* 3.1, 2 (2015), p. 2.

[118] N. Podoliak et al. “Comparative Numerical Studies of Ion Traps with Integrated Optical Cavities”. In: *Phys Rev Applied* 6.4, 044008 (2016), p. 044008.

[119] S. Ragg et al. “Segmented ion-trap fabrication using high precision stacked wafers”. In: *Rev Sci Instrum* 90.10 (2019), p. 103203.

[120] D. Stick et al. “Ion trap in a semiconductor chip”. In: *Nat Phys* 2.1 (2006), pp. 36–39.

[121] M. Brownnutt et al. “Monolithic microfabricated ion trap chip design for scaleable quantum processors”. In: *New J Phys* 8.10 (2006), p. 232.

[122] C. Decaroli et al. “Design, fabrication and characterization of a micro-fabricated stacked-wafer segmented ion trap with two X-junctions”. In: *Quantum Sci Technol* 6.4 (2021), p. 044001.

[123] Z. D. Romaszko et al. “Engineering of microfabricated ion traps and integration of advanced on-chip features”. In: *Nature Reviews Physics* 2.6 (2020), pp. 285–299.

[124] K. K. Mehta et al. “Integrated optical multi-ion quantum logic”. In: *Nature* 586.7830 (2020), pp. 533–537.

[125] K. K. Mehta et al. “Integrated optical addressing of an ion qubit”. In: *Nature nanotechnology* 11.12 (2016), pp. 1066–1070.

[126] M. L. Day et al. “A micro-optical module for multi-wavelength addressing of trapped ions”. In: *Quantum Sci Technol* 6.2 (2021), p. 024007.

[127] S. L. Todaro et al. “State Readout of a Trapped Ion Qubit Using a Trap-Integrated Superconducting Photon Detector”. In: *Phys Rev Lett* 126.1 (2021), p. 010501.

[128] X. Fernandez-Gonzalvo and M. Keller. “A Fully Fiber-Integrated Ion Trap for Portable Optical Atomic Clocks”. In: *arXiv preprint arXiv:2112.06816* (2021).

[129] J. T. Merrill et al. “Demonstration of integrated microscale optics in surface-electrode ion traps”. In: *New J Phys* 13.10 (2011), p. 103005.

[130] P. F. Herskind et al. “Microfabricated surface ion trap on a high-finesse optical mirror”. In: *Opt Lett* 36.16 (2011), pp. 3045–3047.

[131] A. Van Rynbach, P. Maunz, and J. Kim. “An integrated mirror and surface ion trap with a tunable trap location”. In: *Appl Phys Lett* 109.22 (2016), p. 221108.

[132] C. Ospelkaus et al. “Microwave quantum logic gates for trapped ions”. In: *Nature* 476.7359 (2011), pp. 181–184.

[133] J. Stuart et al. “Chip-Integrated Voltage Sources for Control of Trapped Ions”. In: *Phys Rev Applied* 11.2 (2019), p. 024010.

[134] S. Gao et al. “Optimization of Scalable Ion-Cavity Interfaces for Quantum Photonic Networks”. In: *Phys. Rev. Appl.* 19.1 (2023), p. 014033.

[135] S. Gao et al. “An optically heated atomic source for compact ion trap vacuum systems”. In: *Rev Sci Instrum* 92.3 (2021), p. 033205.

[136] E. Kassa et al. “Effects of cavity birefringence in polarisation-encoded quantum networks”. In: *New Journal of Physics* 25.1 (2023), p. 013004.

[137] “Proceedings of the American Physical Society”. In: *Phys Rev* 69.11-12 (1946), pp. 674–674.

[138] H. Kogelnik and T. Li. “Laser Beams and Resonators”. In: *Appl Opt* 5.10 (1966), pp. 1550–1567.

[139] A. Kuhn. “Cavity induced interfacing of atoms and light”. In: *Engineering the Atom-Photon Interaction*. Springer, 2015, pp. 3–38.

[140] R. Lang, M. O. Scully, and W. E. Lamb. “Why is the Laser Line So Narrow? A Theory of Single-Quasimode Laser Operation”. In: *Phys. Rev. A* 7.5 (1973), pp. 1788–1797.

[141] R. Lang and M. O. Scully. "Fluctuations in mode locked "single-mode" laser oscillation". In: *Optics Communications* 9.4 (1973), pp. 331–337.

[142] C. W. Gardiner and M. J. Collett. "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation". In: *Phys. Rev. A* 31.6 (1985), pp. 3761–3774.

[143] R. Bennett, T. M. Barlow, and A. Beige. "A physically motivated quantization of the electromagnetic field". In: *European Journal of Physics* 37.1 (2015), p. 014001.

[144] N. Barré et al. "Waves and rays in plano-concave laser cavities: I. Geometric modes in the paraxial approximation". In: *Eur J Phys* 38.3 (2017), p. 034010.

[145] D. K. Armani et al. "Ultra-high-Q toroid microcavity on a chip". In: *Nature* 421.6926 (2003), pp. 925–928.

[146] M. Arcari et al. "Near-Unity Coupling Efficiency of a Quantum Emitter to a Photonic Crystal Waveguide". In: *Phys Rev Lett* 113.9, 093603 (2014), p. 093603.

[147] C. Vallance et al. "Open-access microcavities for chemical sensing". In: *Nanotechnology* 27.27, 274003 (2016), p. 274003.

[148] M. Harlander et al. "Trapped-ion probing of light-induced charging effects on dielectrics". In: *New J Phys* 12.9 (2010), p. 093035.

[149] F. R. Ong et al. "Probing surface charge densities on optical fibers with a trapped ion". In: *New J Phys* 22.6 (2020), p. 063018.

[150] M. Teller et al. "Heating of a trapped ion induced by dielectric materials". In: *Phys Rev Lett* 126.23 (2021), p. 230505.

[151] M. Kumph et al. "Electric-field noise above a thin dielectric layer on metal electrodes". In: *New J Phys* 18.2, 023020 (2016), p. 023020.

[152] A. Kawasaki et al. "Geometrically asymmetric optical cavity for strong atom-photon coupling". In: *Phys Rev A* 99.1 (2019), p. 013437.

[153] A. Saharyan et al. "Light-matter interaction in open cavities with dielectric stacks". In: *Appl Phys Lett* 118.15 (2021), p. 154002.

[154] S. M. Dutra and G. Nienhuis. "Quantized mode of a leaky cavity". In: *Phys. Rev. A* 62.6 (2000), p. 063805.

[155] P. T. Kristensen et al. "Modeling electromagnetic resonators using quasinormal modes". In: *Adv. Opt. Photon.* 12.3 (2020), pp. 612–708.

[156] W. W. Rigrod. "Isolation of Axi-Symmetrical Optical-Resonator Modes". In: *Appl Phys Lett* 2.3 (1963), pp. 51–53.

[157] D. M. Brink, G. R. Satchler, and M. Danos. "Angular Momentum". In: *Phys Today* 16.6 (1963), p. 80.

[158] C. J. Foot. *Atomic physics*. Vol. 7. OUP Oxford, 2004.

[159] E. Kassa. "Single ion coupled to a high-finesse optical fibre cavity for cQED in the strong coupling regime". PhD thesis. University of Sussex, 2017.

[160] A. Kuhn and D. Ljunggren. "Cavity-based single-photon sources". In: *Contemp Phys* 51.4 (2010), pp. 289–313.

[161] H. Goto et al. "Figure of merit for single-photon generation based on cavity quantum electrodynamics". In: *Phys Rev A* 99.5 (2019), p. 053843.

[162] J. I. Cirac et al. “Quantum state transfer and entanglement distribution among distant nodes in a quantum network”. In: *Phys Rev Lett* 78.16 (1997), p. 3221.

[163] C. Cabrillo et al. “Creation of entangled states of distant atoms by interference”. In: *Phys Rev A* 59.2 (1999), p. 1025.

[164] M. Pompili et al. “Realization of a multinode quantum network of remote solid-state qubits”. In: *Science* 372.6539 (2021), pp. 259–264.

[165] C. W. Chou et al. “Measurement-induced entanglement for excitation stored in remote atomic ensembles”. In: *Nature* 438.7069 (2005), pp. 828–832.

[166] S. Hermans et al. “Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study”. In: *New Journal of Physics* (2023).

[167] D. L. Moehring et al. “Entanglement of single-atom quantum bits at a distance”. In: *Nature* 449.7158 (2007), pp. 68–71.

[168] L. Luo et al. “Protocols and techniques for a scalable atom–photon quantum network”. In: *Fortschr Phys* 57.11-12 (2009), pp. 1133–1152.

[169] A. M. Dyckovsky and S. Olmschenk. “Analysis of photon-mediated entanglement between distinguishable matter qubits”. In: *Phys Rev A* 85.5 (2012), p. 052322.

[170] G. Vittorini et al. “Entanglement of distinguishable quantum memories”. In: *Phys Rev A* 90.4 (2014), p. 040302.

[171] T.-M. Zhao et al. “Entangling Different-Color Photons via Time-Resolved Measurement and Active Feed Forward”. In: *Phys Rev Lett* 112.10 (2014), p. 103602.

[172] S. Olmschenk et al. “Quantum Teleportation Between Distant Matter Qubits”. In: *Science* 323.5913 (2009), pp. 486–489.

[173] D. N. Matsukevich et al. “Bell Inequality Violation with Two Remote Atomic Qubits”. In: *Phys Rev Lett* 100.15 (2008), p. 150404.

[174] A. Mundt et al. “Coupling a single atomic quantum bit to a high finesse optical cavity”. In: *Physical review letters* 89.10 (2002), p. 103001.

[175] J. M. Bennett. “Recent developments in surface roughness characterization”. In: *Meas Sci Technol* 3.12 (1992), pp. 1119–1127.

[176] T. G. Ballance et al. “Cavity-induced backaction in Purcell-enhanced photon emission of a single ion in an ultraviolet fiber cavity”. In: *Phys Rev A* 95.3 (2017), p. 033812.

[177] V. Krutyanskiy et al. “Light-matter entanglement over 50 km of optical fibre”. In: *npj Quantum Inf* 5, 72 (2019), p. 72.

[178] A. Dréau et al. “Quantum Frequency Conversion of Single Photons from a Nitrogen-Vacancy Center in Diamond to Telecommunication Wavelengths”. In: *Phys Rev Applied* 9.6 (2018), p. 064031.

[179] F. Lindenfelser et al. “Cooling atomic ions with visible and infra-red light”. In: *New J Phys* 19.6 (2017), p. 063041.

[180] J. Schmitz, H. M. Meyer, and M. Köhl. “Ultraviolet Fabry-Perot cavity with stable finesse under ultrahigh vacuum conditions”. In: *Rev Sci Instrum* 90.6 (2019), p. 063102.

[181] G. Rempe et al. “Measurement of ultralow losses in an optical interferometer”. In: *Optics letters* 17.5 (1992), pp. 363–365.

[182] C. J. Hood et al. “Real-Time Cavity QED with Single Atoms”. In: *Phys Rev Lett* 80.19 (1998), pp. 4157–4160.

[183] J. Schupp et al. “Interface between Trapped-Ion Qubits and Traveling Photons with Close-to-Optimal Efficiency”. In: *PRX Quantum* 2.2 (2021), p. 020331.

[184] T. Steinmetz et al. “Stable fiber-based Fabry-Pérot cavity”. In: *Appl Phys Lett* 89.11 (2006), p. 111110.

[185] D. Hunger et al. “A fiber Fabry-Perot cavity with high finesse”. In: *New J Phys* 12.6, 065038 (2010), p. 065038.

[186] D. Hunger et al. “Laser micro-fabrication of concave, low-roughness features in silica”. In: *AIP Adv* 2.1 (2012), p. 012119.

[187] F. Rochau et al. “Dynamical Backaction in an Ultrahigh-Finesse Fiber-Based Microcavity”. In: *Phys Rev Applied* 16.1 (2021), p. 014013.

[188] A. Muller et al. “Ultrahigh-finesse, low-mode-volume Fabry-Perot microcavity”. In: *Opt Lett* 35.13 (2010), pp. 2293–2295.

[189] M. Uphoff et al. “Frequency splitting of polarization eigenmodes in microscopic Fabry–Perot cavities”. In: *New J Phys* 17.1 (2015), p. 013053.

[190] H. Takahashi et al. “Novel laser machining of optical fibers for long cavities with low birefringence”. In: *Opt Express* 22.25 (2014), pp. 31317–31328.

[191] K. Ott et al. “Millimeter-long fiber Fabry-Perot cavities”. In: *Optics express* 24.9 (2016), pp. 9839–9853.

[192] B. Brandstätter et al. “Integrated fiber-mirror ion trap for strong ion-cavity coupling”. In: *Rev Sci Instrum* 84.12, 123104-123104-15 (2013), pp. 123104-123104–15.

[193] G. K. Gulati et al. “Fiber cavities with integrated mode matching optics”. In: *Sci Rep* 7, 5556 (2017), p. 5556.

[194] P. R. Dolan et al. “Femtoliter tunable optical cavity arrays”. In: *Opt Lett* 35.21 (2010), pp. 3556–3558.

[195] B. T. Walker et al. “Bespoke mirror fabrication for quantum simulation with light in open-access microcavities”. In: *Opt Express* 29.7 (2021), pp. 10800–10810.

[196] A. A. P. Trichet et al. “Topographic control of open-access microcavities at the nanometer scale”. In: *Opt Express* 23.13 (2015), pp. 17205–17216.

[197] M. Trupke et al. “Microfabricated high-finesse optical cavity with open access and small volume”. In: *Appl Phys Lett* 87.21 (2005), p. 211106.

[198] G. Wachter et al. “Silicon microcavity arrays with open access and a finesse of half a million”. In: *Light: Sci Appl* 8.1 (2019), pp. 1–7.

[199] N. Jin et al. “Scalable fabrication of ultrahigh-finesse micromirrors with variable geometry”. In: *arXiv preprint arXiv:2203.15931* (2022).

[200] G. Biedermann et al. “Ultrasmooth microfabricated mirrors for quantum information”. In: *Appl Phys Lett* 97.18 (2010), p. 181110.

[201] E. Janitz et al. “High mechanical bandwidth fiber-coupled Fabry-Perot cavity”. In: *Opt Express* 25.17 (2017), pp. 20932–20943.

[202] J. Gallego et al. “High-finesse fiber Fabry–Perot cavities: stabilization and mode matching analysis”. In: *Appl Phys B* 122.3 (2016), pp. 1–14.

[203] C. Saavedra et al. “Tunable fiber Fabry-Perot cavities with high passive stability”. In: *Opt Express* 29.2 (2021), pp. 974–982.

[204] M. Teller et al. “Integrating a fiber cavity into a wheel trap for strong ion-cavity coupling”. In: *arXiv preprint arXiv:2207.10500* (2022).

[205] M. Lee et al. “Microelectromechanical-System-Based Design of a High-Finesse Fiber Cavity Integrated with an Ion Trap”. In: *Phys Rev Applied* 12.4 (2019), p. 044052.

[206] J. F. S. Brachmann et al. “Photothermal effects in ultra-precisely stabilized tunable microcavities”. In: *Opt Express* 24.18 (2016), pp. 21205–21215.

[207] B. Casabone et al. “Heralded Entanglement of Two Ions in an Optical Cavity”. In: *Phys Rev Lett* 111.10 (2013), p. 100505.

[208] J. Bochmann et al. “Fast Excitation and Photon Emission of a Single-Atom-Cavity System”. In: *Phys Rev Lett* 101.22 (2008), p. 223601.

[209] T. D. Barrett, T. H. Doherty, and A. Kuhn. “Pushing Purcell enhancement beyond its limits”. In: *New J Phys* 22.6, 063013 (2020), p. 063013.

[210] C. Maurer et al. “A single-photon source based on a single Ca+ ion”. In: *New journal of physics* 6.1 (2004), p. 94.

[211] A. Stute et al. “Toward an ion-photon quantum interface in an optical cavity”. In: *Appl Phys B: Lasers Opt* 107.4 (2012), pp. 1145–1157.

[212] H. G. Barros et al. “Deterministic single-photon source from a single ion”. In: *New J Phys* 11.10 (2009), p. 103004.

[213] J. Kuklinski et al. “Adiabatic population transfer in a three-level system driven by delayed laser pulses”. In: *Phys Rev A* 40.11 (1989), p. 6741.

[214] B. W. Shore. “Picturing stimulated Raman adiabatic passage: a STIRAP tutorial”. In: *Adv. Opt. Photon.* 9.3 (2017), pp. 563–719.

[215] A. Kuhn et al. “Controlled generation of single photons from a strongly coupled atom-cavity system”. In: *Appl Phys B* 69.5 (1999), pp. 373–377.

[216] A. Kuhn, M. Hennrich, and G. Rempe. “Deterministic Single-Photon Source for Distributed Quantum Networking”. In: *Phys Rev Lett* 89.6 (2002), p. 067901.

[217] P. B. Nisbet-Jones et al. “Highly efficient source for indistinguishable single photons of controlled shape”. In: *New J Phys* 13.10 (2011), p. 103036.

[218] H. P. Specht et al. “Phase shaping of single-photon wave packets”. In: *Nat Photonics* 3.8 (2009), pp. 469–472.

[219] B. Weber et al. “Photon-Photon Entanglement with a Single Trapped Atom”. In: *Phys Rev Lett* 102.3 (2009), p. 030501.

[220] P. B. Nisbet-Jones et al. “Photonic qubits, qutrits and ququads accurately prepared and delivered on demand”. In: *New J Phys* 15.5 (2013), p. 053007.

[221] M. Meraner et al. “Indistinguishable photons from a trapped-ion quantum network node”. In: *Phys Rev A* 102.5 (2020), p. 052614.

[222] K. A. Fischer et al. “Signatures of two-photon pulses from a quantum two-level system”. In: *Nat Phys* 13.7 (2017), pp. 649–654.

[223] T. Walker et al. “Improving the indistinguishability of single photons from an ion-cavity system”. In: *Phys Rev A* 102.3 (2020), p. 032616.

[224] G. S. Vasilev, D. Ljunggren, and A. Kuhn. "Single photons made-to-measure". In: *New J Phys* 12.6 (2010), p. 063024.

[225] T. Utsugi et al. "Gaussian-wave-packet model for single-photon generation based on cavity quantum electrodynamics under adiabatic and nonadiabatic conditions". In: *Phys Rev A* 106.2 (2022), p. 023712.

[226] P. P. Rohde, T. C. Ralph, and M. A. Nielsen. "Optimal photons for quantum-information processing". In: *Phys Rev A* 72.5 (2005), p. 052332.

[227] O. Morin et al. "Deterministic shaping and reshaping of single-photon temporal wave functions". In: *Physical review letters* 123.13 (2019), p. 133602.

[228] E. Kassa et al. "Precise positioning of an ion in an integrated Paul trap-cavity system using radiofrequency signals". In: *J Mod Opt* 65.5-6 (2018), pp. 520–528.

[229] T. D. Barrett et al. "Nonlinear Zeeman effects in the cavity-enhanced emission of polarised photons". In: *New J Phys* 20.7 (2018), p. 073030.

[230] T. Doherty. "Optical cavities for quantum information processing with trapped neutral atoms". PhD thesis. University of Oxford, 2021.

[231] D. Clarke and P. Horak. "Alignment requirements of Fabry-Perot microresonators for ion trap quantum information processing (Conference Presentation)". In: *Quantum Technologies 2018*. Ed. by J. Stuhler, A. J. Shields, and M. J. Padgett. Vol. 10674. International Society for Optics and Photonics. SPIE, 2018, 106740P.

[232] A. G. Fox and T. Li. "Resonant modes in a maser interferometer". In: *Bell Syst Tech J* 40.2 (1961), pp. 453–488.

[233] D. Kleckner et al. "Diffraction-limited high-finesse optical cavities". In: *Phys Rev A* 81.4 (2010), p. 043814.

[234] G. Barontini et al. "Deterministic generation of multiparticle entanglement by quantum Zeno dynamics". In: *Science* 349.6254 (2015), pp. 1317–1321.

[235] M. Brekenfeld et al. "A quantum network node with crossed optical fibre cavities". In: *Nat Phys* 16.6 (2020), pp. 647–651.

[236] T. Macha et al. "Nonadiabatic storage of short light pulses in an atom-cavity system". In: *Phys. Rev. A* 101.5 (2020), p. 053406.

[237] S. Garcia et al. "Dual-wavelength fiber Fabry-Perot cavities with engineered birefringence". In: *Opt Express* 26.17 (2018), pp. 22249–22263.

[238] S. Johnson et al. "Tunable cavity coupling of the zero phonon line of a nitrogen-vacancy defect in diamond". In: *New Journal of Physics* 17.12 (2015), p. 122003.

[239] H. Laabs and A. T. Friberg. "Nonparaxial eigenmodes of stable resonators". In: *IEEE J Quantum Electron* 35.2 (1999), pp. 198–207.

[240] J. Benedikter et al. "Transverse-mode coupling and diffraction loss in tunable Fabry–Pérot microcavities". In: *New J Phys* 17.5 (2015), p. 053051.

[241] N. Podoliak et al. "Harnessing the mode mixing in optical fiber-tip cavities". In: *J Phys B: At , Mol Opt Phys* 50.8 (2017), p. 085503.

[242] D. V. Karpov and P. Horak. "Cavities with nonspherical mirrors for enhanced interaction between a quantum emitter and cavity photons". In: *Phys Rev A* 105.2 (2022), p. 023515.

[243] D. V. Karpov and P. Horak. "Evolutionary algorithm to design high-cooperativity optical cavities". In: *New J Phys* 24.7 (2022), p. 073028.

[244] D. V. Karpov, S. Kurdiumov, and P. Horak. "Convolutional neural networks for mode on-demand high finesse optical resonator design". In: (2022).

[245] J. Benedikter et al. "Transverse-mode coupling effects in scanning cavity microscopy". In: *New J Phys* 21.10 (2019), p. 103029.

[246] D. Stoler. "Operator methods in physical optics". In: *J Opt Soc Am* 71.3 (1981), pp. 334–341.

[247] G. Nienhuis and L. Allen. "Paraxial wave optics and harmonic oscillators". In: *Phys Rev A* 48.1 (1993), pp. 656–665.

[248] S. Varró. "Coherent and incoherent superposition of transition matrix elements of the squeezing operator". In: *New J Phys* 24.5 (2022), p. 053035.

[249] K. E. Cahill and R. J. Glauber. "Ordered Expansions in Boson Amplitude Operators". In: *Phys Rev* 177.5 (1969), pp. 1857–1881.

[250] D. Hucul et al. "Modular entanglement of atomic qubits using photons and phonons". In: *Nat Phys* 11.1 (2015), pp. 37–42.

[251] J. L. Hall, J. Ye, and L.-S. Ma. "Measurement of mirror birefringence at the sub-ppm level: Proposed application to a test of QED". In: *Phys Rev A* 62.1, 013815 (2000), p. 013815.

[252] D. Jacob et al. "Supermirror phase anisotropy measurement". In: *Opt Lett* 20.7 (1995), p. 671.

[253] P. Münstermann et al. "Single slow atoms from an atomic fountain observed in a high-finesse optical cavity". In: *Optics communications* 159.1-3 (1999), pp. 63–67.

[254] M. Mader et al. "A scanning cavity microscope". In: *Nat Commun* 6, 7249 (2015), p. 7249.

[255] J. Volz et al. "Measurement of the internal state of a single atom without energy exchange". In: *Nature* 475.7355 (2011), pp. 210–213.

[256] M. Lax, W. H. Louisell, and W. B. McKnight. "From Maxwell to paraxial wave optics". In: *Phys Rev A* 11.4 (1975), pp. 1365–1370.

[257] A. L. Cullen. "On the accuracy of the beam-wave theory of the open resonator". In: *IEEE Transactions on Microwave Theory Techniques* 24 (1976), p. 534.

[258] T. D. Barrett et al. "Polarization Oscillations in Birefringent Emitter-Cavity Systems". In: *Phys Rev Lett* 122.8 (2019), p. 083602.

[259] A. Cadène et al. "Vacuum magnetic linear birefringence using pulsed fields: status of the BMV experiment". In: *Eur Phys J D* 68.1, 16 (2014), p. 16.

[260] P. Asenbaum and M. Arndt. "Cavity stabilization using the weak intrinsic birefringence of dielectric mirrors". In: *Opt Lett* 36.19 (2011), p. 3720.

[261] F. M. Buters et al. "Optomechanics with a polarization nondegenerate cavity". In: *Phys Rev A* 94.6 (2016), p. 063813.

[262] J.-M. Cui et al. "Polarization nondegenerate fiber Fabry-Perot cavities with large tunable splittings". In: *Appl Phys Lett* 112.17 (2018), p. 171105.

[263] D. Deutsch et al. "Quantum Privacy Amplification and the Security of Quantum Cryptography over Noisy Channels". In: *Phys Rev Lett* 77.13 (1996), pp. 2818–2821.

[264] Y. Yu et al. “Entanglement of two quantum memories via fibres over dozens of kilometres”. In: *Nature* 578.7794 (2020), pp. 240–245.

[265] L. Yu et al. “Two-photon interference at telecom wavelengths for time-bin-encoded single photons from quantum-dot spin qubits”. In: *Nat Commun* 6, 8955 (2015), p. 8955.

[266] A. Tchebotareva et al. “Entanglement between a Diamond Spin Qubit and a Photonic Time-Bin Qubit at Telecom Wavelength”. In: *Phys Rev Lett* 123.6 (2019), p. 063601.

[267] P. Farrera, G. Heinze, and H. de Riedmatten. “Entanglement between a Photonic Time-Bin Qubit and a Collective Atomic Spin Excitation”. In: *Phys Rev Lett* 120.10 (2018), p. 100501.

[268] C. Clausen et al. “Quantum storage of photonic entanglement in a crystal”. In: *Nature* 469.7331 (2011), pp. 508–511.

[269] E. Saglamyurek et al. “Broadband waveguide quantum memory for entangled photons”. In: *Nature* 469.7331 (2011), pp. 512–515.

[270] B. Hensen et al. “Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres”. In: *Nature* 526.7575 (2015), pp. 682–686.

[271] N. Gisin et al. “Quantum cryptography”. In: *Reviews of modern physics* 74.1 (2002), p. 145.

[272] C. Antonelli, M. Shtaif, and M. Brodsky. “Sudden Death of Entanglement Induced by Polarization Mode Dispersion”. In: *Phys Rev Lett* 106.8 (2011), p. 080404.

[273] M. Brodsky et al. “Loss of polarization entanglement in a fiber-optic system with polarization mode dispersion in one optical path”. In: *Optics letters* 36.1 (2011), pp. 43–45.

[274] X. Fernandez-Gonzalvo et al. “Quantum frequency conversion of quantum memory compatible photons to telecommunication wavelengths”. In: *Optics express* 21.17 (2013), pp. 19473–19487.

[275] I. Marcikic et al. “Distribution of Time-Bin Entangled Qubits over 50 km of Optical Fiber”. In: *Phys Rev Lett* 93.18 (2004), p. 180502.

[276] J. D. Franson. “Bell inequality for position and time”. In: *Phys Rev Lett* 62.19 (1989), pp. 2205–2208.

[277] D. M. Lucas et al. “Isotope-selective photoionization for calcium ion trapping”. In: *Phys Rev A* 69.1 (2004), p. 012711.

[278] T. G. Ballance et al. “A short response time atomic source for trapped ion experiments”. In: *Rev Sci Instrum* 89.5 (2018), p. 053102.

[279] R. Hendricks et al. “An all-optical ion-loading technique for scalable microtrap architectures”. In: *Appl Phys B* 88 (2007), pp. 507–513.

[280] Q. Wu et al. “Adaptively controlled fast production of defect-free beryllium ion crystals using pulsed laser ablation”. In: *Rev Sci Instrum* 92.6 (2021), p. 063201.

[281] A. Osada and A. Noguchi. “Deterministic loading of a single strontium ion into a surface electrode trap using pulsed laser ablation”. In: *J Phys Commun* 6.1 (2022), p. 015007.

[282] G. Vrijsen et al. “Efficient isotope-selective pulsed laser ablation loading of 174Yb+ ions in a surface electrode trap”. In: *Opt Express* 27.23 (2019), pp. 33907–33914.

[283] K. Sheridan, W. Lange, and M. Keller. “All-optical ion generation for ion trap loading”. In: *Appl Phys B* 104.4 (2011), pp. 755–761.

[284] J. M. Sage, A. J. Kerman, and J. Chiaverini. “Loading of a surface-electrode ion trap from a remote, precooled source”. In: *Phys Rev A* 86.1 (2012), p. 013417.

[285] M. Brandl et al. “Cryogenic setup for trapped ion quantum computing”. In: *Rev Sci Instrum* 87.11 (2016), p. 113103.

[286] M. Weber et al. “Cryogenic ion trap system for high-fidelity near-field microwave-driven quantum logic”. In: *arXiv preprint arXiv:2207.11364* (2022).

[287] F. M. Leupold. “Bang-bang control of a trapped-ion oscillator”. PhD thesis. ETH Zurich, 2015.

[288] P. D. D. Schwindt et al. “A highly miniaturized vacuum package for a trapped ion atomic clock”. In: *Rev Sci Instrum* 87.5 (2016), p. 053112.

[289] K. L. Moore et al. “Collimated, single-pass atom source from a pulsed alkali metal dispenser for laser-cooling experiments”. In: *Rev Sci Instrum* 76.2 (2005), p. 023106.

[290] S. C. Bell et al. “A slow atom source using a collimated effusive oven and a single-layer variable pitch coil Zeeman slower”. In: *Rev Sci Instrum* 81.1 (2010), p. 013105.

[291] D. T. C. Allcock. “Surface-electrode ion traps for scalable quantum computing”. PhD thesis. Hertford College, Oxford, 2011.

[292] C. Alcock, V. Itkin, and M. Horrigan. “Vapour pressure equations for the metallic elements: 298–2500K”. In: *Can Metall Q* 23.3 (1984), pp. 309–313.

[293] C. B. Lucas. *Atomic and molecular beams: production and collimation*. CRC press, 2013.

[294] A. Mohan, R. V. Tompson, and S. K. Loyalka. “Efficient numerical solution of the Clausing problem”. In: *J Vac Sci Technol A: Vac Surf Films* 25.4 (2007), pp. 758–762.

[295] S. Bourdeauducq et al. “ARTIQ”. In: (2021).

[296] Y. Bellouard. “On the bending strength of fused silica flexures fabricated by ultrafast lasers”. In: *Opt Mater Express* 1.5 (2011), pp. 816–831.

A

Calculating Probability Matrices for Multi-Level System Optimisation

All probabilities $P_\zeta(t)$ used in the optimisation can be expressed as expectation values of matrices $\boldsymbol{P}_\zeta(t)$ for Fourier vector $\boldsymbol{\alpha}^{\boldsymbol{F}}_{\boldsymbol{g_1}}$. The simplest example is $P_{R_1}(t)$:

$$\begin{aligned} P_{R_1}(t) &= 2\kappa \int_0^t \alpha^*_{g_1}(t')\alpha_{g_1}(t')\,dt', \\ &= \frac{2\kappa}{T_b}\int_0^t \sum_{n,n'} C^{(1)*}_{n'} C^{(1)}_n e^{i(\omega_n-\omega_{n'})t'}\,dt', \\ &= \sum_{n,n'} C^{(1*)}_{n'}(2\kappa\,\boldsymbol{V}_{n',n}(t))C^{(1)}_n, \\ &= \boldsymbol{\alpha}^{\boldsymbol{F}}_{\boldsymbol{g_1}} \cdot 2\kappa\,\boldsymbol{V} \cdot \boldsymbol{\alpha}^{\boldsymbol{F}}_{\boldsymbol{g_1}}, \\ \boldsymbol{V}_{n',n}(t) &= \frac{1}{T_b}\int_0^t e^{i(\omega_n-\omega_{n'})t'}\,dt'. \end{aligned} \tag{A.1}$$

This can be generalised to other emission probabilities by using the interdependence of the wavefunctions through

$$\begin{aligned} P_{R_J} &= \sum_{n,n'} C^{(j)*}_{n'}(2\kappa\,\boldsymbol{V}_{n',n})C^{(j)}_n, \\ &= \boldsymbol{\alpha}^{\boldsymbol{F}}_{\boldsymbol{g_1}} \cdot \left(2\kappa \boldsymbol{F}^{(j)\dagger}\,\boldsymbol{V}(t)\boldsymbol{F}^{(j)}\right) \cdot \boldsymbol{\alpha}^{\boldsymbol{F}}_{\boldsymbol{g_1}}, \\ \boldsymbol{F}^{(j)}_{n',n} &= f^{(1\to j)}_n \delta_{n',n}, \end{aligned} \tag{A.2}$$

where δ is the Kronecker delta.

For the cavity probabilities, a similar process is used (Eq. A.3). The operators are slightly arbitrary as later in the process an initial condition will be enforced that demands all cavity functions begin from zero. This means that the term from $|\alpha_{g_j}(0)|^2$ is not needed, as it will always evaluate to zero. However, it has been included here so that the written form of the probability matrix has no negative eigenvalues. The probability

$$
\begin{aligned}
P_{C_j}(t) &= |\alpha_{g_j}(t)|^2, \\
&= \int_0^t ((\alpha^*_{g_j}\dot{}(t'))\alpha_{g_j}(t') + \alpha^*_{g_j}(t')(\alpha_{g_j}\dot{}(t'))\, dt' + |\alpha_{g_j}(0)|^2, \\
&= \sum_{n,n'} C^{(1)*}_{n'} \left((f^{(1\to j)}_n)^* (i(\omega_n - \omega_{n'})\, \boldsymbol{V}_{n',n}(t) + 1) f^{(1\to j)}_n \right) C^{(1)}_n, \qquad \text{(A.3)} \\
&= \boldsymbol{\alpha^F_{g_1}} \cdot \boldsymbol{P_{C_j}} \cdot \boldsymbol{\alpha^F_{g_1}}, \\
\boldsymbol{P_{C_j}}_{n,n'}(t) &= \left((f^{(1\to j)}_{n'})^* (i(\omega_n - \omega_{n'})\, \boldsymbol{V}_{n',n}(t) + 1) f^{(1\to j)}_n \right),
\end{aligned}
$$

is most easily expressed through a matrix given in component form.

The calculation for the spontaneous emission probability requires the expression for the excited state component (taken from Eq. 3.16)

$$\alpha_e(t) = -\frac{\kappa + i\delta_1}{g_1}\alpha_{g_1} - \frac{1}{g_1}\dot{\alpha_{g_1}}. \qquad \text{(A.4)}$$

From this, the spontaneous emission probability

$$
\begin{aligned}
P_E(t) &= 2\gamma \int_0^t \alpha^*_e(t')\alpha_e(t')\, dt', \\
&= \boldsymbol{\alpha^F_{g_1}} \cdot \boldsymbol{P_E}(t) \cdot \boldsymbol{\alpha^F_{g_1}}, \qquad \text{(A.5)} \\
\boldsymbol{P_E}_{n',n}(t) &= \frac{2\gamma}{g_1^2}(\kappa^2 + \delta_1^2 + i\kappa(\omega_n - \omega_{n'}) + \delta_1(\omega_n + \omega_{n'}) + \omega_n\omega_{n'})\, \boldsymbol{V}_{n',n}(t),
\end{aligned}
$$

can be expressed. The matrix for probability $P_e(t) = |\alpha_e(t)|^2$ in the excited state at time t can also be written in a similar component form

$$\boldsymbol{P_e}_{n',n}(t) = \frac{1}{g_1^2 T_b} \left(\left(\kappa^2 + \delta_1^2 + +i\kappa\,(\omega_n - \omega_{n'}) + \delta_1\,(\omega_n + \omega_{n'})\,\omega_n\omega_{n'} \right) e^{i(\omega_n - \omega_{n'})t} \right). \qquad \text{(A.6)}$$

B

Choosing the Basis Size for Multi-Level System Optimisation

The numerical procedure conducts calculations of photonic output in a Fourier basis over time $T_b > T$, where T is the extraction time for the photon. The Fourier basis contains positive and negative frequencies, with $2h + 1$ total components, where h is the half basis size. To obtain useful results, T_b and h need to be chosen carefully, avoiding two main issues.

Firstly, the basis must be big enough to capture all of the photon frequencies. For example, suppose one of the cavity outputs had a wavepacket in Fourier space centred on a frequency δ_p. There would be a certain h below which the basis would not include frequency δ_p, and the solution obtained would not be correct. The cavity wavefunction will also have some global shape on the scale of T, that broadens the wavepacket in frequency space. This suggests the half basis size h should satisfy

$$\frac{2\pi}{T_b}h \gg \left(\delta_p + \frac{2\pi}{T}\right), \tag{B.1}$$

so that the central frequency and the width is contained within the basis. This limit is most problematic where $\delta_p T$ becomes large, although for very long times $\kappa T \gg 1$, the system can be driven in the adiabatic regime, and using the methods presented here may not be necessary.

A second issue is most obvious at short times ($T \simeq \frac{2\pi}{\kappa}$ or $T < \frac{2\pi}{\kappa}$). In this case. the optimum photon my leave a large population in the cavity. The Fourier basis for the problem means that the photon function $\alpha_{g_1}(t)$ must return to 0 at T_b (because each photon is constrained to start at zero amplitude). The shape of the function $\alpha_{g_1}(t)$ between T and T_b has not been discussed because it has no physical relevance, however, if the function needs to drop to 0 at T_b from a large value at T, the best solution will have a large frequency component with timescale $T_b - T$. In order to be able to simulate this comfortably, the half basis size h should satisfy

$$h \gg \frac{1}{1 - \frac{T_p}{T_b}}. \tag{B.2}$$

Anecdotally, a factor of 4 has been found to be an acceptable safety margin. The best T_b and h to use is a compromise between the two factors discussed, but understanding these limits makes it easier to choose appropriate parameters for the numerical method.

C

Derivation of Operators in the Hermite Gauss Basis

To derive the operator forms of x and $\frac{\partial}{\partial x}$, start by comparing the wavefunction of the paraxial mode first stated in Eq. 2.7 of Chapter 2 (by convention we will take the outgoing mode)

$$u^{(+)}_{n_x,n_y}(x,y) = aH_{n_x}\left(\frac{\sqrt{2}x}{w}\right)H_{n_x}\left(\frac{\sqrt{2}y}{w}\right)e^{-\frac{x^2+y^2}{w^2}}e^{-ik\frac{x^2+y^2}{2R}}e^{i(n_x+n_y+1)\Psi_G}, \tag{C.1}$$

with the mode of the simple Harmonic oscillator of mass m and resonant frequency Ω

$$\psi^{\text{SHO}}(x,y)_{n_x,n_y} = H_{n_x}\left(\sqrt{\frac{m\Omega}{\hbar}}x\right)H_{n_y}\left(\sqrt{\frac{m\Omega}{\hbar}}y\right)e^{-\frac{m\Omega(x^2+y^2)}{2\hbar}},$$

$$x^{\text{SHO}} = \sqrt{\frac{\hbar}{2m\Omega}}(a_x^{\text{SHO}} + (a_x^{\text{SHO}})^\dagger), \tag{C.2}$$

$$\frac{\partial}{\partial x}^{\text{SHO}} = \sqrt{\frac{m\Omega}{2\hbar}}(a_x^{\text{SHO}} - (a_x^{\text{SHO}})^\dagger),$$

where the operators x^{SHO} and $\frac{\partial}{\partial x}^{\text{SHO}}$ are shown in terms of the simple harmonic annihilation operator a_x^{SHO}. Comparing the two states leads to

$$u^{(+)}_{n_x,n_y}(x,y) = \psi(x,y)^{\text{SHO}}_{n_x,n_y}e^{-ik\frac{x^2+y^2}{2R}}e^{i(n_x+n_y+1)\Psi_G}, \tag{C.3}$$

where the parameters of the simple harmonic oscillator and Gaussian mode are related by $m\Omega/2\hbar = 1/w^2$. The x operator can be found through

$$\begin{aligned}
x_{n'_x,n'_y,n_x,n_y} &= \int_S u^{\pm *}_{n'_x,n'_y}(x,y)xu^{\pm}_{n_x,n_y}(x,y)\,dxdy,\\
x_{n'_x,n'_y,n_x,n_y} &= \exp\left[\pm i\Psi_G(n_x+n_y-n'_x-n'_y)\right]\int_S \psi^{*\text{SHO}}_{n'_x,n'_y}(x,y)x\psi^{\text{SHO}}_{n_x,n_y}(x,y)\,dxdy,\\
x_{n'_x,n'_y,n_x,n_y} &= x^{\text{SHO}}_{n'_x,n'_y,n_x,n_y}\exp\left[\pm i\Psi_G(n_x+n_y-n'_x-n'_y)\right].
\end{aligned} \tag{C.4}$$

This results in the final expression

$$\begin{aligned} x &= U_G^\dagger \frac{1}{2} w (a_x + a_x^\dagger) U_G, \\ (U_G)_{n'_x,n'_y,n_x,n_y} &= \delta_{n'_x,n_x}\delta_{n'_y,n_y} e^{i\Psi_G(n_x+n_y+1)}. \end{aligned} \tag{C.5}$$

A similar approach can be used for the $\frac{\partial}{\partial x}$ operator

$$\begin{aligned} \frac{\partial}{\partial x}_{n'_x,n'_y,n_x,n_y} &= \int_S u^{\pm *}_{n'_x,n'_y}(x,y) \frac{\partial}{\partial x} u^{\pm}_{n_x,n_y}(x,y)\, dxdy, \\ \frac{\partial}{\partial x}_{n'_x,n'_y,n_x,n_y} &= \exp\left[i\Psi_G(n_x+n_y-n'_x-n'_y)\right] \times \\ \int_S \left(\psi^{*\mathrm{SHO}}_{n'_x,n'_y}(x,y) \frac{\partial}{\partial x} \psi^{\mathrm{SHO}}_{n_x,n_y}(x,y)\right) &+ \left(\psi^{*\mathrm{SHO}}_{n'_x,n'_y}(x,y) \left(\mp ik\frac{x}{R}\right) \psi^{\mathrm{SHO}}_{n_x,n_y}(x,y)\right) dxdy, \\ \frac{\partial}{\partial x} &= U_G^\dagger \left(\frac{\partial}{\partial x}^{\mathrm{SHO}} \mp \frac{ik}{R} x^{\mathrm{SHO}}\right) U_G, \end{aligned} \tag{C.6}$$

resulting in the expression

$$\frac{\partial}{\partial x} = U_G^\dagger \left(\frac{1}{w}(a_x - a_x^\dagger) - i\frac{wk}{2R}(a_x + a_x^\dagger)\right) U_G, \tag{C.7}$$

which can be converted algebraically to the more convenient form

$$\frac{\partial}{\partial x} = \left(\frac{1}{w_0}(a_x - a_x^\dagger)\right). \tag{C.8}$$

D

Finding the Gaussian Profile Mirror Matrix in the Hermite Gauss Basis

The formula Eq. 5.12 in Chapter 5 for the Gaussian profile surface matrix in the Hermite-Gauss basis is calculated according to the same method as the appendix A of [248]. However, the derivation is presented below as it differs somewhat.

The Gaussian operator is most simply found through a tensor product of x and y operators. Therefore, start by finding the operator for the one-dimensional function $\exp\left[-x^2/w_e^2\right]$, using the position operator given by Eq. 5.9 in Chapter 5 and derived in App. C. Expanding in terms of ladder operators gives

$$\begin{aligned}\exp\left[-\frac{x^2}{w_e^2}\right] &= U_G^\dagger \exp\left[\chi\left(\frac{1}{2}\left(a^\dagger\right)^2+\frac{1}{2}\left(a^\dagger\right)^2+\frac{1}{2}\left(aa^\dagger+a^\dagger a\right)\right)\right]U_G,\\ &= U_G^\dagger \exp\left[\chi\left(K_+ + K_- + 2K_0\right)\right]U_G, \qquad \text{(D.1)}\\ \chi &= -\frac{w^2}{2w_e^2},\end{aligned}$$

where K_+, K_- and K_0 are assigned to $(1/2)\left(a^\dagger\right)^2$, $(1/4)\left(a^\dagger\right)^2$, and $(1/2)\left(aa^\dagger+a^\dagger a\right)$ respectively, and the annihilation operator a represents a_x (a_y) for finding the x (y) directed Gaussian function. Now use that K_+, K_- and K_0 have the same commutation relations as $-\sigma_+$, σ_- and $(1/2)\sigma_3$ with

$$-\sigma_+ = \begin{pmatrix}0 & -1\\ 0 & 0\end{pmatrix}, \quad \sigma_- = \begin{pmatrix}0 & 0\\ 1 & 0\end{pmatrix}, \quad \frac{1}{2}\sigma_3 = \frac{1}{2}\begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix}. \qquad \text{(D.2)}$$

Now, equate the coefficients of the exponent and normal-ordered exponents of the 2-dimensional matrices, with coefficients ζ, ζ' and η for the normal form

$$\exp\left[\zeta\left(-\sigma_+\right)\right]\exp\left[-\eta\left(\sigma_3\right)\right]\exp\left[\zeta'\left(\sigma_-\right)\right] = \exp\left[\chi\left(-\sigma_+\right)+\chi\left(\sigma_-\right)+2\chi\left(\sigma_3\right)\right]. \qquad \text{(D.3)}$$

Expanding the two sides of this equation gives

$$\begin{aligned}\zeta = \zeta' &= \frac{\chi}{1-\chi},\\ \eta &= 2\ln\left(1-\chi\right).\end{aligned} \qquad \text{(D.4)}$$

The 2-dimensional matrices can then be switched back for creation and annihilation operators to obtain

$$\exp\left[\chi\left(K_{+}+K_{-}+2K_{0}\right)\right]=$$

$$\exp\left[\frac{\chi}{1-\chi}\frac{1}{2}\left(a^{\dagger}\right)^{2}\right]\exp\left(-2\ln\left(1-\chi\right)\frac{1}{4}\left(aa^{\dagger}+a^{\dagger}a\right)\right)\exp\left(\frac{\chi}{1-\chi}\frac{1}{2}\left(a\right)^{2}\right). \tag{D.5}$$

The normal operator form can now be evaluated simply in the Hermite Gauss basis to obtain the result quoted in Sec. 5.3.4 of Chapter 5

$$\exp\left(-\frac{x^{2}}{w_{e}^{2}}\right)_{m',m}=$$

$$U_{G}^{\dagger}\left(1-\chi\right)^{-\left(\frac{m'+m+1}{2}\right)}\left(\frac{\chi}{2}\right)^{\frac{m'-m}{2}}\sqrt{m'!m!}\sum_{k=0}^{\left[\frac{m}{2}\right]}\frac{\left(\frac{\chi^{2}}{4}\right)^{k}}{\left(\frac{m'-m}{2}+k\right)!k!\left(m-2k\right)!}U_{G}, \tag{D.6}$$

$$\chi=-\frac{1}{2}\frac{w\left(z\right)^{2}}{w_{e}^{2}}.$$

E

Mode Transformations in the Hermite Gauss Basis

E.0.1 Changing the Mode Waist

To calculate the coefficients of a mode with a different centre waist, use the property of the transverse basis that, at the centre of the mode (axially), the basis of the Hermite-Gauss modes coincides with the basis of the simple harmonic oscillator. Therefore, the operator required to change the centre waist of the mode is the same as the operator that rescales the coordinate operators of the simple harmonic oscillator: The standard squeeze operators. To change the waist of a mode from w_0 to w_1 in the x-direction requires the application of the operator

$$\begin{aligned} S_x\left(\frac{w_1}{w_0}\right) &= \exp\left[-\frac{1}{2}r\left(a_x^2 - (a_x^\dagger)^2\right)\right], \\ r &= -\log\left[\frac{w_1}{w_0}\right]. \end{aligned} \tag{E.1}$$

The same form of operator applies in the y-direction.

E.0.2 Changing the Mode Angle

To address the impact of rotation on the cavity mode functions, the electric field of a basis state must be rotated and the overlap compared with basis states that have not been rotates. To do this, imagine a coordinate system x, y, z, and construct a new coordinate system with axes x', y' and z', where the axes are rotated about $y = y'$ by an angle θ such that

$$x' = x\cos(\theta) + z\sin(\theta), \quad z' = z\cos(\theta) - x\sin(\theta). \tag{E.2}$$

An unperturbed function in the transformed basis has the same coefficients as the function rotated by an angle θ towards the x axis in the $x - z$ plane in the original basis. Therefore, the Cartesian coordinate transformation can be used to derive the rotation operator. Upon rotating by a small angle $\delta\theta$, the relation between

coordinates in the central axial plane ($z' = 0$) of the new coordinate system can be calculated, using $z = x\tan(\theta)$ and expanding to first order

$$x = x', \quad y = y', \quad z = x'\delta\theta. \tag{E.3}$$

It is important to remember that it is the electric field that is being rotated, not the cavity mode function. For operations in a single transverse plane, such as transverse translation or waist squeezing, it is not necessary to account for this distinction, but for the present operator the dominant contribution arises from this difference. Equality of the electric fields in the two coordinate systems yields

$$\boldsymbol{E}'(x', y', z' = 0, t) = \boldsymbol{E}(x', y', x'\delta\theta, t), \tag{E.4}$$

where $\boldsymbol{E} = \hat{\boldsymbol{\epsilon}} u \exp(-ikz)\exp\{i\omega t\}$ is the field in the non-primed coordinates, and $\boldsymbol{E'} = \hat{\boldsymbol{\epsilon}} u' \exp(-ikz')\exp\{i\omega t\}$ is the field in the primed coordinates, with $\hat{\boldsymbol{\epsilon}}$ the polarisation (as set out in Eq. 2.5 of Chapter 2). Substituting in the cavity mode functions u and expanding to first order in the z-coordinate gives

$$u'(x', y', 0) = \left[1 + x\delta\theta\left(-ik + \frac{\partial}{\partial z}\right)\right] u(x', y', 0), \tag{E.5}$$

where $\partial/\partial z$ is the partial derivative with respect to the axial coordinate.

Note that there is a subtlety in the polarisation that has been glossed over to obtain this formula. It is assumed that the polarisation is unchanged between the two bases, and that it is the scalar function u that is directed along a different axis. However, the polarisation is not independent of the scalar function, as it must lie transverse to the direction of propagation. This is not an issue if the polarisation happened to lie in the y direction, but would be important to consider if there was a component to the polarisation in the x direction. Unfortunately, because the mode mixing method is inherently scalar, these issues cannot be solved mathematically; the only remaining option is to advise the reader that the mode mixing method will ignore polarisation dependent effects.

To evaluate the partial derivative in the last equation, one can use the paraxial equation (Eq. 2.6 of Chapter 2 to express in terms of transverse derivatives, and then App. C to express those transverse derivatives in ladder operators. Applying this infinitesimal result successively generates a matrix exponential as the operator required to rotate the direction of the mode by an angle θ towards x in the xz plane

$$T(\theta_x) = \exp\left[-i\delta\theta\frac{kw_0}{2}\left\{\left(a + a^\dagger\right)\left[1 + \frac{1}{(kw_0)^2}\left(\left(a_x - a_x^\dagger\right)^2 + \left(a_y - a_y^\dagger\right)^2\right)\right]\right\}\right]. \tag{E.6}$$

Some care must be taken extending this formula to general rotation angles, but for the results of this thesis, the direction of mode rotation is assumed without loss of generality to be in the xz plane.

E.0.3 Determining the Fundamental Mode

The mode transformations of Sec. E.0.2 of this appendix are used to determine the mode of interest by finding the mode with the greatest overlap with the geometrical expectation (as described in Sec. 5.2 of Chapter 5). The classical theory generates values for the angle of the mode and its waists (in the two principle directions), whereas the cavity eigenmodes are expressed as vectors in the basis of the on-axis parabolic cavity. To find the overlap of the cavity eigenmodes with the expected mode, the expected mode is expressed in the same basis as the cavity eigenmodes.

To do this, a standard 00 mode of the basis was taken, and then transformed to make the expected mode by first expanding/contracting in the two transverse directions independently to set the waists, and then rotating the mode in the xz plane to set the angle.

E.0.4 Removing Rotation from Mode Occupations

Mode rotations are also used to express the mode coefficients in the basis propagating in the same direction as the mode ('the co-propagating basis' coined in Sec. 5.4.1 of Chapter 5). To do this, the angle of the mode was calculated using the angle operator

$$\theta_x = \frac{-i}{k}\frac{\partial}{\partial x}, \tag{E.7}$$

which is valid in the paraxial approximation. The eigenstates of this operator are plane waves propagating at angle θ_x to the z axis in the x-z plane. Once the angle of the mode is calculated, the mode is rotated back with an applied rotation operator calculated using the methods Sec. E.0.2 of this appendix.

F

Uniqueness of Post-Measurement Corrections to a Single Qubit

In general, birefringent nodes will cause the ion-ion state produced by remote entanglement to differ from the ideal Bell states $\left|\Psi^{\pm}_{\text{Bell}}\right\rangle = (1/\sqrt{2})\left(\left|\uparrow\right\rangle_A \left|\downarrow\right\rangle_B \pm \left|\downarrow\right\rangle_A \left|\uparrow\right\rangle_B\right)$ (with sign + for $n = n'$ and sign − for $n \neq n'$, where $n \in \{c, d\}$ is the detector node of the H detection and $n' \in \{c, d\}$ the node of the V detection). The state produced can be expressed in its Schmidt form

$$|\psi_{nn'}(t_H, t_V)\rangle = \lambda_1 \left|\Uparrow\right\rangle_A \left|\Downarrow\right\rangle_B + \lambda_2 \left|\Downarrow\right\rangle_A \left|\Uparrow\right\rangle_B , \tag{F.1}$$

where $\{\left|\Uparrow\right\rangle_A, \left|\Downarrow\right\rangle_A\}$ and $\{\left|\Uparrow\right\rangle_B, \left|\Downarrow\right\rangle_B\}$ are orthogonal states in the Hilbert spaces of qubit A and qubit B, respectively, and with $\lambda_{1,2} \in \mathbb{R}_+$.

To find the best local corrections (unitaries applied to Nodes A and B) to increase fidelity with the target state, The form of Eq. F.1 suggests applying unitaries U_A^I and U_B^I to $|\psi_{nn'}(t_H, t_V)\rangle$ defined by the following transformations

$$\begin{gathered} \left|\Uparrow\right\rangle_A \to \left|\uparrow\right\rangle_A , \left|\Downarrow\right\rangle_A \to \left|\downarrow\right\rangle_A \\ \left|\Uparrow\right\rangle_B \to \pm\left|\uparrow\right\rangle_B , \left|\Downarrow\right\rangle_B \to \left|\downarrow\right\rangle_B . \end{gathered} \tag{F.2}$$

These transformations result in a state $\left|\psi_I^{\pm}\right\rangle = \lambda_1 \left|\uparrow\right\rangle_A \left|\downarrow\right\rangle_B \pm \lambda_2 \left|\downarrow\right\rangle_A \left|\uparrow\right\rangle_B$. Arbitrary unitaries are then applied to systems A and B (this is equivalent to applying different arbitrary transformations directly to $|\psi_{nn'}(t_H, t_V)\rangle$, but the intermediate step above will be algebraically convenient). The transformations are parametrised according to

$$\begin{gathered} \left|\uparrow\right\rangle_A \to \cos\theta_A \left|\uparrow\right\rangle_A + \sin\theta_A e^{i\phi_A} \left|\downarrow\right\rangle_A , \left|\downarrow\right\rangle_A \to \left(-\sin\theta_A e^{-i\phi_A} \left|\uparrow\right\rangle_A + \cos\theta_A \left|\downarrow\right\rangle_A\right) e^{i\chi_A} \\ \left|\uparrow\right\rangle_B \to \cos\theta_B \left|\uparrow\right\rangle_B + \sin\theta_B e^{i\phi_B} \left|\downarrow\right\rangle_B , \left|\downarrow\right\rangle_B \to \left(-\sin\theta_B e^{-i\phi_B} \left|\uparrow\right\rangle_B + \cos\theta_B \left|\downarrow\right\rangle_B\right) e^{i\chi_B}, \end{gathered} \tag{F.3}$$

where each arbitrary unitary is fully specified by θ, ϕ and χ for each node. After these transformations, the fidelity $\mathcal{F}$ with the target state is given by

$$\begin{aligned}\mathcal{F} = \frac{1}{2}\Big(\left(\lambda_1^2+\lambda_2^2\right)\Big(&\cos^2\theta_A\cos^2\theta_B+\sin^2\theta_A\sin^2\theta_B\\ &-2\,\mathrm{Re}\left(cos\theta_A\cos\theta_B\sin\theta_A\sin\theta_B e^{i(\phi_A-\phi_B)}\right)\Big)\\ \pm\lambda_1\lambda_2\Big(&2\,\mathrm{Re}\left(\cos^2\theta_A\cos^2\theta_B e^{i(\chi_A-\chi_B)}\right)\\ &+2\,\mathrm{Re}\left(\sin^2\theta_A\sin^2\theta_B e^{i(\chi_A-\chi_B)}e^{-2i(\phi_A-\phi_B)}\right)\\ &-4\,\mathrm{Re}\left(\cos\theta_A\cos\theta_B\sin\theta_A\sin\theta_B e^{i(\chi_A-\chi_B)}e^{-i(\phi_A-\phi_B)}\right)\Big)\Big).\end{aligned} \tag{F.4}$$

The maximum fidelity with $\left|\psi^+_{\mathrm{Bell}}\right\rangle$ is obtained for $\theta_A=\theta_B$, $\phi_A=\phi_B+\pi$ and $\chi_A=\chi_B$, and with $\left|\psi^-_{\mathrm{Bell}}\right\rangle$ for $\theta_A=\theta_B$, $\phi_A=\phi_B+\pi$ and $\chi_A=\chi_B+\pi$. Therefore, while there are multiple options for which unitaries to apply to the Nodes A and B to get the maximum fidelity, once one unitary is specified (through its θ, ϕ and χ), the unitary applied to the other qubit is determined. In particular, the arbitrary unitary on a chosen node can be elected to cancel out the initial unitary applied to that node (U_A^I for node A or U_B^I for node B). In this case, one node has no unitary applied, and the unitary for the other node is completely specified. This results in two conclusions:

- To obtain the maximum possible fidelity with local operations, it is sufficient to apply local operations to only one node (this is a property of the entangled nature of the target state).
- In the case that only one node has a unitary applied, that unitary is completely specified.

Finally, the maximum fidelity $F_{\max}=\frac{1}{2}\left(\lambda_1+\lambda_2\right)^2$ is solely a function of the Schmidt values, indicating the recoverable fidelity is limited by the incomplete entanglement of the two nodes after the photonic measurement.

G

Evaluating the Relative Emission Intensity to a Detector of Finite Size

The probability that an atom emits in the direction of the collection system compared to the probability that it would emit in any random region of space with the same solid angle is given the symbol η_p. In the Pico-Oven test, the collection system is oriented orthogonal to the excited dipole transition of the atoms, therefore maximising the collected fluorescence. However, the finite sold angle covered by the collection region makes the calculation of the η_p involved.

An excited dipole decays by coupling to the electric field of travelling vacuum modes. These vacuum modes each have a direction, frequency and polarisation. For determining the spatial profile of spontaneous emission, it is only necessary to know that the total emission rate to both modes of a given direction combined is proportional to the component of the direction that does not lie along the dipole.

For our case, the centre of collection is orthogonal to the dipole to maximise collected counts. A polar coordinate system with angles θ and ϕ for the direction of emission is set up according to Fig. G.1. With this, the relative rate of emission in a particular direction is given by

$$R(\theta, \phi) = 1 - (\sin(\theta)\cos(\phi))^2. \tag{G.1}$$

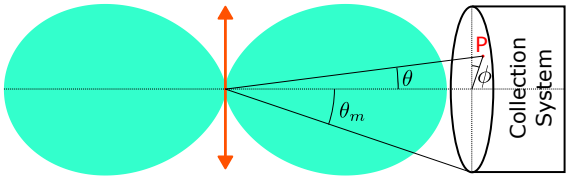


Figure G.1: Schematic diagram for the calculation of η_p for finite detector size. An oscillating dipole (orange) emits photons stochastically, but with statistics determined by the dipole emission pattern (light blue). The direction of emission is parametrised by a polar coordinate system centred on the dipole with angles $0 < \theta < \pi$ and $0 < \phi < 2\pi$ as shown for an example point P on the collection system's surface. Photons are collected from the region $0 < \theta < \theta_m$ and $0 < \phi < 2\pi$, where θ_m is the half angle of the collection lens at the front of the collection system

This leads to

$$\eta_p \left(\frac{1}{4\pi} \int_0^{\theta_m} \int_0^{2\pi} \sin(\theta)\, d\theta\, d\phi \right) = \frac{\int_0^{\theta_m} \int_0^{2\pi} R(\theta,\phi) \sin(\theta)\, d\theta\, d\phi}{\int_0^{\pi} \int_0^{2\pi} R(\theta,\phi) \sin(\theta)\, d\theta\, d\phi}, \tag{G.2}$$

where θ_m is the half-angle of the collection cone. Performing this integration leads to

$$\eta_p = \frac{3}{2} \frac{2\pi(1-\cos(\theta_m)) - \pi\left(\frac{2}{3} - \cos(\theta_m) + \frac{1}{3}\cos^3(\theta_m)\right)}{2\pi(1-\cos(\theta_m))} \tag{G.3}$$

$$= \frac{3}{2}\left(1 - \frac{2\eta_g - \frac{1}{3} + \frac{1}{3}(1-2\eta_g)^3}{4\eta_g}\right), \tag{G.4}$$

where $\eta_g = (1-\cos(\theta_m))/2$ is the geometric efficiency of the collection.