CONTROLLING QUANTUM SYSTEMS AT THE PULSE LEVEL: CAVITY QED AND BEYOND

JAN OLE ERNST

A thesis submitted for the degree of

Doctor of Philosophy

Department of Atomic & Laser Physics

Wolfson College

University of Oxford

"I hold that the mark of a genuine idea is that its possibility can be proved, either a priori by conceiving its cause or reason, or a posteriori when experience teaches us that it is in fact in nature."

Gottfried Leibniz

Dedicated to Kathrin & Arneborg Ernst who always encouraged me to reach for the stars.

ABSTRACT

This thesis addresses challenges in the experimental realisation of optical quantum computing with atom-cavity platforms and develops techniques for experimental quantum control that generalise across a range of quantum modalities. Within the framework of cavity quantum electrodynamics (cQED), methods are developed and numerically validated for fast and precise control of single atoms to generate photonic states of interest. In particular, we consider the generation of high-repetition-rate n-photon bursts and time-bin entangled states. These photonic resources are essential for applications such as boson sampling, multi-photon interference, optical quantum computing, and quantum-enhanced metrology. Experimental improvements in coherent population transfer are demonstrated using adaptive pulse shaping with acousto-optic modulators.

Beyond photon generation, the thesis introduces general-purpose pulse-level techniques, including pulse distortion correction, memory-efficient pulse compression, and machine learning-based optimisation. These methods are shown to be transferable to other platforms, such as neutral atom arrays and superconducting qubits. A particular contribution is the implementation of scalable pulse generation using fixed-point compressed pulse representations on FPGA hardware, achieving substantial memory savings while maintaining high fidelity in simulations of quantum operations. Additionally, reinforcement learning is presented as a powerful tool for optimising complex quantum dynamics, with state-of-the-art results demonstrated across various simulated systems. Throughout the thesis, theoretical developments are supported by numerical simulations and partial experimental validation, providing important insights for the development of future quantum technologies.

PUBLICATIONS AND CONFERENCES

[1] Jan Ole Ernst, Juan-Rafael Alvarez, Thomas D Barrett, and Axel Kuhn. "Bursts of polarised single photons from atom-cavity sources." In: *Journal of Physics B: Atomic, Molecular and Optical Physics* (2023). URL: http://iopscience.iop.org/article/10.1088/1361-6455/acf9d2.

[2] Jan Ole Ernst, Juan-Rafael Alvarez, Thomas D. Barrett, and Axel Kuhn. "Polarized single photons from a cavity-enhanced atom-light interface with coherent re-preparation." In: *Quantum 2.0 Conference and Exhibition*. Optica Publishing Group, 2022, QTh4C.4. DOI: 10.1364/QUANTUM.2022.QTh4C.4.

[3] Jan Ole Ernst, Aniket Chatterjee, Tim Franzmeyer, and Axel Kuhn. "Reinforcement Learning for Quantum Control with Physical Constraints." In: *Proceedings of the 42nd International Conference on Machine Learning (ICML)*. Vancouver, Canada: PMLR, 2025. URL: https://arxiv.org/abs/2501.14372.

[4] Jan Ole Ernst, Matthew King, and Axel Kuhn. *Fast and deterministic production of entangled resource states in a time binned basis*. Invited talk, BQIT Workshop 2024. Upgraded from poster presentation. 2024. URL: https://www.bristol.ac.uk/qet-labs/events/bqit-workshop/bqit-programme/.

[5] Jan Ole Ernst, Jan Snoeijs, Mitchell Peaks, and Jochen Wolf. "Memory-Optimized Cubic Splines for High-Fidelity Quantum Operations." In: *IEEE Transactions on Quantum Engineering* (2025). DOI: 10.1109/TQE.2025.3574463.

[6] Jan Ole Ernst and Felix Tennie. "Mode Entanglement in Fermionic and Bosonic Harmonium." In: *New Journal of Physics* (2024). URL: http://iopscience.iop.org/article/10.1088/1367-2630/ad240f.

ACKNOWLEDGMENTS

Pursuing a DPhil is quite the adventure and would not have been possible without the support of those around me. First and foremost, I have to thank my supervisor Prof. Axel Kuhn for giving me a chance to do research in his group. Axel allowed me to freely explore the things I was interested in and was also a source of great ideas. Moreover, his heart lies in the right place and he knows how to throw a great summer party. I am also grateful for meeting so many other interesting people in the Kuhn group, from whom I could learn a great deal from. Thomas Doherty taught me the basics of cQED as well as some Scottish humour for which I am grateful. Juan Rafael Alvarez Velasquez taught me the basics of laboratory quantum optics for which I also am grateful. Mark IJspeert is a fountain of knowledge. He taught me about electronics, optics, computer graphics, finance, LaTeX and bike maintenance and his attention to detail continues to inspire me every day. Chloe So gave me good insights about oscilloscopes and cavities. I am also grateful for having had a chance to work with two very talented external research students, Greg Ferrero and Marina Llano Piñeiro, who both showed dedication and skill when working with me in the lab. I have no doubt that successful research careers lie ahead of them. Working with Matt King and witnessing his journey from Masters to DPhil was a pleasure. I am convinced that with his deep-seated curiosity for quantum optics, he will do great research and pick up things where I left them. I am also deeply grateful for Felix Tennie's support, particularly his frequent lunch invitations and advice on various research and non-research related topics. He also provided invaluable guidance for writing and publishing my first paper and his mathematical intuition is awe-inspiring. A thank you also goes to Mark, Felix, Juan and Matt who all proofread parts of this thesis, as well as my examiners Stephane Guerin and Alex Lvovsky who had to read all of it!

The research presented here would not have been possible without my external collaborators. I am indebted to all my colleagues I met during my internship at Riverlane who welcomed me with open arms and taught me a lot about engineering, FPGAs and RF electronics, particularly Jan Snoeijs. Working with them beyond my internship and co-authoring a paper was a great pleasure. Moreover, I am very grateful for getting a chance to work with Tim Franzmeyer and Aniket Chatterjee. Tim's drive was always inspiring and I am grateful that he taught me so much about programming and machine learning. Aniket is the most motivated and gifted undergraduate physicist I have encountered and his can-do attitude will take him far. Working at the intersection of physics and machine learning was extremely rewarding and is something that I hope to take into my future endeavours.

I am grateful to the wider research community in Oxford and beyond for their openness in discussing ideas and research directions. Philip Thomas kindly showed me around the MPQ labs in Garching, inspiring me with the possibilities of atomic cQED. I also benefited from Patrick Sund's insights on boson sampling, as well as many fruitful discussions and coding sessions with William Hughes, whose deep understanding of optics and atomic physics was invaluable. Conversations with David Bauch greatly clarified the fidelity calculations of time-bin entangled states, while Sebastian Saner, Matthias Keller, Felix Tennie, and Andrew Daley provided further helpful input on the time-bin scheme. I also appreciated the support of experimental colleagues such as Abel, Gavin, Lorenzo, or Yao-Chih, who would discuss issues and lend me equipment. Beyond my DPhil, I am thankful to the many other researchers I interacted with and to my undergraduate tutors, whose teaching laid the foundations and inspired me to pursue this path.

The experience of doing a DPhil was influenced by all my extracurricular experiences. Rowing provided me with much needed mental distraction from physics and gave me some of the most unforgettable experiences. I had the chance to win summer eights with an Olympic champion, row against Cambridge on the Thames in front of

tens of thousands of people and to face off with the soon-to-be Olympic silver medallists in the pair at Henley Royal Regatta. The rowing community is where I also found a lot of great friends who were always supportive and interested in what I was doing in the lab. Achim and Zach were a source of inspiration for pursuing physics research in the first place and I am extremely grateful for the support and community of Oriel College Boat Club. I will also never forget racing for Oxford University Boat Club and am grateful for all my teammates who stuck with me over the years. After rowing, I also found joy in pursuing triathlon as an alternative endurance sport and am grateful for meeting so many like-minded people. I learnt lots about programming by doing hackathons and remember our trips to Amsterdam and Tokyo fondly, a particular shoutout here goes to Kyle, Keivan, Orlando, Konrad, Fernando and Jessica with whom I share some great experiences. My experience working part-time at Creator-Fund taught me so much about venture capital and commercialising research, and I am grateful to Jamie and Alex for giving me this opportunity. I will fondly miss technical discussions about physics companies with Daniel, Alex, Ben and the rest of my colleagues.

My DPhil would not have happened without the support of my family and friends. I am deeply indebted to my parents who were always there for me when things weren't going well and who were a constant source of support and motivation. While my brother and his wife were often preoccupied with their children, both of whom were born during this period! they were always supportive of my studies. My partner Kat was there for me for most of the DPhil journey and always supported me when things were not going exactly as planned for which I am incredibly grateful.

My experience of Oxford would not have been the same without the many friends I made throughout my time here. Speaking to so many smart and inspiring people kept life fun and exciting. In particular, I am grateful for meeting and getting a chance to interact with Dominik, Ed, Giovanni, Angus, Jonas, Dylan, Orlando, Tom, Shreshth, Fernando, Marcel, Mert, and Madhu (and many others). I also had the pleasure of

living with some great housemates. In particular, I will miss Giovanni's Italian cuisine, philosophical discussions with Ed and Dominik's never-ending arguments. My long-term friends from Germany were also supportive of whatever I was doing and provided some distraction from the Oxford bubble. Big thanks to Linus, Louis, Constantin, Moritz, Jubal, Dennis and Lennart.

CONTENTS

1 INTRODUCTION 1

1.1 Elements of Quantum Mechanics 1

1.2 Quantum Technologies 3

1.2.1 Material Platforms for Quantum Computing 3

1.3 Thesis Outline 6

2 QUANTUM, ATOMIC PHYSICS & CAVITY-QED BACKGROUND 8

2.0.1 Quantum Dynamics & Control 8

2.1 Atomic Physics 14

2.1.1 Coupling of Atomic States with Hyperfine Structure 14

2.1.2 Rubidium-87 17

2.1.3 Stimulated Raman Adiabatic Passage 19

2.2 Cavity QED 21

2.2.1 Optical Cavities 21

2.2.2 The Jaynes-Cummings model 24

2.2.3 V-STIRAP 26

2.2.4 Cavity Polarisation Modes 28

2.2.5 Historical Overview of atomic cQED Experiments for Single Photon Production 32

3 SINGLE PHOTON PRODUCTION SCHEME 35

3.1 Implementation in idealised three level system 37

3.1.1 Atomic Structure 37

3.1.2 Quantisation axis and polarised photon production 38

3.1.3 Adiabatic transfer of population 40

3.1.4 Efficiency 41

3.1.5 STIRAP Re-preparation 45
3.2 Implementation in real atomic system: Rubidium 87 49
3.2.1 Photon Production 50
3.2.2 STIRAP Re-preparation 56
3.2.3 Optical Pumping 59
3.2.4 Full Scheme 61
3.3 Discussion 65
3.3.1 Requirements for Quantum Advantage with Boson Sampling 67
3.3.2 Outlook 72
4 TIME-BIN ENTANGLED PHOTONIC STATE SCHEME 74
4.1 Scheme Overview 76
4.2 Raman Pulses 81
4.2.1 STIRAP Rotations 83
4.2.2 Far Off Resonant Rotations 85
4.3 V-STIRAP Photon Emission 89
4.4 Results 94
4.4.1 Expected Coincidence Rates and Measurement Setup 95
4.4.2 Calculating Photonic State Fidelities 97
4.4.3 Photonic State Fidelity Results 105
4.5 Outlook 108
5 EXPERIMENTAL RESULTS & PROGRESS 110
5.1 MOT 111
5.2 Cavity 113
5.3 Laser System 116
5.3.1 Laser Locking 116
5.3.2 Laser Preparation Table 119
5.4 Electronic Control of Laboratory Equipment 120
5.5 AOM-based Pulse Shaping 122

5.5.1 Time-Dependent Amplitude Shaping 124

5.5.2 Time-Dependent Frequency Chirps 127

5.5.3 Results 128

5.5.4 Power and Rabi Frequency Calibration 130

5.6 Coherent Population Transfer 132

5.7 Outlook 136

6 SCALING PULSE LEVEL CONTROL OF QUANTUM EXPERIMENTS 137

6.1 Introduction 137

6.2 Methods 140

6.2.1 Floating Point Cubic Spline Fitting 140

6.2.2 Implementation on FPGA 142

6.2.3 Quantisation Aware Fit 144

6.3 Benchmarking 148

6.3.1 Coherent Quantum Population Transfer in ^{87}**Rb** Atom 149

6.3.2 Atom Transport 153

6.4 Outlook 157

7 REINFORCEMENT LEARNING FOR CONTROLLING QUANTUM DYNAMICS 159

7.1 Introduction 159

7.2 Related Work 161

7.3 Background and Methods 164

7.3.1 Reinforcement Learning 164

7.3.2 Physics-Constrained Reinforcement Learning 166

7.3.3 Reward Shaping 170

7.3.4 Implementation Details 172

7.4 Numerical Experiments 173

7.4.1 Λ System 174

7.4.2 Rydberg Atom 180

7.4.3 Transmon Reset 183

7.4.4 Multi-Step RL 185

7.5 Physical Experimental Realisation 186

7.6 Outlook 189

8 CONCLUSIONS 191

A APPENDIX 195

A.1 ^{87}**Rb** Transition Line Data 195

A.2 Boson Sampling 195

A.2.1 Connection Between Boson Sampling and Matrix Permanent Calculations 195

A.2.2 Classical Simulation of Boson Sampling 197

A.3 Dynamical Stark Shift Calculation 198

A.4 Far Detuned Raman Pulse Hamiltonian Derivation 199

A.5 Time-Bin V-STIRAP MCWF Error Propagation 201

A.6 Stabiliser expectation values as operator expectation values 203

A.7 Numerical Integration of Photonic Density Matrix Elements 203

A.7.1 Uncertainty Propagation for GHZ Pauli Expectation Values 205

A.8 AOM Background 206

A.9 NLMS Filter Update Rule 208

A.10 Additional Amplitude Optimised Pulses 209

A.11 Frequency Analysis of Pulses 209

A.12 Data Analysis for Experimental Results 210

A.13 RL Background 211

A.13.1 Proximal Policy Optimisation (PPO) 211

A.13.2 Extended Time Horizon in Multi-step RL 212

A.14 RL Implementation Details 213

A.14.1 Training Implementation Details 213

A.14.2 Benchmarking of Simulation Speed 213

A.14.3 Signal Processing & Analysis 214

A.14.4 Noise Model 217

A.15 Lambda System 218

A.16 Rydberg Atoms 219

A.17 Transmon Reset 220

BIBLIOGRAPHY 222

LIST OF FIGURES

Figure 1 Magnetic field dependence of dipole matrix elements and energy levels on the D_2 line in ^{87}**Rb** 18

Figure 2 Comparing the spectral transmission pattern of three cavities with different Finesses at a fixed length of 10^{-3}m. The linewidths (FWHM) are indicated. 22

Figure 3 Cavity used in the current experimental setup with $L_{cav} = 339\mu m$[18] 23

Figure 4 Level Structure for Jaynes-Cummings Hamiltonian. [18]. 27

Figure 5 Cavity coupling components in x, y directions in the top and bottom row respectively, where the columns describe the components in the atomic basis $(\pi, \sigma^+, \sigma^-)$ for arbitrary relative orientation of $\vec{Q}$ and $\vec{C}$ which is parametrised in spherical coordinates θ, ϕ. 29

Figure 6 Idealised 3-level atomic system for photon production 37

Figure 7 Polarisation geometry with quantisation axis perpendicular to the cavity axis 39

Figure 8 V-STIRAP photon production efficiency in an idealised three level system 42

Figure 9 Optimised peak Rabi frequency for V-STIRAP photon production in idealised three level system 43

Figure 10 Cavity design optimisation for maximising photon generation rate 45

Figure 11 Re-preparation efficiency in idealised three-level system 47

Figure 12 Various pulse shapes with optimised parameters used for coherent re-preparation 48

Figure 13 Photon generation sequence in ^{87}Rb using cavity-enhanced emission, STIRAP re-preparation, and optical pumping 51

Figure 14 Variation of relative atomic phase with V-STIRAP driving pulse length 52

Figure 15 V-STIRAP Photon Production efficiency for both D_1 and D_2 lines variation with g and κ 53

Figure 16 V-STIRAP efficiency for various atom cavity coupling strengths g for both the D_1 and D_2 lines, at fixed κ 54

Figure 17 Optimal peak Rabi frequency for V-STIRAP photon production 56

Figure 18 STIRAP re-preparation efficiency for both D_1 and D_2 transitions plotted against the length of the re-preparation pulses Ω 57

Figure 19 STIRAP repumping pulse shapes with optimised parameters for D_1 58

Figure 20 Optical Pumping Efficiency plotted as a function of time 60

Figure 21 Polarised photon stream production likelihood (top) and expected N-photon sequence coincidence rate (bottom) 61

Figure 22 V-STIRAP photon polarisation when the atom is initially falsely prepared 64

Figure 23 Explicit comparison of various single photon sources by repetition rate and brightness. 67

Figure 24 Simplified boson sampling setup with single photon source, demultiplexer, interferometer and detection setup. 68

Figure 25 System efficiency requirements for quantum advantage in boson sampling with realistic atom-cavity photon sources 71

Figure 26 Pulse sequence for generating photonic GHZ states (time-binned) via alternating Raman rotations and V-STIRAP emissions from an atom initially prepared in $|F = 2, m_F = 0\rangle$ 77

Figure 27 Scheme for generating photonic linear cluster states by inserting an $X(\pi/2)$ atomic rotation between emission steps in the GHZ sequence, enabling entanglement between successive photons via a Raman pulse sequence. 79

Figure 28 Total scattering probability and required laser power for a $\sigma^x(\pi)$ pulse with 10, μm waist are plotted versus two-photon detuning, showing both are within experimentally feasible ranges. 86

Figure 29 Comparison of standard flattop and RL-optimised pulses for Rotation 3 shows that only the optimised sequence avoids cross-talk and achieves high-fidelity coherent population transfer. 87

Figure 30 RL-optimised pulses achieve simultaneous high-fidelity (>0.999) coherent transfer between two atomic state pairs with smooth control compatible with standard hardware. 87

Figure 31 Monte Carlo analysis showing that longer, smoothly shaped pulses ($\sin^2$ and Blackman) reduce re-excitation errors compared to fast flattop pulses, balancing emission efficiency with repetition rate constraints. 90

Figure 32 Analysis of different rise times of fixed length flattop Blackman pulses of total duration 1μs on re-excitation rate errors and photon emission probability. 92

Figure 33 The detection setup for routing polarised photons for σ^z or interferometric σ^z measurements, enabling physical detection of photon coherences across time-bins. 95

Figure 34 Photon generation rates with realistic optical losses in Hz (log scale), versus photon number for various different physical platforms. 97

Figure 35 Fidelity of single-photon state $|+\rangle$ with magnetic field and laser phase for far-detuned Raman rotations and STIRAP pulses. 106

Figure 36 Fidelity of n-photon GHZ states is compared between phenomenological and stabiliser methods for STIRAP and far-detuned Raman pulses 107

Figure 37 MOT of approximately 3000μm diameter shown 0.7 cm below the cavity. 113

Figure 38 Diagram showing laser locking and desired transitions one wishes to address for the D_2 line of ^{87}**Rb**. 117

Figure 39 Electronic Control of Experimental System 121

Figure 40 Unoptimised measured STIRAP pulses 123

Figure 41 Comparison of optimised and unoptimised driving on fixed frequency pulses with respect to the desired theoretical shapes (cf. Eq. 40). 129

Figure 42 Fluorescence Readout Pulse Experimental Setup 131

Figure 43 Experimental depopulation of $F = 2$ ground level for different Rabi frequencies and pulses. 134

Figure 44 Quantisation aware fitting method for example Gaussian 145

Figure 45 Benchmarking quantisation aware fit for STIRAP task. 149

Figure 46 Atom Transport Simulation Results with Quantisation Aware Pulse Fit 152

Figure 47 Fidelity and compute time scaling with max solver step constraint 168

Figure 48 Smoothing Parameter Ablation for different environments 169

Figure 49 Instantaneous Frequency Analysis of RL Signals 169

Figure 50 Energy Level Diagram for for four level Λ system 174

Figure 51 Comparison of different RL algorithms for Λ system 175

Figure 52 Quantum state dynamics comparison for RL and OC Solution 175

Figure 53 RL Discovered Pulse Shapes for Λ system 176

Figure 54 Pulse Area Sweep for RL Solutions 178

Figure 55 Robustness to random noise in Λ System 179
Figure 56 RL Discovered Pulses for Rydberg Gate 182
Figure 57 Optimal Transmon Reset Waveforms determined with RL 183
Figure 58 Multi-Step RL for Λ System showing improved fidelity 185
Figure 59 The D_2 transition between the $5^2S_{1/2}$ level and the $5^2P_{3/2}$ levels are shown with hyperfine structure. Diagram taken from [216]. 196
Figure 60 Single-pass AOM configuration 208
Figure 61 Optimising driving of RL determined waveforms 209
Figure 62 Spectral characterisation of the Pump pulse 210
Figure 63 Example fluorescence trace 211
Figure 64 Simulation speedup with GPU parallelisation 214
Figure 65 Comparison of different smoothing functions and learning dynamics 216

LIST OF TABLES

Table 1 Comparison of time evolution methods for quantum systems. Here d is the Hilbert space dimension. 11
Table 2 Probability of emitting a photon of particular polarisation, for a cavity with $C \approx 10$ [67] when applying the optimised driving pulse when correctly initialised (i.e. $m_F = 0$) and when incorrectly initialised into the other sub-states. The purity, i.e. $p(H)/(p(H) + p(\pi))$ is also shown and decreases significantly for the $m_F = \pm 1$ sub-states. 63
Table 3 Population after several repumping cycles. 63

Table 4 Comparison of Far-Off-Resonant Raman Rotations and STIRAP 81

Table 5 Pulse parameters and notes for optimised STIRAP rotations for GHZ state generation. $\Delta_{F'=1}$ denotes a common two photon detuning of the Stokes and Pump Pulses from resonance with respect to the $F' = 1$ line of the D_1 transition. 84

Table 6 Pulse Parameters for far-detuned Raman Pulses 89

Table 7 Explicit Loss Calculation with state of the art components. 96

Table 8 Lasers used in the experiment and how they are locked. 116

Table 9 All AOMs and EOMs used for modulating lasers and their typical operational frequencies. 119

Table 10 Control sequence for shaped pulse experiment with fluorescence readout where Channels correspond to the relevant DAQ Card Channels (Ch.) or AWG Channels (AWG Ch.) which are controlled at various points in time. 133

Table 11 Example coefficients for optimised cubic spline fit 146

Table 12 Comparing RL to other gradient based methods. 164

Table 13 Example timing breakdown for experiment described in Sec. 5.6. Note that one may take many shots per RL policy update. 188

ACRONYMS

MMI Multi Mode Interferometer

AOM Acousto-Optic Modulator

cQED Cavity Quantum Electrodynamics

MBQC Measurement-Based Quantum Computing

MZI Mach–Zehnder Interferometer

STIRAP Stimulated Raman Adiabatic Passage

V-STIRAP Vacuum enhanced STIRAP

NPBS Non Polarising Beam Splitter

PBS Polarising Beam Splitter

RL Reinforcement Learning

MOT Magneto Optical Trap

AWG Arbitrary Waveform Generator

AOD Acousto Optic Deflector

FPGA Field Programmable Gate Array

DAC Digital to Analogue Converter

NCO Numerically Controlled Oscillator

DDS Direct Digital Synthesis

GPU Graphical Processing Unit

PPO Proximal Policy Optimisation

DDPG Deep Deterministic Policy Gradient

TD3 Twin Delayed Deep Deterministic Policy Gradient

1

INTRODUCTION

Quantum mechanics has captivated generations of the world's brightest physicists, challenging our understanding of the most fundamental constituents of the physical world. This introduction begins with a brief conceptual overview of quantum mechanics. It then outlines the translational potential of contemporary quantum technologies, as well as the diverse material platforms being developed for quantum computing. The chapter concludes with a brief overview of the structure and content of the thesis. The first chapters focus on improving the operation of atom-cavity systems for producing photonic states of interest for applications in boson sampling and measurement based quantum computing with some preliminary experimental results. The final two chapters discuss work on memory efficient representations of pulses for controlling quantum systems at scale and reinforcement learning for controlling generalised quantum systems.

1.1 ELEMENTS OF QUANTUM MECHANICS

Quantum mechanics is the fundamental theory governing the behaviour of matter at the subatomic scale, like single photons or electrons. Unlike classical physics, which offers a deterministic and somewhat intuitive description of the macroscopic world, quantum mechanics introduces a radically different framework. It is one characterised by probabilistic measurement outcomes, behaviours associated with both waves and particles, as well as superposition and entanglement which establishes correlations between particles that cannot be explained classically.

Quantum systems are described by wave-functions, which encode probability amplitudes for finding a system in a given state. Measurement causes the wave-function to *collapse* to a definite state, a process which is still subject to much debate in the foundations of the field and often described as the *measurement problem* [200]. The measurement problem is generally understood as the incompatibility between the quantum system description, as given by the wave-function, which permits multiple measurement outcomes, and the observation of a single outcome upon performing a measurement.

Particles such as electrons or photons can exist in a superposition of multiple states. For example, an electron can be a superposition of both "spin-up" and "spin-down" until it is measured. Mathematically, in a Hilbert space spanned by two possible quantum states $\{|\uparrow\rangle, |\downarrow\rangle\}$ a general state can be represented as a superposition $\alpha|\uparrow\rangle + \beta|\downarrow\rangle$, where $|\alpha|^2 + |\beta|^2 = 1$ to conserve probability. Such quantum states provide unique non-classical signatures which manifest in the form of interference phenomena. When two particles are entangled, their quantum states exhibit stronger correlations than those permitted by classical physics. For example, one can construct an entangled pair of states described by the Bell state:

$$\Phi = \frac{1}{\sqrt{2}}\left(|\uparrow\rangle \otimes |\uparrow\rangle + |\downarrow\rangle \otimes |\downarrow\rangle\right), \tag{1}$$

where knowing the outcome of the measurement of one particle determines the possible outcome when measuring the other. Both particles will always be found in the same state, if measured in the same basis, even if they are spatially separated. This "spooky action at a distance" [72], has been experimentally verified in hallmark experiments [14, 103].

Quantum physics underpins much of modern technology, from semiconductors and lasers to MRI machines and atomic clocks. Its more recent applications in quantum computing, quantum cryptography, and quantum sensing promise revolutionary advances in computation, secure communication, and precision measurements. A more

detailed explanation of quantum technologies and their applications, in particular quantum computing will be discussed in Sec. 1.2.

1.2 QUANTUM TECHNOLOGIES

Quantum technologies seek to exploit the fundamental principles of quantum mechanics to surpass what is otherwise possible with classical physics. These technologies have shown promise in revolutionising fields ranging from sensing to cryptography and computing.

Quantum sensing enhances precision in magnetometry [159], gravimetry, and timekeeping [61, 214]. **Quantum cryptography** provides provably secure communication through Quantum Key Distribution (QKD) [27]. **Quantum networking** aims to distribute entanglement over long distances using repeaters [43], entanglement swapping [254], and memories, advancing toward a quantum internet [241]. **Quantum computing** leverages the exponential scaling of the Hilbert space to tackle classically intractable problems. The circuit model [63] and algorithms such as Shor's factoring [209] and Grover's search [94] established the field. Platforms under exploration include superconducting circuits [125], trapped ions [157], spin qubits, photonics [180], and neutral atoms [194]. The threshold theorem [5] and surface codes [79] chart a path to fault-tolerant architectures where logical qubits emerge from imperfect physical ones.

1.2.1 *Material Platforms for Quantum Computing*

Much of this thesis focuses on improving the operation of atomic single-photon sources and photons as a material platform for computing which will be outlined in more detail. Neutral atoms and superconducting qubits are also relevant for the work done in Chapter 6 and Chapter 7 respectively, and are briefly outlined. Other promising mate-

rial platforms, such as trapped ions or spin qubits also show potential for computing but will not be discussed further.

Single photons are promising resources for the realisation of quantum computing and communication; information transfer is quick and can be performed at room temperature. Photons are relatively robust to decoherence, as they interact weakly with other particles. This however poses a challenge for the realisation of two-qubit gates for universal quantum computation. Information can be encoded in discrete degrees of freedom such as photon polarisation, photon number or photon arrival time, but proposals also exist for continuous variable quantum information in the optical domain [40]. Another key goal in quantum technology development is realising quantum networks, where stationary qubits (such as ions or atoms) are connected by photons for distributed computing.

Knill, Laflamme and Milburn demonstrated that with the aid of auxiliary measurements an optical quantum gate could be realised [126]. This requires extremely efficient optical components, as well as photon sources and detectors. If realistic losses are taken into account, it is very challenging to realise quantum circuit operations at reasonable clock rates. There exists another promising proposal for optical quantum computing, namely Measurement-Based Quantum Computing (MBQC) [186]. Sequences of measurements are performed on maximally entangled initial states. Recently, much work has been dedicated to developing formalisms for error correction, as well as designing more hardware efficient implementations [24] and reducing the losses of optical components [180]. This framework requires highly efficient single photon sources for producing entangled states at the outset with high clock rates [87] and high fidelities. This challenge is addressed with a novel encoding in Ch. 4 and it is shown that a time binned basis can provide advantages over the more common polarisation basis.

Another key application of single photons is boson sampling, where photons interfere in different modes of a Multi Mode Interferometer (MMI) and the output distribution is measured. Calculating this distribution is believed to be classically intractable

for large systems [58]. Measuring such distributions with a large number of photons motivates the development of improved single-photon generation schemes; which is discussed in Ch. 3.

Neutral atom quantum computing has advanced rapidly, with trapped atom arrays enabling both error-corrected logical computation [33] and analogue quantum simulation [32]. Single-qubit gates are realised optically or via microwaves, while two-qubit entangling gates exploit the Rydberg blockade to implement controlled-phase operations using phase and amplitude controlled laser pulses. Atoms, confined in optical tweezers that can be dynamically rearranged, providing a reconfigurable architecture with demonstrated arrays of thousands of coherently controlled qubits [151], drawing parallels to the transformative role of FPGAs in classical computing [35]. However, key challenges remain, including probabilistic atom loading requiring rearrangement, optical aberrations, laser power limitations and relatively slow readout. This motivates developing strategies for scalable, high-fidelity control of such systems, as explored in Chapters. 6 and 7.

Superconducting qubits, particularly Transmons, have driven major advances in quantum computing, enabling demonstrations of quantum supremacy [13] and surface-code error correction [4]. Transmons suppress noise while retaining sufficient anharmonicity to encode a qubit in the lowest energy levels which can be controlled with microwave pulses. However, Transmon qubits face challenges including control crosstalk, limited coherence times, spectral crowding in multi-qubit chips, and fabrication-induced variability. High-fidelity qubit reset is critical for optimally using the available coherence time in the context of mid-circuit measurements, feedback, and error correction. Designing optimal reset waveforms is addressed in Ch. 7.

1.3 THESIS OUTLINE

This thesis focuses on improving the fundamental building blocks necessary for experimental tests of quantum mechanics and the realisation of quantum computing, particularly for photonic qubits. The work begins with an exploration through the lens of Cavity Quantum Electrodynamics (cQED), aiming to precisely control single atoms to generate specific photonic states of interest. Tailored and high repetition rate photon bursts can be used to realise boson sampling or to perform other multi-photon interference experiments. Moreover, entangled photonic states are of great interest for realising measurement-based optical quantum computing and also for applications in precision measurement. Many of the insights which we develop in later chapters such as pulse distortion correction, memory compression and machine learning based optimisation, readily apply to other experimental platforms, such as neutral atom arrays or superconducting qubits, as shown in the last two chapters. A unifying theme throughout the thesis is the challenge of achieving dynamic control of quantum experiments. This requires conceptual, theoretical and experimental innovation.

Ch. 2 gives a technical background on quantum physics, atomic physics and cQED. In Ch. 3, the development of novel schemes for experimental generation of n-photon bursts for applications in boson sampling and fundamental tests of the quantum properties of many photon states is described. Thereafter, in Ch. 4 we discuss a novel scheme for generating time-bin entangled photons is extensively analysed and benchmarked with numerical simulations. These entangled states form the bedrock of measurement-based optical realisations of quantum computing and can also be used for quantum metrology. In Ch. 5 we discuss progress towards an experimental realisation of an n-photon experiment with some partial measurements showing optimised coherent population transfer. The results of Ch. 6 show advancements in scaling control systems for larger scale quantum computing experiments by compressing experimental pulse information effectively while correcting for fixed point quantisation errors on

Field Programmable Gate Array (FPGA) platforms. Simulations of population transfer in neutral atoms and atom transport show that optimal fidelities can be attained while significantly compressing pulse information compared to standard implementations with arbitrary waveform generators. Finally, in Ch. 7 the application of Reinforcement Learning (RL) to generalised quantum dynamical control is discussed and state of the art results for various quantum systems are presented with the possibility of extending this approach to real world experiments.

2

QUANTUM, ATOMIC PHYSICS & CAVITY-QED BACKGROUND

This chapter provides the theoretical foundation required to understand the work presented in the following chapters. It begins with an introduction to quantum physics, focusing on the time evolution of quantum systems as well as an explanation of how such systems can be simulated and controlled. Next, essential concepts from atomic physics are reviewed, with particular emphasis on those relevant to the photon production schemes introduced later. This includes a detailed overview of the properties of the hyperfine states of Rubidium (^{87}**Rb**) and an explanation of the Stimulated Raman Adiabatic Passage (STIRAP) technique. Finally, this chapter introduces Cavity Quantum Electrodynamics and the Vacuum enhanced STIRAP (V-STIRAP) method for single-photon generation, along with a historical perspective on key experimental developments in the field.

2.0.1 *Quantum Dynamics & Control*

Quantum dynamics describes the time evolution of quantum systems. The state of a quantum system is represented as a vector in a complex Hilbert space $\mathcal{H}$. The most common representation is the *state vector* $|\psi\rangle \in \mathcal{H}$. A pure quantum state is described by a normalised vector [161] $|\psi\rangle = \begin{pmatrix} \psi_1, & \psi_2, & \dots & \psi_n \end{pmatrix}^\top$, where $\langle\psi|\psi\rangle = 1$. A more general representation is the *density matrix* ρ, a positive semi-definite density operator, which for a pure state reads $\rho = |\psi\rangle\langle\psi|$ [161]. The state of any quantum system can be written as $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ with unit trace $\sum_i p_i = 1$ [161]. In the density

matrix formalism, the state populations are defined as $|\psi_i|^2$ (these are the diagonal terms of the matrix representation of ρ).

Physical observables are associated with Hermitian operators $\hat{O}$ that form an algebra $\mathcal{A}(\mathcal{H})$. Such observables are represented by Hermitian operators $\hat{O} = \hat{O}^\dagger$. The expectation value of any observable $\hat{O}$ with respect to a density matrix ρ reads [161]:

$$\langle \hat{O} \rangle_\rho = tr(\rho \hat{O}). \tag{2}$$

The time evolution of a closed quantum system is described by unitary operators $\hat{U}$, where $\hat{U}^\dagger \hat{U} = \hat{U}\hat{U}^\dagger = \mathcal{I}$ which ensures that such evolution is reversible. The unitary time evolution of $|\psi(t)\rangle$ is governed by the *time-dependent Schrödinger equation* [202]:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle , \tag{3}$$

where $\hbar$ is the reduced Planck constant, and $\hat{H}$ is the Hamiltonian, an observable which describes the system's total energy. The Schrödinger equation captures all closed, quantum mechanical dynamics, but real-world quantum systems are subject to interactions with their environment and decoherence induces irreversible, non-unitary dynamics. While one can, in principle, model environmental decoherence by solving the full Schrödinger equation for the system coupled coherently to a continuum of environmental modes [195], this approach is computationally expensive. In practice, it is far more efficient to describe the reduced system dynamics using effective master equations and their stochastic unravelling, such as the Monte Carlo wave-function method, as will be discussed in what follows.

Such systems are described as open quantum systems, where environmental interactions lead to non-unitary evolution. In the Markov approximation, memory effects in the environment are neglected, meaning the system's future evolution depends only

on its current state, not its history. This assumption allows the system's dynamics to be captured by a time-local master equation [60] which reads:

$$\frac{d}{dt}\rho(t) = \mathcal{L}(t)\rho(t), \tag{4}$$

where the evolution under the super-operator $\mathcal{L}(t)$ can be decomposed into unitary and non-unitary components as follows:

$$\mathcal{L}(t)\rho(t) = -i[\hat{H}, \rho(t)] + \sum_k \left(\hat{L}_k \rho(t) \hat{L}_k^\dagger - \frac{1}{2}\{\hat{L}_k^\dagger \hat{L}_k, \rho(t)\} \right). \tag{5}$$

where the Lindblad operators $\hat{L}_k$ describe non-unitary evolution (e.g. spontaneous emission of an atom, dephasing, cavity decay or other irreversible effects).

The introduction of open system dynamics makes the numerical calculation of a quantum system's time dependence considerably more challenging, as the dimensionality of the system's state space increases from d to d^2 where d is the Hilbert space dimension. Many parts of this thesis rely on numerical simulations of such quantum systems and a detailed comparison of different simulation methods is provided in Tab. 1.

The Monte Carlo wave-function (MCWF) method [59], also known as the quantum jump approach, is a stochastic technique for simulating the dynamics of open quantum systems. It is equivalent to the master equation [59]. Unlike the density matrix formalism, which describes ensemble averages, the MCWF method models individual quantum trajectories, providing insights into the behaviour of single quantum systems under continuous observation. For example, this is useful when analysing so-called re-excitation errors where a coherent process follows an incoherent spontaneous emission process, as will be discussed in Ch. 4. In the Monte Carlo wave-function method, the system's evolution is governed by an effective non-Hermitian Hamiltonian:

$$\hat{H}_{\text{eff}} = \hat{H}_{\text{sys}} - \frac{i\hbar}{2}\sum_n \hat{C}_n^\dagger \hat{C}_n, \tag{6}$$

where $\hat{H}_{\text{sys}}$ is the system Hamiltonian, and $\hat{C}_n$ are the collapse operators representing different dissipative processes. The non-Hermitian term accounts for the probability loss outside of the Hilbert space due to quantum jumps. A computationally efficient implementation is described in [114].

By averaging over many such trajectories, one can reconstruct the ensemble behaviour of the system. The advantages are clear for systems with higher dimensions d, where the linear scaling in d is less computationally demanding than for a full master equation evolution which scales as d^2.

Feature	**Schrödinger Equation [202]**	**Master Equation (Lindblad) [60]**	**MC Master Equation (MCWF) [59]**
System Type	Closed	Open (Markovian)	Open (trajectory-based)
State Representation	Pure state $\|\psi\rangle$	Density matrix ρ	Pure state $\|\psi\rangle$
Scaling with Hilbert Space	$\mathcal{O}(d)$	$\mathcal{O}(d^2)$	$\mathcal{O}(d)$ (per trajectory)
Numerical Accuracy	High	Moderate (approximates ensemble avg.)	$\epsilon \propto 1/N$ (N is no. of traj.)
Parallelisability	Limited	Limited	High (across trajectories)
Includes Decoherence	No	Yes	Yes
Interpretability	Direct state-vector evolution	Mixed state evolution	State-vector trajectories

Table 1: Comparison of time evolution methods for quantum systems. Here d is the Hilbert space dimension.

Tab. 1 provides a comparison between three widespread methods used for calculating quantum dynamics; the Schrödinger equation is most commonly used for computing closed quantum system dynamics, while an ensemble or stochastic open quantum system approach might be used for capturing open system dynamics. This is by no means an exhaustive list, other approximations and methods exist such as the stochas-

tic Schrödinger equation [52] for modelling non-Markovian interactions or tensor network methods which provide computational speed-ups for certain entanglement structures [171].

Quantum control focuses on the design and implementation of tailored external fields, typically laser pulses, microwave signals, or other electromagnetic controls applied to a quantum system to drive its evolution toward a desired target [88, 127, 128]. These targets may include a wide range of objectives, such as quantum state initialisation, state preparation, population transfer, photon emission, quantum gates, and precise quantum measurements. Such tasks are often performed under constraints, including limited time, hardware imperfections, hardware control limits in power, instantaneous bandwidth and response time, as well as environmental decoherence.

Achieving precise quantum control is particularly critical for fault-tolerant quantum computing, where physical operations must be executed with extremely low error rates to allow for scalable quantum error correction [222]. In this context, optimal control techniques tailor the control pulses to reduce sensitivity to noise and control-induced errors, while simultaneously minimising operation times to enhance system performance. Advances in quantum control thus directly impact the feasibility of near-term quantum technology development and fault-tolerant universal quantum computation.

The optimal control of quantum systems is a central challenge in the advancement of quantum technologies, underpinning progress in quantum computing, networking and sensing. Beyond these emerging technologies, quantum control also plays a pivotal role in established fields such as quantum chemistry [45] and solid-state physics [88], where it enables control over molecular dynamics or condensed matter systems at the quantum level.

The manipulation of quantum systems in quantum control with time-dependent control fields is typically represented by the time dependent *control Hamiltonian*. The total Hamiltonian $\hat{H}(t)$ of a controlled system is [86]:

$$\hat{H}(t) = \hat{H}_0 + \sum_i a_i(t)\hat{H}_i, \tag{7}$$

where $\hat{H}_0$ is the drift Hamiltonian of the unperturbed system, $a_i(t)$ are time-dependent control actions, and $\hat{H}_i$ are control Hamiltonians. Real quantum systems are open, so to correctly capture their evolution Lindblad terms of the form $\hat{L}_k$ are added to Eq. 7 (cf. Eq. 5).

Fidelity is a common measure of similarity between quantum states which is often used as a cost function for determining optimal controls. For arbitrary density matrices ρ and σ, the fidelity [116] is defined as:

$$\mathcal{F}(\rho,\sigma) = \left(\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right)^2, \tag{8}$$

where Tr is the trace. A technique which will be used in Chs. 3, 4 and 7 is to evaluate the fidelity between a target state ρ_{des} and the final evolved state $\rho(t_f)$ as a cost function to assess the effectiveness of the applied controls $a_i(t)$. Quantum control for larger more complex systems is challenging due to the explosion in dimensionality of the Hilbert space, which makes numerically simulating larger systems very challenging. Moreover, real world experimental control is challenging, due to control parameter drifts, noise (at the control system and quantum system level) and other experimental imperfections. One promising approach, which will be discussed in Ch. 7, is the use of Reinforcement Learning to devise optimal control solutions by interacting with and sampling data from realistic simulations or real experiments.

2.1 ATOMIC PHYSICS

Alkali atoms are widely used in quantum information processing due to their simple electronic structure, consisting of a single valence electron. This allows precise control using laser fields for a variety of quantum information processing tasks. In this section, we briefly review key concepts from atomic physics, followed by an analysis of the relevant transition lines of $^{87}\mathbf{Rb}$ in the presence of external magnetic fields which is important for later experimental work. We also outline a standard method for coupling ground-state levels via a Λ-type configuration.

2.1.1 *Coupling of Atomic States with Hyperfine Structure*

To realise an atomic single-photon source, specific dipole transitions between atomic energy levels are utilised. These transitions and their properties are best understood by analysing the dipole coupling strengths between eigenstates of an atom in an external magnetic field.

The simplest atomic model is the hydrogen atom, which includes only the kinetic energy of the electron and the Coulomb potential. However, this idealised model fails to capture the behaviour of multi-electron atoms. To account for additional effects, perturbations are added to the hydrogenic Hamiltonian, including the central field and residual electrostatic interactions, which model electron shielding and spin effects. Spin-orbit coupling or fine structure, is also included to account for the interaction between the electron's spin and its motion in the attractive potential. The total electronic angular momentum $\vec{J} = \vec{L} + \vec{S}$ is a sum of the orbital and spin angular momentum. J thus lies in the range $|L - S| \leq J \leq L + S$.

The Hamiltonian most relevant to the work that follows is the hyperfine structure interaction which is a result of the coupling of the total angular momentum J with the

total nuclear angular momentum I giving rise to the total atomic angular momentum F which is the sum of I and J:

$$\vec{F} = \vec{I} + \vec{J}, \text{where } |J - I| \leq F \leq J + I \tag{9}$$

The Hamiltonian describing this interaction reads:

$$\hat{H}_{\text{hfs}} = A_{\text{hfs}}\mathbf{I} \cdot \mathbf{J} + B_{\text{hfs}} \frac{3(\mathbf{I} \cdot \mathbf{J})^2 + \frac{3}{2}\mathbf{I} \cdot \mathbf{J} - I(I+1)J(J+1)}{2I(2I-1)J(2J-1)}, \tag{10}$$

where A_{hfs} and B_{hfs} are the magnetic dipole and quadrupole constants respectively, the latter only applying to the D_2 transitions. Every hyperfine level F contains $2F + 1$ sub-levels with quantum numbers m_F which are degenerate in the absence of an external magnetic field. In the limit of a weak external magnetic field, such that $A_{\text{hfs}} >> \mu_B|B|$, an external field perturbs the energy levels with respect to the energy of a particular state $|F, m_F = 0\rangle$ such that [1]:

$$\Delta E(m_F) = \mu_B g_F m_F B_z \tag{11}$$

This is also called the *linear Zeeman regime* as the energy splitting is linear in B_z. For strong external magnetic fields this approximation starts to break down as the total angular and total nuclear angular momenta decouple and couple individually to the external magnetic field. The interaction is described by the following Hamiltonian:

$$\hat{H}_B = \frac{\mu_B}{\hbar}(g_s\vec{S} + g_I\vec{I} + g_L\vec{L}) \cdot \vec{B}. \tag{12}$$

Eigenstates in the regime of $A_{hfs} << \mu_B|B|$ are given by: $|I, J, m_I, m_J\rangle$. This is also called the "Paschen-Back regime". In general, it is desirable to operate either in the linear Zeeman or strong field regime such that there are well defined eigenstates which can be efficiently coupled.

1 The quantisation axis is picked arbitrarily to lie along $\vec{B}$, without loss of generality.

In the intermediate field regime, one has to numerically diagonalise the total Hamiltonian which contains both contributions from (12) and (10) to find the atomic eigenstates. These are linear combinations of various $\left|\vec{F}, m_F\right\rangle$ [2] with mixing coefficients $c_{FF_\dagger}(\left|\vec{B}\right|)$ [20]:

$$|\psi(F, m_F)\rangle = \sum_{F_\dagger} c_{FF_\dagger}\left(\left|\vec{B}\right|\right)|F_\dagger, m_{F_\dagger}\rangle\,, \tag{13}$$

where $m_F = m_{F_\dagger}$ due to the conservation of angular momentum [20]. The dependence of the mixing coefficients on the magnetic field make it impossible to write down a general solution.

The interaction strength between light and an atom is described by the dipole matrix, where the individual matrix elements describe the coupling between individual states. In particular we are interested in the coupling of two hyperfine states of the form $|F, m_F\rangle$ and $|F', m'_F\rangle$ which denote a ground and excited state respectively. One can use the Wigner-Eckart theorem to express individual matrix elements as products of the reduced matrix elements and the Clebsch-Gordan coefficient:

$$\left\langle Fm_F\left|er_q\right|F'm'_F\right\rangle = \left\langle F\|e\mathbf{r}\|F'\right\rangle\left\langle Fm_F \mid F'1m'_Fq\right\rangle. \tag{14}$$

Here q denotes the component of r in the spherical basis which also corresponds to the polarisation in the ordered basis $(\sigma_-, \pi, \sigma_+)$. This Clebsch-Gordan coefficient can be equivalently expressed in terms of a Wigner 3-j symbol:

$$\left\langle Fm_F\left|er_q\right|F'm'_F\right\rangle = \left\langle F\|e\mathbf{r}\|F'\right\rangle(-1)^{F'-1+m_F}\sqrt{2F+1}\begin{pmatrix} F' & 1 & F \\ m'_F & q & -m_F \end{pmatrix} \tag{15}$$

Notice that the 3-j symbol (or, equivalently, the Clebsch-Gordan coefficient) vanishes unless $m_F = m'_F + q$. The reduced matrix element can be further simplified by factoring out $\langle F\|e\mathbf{r}\|F'\rangle$ into a Wigner 6-j symbol. This gives rise to a final reduced matrix element

2 or equivalently of various $\left|\vec{I}, \vec{J}, m_I, m_J\right\rangle$

which is proportional to the dipole matrix element of a particular transition (D_1 or D_2 for $^{87}\mathbf{Rb}$) $d_{JJ'} = \langle J \| e\mathbf{r} \| J' \rangle$, which is well characterised, and a prefactor which captures the angular dependence of the transition:

$$\begin{aligned} \langle F \| e\mathbf{r} \| F' \rangle &\equiv \langle JIF \| e\mathbf{r} \| J'I'F' \rangle \\ &= \langle J \| e\mathbf{r} \| J' \rangle \, (-1)^{F'+J+1+I} \sqrt{(2F'+1)\,(2J+1)} \begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix}. \end{aligned} \tag{16}$$

Putting these two results (cf. Eq. 15 and Eq. 16) together one obtains the following expression for the matrix elements at zero external field:

$$\langle F m_F \left| e r_q \right| F' m'_F \rangle = d_{JJ'} \cdot A(F, F', m_F, m'_F, J, J'), \tag{17}$$

where $A(F, F', m_F, m'_F, J, J')$ is a combination of a Wigner 3−j, 6−j and various prefactors. Substituting the field dependent eigenstates of the Hamiltonian together allows for an expression of the field dependent coupling strength of individual states :

$$\langle F m_F \left| e r_q \right| F' m'_F \rangle_{(B_{ext}=\vec{B})} / d_{JJ'} = \sum_{F_\dagger, F'_\dagger} c_{F'F'_\dagger} \left(\left| \vec{B} \right| \right) A(F, F', m_F, m'_F, J, J') c_{FF_\dagger} \left(\left| \vec{B} \right| \right) \tag{18}$$

The normalised element $\langle F m_F \left| e r_q \right| F' m'_F \rangle_{B_{ext}=\vec{B}} / d_{JJ'}$ (i.e. no JJ′ dependence) is also often abbreviated as d_{ij} where indices i, j each correspond to a pair of F, m_F quantum numbers.

2.1.2 *Rubidium-87*

Rubidium-87 ($^{87}\mathbf{Rb}$) is one of the most widely used atoms in quantum optics due to its simple electronic structure, well-characterised energy levels, and ease of experimental control with off-the-shelf hardware. Its ground-state hyperfine splitting between the $F = 1$ and $F = 2$ levels of the $5S_{1/2}$ manifold makes it ideal for implementing

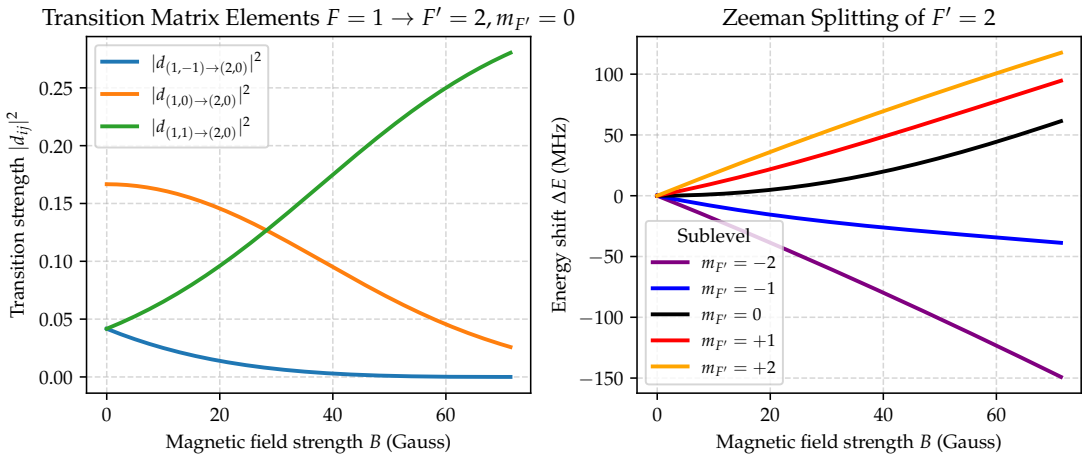

Figure 1: Using the D_2 line parameters, shown is the the variation of absolute dipole matrix elements $|d_{i,j}|^2$ (cf. Eq. 18) where i, j are indices corresponding to particular eigenstates of F and m_F quantum numbers, with respect to an external bias field (left plot). Using the mixing coefficients described in Eq. 13 and the hyperfine Hamiltonian (cf. Eq. 10) also allows for an expression of the variation in energy with magnetic field (right plot). Moderate fields of over 10G already induce strong non-linearities in energies and dipole coupling strengths.

effective two-level or Λ-type three-level systems. The D_1 (which couples to the $J' = \frac{1}{2}$ excited level) and D_2 (which couples to the $J' = \frac{3}{2}$ level) transitions near 795nm and 780nm are easily addressed with diode lasers, and well-established cooling and trapping techniques enable state preparation and coherent manipulation. These properties, along with access to highly excited Rydberg states, make ^{87}**Rb** a cornerstone of many quantum optics and quantum information experiments and it will be used for the simulation and experimental work in this thesis. The full structure of the D_2 line is given in App. A.1 for reference.

Elaborating on the derivations discussed in the previous section, Fig. 1 shows the variation of dipole matrix elements $d_{i,j}$ of the transitions between $F = 1$ and $F' = 2, m_F = 0$, as well as the Zeeman splitting of the $F' = 2$ excited level. This is a transition of interest as it is discussed for single photon production in Ch 3. We observe that moderate fields of 10G and above break the dipole matrix element symmetry in $B = 0$ field case due to the onset of state mixing. This precluded the efficient generation of chains of photons of alternating polarisations [20] as different transitions corresponding to different polarisations were driven with different effective couplings.

Fig. 1 shows there exists a narrow window for which an atom-cavity single photon source can be operated without parasitic background field noise which induces a time evolving quantisation axis, or the onset of non-linear Zeeman effects [20], at approximately 1G and below. This is particularly relevant for the schemes discussed in Chapter 3 and 4.

2.1.3 *Stimulated Raman Adiabatic Passage*

It is instructive to review a foundational experimental technique for establishing a direct coupling between two atomic ground states $|u\rangle$ and $|g\rangle$ with two fields via a shared coupling to an excited atomic state $|e\rangle$. First introduced in [132], the original Stimulated Raman Adiabatic Passage proposal enables robust and complete population transfer in a three-level Λ system—from the initial state $|u\rangle$ to the target state $|g\rangle$, without significant occupation of the intermediate excited state $|e\rangle$. The process is mediated by two coherent laser fields: the Pump field, which couples $|u\rangle \leftrightarrow |e\rangle$, and the Stokes field, which couples $|g\rangle \leftrightarrow |e\rangle$ [238].

The dynamics of this system is described with the time-dependent Schrödinger equation (cf. Eq. 3). The state-vector $\Psi(t)$ is a three-component column-vector spanned by $\{|u\rangle, |e\rangle, |g\rangle\}$. In the RWA, the Hamiltonian is expressed as [238]:

$$\hat{H}(t) = \frac{\hbar}{2}\begin{pmatrix} 0 & \Omega_P(t) & 0 \\ \Omega_P(t) & 2\Delta & \Omega_S(t) \\ 0 & \Omega_S(t) & 2\delta \end{pmatrix}, \tag{19}$$

where Δ describes the single photon detuning $\omega_{ue} - \omega_P$ from resonance and δ describes the two photon detuning $\delta = \omega_{eg} - \omega_S - \Delta$, which is the difference of the Stokes laser from the single photon resonance Δ. For atomic excitations, the Rabi frequencies $\Omega_P(t)$ and $\Omega_S(t)$ are calculated from the dipole energy $-\mathbf{d} \cdot \mathbf{E}(t)$. The time-dependent Rabi

frequencies $\Omega_P(t)$ and $\Omega_S(t)$ [3] arise from the dipole interaction $-\mathbf{d}\cdot\mathbf{E}(t)$ and are given by [238]:

$$\Omega_P(t) = -\frac{d_{u,e}\mathcal{E}_P(t)}{\hbar}, \qquad \Omega_S(t) = -\frac{d_{e,g}\mathcal{E}_S(t)}{\hbar}, \tag{20}$$

where $d_{u,e}$ and $d_{e,g}$ [4] are the dipole matrix elements for their respective transitions, and $\mathcal{E}_P(t)$ and $\mathcal{E}_S(t)$ are the time evolving electric field envelopes in the rotating wave approximation (RWA). In a multilevel atom like ^{87}**Rb**, excited states can undergo spontaneous emission to lower-lying states. Emission can occur to a state outside the three state manifold described above which manifests as a loss of probability, or as an incoherent loss to either of the ground states $|u\rangle$ or $|g\rangle$ which is also undesirable. This can be addressed phenomenologically by adding an imaginary component on the diagonal of the Hamiltonian described in Eq. 19, or by evolving the system with the master equation described in Eq. 5.

One of the eigenstates of the Hamiltonian (setting $\delta = 0$) described in Eq. 19 has zero eigenenergy and reads [238]:

$$|\Psi_D\rangle = \cos(\theta)\,|u\rangle - \sin(\theta)\,|g\rangle\,, \tag{21}$$

where the mixing angle $\theta(t)$ is given by $\tan(\theta(t)) = \Omega_P(t)/\Omega_S(t)$ [238]. This state contains no excited state contribution and is therefore resilient against loss of population from the Λ system through spontaneous emission from state $|e\rangle$. To ensure that the system initially follows the dark state Ψ_D, the mixing angle must satisfy $\theta(t) = 0$, which implies $\Omega_P/\Omega_S \to 0$. This condition ensures that the dark state Ψ_D coincides with the initial state $\Psi(t) = |u\rangle$. The population transfer is completed when Ψ_D coincides with the target state $|g\rangle$, which corresponds to a mixing angle of $\theta(t) = \pi$ and

3 It should be noted that these are defined here as angular frequencies.
4 For ^{87}**Rb** these dipole matrix elements consist of products of a reduced dipole matrix element and a Clebsch-Gordan coefficient (cf. Eq. 14.)

requires $\Omega_S/\Omega_P \to 0$. These boundary conditions define the counter-intuitive pulse sequence [132, 238], in which the Stokes pulse $\Omega_S(t)$ precedes the Pump pulse $\Omega_P(t)$. As a result, the initially populated state $|u\rangle$ remains uncoupled at early times, enabling adiabatic following of the dark state throughout the evolution. A suitable temporal overlap between the pulses is essential to ensure adiabaticity. More precisely, the global adiabaticity condition [28], $\Omega_{\text{rms}}T \gg 1$, where $\Omega_{\text{rms}} = \sqrt{\hat{\Omega}_S^2 + \hat{\Omega}_P^2}$, must be satisfied, where T denotes the total pulse time and $\hat{\Omega}$ describes the time average amplitudes for Ω_S and Ω_P. In the original STIRAP experiments, this was achieved using temporally interleaved Gaussian pulses [84].

2.2 CAVITY QED

2.2.1 *Optical Cavities*

A widespread method for single photon production is the controlled interfacing of light and matter with an optical cavity. Photons can be emitted into the cavity mode allowing for precise control over the photon that is produced. In what follows, the fundamental principles of Cavity Quantum Electrodynamics will be explained.

An optical cavity denotes a structure formed by two mirrors whose optical surfaces are oriented towards each other. The interaction of light resonating between the mirror surfaces leads to the formation of a particular interference pattern. The simplest and most well-known cavity is a Fabry-Perot cavity, comprising of two identical, plane parallel mirrors.

At certain frequencies, the interaction of light with the cavity results in a standing wave. At these resonant frequencies, the power of the light inside the cavity is enhanced by a factor $\mathcal{F}$, known as the *finesse* of the cavity. To model this behaviour, each mirror is characterised by its electric field reflection, transmission, and loss coefficients, denoted

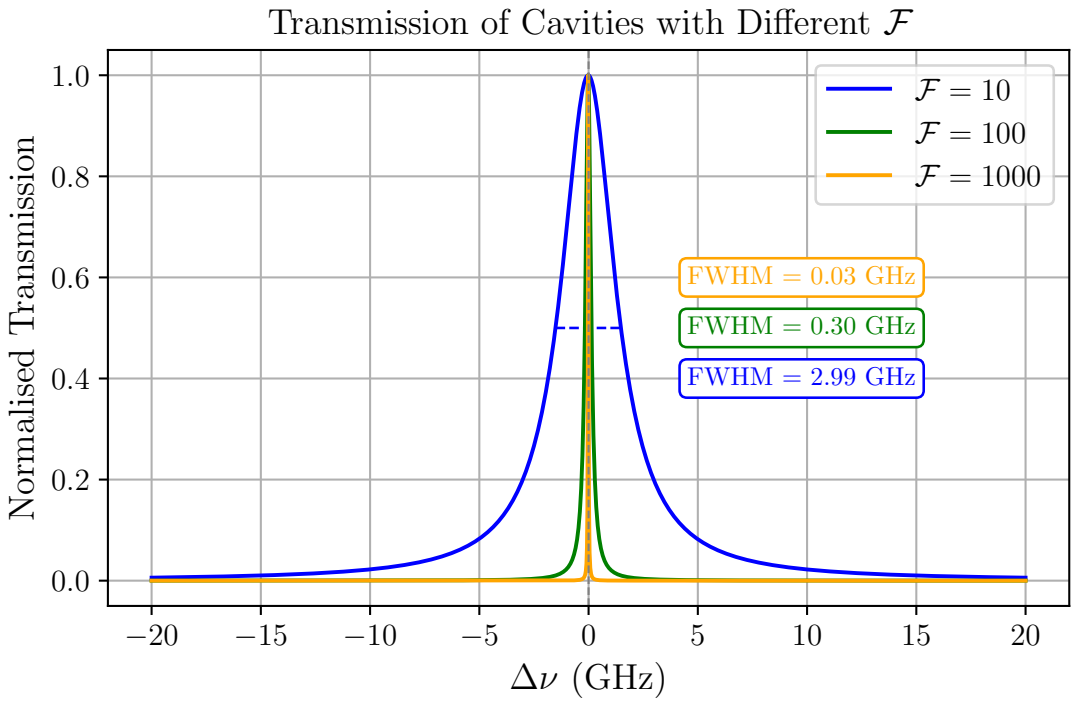

Figure 2: Comparing the spectral transmission pattern of three cavities with different Finesses at a fixed length of 10^{-3}m. The line-widths (FWHM) are indicated.

r_i, t_i, and l_i, respectively, where the index i refers to each mirror. These coefficients satisfy the relation $|r_i|^2 + |t_i|^2 + |l_i|^2 = 1$. One can calculate the effect of the cavity of length L_{cav} on the incident field E_i, to give a transmitted field E_t:

$$E_t = \frac{t}{1 - re^{i\phi}} E_i \tag{22}$$

where $r = \sqrt{r_1 r_2}, t = \sqrt{t_1 t_2}$ and $\phi = 2L_{\text{cav}} \cdot \frac{2\pi}{\lambda}$ is the phase acquired by the light on a single round trip. The transmitted intensity profile of a cavity reads $I_t \propto |E_t|^2$ and this leads to a characteristic profile of spectrally distributed transmission peaks as shown in Fig. 2. When moving to intensities, the intensity reflection and transmission coefficients are described as $R_i = |r_i|^2$ and $T_i = |t_i|^2$.

The properties of a cavity are characterised by its line-width[5] $\Delta\nu_{\text{FWHM}}$ and the free spectral range (FSR) $\Delta\nu_{\text{FSR}}$. These determine the widths of the spectral transmission peaks and the spacing between them. The cavity FSR is calculated by considering the

5 This corresponds to the full width at half maximum of the intensity of the cavity transmission.

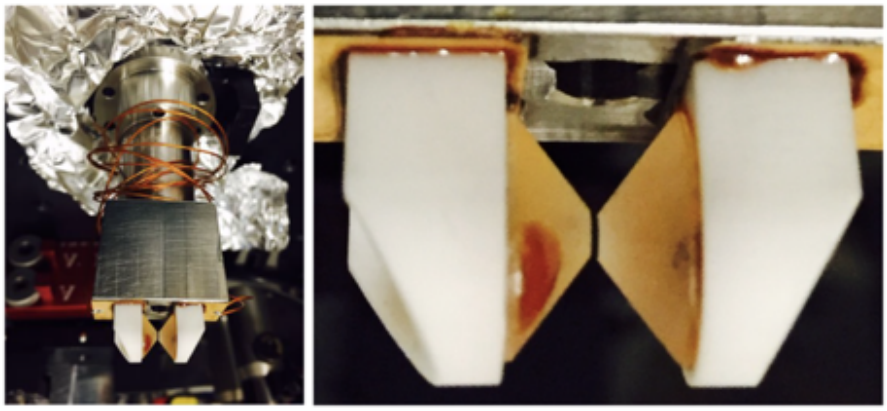

Figure 3: Cavity used in the current experimental setup with $L_{cav} = 339\mu m$[18]

constructive interference of light passing through the cavity, i.e. $\phi = 2n\pi$, where $n \in \mathcal{N}$, such that $\Delta\nu_{\text{FSR}} = \frac{c}{2L_{\text{cav}}}$. The cavity Finesse is another important metric for cavity characterisation and it expresses the ratio of distance between transmission peaks and their widths. In the regime of identical mirrors and high reflectivities, the Finesse can be related to the reflectivity of the cavity mirrors and is approximated [6] as:

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1-R} = \frac{\Delta\nu_{\text{FSR}}}{\Delta\nu_{\text{FWHM}}}. \tag{23}$$

The cavity being used for the current experimental setup has a Finesse of $\mathcal{F} = 117800 \pm 200$. [18] and an image is shown in Figure 3.

In general, our cavities are fabricated asymmetrically, such that one mirror is more transmissive and the emission of a photon only occurs in that preferential direction. The cavity decay rate:

$$2\kappa = 2\pi\Delta\nu_{\text{FWHM}}, \tag{24}$$

is the characteristic time scale at which light decays from the cavity, where κ is the decay rate of the electric field amplitude.

6 When $R \approx 1$, one can use the small angle approximation $sin(\phi) \approx \phi$ in the normalised transmitted intensity distribution and arrives as the following expression.

2.2.2 *The Jaynes-Cummings model*

The behaviour of a single quantum emitter coupled to a cavity is described by the Jaynes Cummings Hamiltonian. The dynamics are determined by three key quantities: the spontaneous decay rate, γ[7], the cavity field decay rate, κ, and the coupling between emitter and the cavity, g.

The Jaynes-Cummings model [80] is described by a Hamiltonian with three terms, namely a two level atom with one ground and one excited state $|g\rangle, |e\rangle$:

$$\hat{H}_{JC} = \hat{H}_{\text{cav}} + \hat{H}_{\text{atom}} + \hat{H}_{\text{int}}, \tag{25}$$

$\hat{H}_{\text{cav}} = \hbar\omega_C \hat{a}^\dagger \hat{a}$ is the cavity field, where $\hat{a}^\dagger, \hat{a}$ denote creation and annihilation operators, respectively. [8]

$\hat{H}_{\text{atom}} = \hbar\omega_g \hat{\sigma}\hat{\sigma}^\dagger + \hbar\omega_e \hat{\sigma}^\dagger\hat{\sigma}$ is associated with the atomic energy. Here $\hat{\sigma} = |g\rangle\langle e|$ is the lowering operator and $\hat{\sigma}^\dagger = |e\rangle\langle g|$ corresponds to the atomic raising operator.

$\hat{H}_{\text{int}} = -\hbar g_0 \left(\hat{a}^\dagger\hat{\sigma} + \hat{a}\hat{\sigma}^\dagger\right)$ is the interaction Hamiltonian, which describes the interaction between the atom and the external electric field, mediated by the cavity.

The cavity is tuned close to the atomic resonance ($\omega_C \approx \omega_e - \omega_g$) and thus, by invoking the rotating wave approximation, one can neglect the terms rotating very rapidly at twice the optical frequency ($\hat{a}^\dagger\hat{\sigma}^\dagger$ and $\hat{a}\hat{\sigma}$), as they average to 0. This only leaves the terms corresponding to an atomic excitation and photon destruction, or photonic creation and atomic de-excitation. The atom-cavity coupling is given by $g_0 = \frac{\vec{\mu}\cdot\vec{E}}{\hbar}$. The electric field strength for the vacuum (i.e. an empty cavity) is given by $\left|\vec{E}\right| = \sqrt{\frac{\hbar\omega_C}{2\epsilon_0 V}}$ where $V = \frac{\pi w_0^2 L_{\text{cav}}}{2}$ is the mode volume of the cavity. There is likely going to be some

7 It is important to note that I have adopted the convention $\gamma = \frac{1}{2\tau}$ such that γ describes the spontaneous decay of the electric field and not the intensity. The same convention is adopted for κ.

8 Strictly speaking n photons have an eigen-energy of $\hbar\omega_C(n + \frac{1}{2})$ where ω_C denotes the resonance frequency of the cavity mode with which the atom interacts, but one drops the zero point energy because it does not alter the dynamics.

variation in the coupling of the atom to the cavity as it moves so g_0 expresses the maximum coupling, whereas in reality the atom may experience a reduced coupling.

The eigenstates of this Hamiltonian are expressed as tensor products of the state of the atom and the photons in the cavity mode, such that $|e,0\rangle$ and $|g,1\rangle$ express the excited state with no photons in the cavity and the ground state with one photon in the cavity respectively. They are also known as dressed states.

The dynamics of the system are also determined by the non-unitary processes of spontaneous emission and photonic field decay. A photon emitted spontaneously by the emitter inside the cavity will be radiated in an arbitrary direction into free space. This is undesirable for quasi-deterministic single photon production. It is desirable for the atom-cavity interaction mediated by g to dominate compared to the incoherent decay of photons from either the atom or the cavity mode. The interplay between atom-cavity coupling, spontaneous decay and field decay is captured by the cooperativity [80]:

$$C = \frac{g^2}{2\kappa\gamma} \tag{26}$$

There are various coupling regimes, dependent on the relative strength of the three parameters g, γ and κ. A desirable regime for efficient single-photon production is known as the strong coupling regime, characterised by $C > 1$, such that the cavity-emitter interaction dominates and $g > \kappa, \gamma$. This usually requires a small mode-volume as $g \propto \sqrt{\frac{1}{V}}$. The strong coupling regime ensures that the dominant time scale is determined by the coherent process of coupling the atoms to the cavity, i.e. g is faster than that which determines the spontaneous decay of the atom γ or the photonic field decay κ.

An important effect relevant to single-photon generation is the *Purcell effect*, which describes the enhancement of spontaneous emission due to a modified photonic en-

vironment, such as a resonant cavity [181]. The Purcell factor quantifies this enhancement:

$$F_P = \frac{3}{4\pi^2}\left(\frac{\lambda}{n}\right)^3 \cdot \frac{Q}{V}, \tag{27}$$

where λ is the transition wavelength, n the refractive index, $Q = \omega/\kappa$ the cavity quality factor, and V the mode volume. In cQED, this enhancement relates directly to the cooperativity $C = \frac{g^2}{2\kappa\gamma}$:

$$F_P \approx \frac{g^2}{\kappa\gamma} = 2C. \tag{28}$$

The fraction of emission coupled into the cavity mode reads $\frac{F_P}{F_P+1}$ which is equivalent to the upper bound in cavity emission efficiency that can be derived from the time dynamics of the dissipative Schrödinger equation described in the next Section.

2.2.3 *V-STIRAP*

To implement single-photon generation, we consider a three-level atomic system in the Λ-configuration. This consists of two long-lived ground states, $|u\rangle$ and $|g\rangle$, which in the context of this thesis correspond to two hyperfine levels of Rubidium, and an excited state $|e\rangle$. The goal is to transfer the atom from state $|u\rangle$ to $|g\rangle$, a transition that is dipole-forbidden and cannot be addressed directly with optical fields. Instead, this population transfer is achieved adiabatically via the intermediate state $|e\rangle$, using a technique known as Vacuum-induced Stimulated Raman Adiabatic Passage (V-STIRAP). The setup is very similar to STIRAP except that the radiation field coupling the $|e\rangle$ to $|g\rangle$ transition is replaced with atom-cavity coupling (cf. Eq. 19). This particular photon production process is appealing insofar as one can shape the temporal wave-function of the produced photons with an analytically computable driving shape $\Omega(t)$ as described in [234].

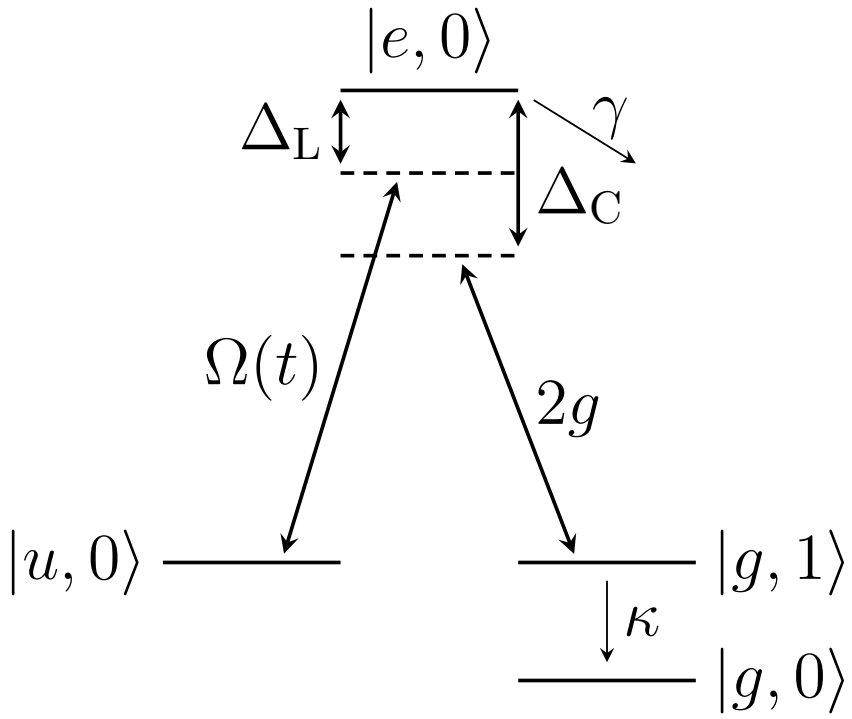

Figure 4: Level Structure for Jaynes-Cummings Hamiltonian. [18].

An external laser with Rabi frequency Ω drives the $|u\rangle \leftrightarrow |e\rangle$ transition, detuned from resonance by Δ_L. Simultaneously, a cavity mode couples the $|e\rangle \leftrightarrow |g\rangle$ transition with detuning Δ_C. This configuration is illustrated in Figure 4.

The Hamiltonian for this system is an extension of the two-level Jaynes Cummings Hamiltonian. The basis states $\{|u,0\rangle, |e,0\rangle, |g,1\rangle\}$ with corresponding probability amplitudes $\mathbf{c}(t) = [c_u(t), c_e(t), c_g(t)]^T$ evolve according to the Schrödinger equation [234]:

$$i\hbar \frac{d}{dt}\mathbf{c}(t) = -\frac{\hbar}{2}\begin{pmatrix} 0 & \Omega(t) & 0 \\ \Omega(t) & 2i\gamma & 2g \\ 0 & 2g & 2i\kappa \end{pmatrix}\mathbf{c}(t), \tag{29}$$

where one neglects higher photon number states and invokes the rotating wave approximation. The non-unitary dynamics are incorporated phenomenologically with imaginary diagonal elements, corresponding to a loss of population from the system. One can show that for a Raman excitation which satisfies the two-photon resonance

condition [9], i.e. $\Delta_L = \Delta_C = 0$, this system has a dark eigenstate with no contribution from the excited state $|e\rangle$:

$$|\phi_D\rangle = \cos(\Theta)\,|u,0\rangle - \sin(\Theta)\,|g,1\rangle\,, \tag{30}$$

where $\tan(\Theta) = \frac{\Omega}{2g_0}$. Thus, slowly modulating the Rabi frequency Ω of the external driving laser allows for an adiabatic transfer of population from $|u,0\rangle$ to $|g,1\rangle$, by avoiding a population of the excited state, which would lead to undesirable spontaneous emission of a photon into free space as parametrised by γ.

This adiabatic process ensures that the collection of photons in the cavity exceeds free-space emission and photons accrued in $|g,1\rangle$ leak out of the cavity due to its decay as given by κ. An upper bound on the efficiency of photon generation with V-STIRAP is given by [234]:

$$\eta_{\max} = \frac{2C}{2C+1}, \tag{31}$$

where C denotes the cooperativity defined in (26) and it is evident that stronger atom-cavity coupling enhances photon production efficiency, all else staying constant.

2.2.4 *Cavity Polarisation Modes*

The polarisation modes supported by a cavity, depend on the relative orientation of the cavity and the magnetic field direction which coincides with the preferred quantisation axis [10]. In practice, photons produced by a V-STIRAP process as described in Sec. 2.2.3 have polarisations determined by the particular atomic transitions which are close to resonance with the cavity.

9 Solving the system for other off-resonant excitations becomes rather involved [229] and no simple, general and closed-form solution exists.

10 The choice of quantisation axis is arbitrary, but it is preferable to pick it along the external magnetic field direction to simplify the calculation of the atomic eigenstates.

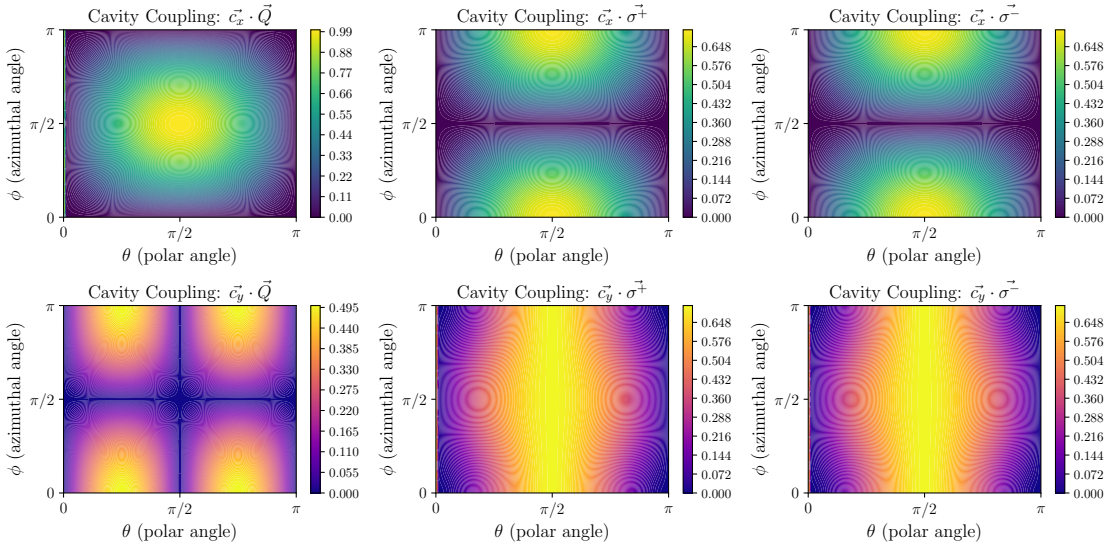

Figure 5: Cavity coupling components in x, y directions in the top and bottom row respectively, where the columns describe the components in the atomic basis $(\pi, \sigma^+, \sigma^-)$ for arbitrary relative orientation of $\vec{Q}$ and $\vec{C}$ which is parametrised in spherical coordinates θ, ϕ.

Consider a cavity axis vector $\vec{C}$ and magnetic field (i.e. quantisation axis) axis vector $\vec{Q}$. The cavity polarisation plane is spanned by two vectors $(\vec{c_x}, \vec{c_y})$ perpendicular to the cavity axis. In general, the basis can be an arbitrary orthogonal basis set, a particularly simple choice being two linear polarisation directions which form a right-handed orthogonal basis with the cavity axis vector. Photons must be polarised in this plane since they are out-coupled through one of the cavity mirrors. A photon polarised in the direction of the cavity axis cannot propagate in this same direction, as the oscillation of its electric field must be orthogonal to its wave-vector which points along the direction of propagation.

The mapping from the cavity polarisation vectors $(\vec{c}_x, \vec{c}_y)$ into the atomic polarisation basis $(\vec{Q}, \vec{\sigma}^+, \vec{\sigma}^-)$ is obtained by projecting each cavity vector onto the atomic basis. The resulting coefficient vectors:

$$c_x^{atom} = \begin{pmatrix} \vec{c_x} \cdot \vec{Q} \\ \vec{c_x} \cdot \vec{\sigma^+} \\ \vec{c_x} \cdot \vec{\sigma^-} \end{pmatrix}, c_y^{atom} = \begin{pmatrix} \vec{c_y} \cdot \vec{Q} \\ \vec{c_y} \cdot \vec{\sigma^+} \\ \vec{c_y} \cdot \vec{\sigma^-} \end{pmatrix}, \tag{32}$$

give the representation of the cavity polarisations in the atomic basis.

Assigning $\vec{Q} = (1,0,0)$ without loss of generality, the cavity axis vector can be expressed in spherical coordinates using an azimuthal angle ϕ and a polar angle θ:

$$\vec{C} = \begin{pmatrix} \cos\theta \\ \sin\theta\cos\phi \\ \sin\theta\sin\phi \end{pmatrix}, \tag{33}$$

θ (polar angle) represents the angle between $\vec{C}$ and the x-axis and ϕ (azimuthal angle) represents the angle between the projection of $\vec{C}$ onto the yz-plane and the y-axis. Fig. 5 show the variation of coupling strengths of the cavity polarisation vectors $\vec{c_x}$ and $\vec{c_y}$ expressed in the atomic polarisation basis given in Eq. (32). The orientation of the quantisation axis (i.e., the magnetic field direction) plays an important role in atom–cavity coupling and should be optimised for a given experimental setup and atomic structure. In the following chapter, we examine the generation of bursts of linearly polarised photons, which requires coupling the atom to a single cavity polarisation mode (scenario I) [74]. A more general scenario (scenario II) involves coupling to two orthogonal polarisation modes, relevant for generating polarisation-encoded entangled states [223].

For atoms with fine or hyperfine structure, there are seven distinct coupling configurations: three involving a single atomic transition driven by a specific polarisation (π, σ^+, or σ^-), three involving two such transitions, and one involving all three. When only a single dipole-allowed transition is used, the quantisation axis can be chosen to maximise coupling into the desired cavity mode. For instance, to drive a σ^+ transition, the magnetic field should be aligned parallel to the cavity axis; for a π transition, it should be perpendicular.

The situation becomes more nuanced when multiple atomic transitions corresponding to orthogonal polarisations (two or all three of π, σ^+, σ^-) are allowed. These transitions generally have different dipole coupling strengths (see Eq. 18 for the hyperfine case). It is important to distinguish the optimal quantisation axis orientation depend-

ing on whether the experimental goal is scenario I (emission of photons into a single polarisation mode) or scenario II (emission into two modes).

In scenario I, the magnetic field should be aligned to maximise coupling of the strongest atomic transition(s) into one cavity polarisation mode (e.g., $\vec{c_x}$), while suppressing undesired emission into the orthogonal mode (e.g., $\vec{c_y}$). This suppression can be achieved by ensuring that the weakest dipole coupled transition aligns with the unwanted mode or by eliminating its contribution via the basis transformation in Eq. 32. For example, if the π transition is weak, aligning the quantisation axis with the cavity axis couples only the strongest dipole coupled mode. Conversely, if the π transition is strong and the $\sigma^{\pm}$ transitions are weak, a perpendicular magnetic field is preferable to achieve stronger mode coupling.

A particularly useful case arises when the $\sigma^{\pm}$ transitions are strongly dipole coupled. Here, choosing the quantisation axis perpendicular to the cavity axis ensures that the cavity supports a single linear polarisation mode that couples to a specific superposition of the $\sigma^{\pm}$ components. This configuration is used in the photon emission scheme discussed in the following chapter.

For schemes of type II, which require emission into two orthogonal polarisation modes, the optimal configuration involves maximising coupling from all available atomic transitions into the two cavity modes. For example, if the $\sigma^{\pm}$ transitions are strongly coupled, aligning the quantisation axis with the cavity axis ensures that each transition couples efficiently to one of the orthogonal polarisation modes. If all three transitions (π, σ^{+}, σ^{-}) are strongly coupled, it is preferable to align the quantisation axis perpendicular to the cavity axis. This enables emission of a π-polarised photon and a linearly polarised photon orthogonal to it.

The ideal quantisation axis depends on both the dipole matrix elements and the energy level structure of the atom. However, in practice, orienting the magnetic field along the cavity axis or one of the two orthogonal directions perpendicular to it generally yields the highest photon production efficiency if one is after a particular polari-

sation. Some interesting enhancements of coupling when the cavity modes are split in energy due to cavity mirror birefringence are described in [19].

Although it is possible to tilt the magnetic field away from the principal axes in the lab frame, this introduces mixing of different cavity modes being coupled by multiple atomic transitions. Specifically, it causes a single atomic transition to couple into both cavity modes, breaking the unique correspondence between photon polarisation and coupled atomic transition. Such mixing can critically undermine the fidelity of atom-photon states [223].

2.2.5 *Historical Overview of atomic cQED Experiments for Single Photon Production*

The generation of single photons using atom–cavity systems has been a central pursuit in quantum optics. As discussed in detail in Sec. 2.1.3, STIRAP laid the groundwork for much of the following hallmark experiments, initially demonstrated by Gaubatz et al. [84]. Originally developed in physical chemistry for efficient population transfer in molecular systems, it was later adapted for coherent control in atomic systems.

In 2000, Hennrich et al. [101] demonstrated vacuum-stimulated Raman scattering in a high-finesse optical cavity using a single ^{85}**Rb** atom. This experiment showcased the potential of combining STIRAP with cavity enhancement to achieve controlled photon emission which marked a significant milestone.

Building upon this, Kuhn et al. [131] in 2002 realised a deterministic single-photon source employing a V-STIRAP scheme, where the atomic population was coherently transferred between two hyperfine ground states, emitting a photon into the cavity mode. This setup introduced an incoherent repumping mechanism to reset the atomic state after each emission cycle.

The coherence properties of such single-photon sources were further explored by Legero et al. [136], who observed quantum interference of photons emitted from the atom–cavity system. Their time-resolved measurements revealed quantum beats, un-

derscoring the high degree of indistinguishability achievable with cavity-enhanced emission.

In parallel, Kimble and colleagues advanced the field by demonstrating a deterministic single-photon source using a single Caesium atom trapped in an optical cavity [152].

Further advancements included polarisation control of emitted photons. Wilk et al. [242, 244] in 2007 achieved polarisation-controlled single-photon emission by driving Raman transitions between magnetic sublevels of an atom in a cavity, which formed the conceptual basis for many photon production schemes that followed, including that described in Ch. 3. This experiment had the external magnetic field direction, i.e. the quantisation axes, lying along the cavity axis which defined a convention followed by many future experiments.

The manipulation of photon wavepackets was demonstrated by Nisbet-Jones et al. [164], who shaped the temporal profile of single photons emitted from an atom–cavity system. Dynamic control over quantum states was first achieved by Barter et al. [23], who implemented a feedback mechanism to influence the collapse of the wave-function during photon emission [7]. This real-time control opens avenues for adaptive quantum measurements and error correction.

Ritter et al. [191] pursued important work towards the realisation of the quantum internet, by creating an elementary quantum network of single atoms in optical cavities, achieving entanglement between distant nodes. The research group of Gerhard Rempe has since been at the forefront of this, demonstrating efficient quantum repeaters for networking [134] and a quantum register of multiple atoms in a single cavity [97].

cQED platforms have also been used for elementary demonstrations of quantum computing. Holleczek et al. [105] utilised atom-cavity photons interfaced with photonic chips to perform on-chip quantum logic operations, such as a controlled-NOT gate, highlighting the potential for integrated quantum photonics and Thomas et. al. demonstrated efficient generation of multi-photon graph states for measurement based quantum computing using a cQED platform in 2022 [223]. This work was further ex-

tended by using measurements to *fuse* together smaller entangled photon states to produce a larger one [224].

Beyond neutral atoms, ion-based cQED systems have also been developed extensively. Keller et al. [120] demonstrated continuous single-photon generation from a single ion in a cavity, achieving high purity and control. Similarly, Blatt and colleagues [21] realised a deterministic photon source from a single trapped ion, emphasising the versatility of ion–cavity platforms. The charge of the ion poses significant challenges for cavity integration. Thus, much recent work with single photons emitted from ions has been done in free space with large collection lenses, such as in the group of David Lucas and colleagues [214, 217], who have demonstrated many important building blocks for networked quantum computation.

These developments collectively underscore the progress in atom–cavity QED systems, transitioning from fundamental research platforms to more practical components for quantum technologies.

3

SINGLE PHOTON PRODUCTION SCHEME

This chapter is concerned with the development of a novel scheme for efficiently producing bursts of polarised photons from a single atomic emitter placed in an optical resonator and closely follows the description in [74]. Photons are key to unlocking a number of quantum technologies such as measurement-based [42] quantum computing, as well as boson sampling [2]. Moreover, photons are important for the development of quantum networks [191], where information from stationary quantum bits is coherently mapped to a photonic qubit for communication and entanglement distribution. Single photons are also ubiquitous in quantum metrology [172]. What makes photons appealing is the ease with which they can be manipulated and transmitted with standard optical components, as well as their robustness to decoherence.

The realisation of the aforementioned applications requires highly efficient single-photon sources. Spontaneous parametric down conversion sources have been used widely and successfully in this regard [250], yet their probabilistic nature makes generating larger photon sequences and larger entangled states [246], required for scalable quantum computing architectures, challenging. To overcome the scalability challenges associated with an inherently probabilistic source, deterministic schemes for single photon generation have been proposed and shown promise, such as quantum dots [111], ions [120, 153] or neutral atoms [131].

Many deterministic single photon production schemes exploit the coupling of a quantum emitter to an optical resonator for the Purcell enhanced emission of photons into a well defined mode of the radiation field. In particular, we consider a process known as Vacuum enhanced STIRAP [131] which is described extensively in Sec. 2.2.3. Photons produced coherently with Raman pulses and optical cavities can have very

long coherence times, on the order of a few hundreds of nanoseconds [164]. The advantages of this technique are the precise control over all photonic degrees of freedom, including polarisation, frequency and wave-packet shape [234]. Cavity-QED platforms have been used successfully for quantum information processing, such as the demonstration of an efficient quantum repeater [134] or the efficient generation of large entangled cluster states [223]. Moreover, they can also be used to explore foundational questions in quantum optics such as the nature of the photon as a fundamental excitation [7].

Nevertheless, atom-cavity architectures are subject to various limitations. Often, the rate at which single photons are produced coherently from a deterministic atomic source is far lower than that of many non-deterministic sources, making it challenging to implement quantum information processing protocols that require large photon numbers or reasonable clock rates. Moreover, polarisation control is hard to achieve due to a manifold of degenerate magnetic sub-states present in most atoms. To overcome the latter, some schemes for producing polarised photons rely on the application of external magnetic fields to individually address a single Zeeman-state [242, 244]. However, finite cavity line-widths require a relatively strong magnetic field to lift the degeneracy of states, which leads to the apparition of non-linear Zeeman effects [20], strongly curtailing the achievable efficiencies.

In this chapter, a novel scheme for the efficient production of bursts of polarised photons is proposed, which does not require the application of a strong external magnetic field and is extensively explored in numerical simulations [1] of a realistic atomic cavity-QED system. First, in Sec. 3.1 the implementation of this scheme in an idealised atomic system is outlined, which encompasses the production of polarised photons, as well as the coherent re-preparation of the atom into the desired initial state. Thereafter, in Sec. 3.2 this scheme is applied to ^{87}**Rb** which exhibits two particular transitions that have the required energy level structure for photon production. Various trade-offs and limitations with respect to a variety of laser and cavity parameters are discussed.

1 All relevant code is available at https://github.com/jan-o-e/rb_photon_prod.

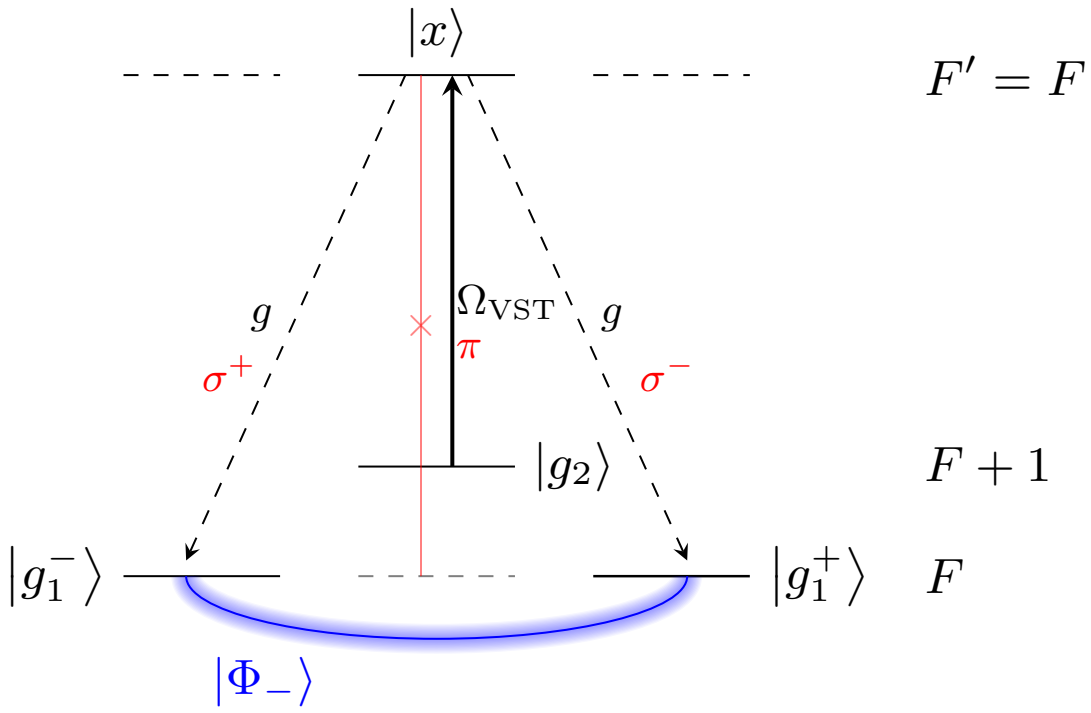

Figure 6: Idealised system with three atomic energy levels: two stable ground levels with total angular momenta $F, F+1$ and one excited level with total angular momentum $F' = F$. The individual states $|g_1^-\rangle, |g_1^0\rangle, |g_1^+\rangle$ are degenerate. The excited state $|x\rangle$ has dipole allowed transitions to $|g_1^\pm\rangle$ and $|g_2\rangle$, and the π transition to $|g_1^0\rangle$ is dipole forbidden. The same configuration can of course arise for atoms with no hyperfine structure and $F = J$. This configuration of levels is required for the polarised photon production scheme.

Additionally, the efficiency limitations are analysed and an optimal sequence which uses a combination of incoherent and coherent state preparation methods is presented. Finally, the results are contextualised, the requirements for demonstrating quantum supremacy [253] are analysed and future directions are illustrated.

3.1 IMPLEMENTATION IN IDEALISED THREE LEVEL SYSTEM

3.1.1 *Atomic Structure*

The proposal for polarisation controlled single photon production is first described in an idealised three level atomic system illustrated in Fig.6.

The atomic structure required for the production of polarised photons comprises two ground levels with total angular momenta F and $F+1$, as well as one excited level with total angular momentum $F' = F$ [2]. Three of the individual sub-states of the two ground levels given by F and $F+1$ are required for photon production, which are labelled as $\{|g_2\rangle, |g_1^+\rangle, |g_1^-\rangle\}$. The latter two states can be considered to be degenerate in the absence of a magnetic field. The excited state, $|x\rangle$, has dipole-allowed transitions to all three ground states (cf. Fig 6), yet the transition to $|g_1^0\rangle$ is dipole forbidden. A transition from $|x\rangle$ to either of the degenerate ground states $|g_1^\pm\rangle$ corresponds to the emission of photons of orthogonal circular polarisations $\sigma^\mp$ with respect to the chosen quantisation axis. On the other hand, the transition $|g_2\rangle \rightarrow |x\rangle$ corresponds to the absorption of a photon of π-polarisation. This configuration arises naturally in alkali atoms and ions, both with and without hyperfine structure (for example **Rb**, **Cs**, $\mathbf{Ca}^+$, $\mathbf{Sr}^+$).

To best illustrate this scheme, a proper choice of quantisation axis is made and polarisation directions for the atomic driving, all combined with a fast coherent re-preparation of the atom back to its original state. This enables the production of a burst of polarised single photons emitted with repetition rates on the MHz scale. Optical pumping can be used to initialise the atom at the outset and to re-initialise it after failed photon production. In what follows, these cornerstones, as well as the implementation of the overall sequence are described.

3.1.2 *Quantisation axis and polarised photon production*

The scheme is best explained with a choice of quantisation axis which is perpendicular to the cavity axis as shown in Fig. 7. In combination with the atomic selection

2 The scheme also works for atoms with no nuclear angular momentum and only fine structure, s.t. F=J.

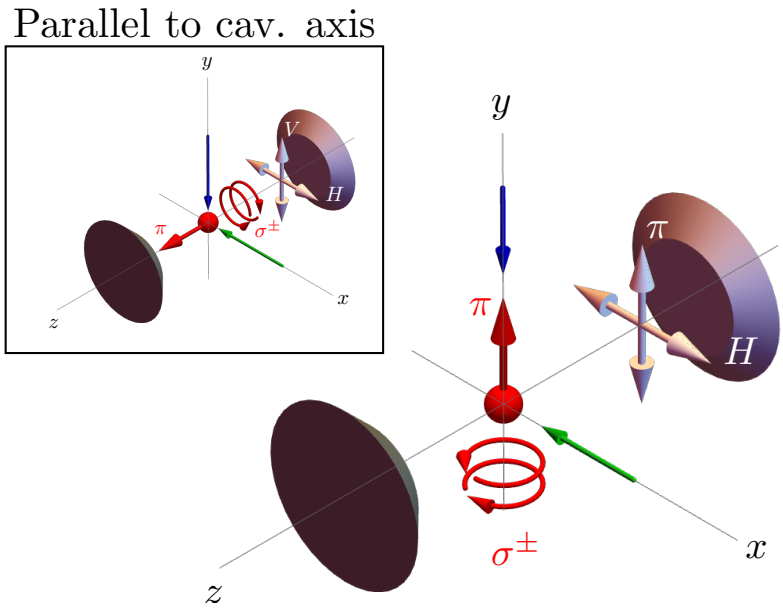

Figure 7: Atomic quantisation axis perpendicular to the cavity axis. In red the π, σ^+ and σ^- polarisation basis with respect to this particular quantisation axis is shown. In green, the light addressing the atom for V-STIRAP is shown. In white, one sees the two linear, orthogonal polarisation modes that span the cavity basis. The inset shows the traditional choice of quantisation axis, parallel to the cavity axis, where the cavity supports either a circularly polarised basis $\{\sigma^+, \sigma^-\}$ or a linear basis $\{H, V\}$.

rules, this uniquely defines the polarisation of the photon produced by the atom-cavity interaction.

The cavity supports two degenerate polarisation modes, one is π-polarised and the other a linear superposition of σ^+ and σ^-. Without loss of generality, for a quantisation axis along the y axis, the polarisation modes supported by the cavity correspond to $\{\pi, H\}$ (in the lab frame, these correspond to polarisations parallel to the $\{y, x\}$ axes, cf. Fig. 7). One aims to emit only into the linear polarisation modes supported by the cavity by enabling the atom to emit photons in a superposition state with respect to the quantisation axis, following ideas first introduced in [243] for H polarised photons from atom-photon entanglement.

A single atom is initially prepared in the state $|g_2\rangle$ and addressed by a laser with $\pi-$polarisation and total pulse duration T (polarised in the y direction in the lab frame, propagating along x cf. Fig. 7), resonant with the transition from $|g_2\rangle$ to $|x\rangle$. The cavity is tuned to be resonant with the $|g_1^{\pm}\rangle \rightarrow |x\rangle$ transitions, supporting the dipole allowed $\sigma^{\mp}$ transitions while there is no dipole allowed π transition. The cavity therefore only

supports the $H = (\sigma^+ + \sigma^-)/\sqrt{2}$ polarisation mode. This paves the way for efficiently producing linearly polarised photons without a strong external magnetic field to address individual hyperfine states.

3.1.3 *Adiabatic transfer of population*

A process known as V-STIRAP [101] is considered for photon production, whereupon the atomic population is adiabatically transferred from the initial state $|g_2\rangle$ to a superposition of states $|g_1^{\pm}\rangle$, leading to the following entangled atom-photon state, $|\Psi\rangle$:

$$\begin{aligned} |\Psi\rangle &= \frac{1}{\sqrt{2}} \left(|g_1^+, \sigma^-\rangle - |g_1^-, \sigma^+\rangle \right) \\ &= \frac{1}{\sqrt{2}} \left(|\Phi_{\phi=\pi}\rangle \otimes |H\rangle - i\, |\Phi_{\phi=0}\rangle \otimes |V\rangle \right), \end{aligned} \tag{34}$$

where $|\Phi_\phi\rangle = \frac{1}{\sqrt{2}} \left(|g_1^+\rangle + e^{i\phi} |g_1^-\rangle \right)$ [3]. However, the cavity only supports one linear polarisation mode $H = (\sigma^+ + \sigma^-)/\sqrt{2}$, where H is perpendicular to the cavity axis. V would correspond to an oscillation along the cavity axis and photons propagating perpendicular to the cavity axis. Therefore, the cavity acts like a polarisation filter, only supporting H-polarised light. As a consequence, the atom-cavity state collapses to the product state:

$$|\Psi_H\rangle = |\Phi_{\phi=\pi}\rangle \otimes |H\rangle, \tag{35}$$

which, upon photon leakage from the cavity, evolves to

$$|\Psi_f\rangle = |\Phi_{\phi=\pi}\rangle \otimes |0\rangle. \tag{36}$$

3 In general, the phase ϕ depends on the choice of quantisation axis and it should be noted that it is an important aspect of the coherent re-preparation scheme described in Sec. 3.1.5.

3.1.4 *Efficiency*

The efficiency of the V-STIRAP process is determined by estimating the cavity field decay rate with respect to the desired final state, given by an integral over the photon production time interval $[t_i, t_f]$:

$$\eta_1 = 2\kappa \int_{t_i}^{t_f} \langle \Phi_\pi, 1|\rho(t)\rangle \langle \rho(t)|\Phi_\pi, 1\rangle \, dt, \tag{37}$$

where the quantity $\langle \Phi_\pi, 1|\rho(t)\rangle \langle \rho(t)|\Phi_\pi, 1\rangle$ describes the instantaneous population of the photon emitting state and the pre-factor accounts for the cavity decay rate out of the cavity mode giving the total efficiency of emitting a photon of the desired polarisation in the finite time interval $[t_i, t_f]$. In the idealised state configuration, there is an equivalent analytic formulation which does not require an integral. This is because, the desired coherence between $\left|g_1^+\right\rangle$ and $\left|g_1^-\right\rangle$ can only be attained from the initial state $|g_2\rangle$ by adiabatically transferring population through the dark state with the coherent system dynamics. Spontaneous emission may populate a statistical mixture of the two ground-states $\left|g_1^+\right\rangle$ and $\left|g_1^-\right\rangle$, yet an incoherent process cannot lead to the creation of coherence between the ground states when the system is initialised in $|g_2\rangle$. Thus, one can also estimate the efficiency as:

$$\eta_2 = \langle \Phi_\pi, 0| \, \rho(t_f) \, |\Phi_\pi, 0\rangle - \langle \Phi_0, 0| \, \rho(t_f) \, |\Phi_0, 0\rangle \underset{\textit{ideal sys.}}{=} \eta_1, \tag{38}$$

where $\rho(t_f)$ denotes the density matrix at the end of the simulation and $|\Phi, n\rangle$ represents a tensor product of the atomic and Fock state, with $\Phi_\phi = \frac{1}{\sqrt{2}} \left(\left|g_1^+\right\rangle + e^{i\phi} \left|g_1^-\right\rangle\right)$.

Various numerical photon production efficiencies are presented for an idealised state configuration in Fig. 8, by assuming an atom is successfully initialised in $|g_2\rangle$ and addressed with a sin-squared laser pulse of $\Omega_{\text{VST}} = \Omega_{\max} \sin^2(\pi t/T)$, where the total pulse duration $T = 10/\gamma$. The variation as a function of κ/γ, as well as g/γ, where κ

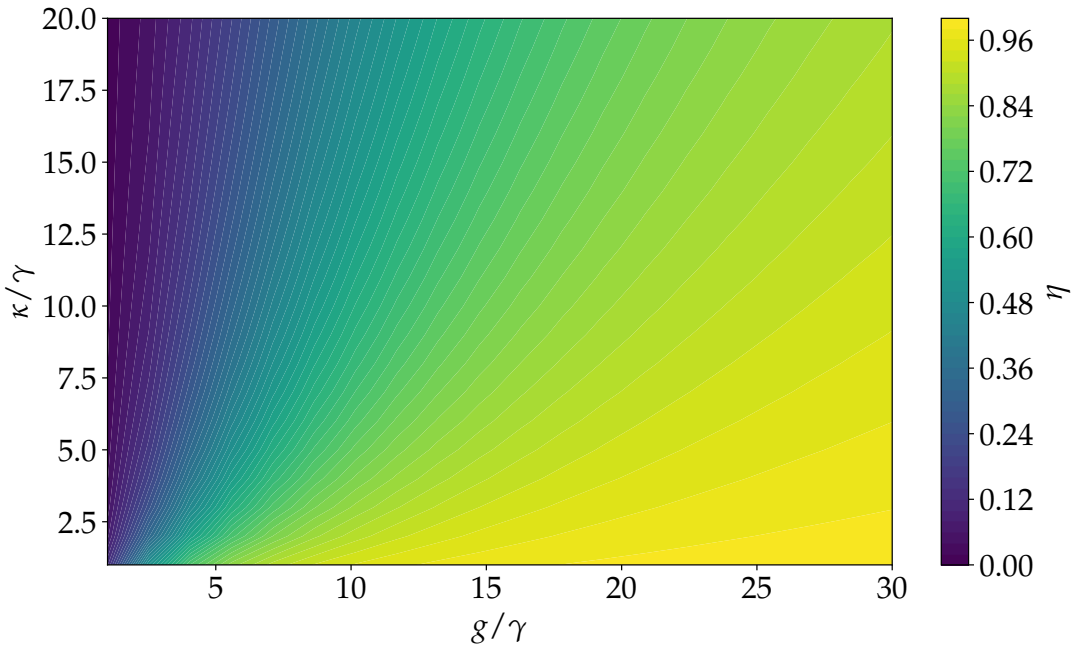

Figure 8: V-STIRAP photon production efficiency in an idealised system, shown for a fixed $\sin^2$ driving pulse of duration $T = 10/\gamma$, plotted as a function of the normalised cavity field decay rate κ/γ. The efficiency increases with both stronger atom-cavity coupling, reflecting improved coherent transfer. In contrast, weakly coupled systems with large κ exhibit reduced efficiency when driven over a finite duration, as the cooperativity which is $\propto \frac{1}{\kappa}$ is reduced.

is the cavity field decay rate, g the atom cavity coupling strength [4] and γ the atomic depolarisation rate are shown in Fig. 8. Each data-point corresponds to a local optimisation problem for finding the optimal peak Rabi frequency Ω_{max} which maximises the photon emission efficiency.

The peak Rabi frequencies were numerically optimised for maximising photon production probability. Since the system is idealised, everything was normalised with respect to γ and then g and κ, the atom cavity coupling, as well as the photonic field decay rate were varied. As evident in Fig. 9, the optimal peak Rabi frequencies for these atom-cavity coupling strengths increase drastically as a function of g, especially for fast field decay. Changing the detuning from the state $|x\rangle$ has close to no effect on the efficiency calculation and hence all V-STIRAP processes considered here are resonant. It should also be noted that one only approaches the theoretical upper limit in photon production efficiency (cf. Fig. 9) for higher cavity cooperativities in the regimes

4 $g = g_{\text{max}} * d_{\text{i,f}}$, where g_{max} is the maximal atom cavity coupling $\propto 1/\sqrt{V}$ where V is the cavity mode volume and $d_{\text{i,f}}$ is the Clebsch-Gordan coefficient for the $\sigma^{\mp}$ transitions from $|g_1^{\pm}\rangle$ to $|x\rangle$.

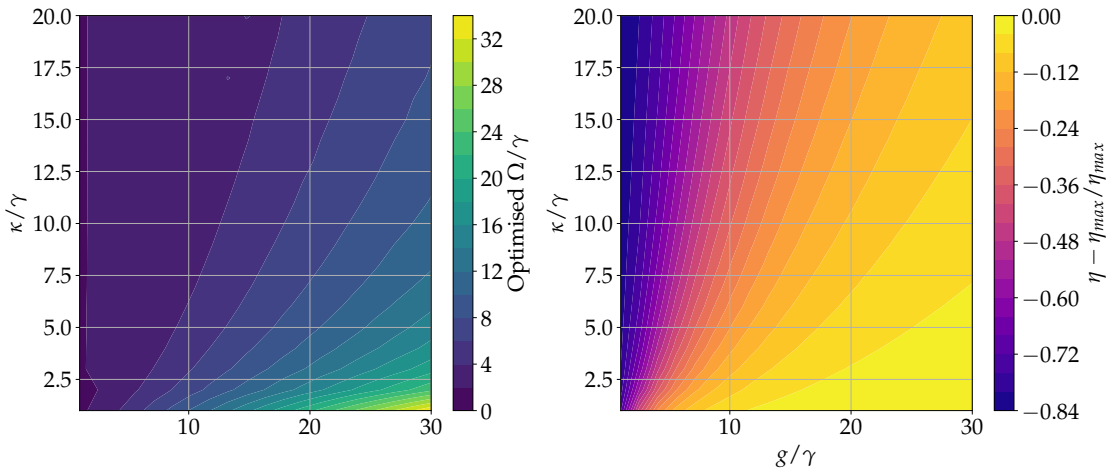

Figure 9: Optimised peak Rabi frequency for V-STIRAP photon production in idealised three level system, at the corresponding data points $g/\gamma, \kappa/\gamma$) (left) and difference from the theoretical upper limit in efficiency $\eta_{\max} = 2C/(2C+1)$, where C is the cavity cooperativity (right). For finite-time driving, the maximal achievable efficiency can only be approached when the critical system time is sufficiently short (cf. Eq. 39) relative to photon driving pulse time $T = \frac{10}{\gamma}$. In this case, this requires κ to be comparable to g in magnitude or smaller. In this regime, the coherent atom-cavity coupling is strong enough to transfer population into the cavity mode on a timescale that matches or exceeds the photon extraction rate as shown for larger g. For smaller g and high κ we observe a strong deviation from maximal attainable efficiency as the critical system time becomes comparable to the photon driving time implying that field decay is too fast in relation to the coherent photon production process.

of high g and low κ or for very large κ, i.e. in the fast cavity regime.

The efficiency of cavity-QED photon production is limited by two competing physical constraints: the requirement of adiabatic population transfer and the ability to efficiently extract the generated photon from the cavity. Adiabaticity is governed by the relative strengths of the atom-cavity coupling g, the cavity field decay rate κ, and the timescale of the driving pulse $\Omega(t)$, while photon extraction is directly controlled by the cavity decay rate κ. As shown in Fig. 8, the efficiency initially increases with κ and approaches unity for large g once $T\kappa > 1$, indicating that cavity extraction no longer limits the finite-time efficiency of the process[5].

These constraints are compactly summarised by the critical system time [229]

$$t_{\text{crit}} = \max\left(\frac{1}{\kappa}, \frac{\kappa}{g^2}\right), \tag{39}$$

5 This behavior is discussed in more detail in [229].

which sets the minimum pulse duration required for efficient photon generation. The timescale $1/\kappa$ corresponds to the cavity lifetime and determines how quickly a photon can be extracted from the cavity. κ/g^2 characterises the effective coherent transfer rate between the atomic excitation and the output mode and reflects the need for adiabatic following during emission. Pulses shorter than either limit result in incomplete extraction or non-adiabatic population of the lossy excited state, respectively. Consequently, efficient photon production requires pulse durations exceeding both timescales, with the slower process defining the system's critical time. This leads to an intrinsic trade-off between repetition rate and efficiency: increasing κ facilitates photon extraction but simultaneously demands stronger coupling g to maintain adiabaticity. The implications of this trade-off for a realistic atomic system are discussed further in Sec. 3.3.

A fundamental upper bound in the photon production efficiency obtained for V-STIRAP arises from the atom-cavity parameters. Controlling the time evolution of the Rabi frequency Ω enables the precise shaping of the time-dependent photon amplitude [234], as well as the efficiency with which it is produced. Even with lossless mirrors, the production of a temporally shaped photon is limited in efficiency by a theoretical maximum of $\frac{2C}{2C+1}$ (cf. Eq. 31). Further losses arise from photonic field decay which does not occur to an external mode, but which are internally lost due to e.g. mirror absorption or scattering. This introduces an additional prefactor called the out-coupling efficiency [229] $\eta_{esc} = \frac{\kappa_e}{\kappa_e+\kappa_i}$, where the subscripts e and i denote external field decay and internal field decay. This loss can be optimised by improving the mirror surface quality with low absorption and scattering losses. This is discussed more extensively in [66].

The previous analysis begs the question as to what the ideal cavity parameters are for maximising photon coincidences with a particular cavity design. Combining the results from Eq. 39 and 31 allows for the definition of a cost function which takes into account both photon production efficiency (assuming perfect out-coupling efficiency) as the ratio of the photon production efficiency and critical time. This leads to

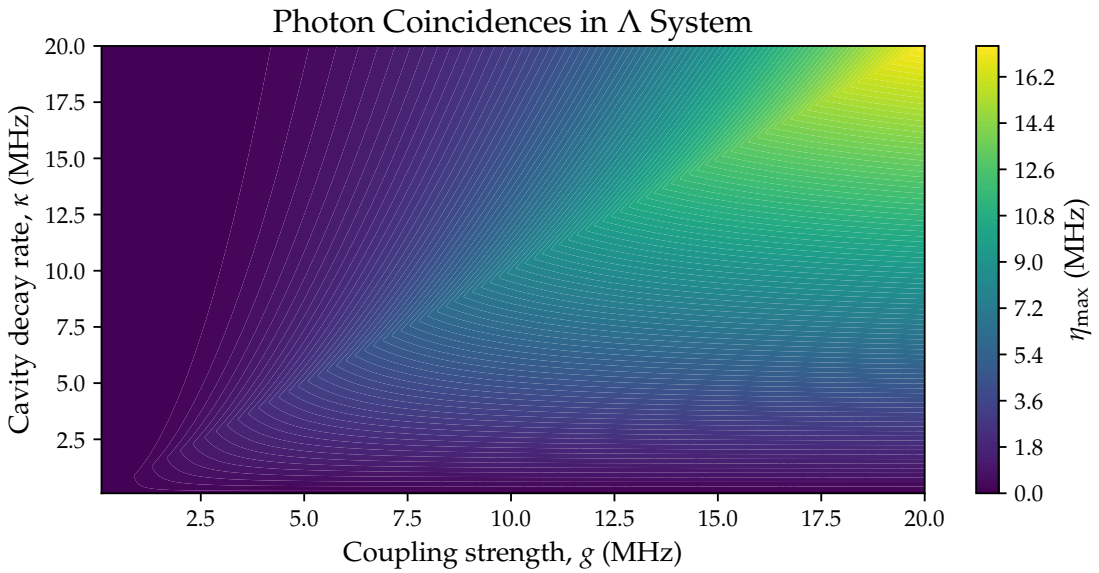

Figure 10: Cavity design optimisation for maximising photon generation rate $\eta_{max} = \eta / t_{\mathbf{crit}}$ at fixed spontaneous emission rate $\gamma/2\pi = 3$MHz. One can see that for achieving the optimal trade-off between high photon production efficiency and repetition rate one should maximise $\kappa = g$ to balance the coherent production and photon extraction.

contours visible in Fig. 10 where maximal photon production rate is achieved in the regime where $\kappa = g$. This is interesting insofar one is neither in the strong coupling regime where $g > \kappa$, nor in the bad-cavity (or Purcell) regime where $\kappa > g^2/\kappa > \gamma$ [181]. Rather the intermediate regime produces the highest expected photon coincidences. Further experimental work is needed to understand how these results relate to practical considerations such as mirror fabrication and coating, feasible cavity lengths, cavity mode structure, and overall system stability. Nevertheless, this provides a useful insight, especially as much of the prior work has focused primarily on optimising production and escape efficiencies.

3.1.5 *STIRAP Re-preparation*

Coherent STIRAP re-preparation [28] [6] enables fast and efficient re-preparation of the atom back into the desired initial state $|g_2\rangle$ with high re-preparation efficiencies. Unlike incoherent methods, which rely on eventual spontaneous emission into the desired

6 This is introduced more completely in Sec. 2.1.3

state and are bounded by the atomic decay rate, coherent repumping over a time-scale T must satisfy the constraint of global adiabaticity [28], $\Omega_{\text{rms}}T \gg 1$, where $\Omega_{\text{rms}} = \sqrt{\hat{\Omega}_S^2 + \hat{\Omega}_P^2}$, such that $\hat{\Omega}$ describes the time average amplitude over the pulse duration for the Stokes pulse Ω_S and the Pump pulse Ω_P. The desire for short pulses with high repetition rates can be achieved with intense pulses with correspondingly high Rabi frequencies. Moreover, fully coherent atomic population transfer is reversible, essential for the implementation of quantum networks [56, 65] and for the preparation of highly entangled optical states with a single atomic memory [223].

After a photon has been successfully emitted, the atom, which is left in the state $|\Phi_\pi\rangle$, can be prepared back into the desired initial state $|g_2\rangle$ using STIRAP [238] which involves using two laser pulses (Stokes and Pump) to transfer the atomic population back to its initial state. Coherent methods like STIRAP are generally extremely good for achieving high population transfer efficiencies [133, 233] and are also fast compared to incoherent optical pumping.

To perform coherent repumping, the pump pulse must be resonant with the transitions from the groundstates $|g_1^+\rangle$ and $|g_1^-\rangle$, to $|x\rangle$ and a π polarised Stokes pulse must be resonant with respect to the $|g_2\rangle \rightarrow |x\rangle$ transition [7]. The pump pulse must be polarised in a linear combination of σ^+ and σ^- polarisation with respect to the quantisation axis, maintaining a particular phase relation between the circularly polarised components (cf. Eq. 36), such that it is linearly polarised perpendicularly both to the cavity and the quantisation axis [8]. As a result, the polarisation of the laser light and thus the phase between the two circularly polarised components of the light must match the phase of photon polarisation derived from Eq. (36). Both laser pulses must have equal effective Rabi frequencies [9].

By carefully shaping the time dependent amplitudes of two pulses which address the transitions $(|g_1^+\rangle + |g_1^-\rangle) \rightarrow |x\rangle$ and $|x\rangle \rightarrow |g_2\rangle$, it is possible to maximise re-

7 Compared to the V-STIRAP process, the initial and final states swap roles and a pump laser and stimulating cavity mode are replaced by a stokes and pump laser respectively.

8 This requires the laser to be pointing along the quantisation axis, as explained in Fig. 13

9 It is the effective peak Rabi frequencies which are equal. The absolute laser intensities may differ if the Clebsch-Gordan coefficients of the individual transitions vary.

preparation efficiency between the states $|g_1^{\pm}\rangle$ and $|g_2\rangle$. Various pulse-shaping techniques have been developed extensively, such as optimal control methods [122, 124, 133], reinforcement learning [86] or other numerical methods to minimise the non-adiabatic transitions [233]. It should be noted that for experimentally feasible pulse shapes, one requires $\Omega_S(t=0) = \Omega_P(t=0) = \Omega_S(t=T) = \Omega_P(t=T) = 0$, where T is the total pulse sequence duration [10]. With this in mind, the highest population transfer efficiency was attained by adopting a technique that involved employing two interleaved pulses of constant rms Rabi frequency Ω_{rms}, which were applied with a hyper-Gaussian mask [233]. The simulated STIRAP re-preparation efficiency approach unity in the idealised system configuration ($\eta_{\max} \approx 0.997$ in Fig. 11), with losses parameterised by γ, as shown in Fig. 11.

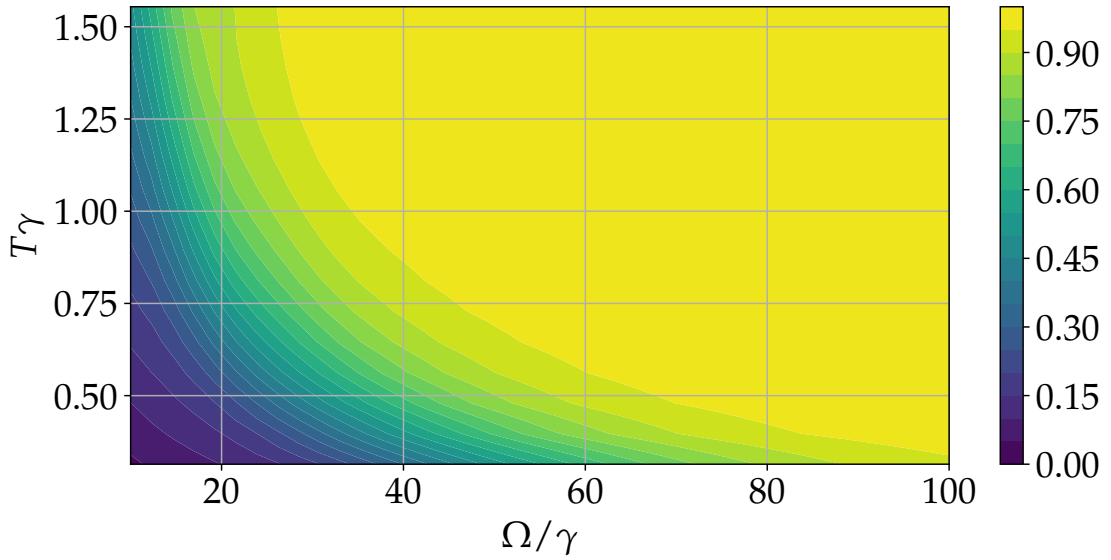

Figure 11: Repumping efficiency in idealised system plotted for length of the repumping pulses T and peak Rabi frequency Ω, which are both normalised with respect to the atomic depolarisation rate γ. All pulses are resonant.

There is a trade-off between efficiency, repetition rate and driving laser intensity. The variation of the repumping efficiency with re-preparation time and peak Rabi frequencies is shown in Fig. 11. In general, faster re-preparation requires larger peak Rabi frequencies and in the idealised system configuration one can in principle arbitrarily

10 It is worth noting that if the transition used for V-STIRAP photon production corresponds to that of the Stokes laser (i.e. Ω_S), the laser can be left on after the photon has been successfully emitted from the cavity and turned off once the Pump pulse begins to increase in amplitude.

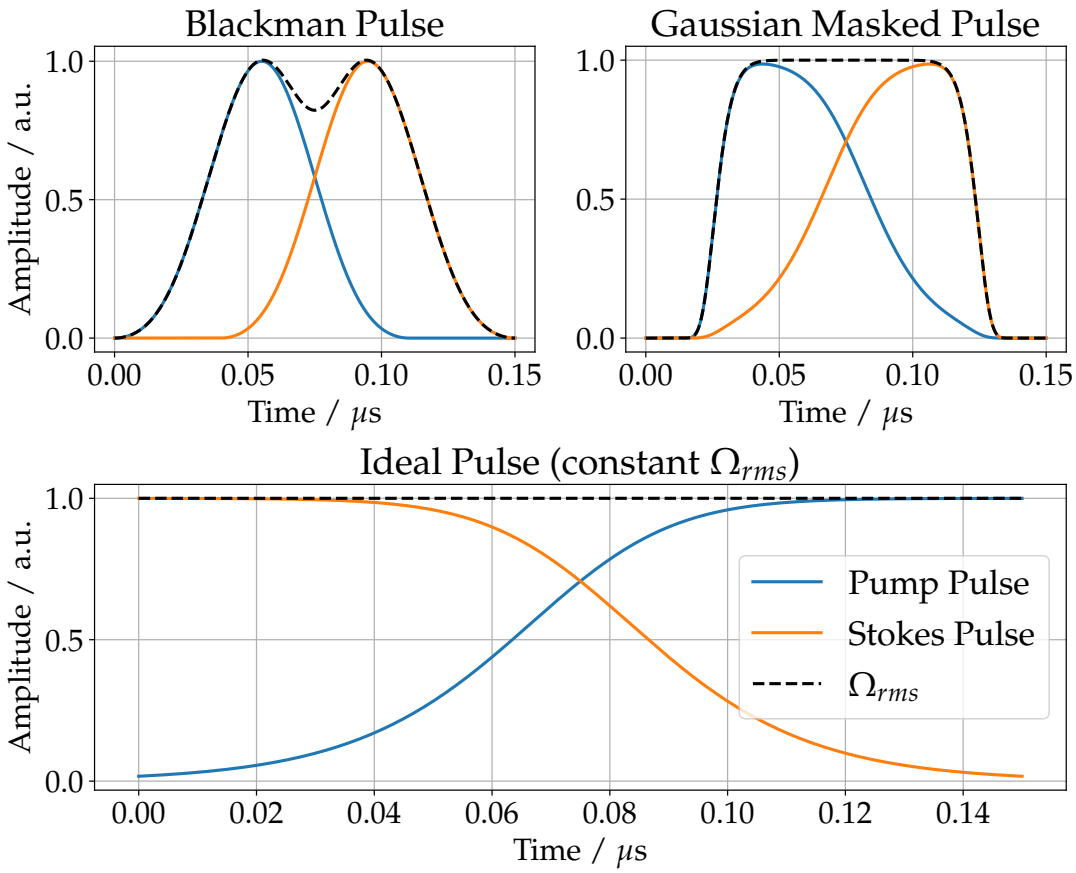

Figure 12: Various pulse shapes with optimised parameters used for coherent re-preparation, with the ideal (yet unphysical) shape with constant Ω_{rms} being the most efficient and also exhibiting the ideal case of constant root mean square Rabi frequency (bottom). The top left is a common pulse form used because of its optimal frequency window, yet no matter what pulse delay is chosen, it does not give constant Ω_{rms}. The top right pulse is a physically realistic implementation of the ideal pulse (bottom) as it starts and ends at zero amplitude.

increase the peak Rabi frequency to increase the efficiency, this is in stark contrast to many real atomic systems where these are limited. The pulses are shaped with a technique introduced in [233] by considering two pulses of constant rms Rabi frequency $\sqrt{(\hat{\Omega_S}^2 + \hat{\Omega_P}^2)}$. These correspond to the ideal adiabatic trajectory in a three level system, but are unphysical experimentally as they do not start or end with zero amplitude (confer bottom pulses in Fig. 12). Thus, this optimal shape is scaled with a hyper-Gaussian mask with numerically optimised envelope. The analytic expressions for the pulse shapes of total duration T and peak Rabi frequency Ω_{max} are as follows:

$$\begin{aligned}\Omega_S(t) &= \Omega_{\max}\cdot e^{\left(-\left(\frac{t-T/2}{c}\right)^{2n}\right)}\sin\left(\frac{\pi}{2\cdot\left(1+e^{-\left(\frac{a(t-T/2)}{T}\right)}\right)}\right)\\ \Omega_P(t) &= \Omega_{\max}\cdot e^{\left(-\left(\frac{t-T/2}{c}\right)^{2n}\right)}\cos\left(\frac{\pi}{2\cdot\left(1+e^{-\left(\frac{a(t-T/2)}{T}\right)}\right)}\right)\end{aligned} \tag{40}$$

where (n, a) are control parameters and $c = 0.05$ is chosen such that both pulses exhibit (within the numerical limit of the solver) zero amplitude at $t = 0$ and $t = T$. This is a desiderata for making the simulation more realistic. The optimal values of (n, a) for the idealised system are $(6, 14)$. The pulses are exhibited in Fig. 12 and lead to re-preparation efficiencies which approach unity at large peak Rabi frequencies. For completeness, a Blackman amplitude envelope is also shown, as it has an *optimal* frequency transformed window and is common in experiments as the smoothness of the amplitude modulation makes it straightforward to implement with bandwidth limited electronics. However, the population transfer efficiency achieved with this method is below that achieved by the hyper-Gaussian mask, as the r.m.s. Rabi frequency exhibits some variation around the centre of the plotted time window. It shall also be stressed that the discussion in this section did not take into account the entire structure of a real atom where further losses arise from off-resonant coupling to undesired levels which will be explored in what follows.

3.2 IMPLEMENTATION IN REAL ATOMIC SYSTEM: RUBIDIUM 87

Many alkali atoms and earth alkaline ions exhibit the desired atomic structure. A particular choice for an atomic emitter is $^{87}\mathbf{Rb}$, as its properties and atomic structure are well understood [216] and it exhibits a hyperfine structure suitable for the proposal. More details on the structure of $^{87}\mathbf{Rb}$ are provided in Sec. 2.1.3 Following the exposition in Sec. 3.1.3, one picks $|g_1^{\pm}\rangle = |F = 1, m_F = \pm 1\rangle$, $|g_2\rangle = |F = 2, m_F = 0\rangle$, and

$|x\rangle = |F' = 1, m_F = 0\rangle$. This configuration of states in Rubidium exists for two different transitions, one has the choice of either the D_2 transitions defined with respect to $5^2S_{\frac{1}{2}} \Longleftrightarrow 5^2P_{\frac{3}{2}}$ or the D_1 transitions, defined with respect to $5^2S_{\frac{1}{2}} \Longleftrightarrow 5^2P_{\frac{1}{2}}$. In what ensues, all components of the photon production sequence are analysed and the ideal parameters are presented. An important trade-off with respect to photon production efficiency and the explicit cavity and laser parameters is discussed in Sec 3.2.1 and the results are discussed and presented in Sec. 3.3.

3.2.1 *Photon Production*

3.2.1.1 *Quantisation Axis & External Magnetic Field*

The proposed photon production scheme does not require a strong external magnetic field to significantly split the hyperfine states, because there is no need to address them individually. As shown in [20], in ^{87}**Rb** a too strong field leads to the apparition of non-linear Zeeman effects which strongly curtail achievable efficiencies. To ensure that the atomic quantisation axis does not vary erratically in time, one does however seek to impose a weak external magnetic field, otherwise background field fluctuations would adversely affect photon polarisation purity within a particular burst. The achievable simulated efficiencies in Sec. 3.2 all assume perfectly degenerate ground-states, yet in reality one seeks to impose a small external magnetic field, aligned with the quantisation axis which in turn imposes a small hyperfine state splitting and induces a Larmor precession. As demonstrated in Fig. 14, increasing ground state splitting affects the phase $\phi(t)$ of the atomic superposition state Φ:

$$\left|\Phi_\phi\right\rangle = \frac{1}{\sqrt{2}}(|F = 1, m_F = 1\rangle + e^{i\phi(t)} |F = 1, m_F = -1\rangle). \tag{41}$$

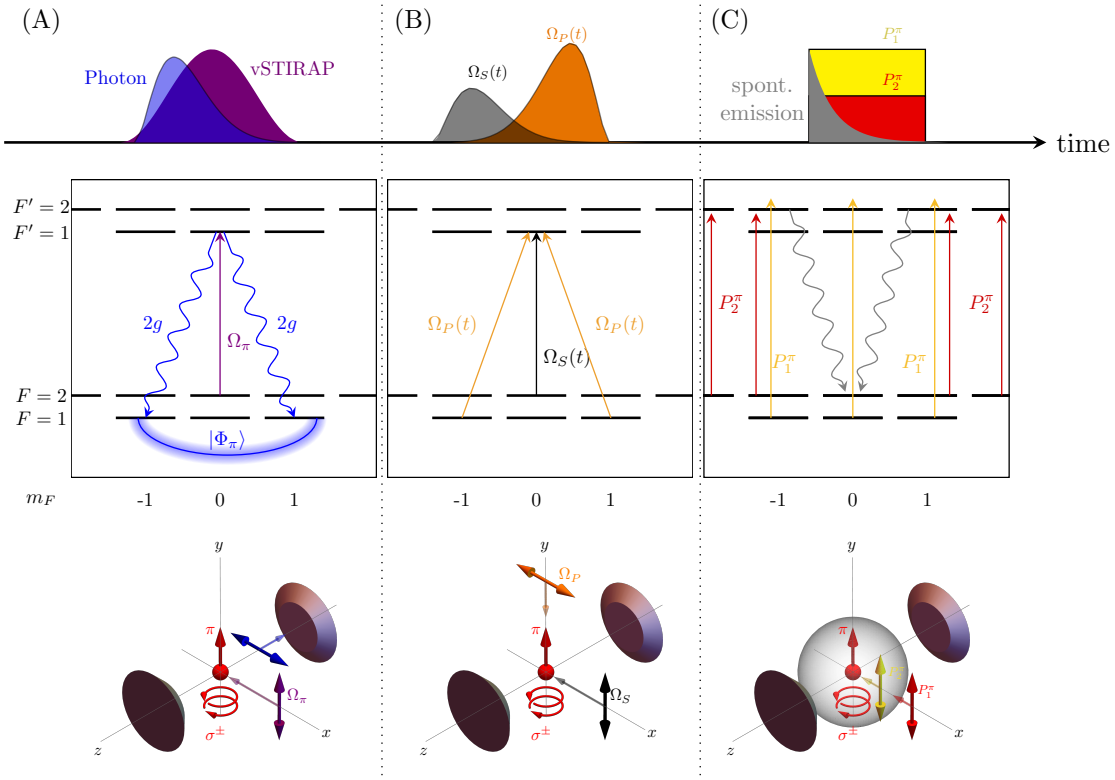

Figure 13: Sequence for photon generation with ^{87}**Rb**. (A) Starting with an atom in $|F = 2, m_F = 0\rangle$, a π polarised driving pulse Ω_π, resonant with the transition to $|F' = 1, m_{F'} = 0\rangle$, addresses the atom. The cavity is resonant with the transitions $F = 1 \leftrightarrow F' = 1$. Importantly, the π-transition from $|F' = 1, m_{F'} = 0\rangle \rightarrow |F = 1, m_F = 0\rangle$ is dipole forbidden such that the atomic population is transferred into the superposition state $|\Phi_{\phi=\pi}\rangle = \frac{1}{\sqrt{2}}(|F = 1, m_F = 1\rangle - |F = 1, m_F = -1\rangle)$. A linearly polarised photon is then emitted. (B) This is followed by fast, coherent STIRAP re-preparation. A driving laser of peak Rabi frequency Ω_P is polarised linearly and parallel to the emitted photon and a stimulating laser Ω_S that is π polarised. Together, they adiabatically re-prepare the atom back into the desired initial state $|F = 2, m_F = 0\rangle$. These two processes can be repeated n times to produce the desired number of photons sequentially. (C) Thereafter, the atom is re-prepared incoherently into the desired initial state with optical pumping. One laser is blue detuned from the $F = 1 \rightarrow F' = 2$ transition and the other is red detuned from the $F = 2 \rightarrow F' = 2$ transition. All transitions from the excited $F' = 2$ line are dipole allowed except the $|F' = 2, m_{F'} = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$ transition. The lasers are two overlapping top-hat pulses, thus eventually the population is continually re-excited out of the undesired Zeeman ground-states and will spontaneously decay (isotropically, shown as a grey sphere) to the desired $|F = 2, m_F\rangle = 0$ state (cf. grey lines). This is also used to initialise the atom into the desired Zeeman state in the first place.

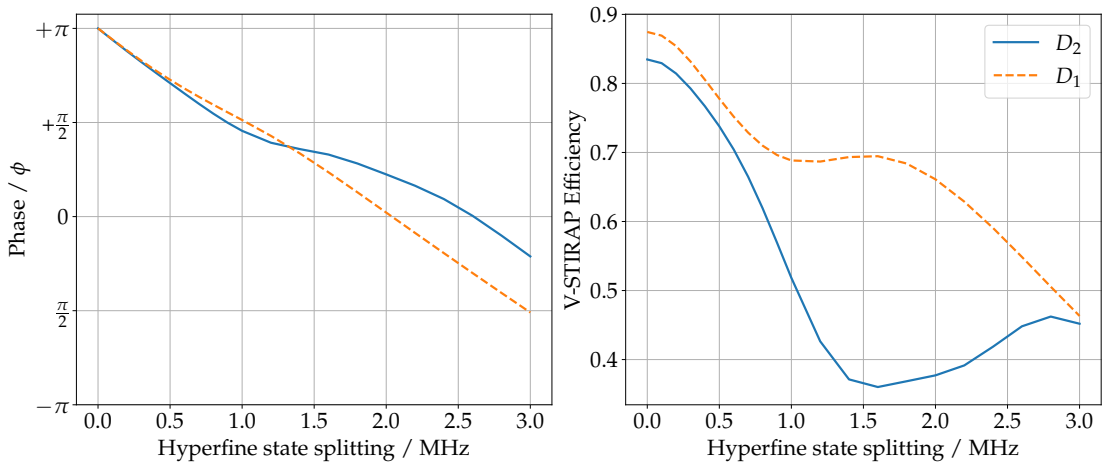

Figure 14: Variation of the phase $\phi(t = T)$ (y) (where $T = 500$ns is the V-STIRAP driving pulse length) between the two hyperfine groundstates $F = 1, m_F = \pm 1$ against hyperfine state splitting (x) by imposing a weak external magnetic field aligned with the preferred quantisation axis (left) and variation of the V-STIRAP photon production efficiencies for state of the art optical cavity [67] (right). The phase rotates linearly for D_1, yet varies slightly erratically for D_2 from about 1 MHz splitting onwards, as the D_2 transitions experience an asymmetric variation in the dipole matrix elements for the $\sigma^{\pm}$ transitions which couple to the individual states in Eq. 41 [20]. This is also reflected in the sharp decline in efficiency for the D_2 transitions, which eventually increases again at larger splittings.

Stronger hyperfine splitting eventually rotates the phase by an entire revolution over the V-STIRAP pulse duration T and this induces a time dependent photon polarisation drift and also adversely affects the achievable photon production efficiencies as shown in Fig. 14. A ground state splitting of 1 MHz corresponds to a field of 1.43 G, therefore, in the mG regime, as demonstrated in [223], the phase evolution over the V-STIRAP pulse duration is close to the desired phase of π and the efficiencies are very close to the limit achieved in a perfectly degenerate state configuration. In this regime, there is a fixed quantisation axis and photon polarisation, but close to no decline in efficiency from the perfectly degenerate state configuration, as well as no fast Larmor precession, with respect to the photon production timescale. Stronger hyperfine splitting would also require a modification of the polarisation of the STIRAP pump beam, such that its polarisation matches the polarisation of the emitted photon at the end of the driving pulse (cf. Eq. 41), where $\phi(t = T)$ corresponds to the relative phase of the $\sigma^{\pm}$ beams. If $\phi(t = T)$ changes quickly on the photon production timescale $[0, T]$, synchronising

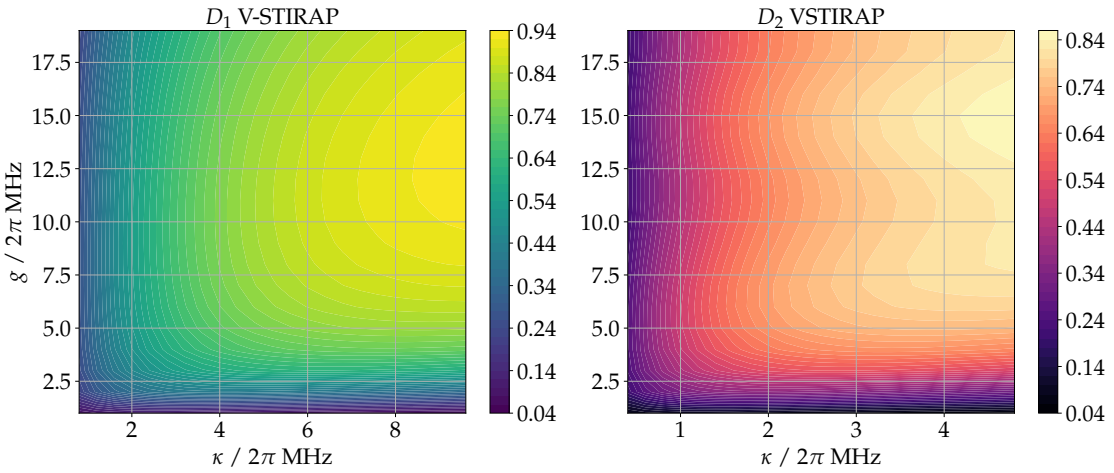

Figure 15: V-STIRAP Photon Production efficiency for both D_1 and D_2 lines plotted for various atom cavity coupling strengths g and photonic field decay rates κ. The laser pulse length is fixed at 500 ns and it takes a $\sin^2$ shape. Each point in the plot corresponds to a local optimisation problem to find the optimal detuning and peak Rabi frequency to maximise photon production efficiency. For the same photonic field decay rate κ and for the same atom cavity coupling g the D_1 scheme is more efficient than the D_2 scheme. Note however, that g for different atomic transitions does not correspond to the same physical mirror parameters, since the Clebsch-Gordan coefficients for the V-STIRAP transition differ.

repreparation pump polarisation becomes challenging and it is therefore important to remain in the low B-field regime.

3.2.1.2 *V-STIRAP*

The V-STIRAP process lies at the heart of the photon production scheme and its merits have been demonstrated for an idealised three level system in Sec. 3.1.3. The full atomic structure of ^{87}**Rb** is now considered and the optimal parameters and efficiency trade-offs exemplified by this real atomic quantum system are carefully determined.

The relevant atomic states are shown in Fig 13. After initialising the atom into the $|F = 2, m_F = 0\rangle$ ground state, a π polarised driving pulse, near resonant with the transition to $|F' = 1, m_{F'} = 0\rangle$ addresses the atom and the cavity is tuned to be resonant, such that one transfers the atom into the superposition state:

$$\left|\Phi_{\phi=\pi}\right\rangle = \frac{1}{\sqrt{2}}(|F = 1, m_F = 1\rangle - |F = 1, m_F = -1\rangle). \tag{42}$$

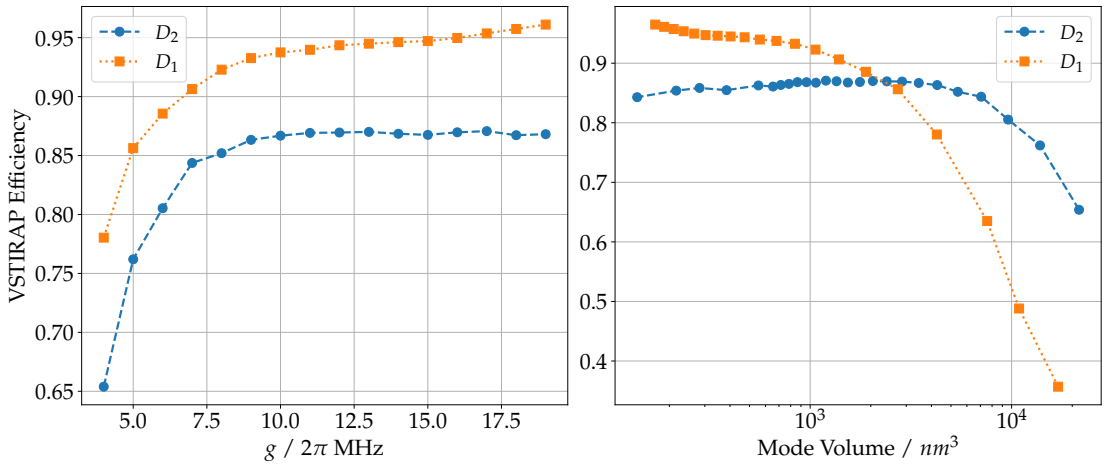

Figure 16: Efficiency of the V-STIRAP photon production procedure for various atom cavity coupling strengths g for both the D_1 and D_2 lines (left), at fixed κ. Given the desire to reasonably compare physical cavity parameters across possible schemes the efficiency is also plotted against mode volume which is a function of the cavity mirror properties (right). It is clearly visible that for most feasible cavity mode volumes $> 2 \times 10^3 \text{nm}^3$, the D_2 V-STIRAP process is more efficient than the D_1 V-STIRAP process.

A linearly polarised photon is then emitted preferentially through one of the cavity mirrors.

It was shown in Sec. 3.1.3 that in order to push the efficiencies near unity, one requires strong atom cavity coupling with respect to the atomic depolarisation rate and long photons or, alternatively, a fast photonic field decay. Some specific cavity regimes are explored in Fig. 15. A fixed pulse of length 500ns which takes the following form: $\Omega_{\text{VST}} = \Omega_{\text{max}} \sin^2(\pi t/T)$ is chosen in what follows. Every particular cavity parameter configuration corresponds to a local optimisation problem with regards to the detuning $\Delta = \Delta_{cav} = \Delta_{laser}$ (only two photon detuning was considered) and peak Rabi frequency of the laser pulse Ω_{VST}. These are numerically optimised for maximising the photon production probability (cf. Eq. 37). In general, the optimal peak Rabi frequency increases with regards to an increase in atom-cavity coupling as shown in Fig. 17.

Fig. 15 shows that the D_1 transitions lead to higher efficiencies than the D_2 transitions. To compare the two possible transition lines more explicitly, a realistic photonic field decay rate of $\kappa = 2 \times 2\pi$ MHz is imposed and only the atom-cavity coupling strength g is varied, as exemplified in Fig. 16. Importantly, unlike in an idealised three level sys-

tem (cf. Fig. 8), the efficiency does not asymptotically approach a theoretical maximum, as the atom cavity coupling strength (and cooperativity) further increase (cf. Fig. 16). This can be attributed to the fact that, stronger coupling necessitates higher peak Rabi frequencies Ω_{VST} which lead to undesirable off-resonant couplings to other excited levels. For the D_2 line, the excited levels $F' = 1$ and $F' = 2$ are split only by ≈ 160 MHz [216]. As the strength of the atom-cavity coupling increases beyond $\approx 10 \times 2\pi$ MHz, no further increases in V-STIRAP photon production efficiency on D_2 is observed, as shown in Fig. 16. The excited levels $F' = 1, 2$ for the D_1 transitions are significantly further detuned at ≈ 816 MHz. Thus, for the D_1 scheme, a decline in efficiency is not visible up to the couplings shown in Fig. 16.

It shall also be stressed that the same physical cavity mirror parameters do not give rise to the same atom-cavity coupling strength for different atomic transitions, because the atom cavity coupling g is a function of the Clebsch-Gordan coefficients of the particular transitions between the states $\left|g_1^{\pm}\right\rangle$, $|x\rangle$ and $|g_2\rangle$. For the same physical cavity parameters, the coupling strengths scale as $g_{D_2} \approx g_{D_1} \times 2.3$. The trade-off between cavity mode volume and single photon production efficiency is also shown in Fig. 16 to compare existing cavity designs. There exists a cross-over point at a mode volume of around $2 \times 10^3 \mathrm{nm}^3$ below which it becomes advantageous to use the D_1 line for photon production, but for most existing physical cavities [67], the D_2 line is preferential.

As for the idealised system, there exists an optimisation problem at each particular cavity parameter configuration (g, κ), to find the optimal peak Rabi frequency and detuning for maximising polarised photon production probability. The optimal peak Rabi frequencies at various atom cavity coupling strengths g and fixed $\kappa = 2 \times 2\pi$ MHz are given in Fig. 17. One generally observes an increase with g and Ω varies approximately linearly for low coupling strengths. Thereafter, Ω varies more erratically for higher coupling strengths as competing off-resonance effects adversely affect the photon production probability. Detuning the V-STIRAP process further from the atomic resonance had little effect on the efficiency, although detuning the cavity is im-

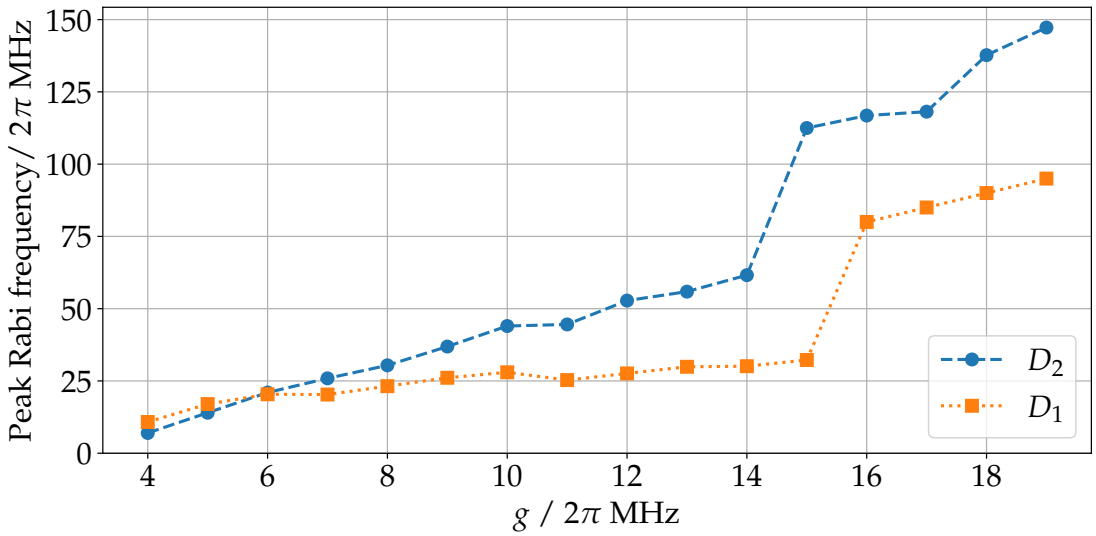

Figure 17: Optimal peak Rabi frequency for the V-STIRAP photon production procedure for various cavity coupling strengths for both the D_1 and D_2 transition lines.

portant insofar as STIRAP re-preparation is performed on the same transitions used for photon production.

3.2.2 *STIRAP Re-preparation*

The coherent re-preparation scheme described in Sec. 3.1.5 can be applied to the real atomic system of ^{87}**Rb**, which offers either the D_1 or D_2 transitions for coherently re-preparing the atom back into the desired initial state with two counter-intuitively sequenced laser pulses of appropriate polarisations. In general, when simulating the full atomic structure one observes similar trade-offs noted in the previous section for polarised photon production with V-STIRAP.

The STIRAP re-preparation on either D_1 or D_2 transitions comprises of a Stokes pulse which is π polarised and resonant with the $F = 2$ to $F' = 1$ transition, as well as a pump pulse which comprises of two $\sigma_\pm$ polarised pulses which are resonant with the $F = 1$ to $F' = 1$ transitions. Following the same principles introduced in Sec. 3.1 we consider two pulses of constant rms Rabi frequency with a hyper-Gaussian mask and numerically optimised Gaussian shape/peak Rabi frequencies (cf. Eq. 40).

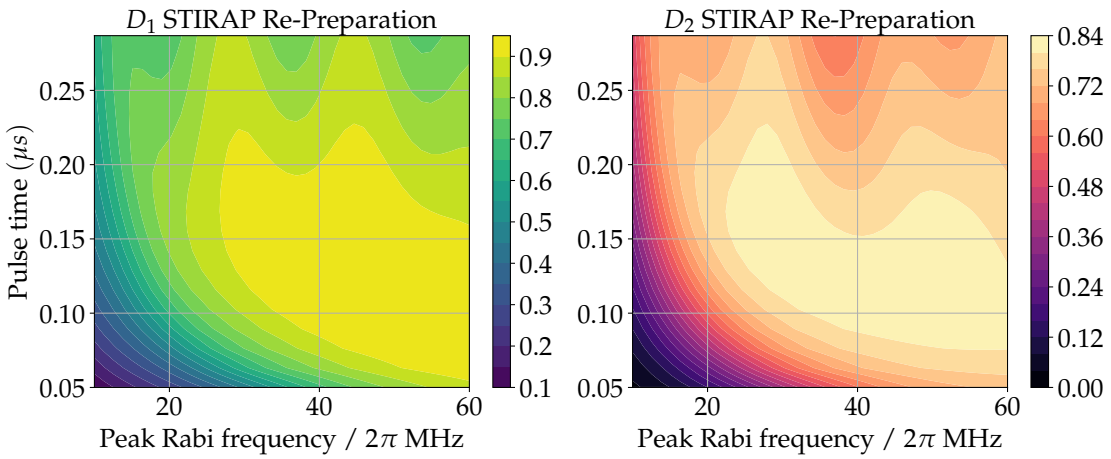

Figure 18: STIRAP re-preparation efficiency for both the D_1 and D_2 transitions plotted against the length of the re-preparation pulses for various peak Rabi frequencies Ω. STIRAP re-preparation is generally more efficient on D_1 and, unlike in an idealised system, arbitrarily large peak Rabi frequencies do not further increase the re-preparation efficiency as undesired off-resonance effects adversely affect the re-preparation efficiency.

Thus, $(n, a, \Omega_{\{max})$ are free parameters and $c = 0.05$ is chosen such that both pulses exhibit (within the numerical limit of the solver[11]) zero amplitude at $t = 0$ and $t = T$. The optimal parameters for maximal population transfer efficiency are found to be as follows: $(n, a) = (6, 11)$ and $(\Omega^{D_1}_{\text{max}}, \Omega^{D_2}_{\text{max}}) \approx (41, 49) \times 2\pi$. At a fixed re-preparation time of 150 ns, the maximal population transfer efficiencies with these pulse shapes are $\approx (0.95, 0.83)$. The ideal, but un-physical constant Ω_{rms} pulses yield a re-preparation efficiency of $(0.95, 0.85)$. The population transfer efficiencies could be further improved by introducing longer pulses which however, detrimentally affect the produced photon repetition rate. The re-preparation pulse shapes for D_1 are depicted in Fig. 19. It was found that for a time independent pulse detuning, resonant pulses gave the highest population transfer efficiency, within the constraint of a reasonable [12] laser intensity. A large detuning required significantly higher peak Rabi frequencies at this pulse time scale which is experimentally challenging due to the need to introduce a fast frequency variation to counteract the differential Stark shift introduced by the pulse. However, if

11 $\approx 10^{-}9$

12 Reasonable here implies experimentally viable with standard optical components and diode lasers and it also implies that the lasers do not induce differential Stark shifts which significantly shift the relative energies of the atomic states.

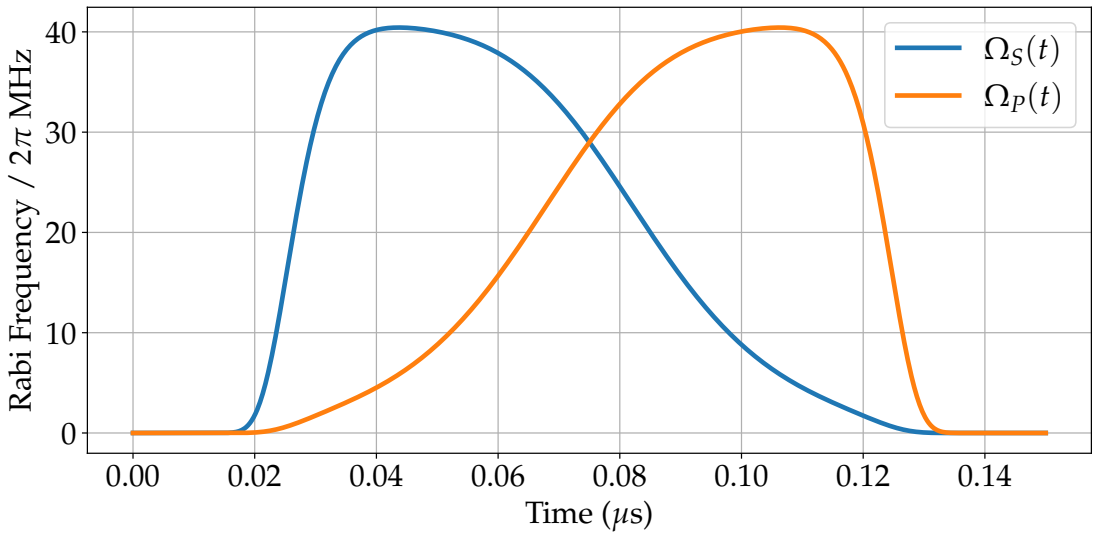

Figure 19: STIRAP repumping pulses with optimised parameters shown for D1 repumping scheme. Note that although the absolute laser pulse amplitudes differ, the effective Rabi frequencies of the Pump and Stokes pulses for the respective transitions $|\Phi_\pi\rangle \rightarrow |F' = 1, m_{F'} = 0\rangle$ and $|F = 2, m_F = 0\rangle \rightarrow |F' = 1, m_{F'} = 0\rangle$ are equal at $41 \times 2\pi$ MHz.

one seeks to further improve the repreparation efficiency, a possible extension is to also optimise the frequencies of the pulses in time.

As for the photon production process, off-resonant coupling to other excited levels is the main source of depolarisation and adversely affects the re-preparation efficiency, particularly for extremely fast re-preparation which requires high peak Rabi frequencies, as exemplified in Fig. 18 for both the D_1 and D_2 transitions. One observes that for re-preparation times of up to 150 ns the maximal population transfer efficiency for re-preparation on these lines is 0.95 and 0.83 respectively. Like in the V-STIRAP process, the relative proximity of the higher lying excited levels detrimentally affects the efficiency of the re-preparation with respect to the D_2 transitions. It should be noted, that the presence of the cavity may also adversely affect the re-preparation efficiency, as it produces photons on the same transitions used to coherently re-prepare the atom back into the initial state. However, significantly detuning the cavity from the resonantly driven transition to $F' = 1, m_F = 0$ ensures that there is no such effect.

Due to the reduced achievable efficiencies with the D_2 line (cf. Fig. 18) and the adverse affect of a near-resonant cavity with respect to the re-preparation efficiency given

current state of the art cavity parameters [67], the optimal photon production scheme uses a combination of V-STIRAP with respect to the D_2 transitions for single photon production, as well as STIRAP re-preparation via the D_1 transitions. It exploits the optimal efficiency of the underlying processes and avoids the introduction of further losses in efficiency, as the cavity is now extremely far detuned from the transitions used for quickly re-preparing the atom back into the desired initial state.

3.2.3 *Optical Pumping*

An ^{87}**Rb** atom loaded into an optical cavity is initially populated in a random ground state. Moreover, during the photon production cycle consisting of V-STIRAP and coherent re-preparation, the atom may spontaneously decay into any Zeeman state. For these reasons, it is necessary to use an optical pumping process to efficiently prepare the atom into the desired initial $|F = 2, m_F = 0\rangle$ state.

Optical pumping for the D_1 and D_2 transitions requires two π polarised lasers (with respect to the y axis, traveling along the x axis, as shown in Fig. 13) which are nearly resonant with the $F = 1 \rightarrow F' = 2$ and $F = 2 \rightarrow F' = 2$ transitions. Both lasers act simultaneously and through spontaneous emission all population eventually accumulates in the $F = 2, m_F = 0$ groundstate, because the transition from $F = 2, m_F = 0 \rightarrow F' = 2, m'_F = 0$ is dipole forbidden. This is shown in Fig 13.

Optical Pumping proceeds with two π polarised top-hat pulses which have different detunings from the $F = 1 \rightarrow F' = 2$ and $F = 2 \rightarrow F' = 2$ transitions, as well as different peak Rabi frequencies. Numerical optimisation with these control parameters was performed to determine the optimal population efficiency to the desired initial state $F = 2, m_F = 0$. The simulation results exhibited low pumping efficiencies for resonant pulses due to the creation of two-photon resonances which would repopulate other states of the $F = 2$ level, due to the formation of unwanted dark states formed by indi-

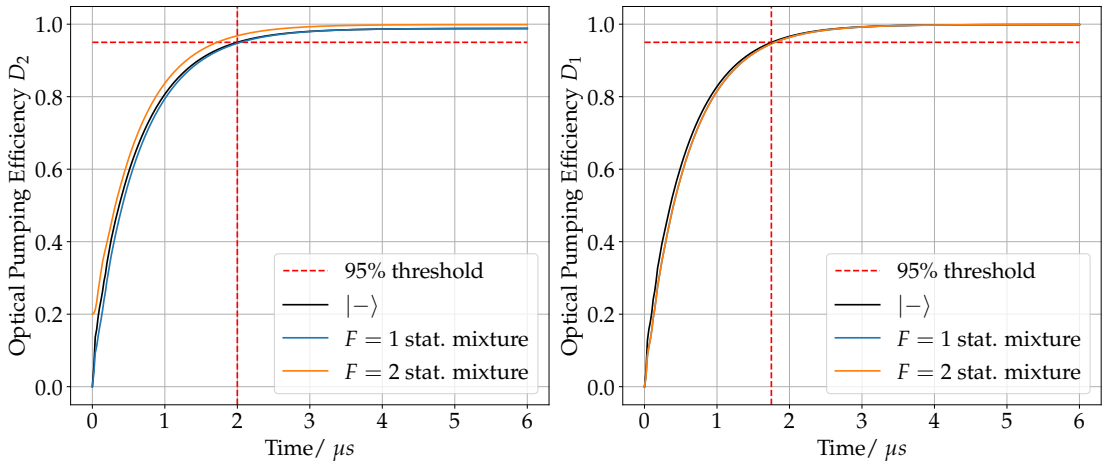

Figure 20: Optical Pumping Efficiency plotted as a function of time for various initial conditions for the D_2 and D_1 transitions. The black trace corresponds to an initial state corresponding to the coherent superposition of the $F = 1$ groundstates $|\Phi_\pi\rangle$(cf. Eq. 42), the green and blue traces correspond to statistical mixtures of the states of the $F = 1$ and $F = 2$ levels respectively. Re-preparation efficiency depends on the cost one is willing to incur with respect to the repetition rate, the 95% threshold is indicated with the dotted red line, and it is approached for times just below 2μs for D_1 which is slightly faster than D_2 due to a reduced number of excited levels reducing the potential for crosstalk.

vidual sub states of $F = 1$ and $F = 2$. Optimal detunings are determined to be $\Delta_1 = 4$ MHz and $\Delta_2 = -7.5$ MHz, for the $F = 1$ and $F = 2$ transitions respectively.

Both the D_1 and D_2 lines were considered for repumping. The optimal peak Rabi frequencies for the D_1 transitions were $\Omega_1/2\pi = d_{(1,m_i)\to(2',m_i)} \times 34$ MHz which addresses $F = 1$ and $\Omega_2/(2\pi) = d_{(2,m_i)\to(2',m_i)} \times 24$ MHz which addresses $F = 2$. The optimal peak Rabi frequencies for the D_2 transitions are $\Omega_1/2\pi = d_{(1,m_i)\to(2',m_i)} \times 57.5$ MHz which addresses $F = 1$ and $\Omega_2/2\pi = d_{(2,m_i)\to(2',m_i)} \times 25.5$ MHz which addresses $F = 2$ and where $d_{i,j}$ describe the Clebsch-Gordan coefficients of the transition. The difference in the numerical factors can be attributed to the difference in Clebsch-Gordan coefficients $d_{i,j}$ of the two F' levels.

The efficiency with which initialisation or re-preparation with optical pumping occurs over different time scales and various initial conditions is shown in Fig. 20 for the D_1 line. This is more efficient than using the D_2 transitions, which has additional excited atomic levels that adversely affect the population transfer efficiency. In principle, one can achieve arbitrarily high efficiencies with extremely long pulses. It should

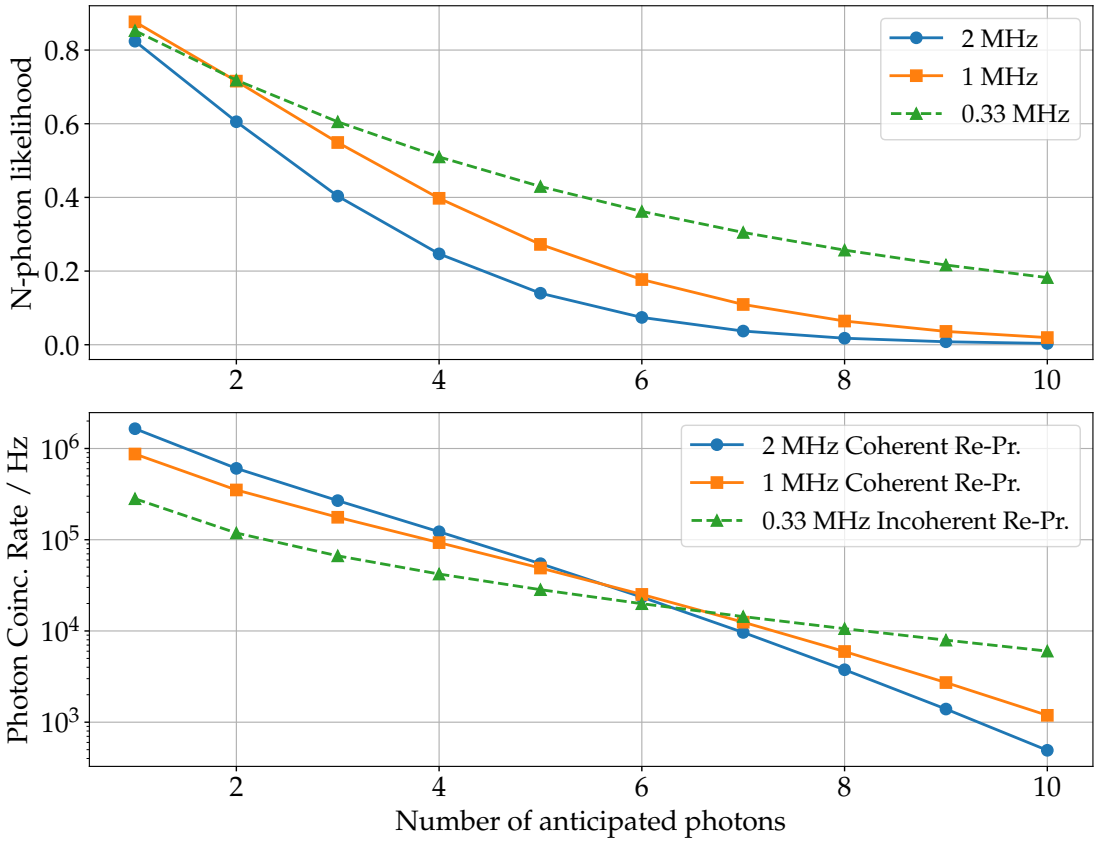

Figure 21: Polarised photon stream production likelihood (top) and expected N-photon sequence coincidence rate (Efficiency × Single Photon Repetition Rate/N), where N denotes the number of photons (bottom) for existing cavity design [67]. Slower driving procedures with longer photons or incoherent re-preparation methods lead to higher coincidence counts for longer photon sequences as the efficiencies are higher and thus the N-photon production efficiencies scale more favourably.

be stressed that the time scale at which optical pumping occurs is about an order of magnitude longer than that of coherent re-preparation because of the random walk of the population back into the desired initial state.

3.2.4 *Full Scheme*

Following the analysis of the previous two sections, the combination of optical pumping, V-STIRAP, and STIRAP with respect to the D_2 and D_1 transitions, respectively, for polarised photon stream production is considered. Simulated efficiencies for generating bursts of polarised photons with recurring cycles of D_2 V-STIRAP photon production and D_1 STIRAP re-preparation from a state of the art optical cavity [67] and ^{87}**Rb**

atoms - which for simplicity are assumed to be stationary [13] - in Fig. 21 are shown. It should be noted that the proposal outlined in Sec. 3.1 and applied to ^{87}**Rb** describes a perfectly degenerate state configuration, yet perfect degeneracy is nearly impossible to achieve due to small background B-field fluctuations. One can mitigate this by imposing a small constant external field along the preferred quantisation axis (i.e. y) which induces a hyperfine state splitting smaller than the line-width of the transition or the cavity and inducing a Larmor precession the timescale of which is $<< T$ where T defines the photon duration. This makes the scheme relatively robust to small magnetic field fluctuations and a detailed analysis is provided in Fig. 14.

The repetition rate of the V-STIRAP and STIRAP sequence can be altered by modifying the duration of either or both. In what follows the coherent STIRAP re-preparation duration is fixed to 140ns. Shorter re-preparation pulses lead to reduced efficiencies without significantly improving the overall photon production repetition rate and longer pulses only marginally increase efficiency. However, two different V-STIRAP photon driving pulses of length 360 and 860 ns are considered, which lead to overall repetition rates of 2 and 1 MHz respectively. Longer driving leads to higher single photon production efficiencies, as evident in Fig. 21, yet the expected number of single photon counts in a particular time interval is lower. However, when the photon burst length exceeds six or more consecutive photons, the expected count rate is larger for the photons which are driven more slowly at a repetition rate of 1 MHz.

Neither the V-STIRAP photon production, nor the STIRAP re-preparation are perfectly efficient. Thus, it is important to characterise the effects of decoherence on the produced photon stream. Atomic depolarisation during the V-STIRAP procedure may

13 For the sake of simplicity, atoms in the simulation are assumed to be stationary. In the case of atoms in free flight or trapped in a conservative trap, noting that each V-STIRAP/STIRAP cycle imparts a maximal momentum change of $3\hbar k$ on each atom (in reality this will be much lower), thus increasing the temperature of a trapped ^{87}**Rb** atom by approximately 538nK and leading to a corresponding Doppler shift of approximately 22.48 kHz. This might reduce the efficiency of the process if the burst length is of the order of hundreds of photons, but it is negligible for the photon burst lengths described here.

$F=2, m_F=$	-2	-1	0	1	2
$p(H)$	0.188	0.158	0.853	0.158	0.188
$p(\pi)$	0.049	0.710	0.009	0.710	0.049
Purity	0.8	0.182	0.99	0.182	0.8

Table 2: Probability of emitting a photon of particular polarisation, for a cavity with $C \approx 10$ [67] when applying the optimised driving pulse when correctly initialised (i.e. $m_F = 0$) and when incorrectly initialised into the other sub-states. The purity, i.e. $p(H)/(p(H)+p(\pi))$ is also shown and decreases significantly for the $m_F = \pm 1$ sub-states.

manifest itself, such that the atom is not prepared into a superposition of the states of the $F = 1$ level. This adversely affects the coherent STIRAP re-preparation efficiency which only correctly re-prepares the atom into the initial state if starting from the superposition state $|\Phi_\pi\rangle$. Moreover, atomic depolarisation during STIRAP re-preparation can similarly lead to incoherent decay of the atom into the wrong Zeeman state. Simulations for a state of the art optical cavity [67] and optimal pulse parameters show that for 10 consecutive photon production attempts with STIRAP coherent re-pumping, a non-negligible accrual of population in other states of the $F = 2$ manifold is unavoidable. The population of undesired Zeeman states however leads to a small efficiency gain with respect to the production of photon sequences with fixed linear polarisation. Even when the atomic population falls to the wrong Zeeman state after re-preparation, one obtains a useful photon (of desired H polarisation). This is exemplified in Tab. 2 as for instance when starting in the state $F = 2, m_F = -2$ the probability of emitting an H-polarised photon is 0.188 as exemplified in Fig. 22.

$F=2, m_F=$	-2	-1	0	1	2
$n=1$	0	0.01	0.878	0.01	0
$n=10$	0.121	0.078	0.44	0.078	0.121
$n=100$	0.129	0.109	0.358	0.109	0.129

Table 3: Population after several repumping cycles.

It was important to analyse the distribution of of population after several V-STIRAP & STIRAP repumping cycles to consider the resulting photon polarisation when a V-STIRAP pulse adresses the atom, after it has been re-prepared to the wrong initial

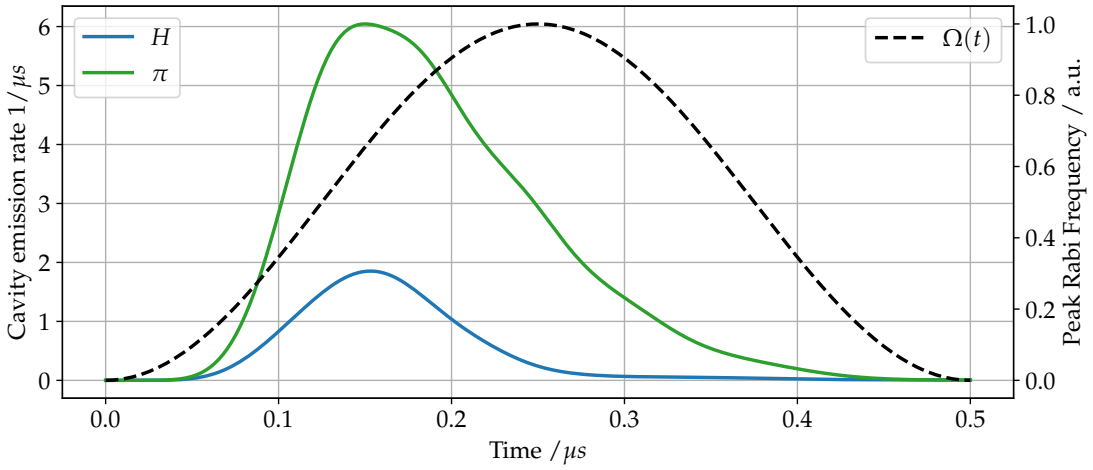

Figure 22: V-STIRAP photon polarisation when the atom is initially falsely prepared into the state $F = 2, m_F = -1$, since the π transition is no longer dipole forbidden.

state. Tab. 3 shows the distribution of atomic population in the $F = 2$ level after several V-STIRAP and STIRAP cycles. Significant accrual in undesired states adversely affects the efficiency for larger photon chains and since the π transitions are no longer dipole forbidden for the $m_F = \pm 1$ states, polarisation errors are introduced, as exemplified in Fig. 22. This can however be experimentally mitigated by filtering conditional on photon polarisation at the cavity output. The optimal parameters for the simulations showcasing expected generation efficiencies for various photon chains, as well as the expected count rates in Fig. 21 for an existing cavity design [67], are as follows. $\Omega_{\text{vst}} \approx 105$ and $69 \times 2\pi$ MHz for V-STIRAP pulses of duration 0.36 μs and 0.86 μs respectively, with no detuning with respect to the $|F = 2, m_F = 0\rangle \rightarrow |F' = 1, m_{F'} = 0\rangle$ transition and the cavity is resonant with the $|\Phi_\pi\rangle \rightarrow |F' = 1, m_{F'} = 0\rangle$. The optimal pulses are as shown in Fig. 19 for re-preparation via the D_1 transition.

Practically, undesired π polarised photons would be detected at the cavity output with a polarising beam splitter, indicating the failure of a particular coherent photon productions sequence. This could be used to immediately trigger incoherent re-preparation with optical pumping before initialising another coherent photon production sequence. Thus, the detection of a photon of the wrong polarisation would enable real time error correction on the atomic single photon source. This is a promising insight beyond the use of atoms as sources of highly indistinguishable photons. In

neutral atom quantum computing, fast and efficient atomic state readout is required for error correction and remains as one of the main challenges with regards to scaling neutral atom arrays [102, 193] and improving the clock times of logical quantum operations [92, 168]. Mid circuit measurements have been demonstrated, including cavity assisted readout, [62] but error correction with cavity assisted readout has not been demonstrated in neutral atom arrays.

One can also consider a possible sequence where the atom is re-prepared entirely incoherently with optical pumping lasting for 2.5 μs and with the V-STIRAP process lasting for 500ns, yielding a single photon repetition rate of 0.33 MHz. This has the advantage of achieving repumping efficiencies close to unity for longer time scales, regardless of the initial atomic state after photon production. This of course incurs other costs; the repetition rate is significantly reduced and the atom does not maintain coherence throughout the process which precludes the performance of quantum information processing tasks with the atom as an entanglement mediator (cf. [223]). However, the likelihood to obtain counts for longer photon sequences (> 7) exceeds those achievable with coherent re-preparation.

3.3 DISCUSSION

It is instructive to contextualise the results demonstrated in this section in the existing literature and to compare the achievable photon generation rates to other physical platforms, these could be other deterministic sources such as ions or quantum dots, or non deterministic sources like spontaneous parametric down conversion. Two important metrics are the repetition rate and the brightness respectively. The single photon repetition rate is defined as the frequency at which single photons can be produced from a single source[14], this includes the photon duration and hence typically those

14 With a caveat for multiplexed sources, see [117] where several sources are combined to effectively form a single source.

platforms which emit photons which are shorter in duration also reach much higher repetition rates. The brightness is defined as the conditional probability of measuring a single photon, given that an emitter has been appropriately excited. Compared to non-deterministic physical platforms like SPDC sources [252], atom-cavity sources do not reach comparable repetition rates. Nevertheless, the proposal demonstrates the ability to efficiently generate highly indistinguishable photons at the MHz scale. The high single photon production efficiencies of deterministic sources lead to a favourable scaling behaviour for longer photon sequences. The use of a single photon source for n-photon burst production is limited to a scaling of p_1^n at best, where p_1 represents the single photon production efficiency and n represents the photon burst length. Consequently, when dealing with non-deterministic sources, which frequently exhibit low p_1, the probability of larger n-photon burst production rapidly decreases. Besides atomic sources, other deterministic single photon sources offer similar advantages. Recently, much progress has been made in demonstrating efficient and fast photon production with quantum dots [147, 213, 227] which have shown the best trade-off of high brightness as well as repetition rate. A comprehensive comparison of different sources is provided in Fig. 23.

The scheme described in this chapter shows that bursts of 10 consecutive photons can be generated with a production efficiency of up to around 2% and 20%, respectively, for coherent and incoherent re-preparation methods. Moreover, 10 photon counts rates on order of kHz represent a significant improvement from any previous photon generation scheme based on atom cavity sources [8, 34, 131, 165].

It shall be noted that the source comparison does not take into account the photon indistinguishability, which is indispensable when using single photons for quantum information processing tasks. For example to ensure computational hardness of boson sampling, one requires near unity photon indistinguishability and lower photon indistinguishability imposes stricter constrains on overall optical loss tolerances and thus maximal photon indistinguishability is desired [219]. This further complicates the

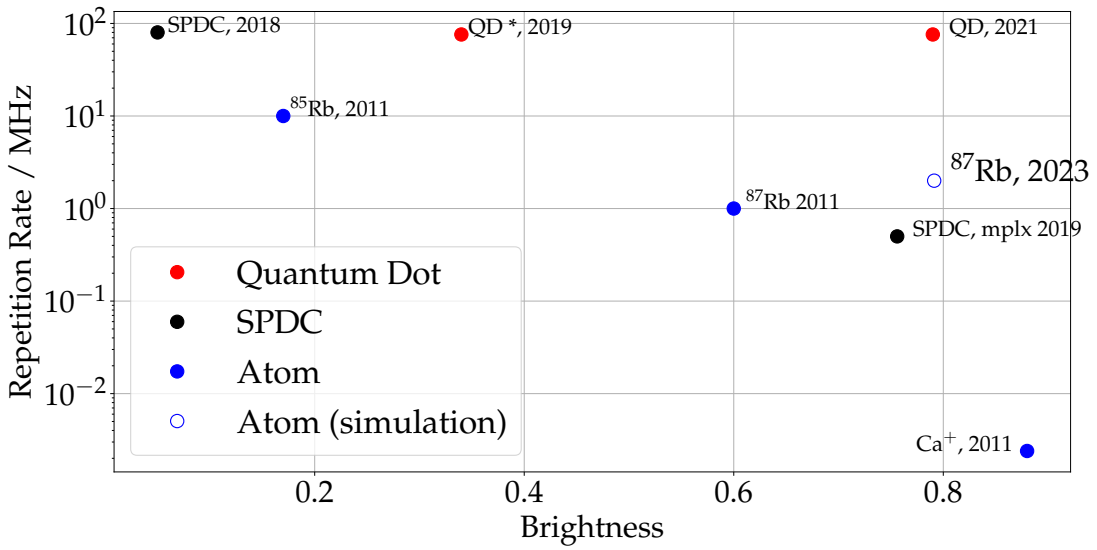

Figure 23: Explicit comparison of various single photon sources by repetition rate and brightness. The brightness is defined as the conditional probability of measuring a single photon when an excitation pulse addresses the source, but neglecting all optical losses and detection inefficiency (except *) to warrant a comparison with simulation. Benchmarks against some state of the art single photon sources are shown; including spontaneous parametric down conversion ([252] and [117] for a multiplexed SPDC source), quantum dots ([147] (*with optical losses), [227]), as well as atomic (ionic) sources ([21] for an ionic source, [34] for ^{87}**Rb** and [8] for ^{85}**Rb**).

source comparison given in Fig. 23; atom cavity sources tend to offer extremely good indistinguishability of up to 98% reported in [18], while quantum dots have historically struggled with spectral drifts. However, recently a similar indistinguishability value of 97.5% was reported in [227].

3.3.1 *Requirements for Quantum Advantage with Boson Sampling*

Boson sampling [2] is a particularly interesting application to consider for the single photon production scheme developed in this chapter because it offers a feasible demonstration of *quantum advantage* [253] using current photonic technology, and advances experimental techniques with potentially broader applications in cryptography, graph theory and drug discovery [12, 16, 162, 163, 210, 248]. It entails sampling the output photon distribution after photons interfere on an optical multimode interferometer. In what follows, we follow the detailed analysis provided in [219], but consider an atomic

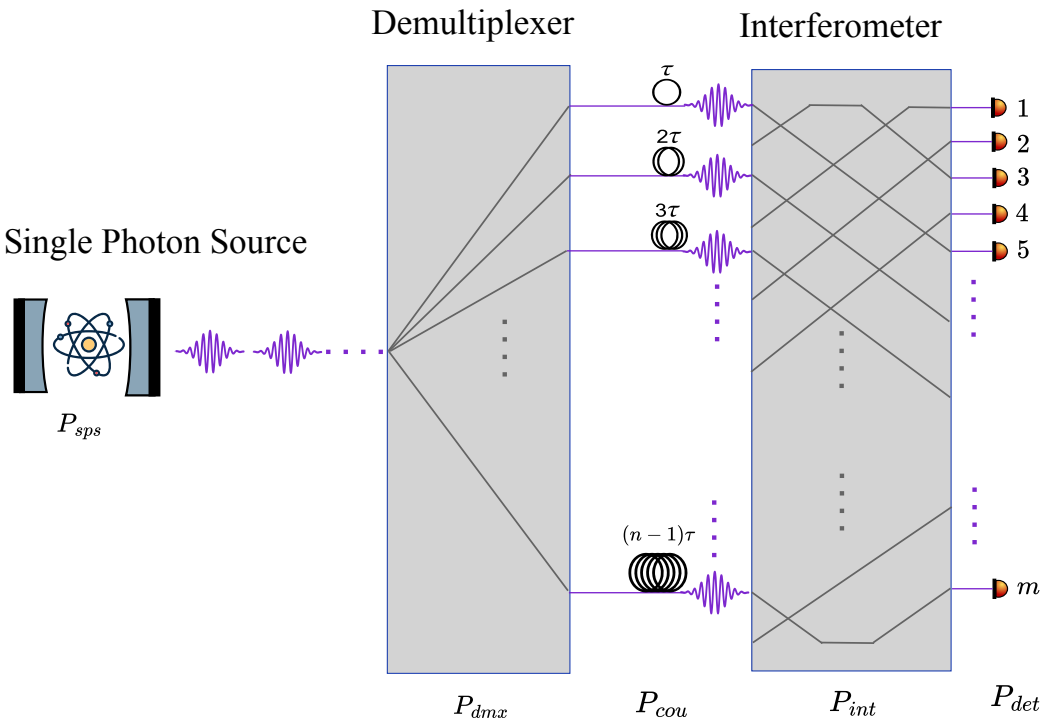

Figure 24: Simplified boson sampling setup with single photon source, demultiplexer, interferometer and detection setup. All efficiencies P are given the efficiency per photon (cf. Eq. 43). This includes a single photon source efficiency of P_{sps}, demultiplexer efficiency of P_{dmx}, coupling efficiency into and out of the demultiplexer of P_{cou}. A universal interferometer [57] is depicted with $n-$input photons and m output modes, which in this figure is equal to the number of input photons n (this may be different for other interferometer architectures). The efficiency through the m interferometers is $P_{int} = mP_{mzi}$, and the detection efficiency P_{det}. Delay lines of variable length are included to ensure simultaneous photon arrival at the interferometer input as a stream of n photons is produced sequentially from a single atomic source.

photon source instead of a quantum dot. A typical boson sampling setup with a single photon source comprises of a multimode interferometer with multi-photon input, preceded by a photon demultiplexer, followed by single photon detectors, as shown in Fig. 24. The computational hardness of simulating boson sampling arises from the connection between multi-photon interference and the computation of matrix permanents (cf. App. Sec. A.2 for more details). However, this complexity diminishes in instances where duplicate rows or columns emerge due to multiple photons occupying identical input or output modes. To maintain computational complexity, it is crucial to ensure collision-free states at both the input and output stages of the boson sampling protocol, meaning each mode accommodates no more than one photon. In practice, especially with imperfect photon sources, this requires a significant number of modes per photon, scaling at least quadratically with the number of photons. This phenomenon is often referred to as the bosonic birthday paradox [11].

To demonstrate quantum advantage, it is necessary to detect a sufficient number of output samples that can be validated as computationally hard to generate classically [46]. This requires considering the classical algorithms available for simulating the interference process. In the proceeding analysis, a particular variant of boson sampling known as Aaronson-Brod sampling [3] is focused on, as discussed in [219]. This variant is designed to be more robust against photon losses, increasing the rate at which collision-free boson sampling events are experimentally generated by raising the number of input photons beyond the number one seeks to detect and post-selecting on all detection events. This approach reduces computational complexity. However, the classical algorithm required to simulate noisy boson sampling (see App. A.2) depends on the photon indistinguishability, denoted x^2. As long as this indistinguishability approaches unity and the post-selected photon number is large enough relative to the additional input photons, the classical algorithm's approximation error remains bounded. This motivates the use of Aaronson-Brod sampling for demonstrating quantum advantage in realistic, lossy experiments. Additionally, sufficient collision-free samples need to be

acquired from the boson sampling experiment to allow statistical tests to verify that the experimentally measured distribution differs from a set of classical distributions [219]. Following [219] the desired number of samples is set to 100. The sample acquisition rate can be defined as the product of the multi-photon input rate $r_n = r_1/n$ (which is determined by the repetition rate of the source r_1 for n photons) and the single photon system efficiency P_{ph}. The single photon system efficiency is defined as [219]:

$$P_{ph} = P_{src}P_{cou}P_{dmx}P_{int}P_{det}, \tag{43}$$

comprised of single photon source source efficiency, coupling efficiency (both in and out of the demultiplexer), demultiplexing efficiency, interferometer efficiency, as well as single photon detection efficiency. These numbers can also be expressed in decibles as losses where $\mathcal{L} = log_{10}(P)$. Limiting sample acquisition to a day for 100 total samples gives a desired acquisition rate of $r_{sample} = \frac{r_1}{n}P_{ph} \approx 0.00116$Hz [15]. This sets a constraint which will be used in an explicit hardware requirement calculation exhibited in Fig. 25. In this analysis, a universal architecture introduced in [57] is considered. The optical depth of this architecture is equal to the number of input and output modes, such that each photon in one of the m-input modes passes through m Mach–Zehnder Interferometer. The optical loss thereafter scales as $P_{int} = mP_{mzi}$ where P_{mzi} is the loss associated with a single interferometer. There exist a variety of interferometer architectures which are non-universal and therefore offer interferometer losses $< mP_{mzi}$, but this goes beyond the superficial analysis provided in this section. A detailed comparison of various interferometer architectures is provided in [219].

Using the two constraints of a sample acquisition of 100 samples per day, as well as a lower bound on the error of a classical approximation [188] of 10^{-3} one can calculate the per photon efficiency (cf. Eq. 43) required for a demonstration of quantum advantage which is shown in Fig. 25 [16]. Assuming an optimistic repetition rate of 2MHz for

15 This is also broadly consistent with the experimental sample acquisition rate demonstrated in [240].

16 The code for the calculation of the quantum advantage bounds was originally written by Patrick Sund [219] for quantum dot sources and it was adapted to the photon production scheme described in this

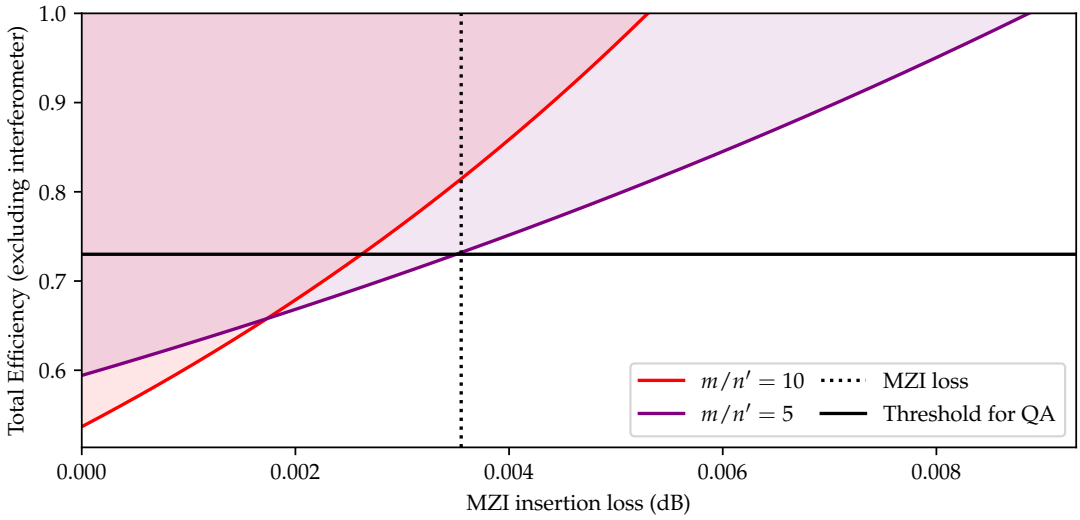

Figure 25: Using an indistinguishability of $x^2 = 0.98$ and a repetition rate of 2MHz for the atom-cavity scheme the requirements for the total system efficiency P_{sys} comprised of source, demultiplexer, coupling and detection for realising quantum advantage are plotted for a universal interferometer architecture, against the insertion loss of the Mach–Zehnder Interferometer (MZI). The plot shows different interferometer dimensions m for an input photon number n of 59 and a post-selected photon number n' of 50. This number n is chosen, such that the error of the classical noisy boson sampling simulation algorithm has an error which exceeds 10^{-3}. (cf. App. A.2). The state of the art MZI insertion loss (black dotted line) is plotted vertically and one can infer that the total source efficiency required is ≈ 0.73.

the scheme described in this chapter (photon duration 360 ns and re-preparation duration of 140 ns), as well as an experimentally demonstrated indistinguishability of 0.98 [18] with an atom cavity source and state of the art MZI loss of $l \approx 0.00355$ dB [240] the requirements for total system efficiency per photon are $P_{ph}/P_{int} \approx 0.73$ (excluding interferometer losses). Currently, this requirement exceeds what is achievable with the scheme outlined in this chapter, which demonstrates $P_{sps} \approx 0.79$ [17] at a repetition rate of 2 MHz, or any other known atom-cavity scheme, as state-of-the-art demultiplexing, coupling, and detection efficiencies amount to approximately 0.711 [219]. One can reduce the required system efficiency threshold with a non-universal interferometer architecture (see. [219]), reduce the number of desired daily samples or consider multiplexed sources with multiple atom cavity sources operating in parallel.

chapter. I am grateful for our discussions and his openness which proved very helpful for writing this section.

17 This includes coherent re-preparation efficiency.

However, there is reason to be hopeful, as the dynamics of the atomic photon source are not a fundamental bottleneck. Rather, the limiting factor are rather the demultiplexer, coupling, detection and interferometer losses put together which are likely to improve with improved engineering of low loss photonic components as demonstrated in [180]. Moreover, if one reduces the number of desired samples per day from 100 to 10 in the aforementioned calculation, one requires $P_{ph}/P_{int} > 0.695$ for $m/n' = 5$ which is achievable with current state of the art components and the single photon production scheme described in this chapter (≈ 0.711).

3.3.2 *Outlook*

In this chapter, a novel scheme for polarisation controlled single-photon emission from atom-cavity systems [74] was described and it was shown through extensive numerical simulations that it might be used for the production of n-photon bursts in a deterministic manner. Combined with coherent re-preparation, expected success and coincidence rates are up to around 2% and 20% respectively for coherent and incoherent re-preparation methods, as well as $\mathcal{O}$(kHz) for 10 photon bursts. This represents a significant improvement over any previous atom cavity schemes. The simulated single photon production efficiency with a state of the art optical cavity [67] reaches $\approx 85\%$ [18] which to the knowledge of the author lies beyond any existing single photon production efficiencies achieved to date with atoms at a comparable repetition rate (similar efficiencies are demonstrated with ions [21], yet at a repetition rate several order of magnitudes lower).

This is essential for demonstrating multi-photon coincidences with high clockrates [158] for quantum information processing applications. It follows from the discussion of section 3.3.1 that with optimistic assumptions and state of the art parameters, a demonstration of quantum advantage with boson sampling with an atomic single pho-

18 This does not include coherent re-preparation efficiency.

ton source is challenging but possible, if only few sample measurements demonstrating quantum advantage are required. The comparatively slow repetition rate imposes stringent requirements on the per photon efficiency.

Preserving coherence of the emitter throughout the entire process is an extremely appealing feature. It has been demonstrated in [223] that for optical GHZ and cluster state generation with more than 10 photons, atoms outperform quantum dots [140, 204] or state of the art spontaneous parametric down conversion sources [252]. Creating large highly entangled optical states is instrumental for the realisation of measurement based quantum computing [42]. This will be explored more extensively in the next chapter.

Moreover, improved cavity designs with higher atom-cavity coupling (e.g. fibre tip cavities [19] cf. [110]) would allow for single photon production efficiencies $> 90\%$. It should however be noted that, even with increased atom-cavity coupling strength, the high peak Rabi frequencies would lead to increasingly severe off-resonant coupling errors in ^{87}**Rb**, especially for the D_2 line which have been discussed in Sec. 3.2.4. In addition, the increased laser intensity of the V-STIRAP driving pulse could also induce differential Stark shifts which would have to be compensated with frequency chirping. Thus, even if $\eta = \frac{2C}{2C+1} \to 1$, other challenges are likely to materialise. In addition to the extensive exploration of ^{87}**Rb**, it would be interesting to consider applying the scheme to other neutral atoms (e.g. Caesium [152]) or ions (e.g. Calcium [21], Strontium [217]), which could prove more favourable for a cavity based single photon production scheme.

19 There are significant practical challenges like mechanical instability and mode matching issues associated with realising such resonator designs and this is an ongoing research effort.

4

TIME-BIN ENTANGLED PHOTONIC STATE SCHEME

Entangled resource states, such as Greenberger–Horne–Zeilinger (GHZ) states [93] and cluster states, are central to many quantum information processing tasks. GHZ states provide a striking demonstration of quantum non-locality and have applications in quantum metrology and secret sharing protocols [104]. Cluster states are highly entangled multipartite states that serve as the foundational resource for measurement-based quantum computation (MBQC) [44]. One way quantum computing [186] or measurement based quantum computing are appealing alternative paradigms to the more widespread gate based model. In the one way quantum computing model, quantum operations are implemented as sequences of adaptive single qubit measurements on an entangled resource state prepared in advance. This has advantages, particularly for photonic architectures, as it more efficiently uses entanglement in computation, as the implementation of high fidelity two qubit gates in the optical domain is an open challenge and heralded two qubit gates lead to a large clock cycle overhead for larger systems. The high rate and high fidelity preparation of large entangled resources, however, remains a central challenge to realise measurement based quantum computing.

Probabilistic sources, like spontaneous parametric down conversion have shown promise [252] but show severe scaling limitations for preparing larger resource states due to the limited intrinsic entanglement generation efficiencies of a single entangled pair. Deterministic schemes have emerged with two particularly promising physical platforms; quantum dots [111] and single atoms [223].

Atom-cavity sources are versatile platforms for producing photons with long coherence times (up to hundreds of nanoseconds [164]) and temporally shaped envelopes. However, typical atom-cavity experiments cannot compete with quantum dots as sources

of photons streams with regard to repetition rates or in their potential for miniaturisation and scalability. Each atomic source requires a vacuum chamber as well as a substantial electronic and optical control overhead. Proposals exist for combining several sources inside a single vacuum chamber, but this remains an experimental challenge [67]. Remarkable progress has been made in integrating quantum dots with micro-pillar cavities as exemplified in [147, 227]. Previously, spectral drift precluded indistinguishability across individual sources, but theoretical work has shown that an indistinguishability of 97.5% can be combined with near unity efficiency [232], all at a repetition rate of several tens of MHz. This raises the question of what advantages atom-cavity systems still offer over quantum dots as deterministic single-photon sources. The answer lies in the fact that the coherence times of atoms exceed those of electron spins or holes by many orders of magnitude. Such emitters act as entanglement mediators in the generation of entangled photons and thus directly affect their properties. It was shown in [223] that for large entangled states ($n > 10$), atom-cavity sources are superior to any other physical platform in terms of the measured coincidences and achieved fidelities.

The scheme realised in [223] used a $\sigma^{\pm}$ polarisation encoding for GHZ and linear cluster states. This is the natural choice of encoding for a quantisation axis which lies along the cavity axis. The drawbacks of a polarisation encoding are the dynamic decoupling of the atomic states which introduce a time-dependent polarisation rotation on the emitted photons, as well as the possibility of cavity birefringence affecting the fidelity of the produced states as it splits the polarisation modes in energy. It is shown in [119] that even modest cavity birefringence can be detrimental to remote entanglement performance for integrating multiple sources to form a networked architecture. This strongly motivates the search for an alternative to polarisation-encoding for distributed entanglement generation. Therefore, time-bin entanglement is explored as a viable alternative to more widespread and conventional polarisation encoding.

In this chapter, it is shown the reported experimental fidelities and measured repetition rates in [223] could be further improved by introducing a time-bin encoded photon entanglement scheme which relies on a different quantisation axis, as well as a modified re-preparation sequence. Extensive simulations [1] with realistic experimental parameters and limitations underpin the merits of this approach. For the remainder of this chapter, all simulations use parameters from an existing cavity design described in [67] and use the full structure of the D_1 and D_2 lines of ^{87}**Rb**.

4.1 SCHEME OVERVIEW

As in Ch. 3 we consider an atom-photon system described by the following Hilbert space:

$$\mathcal{H} = \mathcal{H}_{\text{atom}} \otimes \mathcal{H}_{\text{photon,}} \tag{44}$$

where there are two orthogonal photon polarisation modes π, H which are Fock states with number occupation n_π, n_H. The atomic Hilbert space is comprised of several ground and excited states and it is the atom which mediates the entanglement between subsequently emitted photons. The following analysis considers a single ^{87}**Rb** atom assumed to be trapped by an optical dipole trap inside an optical cavity as in Ref [223]. A weak bias field is assumed to be set up which lies perpendicular to the cavity axis to fix the preferred quantisation axis. As outlined in Fig. 7 in Ch. 2, this describes the basis transformation from the atomic polarisation basis to the Cartesian lab basis in which photons are measured. The field and the driving pulse are chosen such that all emission occurs into the π polarisation mode. This is akin to the choice of quantisation axis described in Ch. 3. As for the scheme in Ch. 3, a reasonable external field strength induces a 0.07 MHz ground state splitting to provide a fixed quantisation axis, larger

1 All relevant code is available at https://github.com/jan-o-e/rb_photon_prod

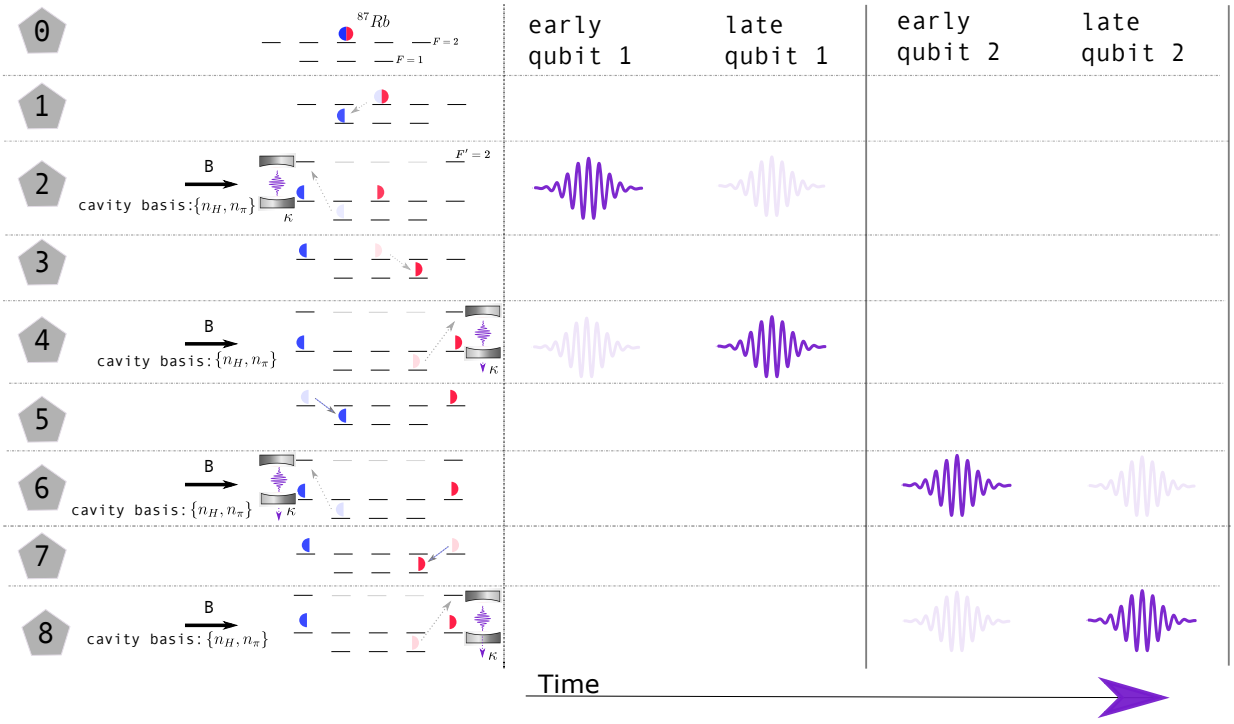

Figure 26: Generalised Driving Scheme for producing photonic GHZ states. The atom is initialised in $|F = 2, m_F = 0\rangle$ (0) with optical pumping and then a series of alternating Raman rotations $(1, 3, 5, 7)$ and V-STIRAP photon emissions $(2, 4, 6, 8)$ create the desired entangled photonic state in time. The first rotation creates the ideal superposition $1/\sqrt{2}(|F = 1, m_F = -1\rangle + |F = 2, m_F = 0\rangle)$, i.e. a $\pi/2$ pulse and subsequent rotations $3, 5, 7$ maintain that phase relationship and implement π pulses to completely transfer population between the desired sub-states. Photons, produced during steps $2, 4, 6$ and 8 are coupled into the same polarisation mode of the cavity (π) and the pulses are sequenced such that early and late emissions always occur in the time intervals $[(2n-2)T, (2n-1)T)$ and $[(2n-1)T, (2n)T)$ where $T = t_{\text{vst}} + t_{\text{raman}}$ and n the number of photons. Here the generation of a two photon GHZ state is shown, but the steps $5 - 8$ can be repeated $n - 1$ times for an n-photon GHZ state. The pulses are sequenced carefully so there is no possibility of scattering on the idling half of the atomic superposition (alternating blue and red) by a laser pulse pair acting on the other half of the atomic superposition (red and blue respectively).

than any background field noise, but weak enough to avoid breaking any symmetries in the ground state manifold for the early and late photon emissions.

Following the insights of Ch. 3, re-preparation pulses are implemented via the D_1 line, due to its simplified level structure, and V-STIRAP photon emission is implemented via the D_2 line due to its stronger dipole coupling. Moreover, all the analysis and the simulations in this chapter use the parameters of a state of the art optical micro-cavity [67].

The qubit basis encoding is picked in time, following work with quantum dots [10, 25, 135], such that the emission time encodes the qubit basis state. The notation used for the computational qubit states is $|0\rangle \Leftrightarrow |e\rangle$, for an early emission and $|1\rangle \Leftrightarrow |l\rangle$, for a late emission. A running clock, with respect to the pulses driving the production of single photons is set up, which partitions emissions into early $[0, T)$ or late $[T, 2T)$ time intervals for a single qubit, where T is the time-bin width. For n-qubits or photons this generalises to $[(2n-2)T, (2n-1)T) \Leftrightarrow |e\rangle$ and $[(2n-1)T, (2n)T) \Leftrightarrow |l\rangle$.

Entanglement generation is mediated by the atomic emitter. Raman pulses can be used to (re)-prepare the atom in the desired $F = 1, 2, m_F = \pm 1, \pm 2$ ground-states (cf. Fig. 26) and a series of these Raman pulses and V-STIRAP photon production attempts are used to create a photonic GHZ state of the form:

$$|GHZ\rangle_n = \frac{1}{\sqrt{2}}(|g_e\rangle\,|e\rangle^{\otimes n} + |g_l\rangle\,|l\rangle^{\otimes n}), \tag{45}$$

where the $|g_{e,l}\rangle$ correspond to the atomic states associated with the emission of early ($|e\rangle$) or late ($|l\rangle$) photons. Each of the n photons is emitted at a well-defined time corresponding to either the early or late component of a time-bin qubit. Care must be taken in interpreting this state: the photons are not emitted simultaneously, but rather sequentially in time. Thus, the state (45) is not defined at a single instant in any global reference frame. Instead, it is a history-dependent entangled state constructed with

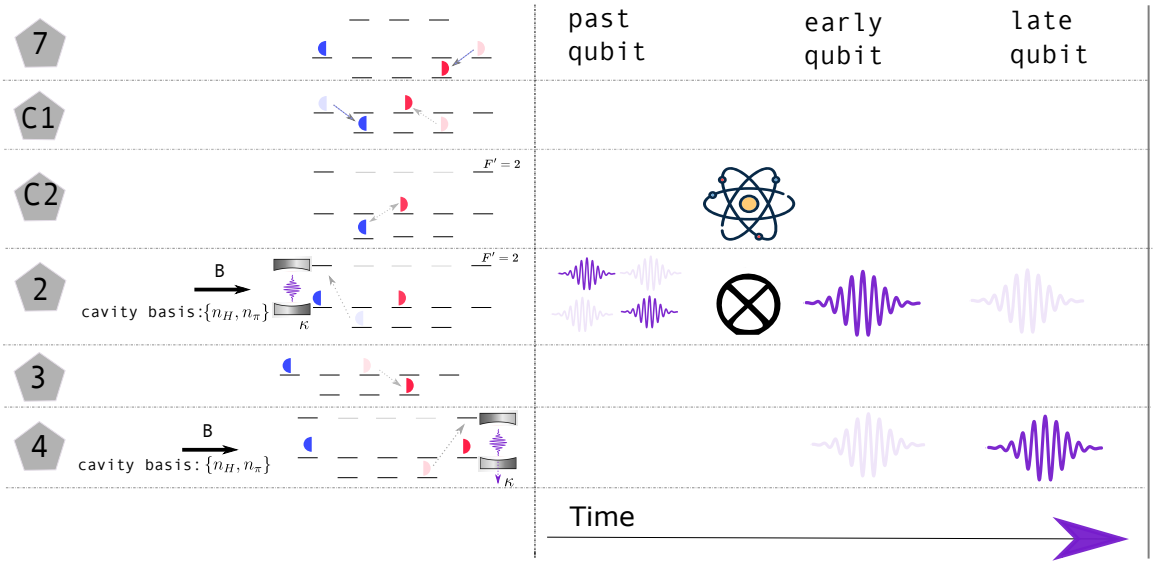

Figure 27: The scheme for realising photonic linear cluster states is shown (left) as an extension to the GHZ state scheme shown in Fig. 26. Starting from after step 4 in Fig. 26, an $X(\pi/2)$ rotation is realised between the atomic states corresponding to early and late emission, after first applying a Raman pulse (7 as in Fig. 26) which transfers population from $|F = 2, m_F = 2\rangle$ to $|F = 1, m_F = 1\rangle$, another rotation (C1) transfers both superposition states to states which can be directly jointly addressed. (C2) then realises a $\pi/2$ rotation to implement the desired relative phase change between subsequently emitted photons. The pulse durations are synchronised such that $t_7 + t_{C1} + t_{C2} = t_3$. The following pulses (2,3,4) are the same as in the GHZ state scheme. The right part of the plot exemplifies how the atomic emitter mediates entanglement between past and future photonic qubits.

respect to a particular time trajectory (e.g., that of the atomic emitter or the detector). This makes the interpretation of entanglement present in this state frame-dependent[2].

The atom mediates entanglement in this context and the photon–photon entanglement and their relative phase are determined by the interaction with the atomic emitter. Measuring the photonic GHZ state requires disentangling the atom from the photonic state in the last step, such that the following state is obtained:

$$|GHZ\rangle_{n,mx} = |g_e\rangle\,|e\rangle\,(|e\rangle^{\otimes n} + |l\rangle^{\otimes n})/\sqrt{2}. \tag{46}$$

Starting from step 5. in Fig 26, the rotation described in step 1. is driven in reverse, transferring population from $F = 1, m_F = -1$ to $F = 2, m_F = 0$, followed by a complete population transfer from $F = 2, m_F = 0$ to $F = 1, m_F = 1$ which is the transition addressed in step 3. The final photon will thus be emitted only into an early time-bin and one attains the desired n photon state with no remaining atom-photon entanglement.

A linear cluster state is prepared in similar fashion but requires a single qubit rotation $\sigma^x(\pi/2)$ in the atomic basis $(|g_e\rangle, |g_l\rangle)$ between subsequent photon emissions [223]. This allows for the preparation of linear clusters where the initial photon is prepared in the $|+\rangle = |GHZ\rangle_{n=1}$ one photon GHZ state and the $\sigma^x(\pi/2)$ rotation is applied between subsequent photon emissions to realise an effective CZ gate and to prepare an n-dimensional 1-d linear cluster state which is defined as:

$$|C\rangle_n = \prod_{j=1}^{n-1} CZ_{j,j+1}(|g_e\rangle\,|e\rangle + |g_l\rangle\,|l\rangle)^{\otimes n}, \tag{47}$$

where $CZ_{j,j+1}$ applies a relative phase change (-1) to the state $|l_j l_{j+1}\rangle$, leaving other basis states unchanged.

2 The coherent superposition of time-separated events implies that the relevant degree of freedom is not the photon's spatial wave-function at a fixed time, but its temporal mode structure. Consequently, relativistic effects such as Lorentz boosts can alter the interpretation of what is an early and a late time-bin, which are tied to specific detection events in a particular frame. For a more detailed discussion on relativistic invariance and entanglement see [53, 178]

Unlike in the GHZ state case, for the cluster state, a measurement of the first and last photon in the σ^z basis disentangles the atom from the photonic state [143].

4.2 RAMAN PULSES

Feature	**Far-Off-Resonant Raman Rotations** [115, 118, 247]	**STIRAP** [132]
Mechanism	Two-photon transition directly coupling ground states via virtual excited state	Uses temporally interleaved near-resonant pulses to couple ground states via a dark state
Robustness	Sensitive to fluctuations in timing, detuning and laser intensity	More robust to pulse imperfections and amplitude variations due to adiabaticity
Laser Power Requirement	Requires high laser power to suppress spontaneous scattering	Can operate near resonance; lower power required
Excited State Population	Non-zero population in excited state leads to decoherence via spontaneous emission	Dark-state evolution suppresses population in excited state
Addressability in Multi-level Systems	Can be tailored with frequency and polarisation, but challenging in closely spaced multi-level systems	Same as for Far-Off-Resonant Raman Rotation
Gate Symmetry and Flexibility	Can implement arbitrary single-qubit rotations (e.g., $\sigma^x(\theta)$, $\sigma^y(\theta)$ pulses)	Limited operations; cannot perform all universal single-qubit gates needed for cluster state generation

Table 4: Comparison of Far-Off-Resonant Raman Rotations and STIRAP

Producing entangled photonic states requires the preparation of atomic states of the following form (in the hyperfine basis for $^{87}\mathbf{Rb}$):

$$\frac{1}{\sqrt{2}}\left(|F_1, m_{F_1}\rangle + e^{-i\phi_r(t)}\,|F_2, m_{F_2}\rangle\right). \tag{48}$$

Many approaches exist for high fidelity ground state rotations in generalised Λ systems. In general, there is a trade-off between the speed at which such rotations can be carried out and the fidelity with which the desired states can be prepared. Faster rotations require higher peak Rabi frequencies which require higher laser (or microwave) power, increased scattering rates and cross-talk to undesired transitions in multilevel systems. Two particular approaches will be analysed and contrasted in what follows. One follows the original STIRAP proposal [132], as explained in the previous Ch. in Sec.

3.1.5 and considers the application of two near resonant, but temporally interleaved pulses which couples a so called "dark-state", exhibiting no excited state component (cf. Sec. 2.1.3)). Such a technique is robust with respect to parameter fluctuations and since it can be performed close to resonance, it does not require high laser powers. However, it is challenging to address only one particular transition in a multilevel system with closely spaced transitions frequencies, and it is not symmetric with respect to the coupled states so it cannot perform an arbitrary single qubit gate rotation, which is required for the production of cluster states.

A second approach considered here is a far-detuned off-resonant Raman Rotation which was first introduced in [118] for atom interferometry. This relies on the application of two far-detuned beams Ω_1 and Ω_2 with a common two photon detuning of Δ forming a direct coupling between two ground states with an effective Rabi frequency of $\frac{|\Omega_1||\Omega_2|}{2\Delta}$. A detailed derivation of the effective Hamiltonian is provided in App. Sec. A.4. The use of far-detuned beams to implement Raman rotations was demonstrated experimentally in ^{87}**Rb** in [115, 247]. Although this approach is less robust in the face of experimental imperfections, such as fluctuations in amplitude or pulse duration, it allows for the implementation of an arbitrary single qubit gate, which is important for the realisation of cluster states as described in Fig 27. Moreover, composite pulse sequences [49] that decompose a single qubit rotation into many individual rotations can also aid in improving robustness, at the expense of an increased overall rotation duration.

Building upon the lessons learned about high fidelity ground state STIRAP re-preparation pulses from Ch. 3, only the D_1 line will be considered for implementing any of the Raman rotations, due to its comparatively simpler structure compared to the D_2 line.

4.2.1 *STIRAP Rotations*

In this subsection, the implementations of the rotations numbered 1,3,5,7 in Fig. 26 are described. These rotations shall be performed as fast as possible without detrimentally affecting the graph state preparation fidelity. An analytically optimal amplitude evolution with Gaussian masks is described in [233] which is also used in Ch. 3. However, at the timescales of interest (100's of nanoseconds), numerically optimised shapes of this form and correspondingly high peak Rabi frequencies induce non-negligible differential Stark shifts. This does not affect the state populations, but adversely affects the relative phase ϕ for superposition states of the form described in Eq. 48. Maintaining fine control over the phase relationship between the atomic superpositions is important for the preparation of both GHZ and linear cluster states as will be described extensively in the results section. A potential solution is to apply a frequency chirp to the Stokes and Pump pulses to compensate for the differential Stark shift caused by amplitude modulation, which arises from interactions with off-resonantly coupled states in the multi-level atomic structure. A similar Ansatz is presented in [235]. However, the authors consider a 4 level system with a single off-resonant excited level, which does not capture the complexity of a multilevel atom with several additional states that can be coupled to. A detailed calculation is provided in App. Sec. A.3 establishes expressions for the required frequency chirps of every pulse. It should be noted that the chirps are smooth and follow structural similarities to the amplitude modulation guaranteeing that such signals can actually be generated by off the shelf Arbitrary Waveform Generator (AWG)'s and modulators. The required chirp is a sum of the dynamic Stark shifts induced by both the Pump and Stokes beam and ensures that the pulses follow the dark state which has a time evolving eigenenergy. This phenomenon does not occur in isolated three level Λ systems but is a result of the complex multilevel structure of ^{87}**Rb**.

Parameter	**Rot. 1**	**Rot. 3**	**Rot. 5**	**Rot. 7**
Pulse Type	f-STIRAP	STIRAP	STIRAP	STIRAP
Ω_0 / MHz	573.4	754.1	1459	1139
Polarisations (Ω_S, Ω_P)	σ^-, π	σ^+, π	π, σ^+	π, σ^-
Duration / μs	0.163	0.25	0.25	0.25
$\Delta_{F'=1}$ / MHz	-37.9	-18.7	-33.82	-40.28
Pulse Parameters	0.7293	10.32, 2	14.92,6	14.53,6
Fidelity	0.989	0.985	0.99044	0.9948

Table 5: Pulse parameters and notes for optimised STIRAP rotations for GHZ state generation. $\Delta_{F'=1}$ denotes a common two photon detuning of the Stokes and Pump Pulses from resonance with respect to the $F' = 1$ line of the D_1 transition.

The atom is initialised in the $|F = 2, m_F = 0\rangle$ ground state of the $^2S_{1/2}$ hyperfine level with optical pumping, as described in [74] and Ch. 3, as it initially populates a random ground state. After initial state preparation with optical pumping, the sequence proceeds with a carefully timed f-STIRAP [237] (fractional Stimulated Raman Adiabatic Passage) pulse sequence which allows for the coherent preparation of superposition states by inducing only a partial dark state population transfer unlike the typical STIRAP pulses which transfer population entirely from one groundstate to the other. The aim is to go from $|F = 2, m_F = 0\rangle \rightarrow \frac{1}{\sqrt{2}}(|F = 2, m_F = 0\rangle + |F = 1, m_F = -1\rangle)$. For a generalised dark state in an idealised three level system:

$$|\psi_d\rangle = cos(\theta)\,|F = 2, m_F = 0\rangle + e^{i\phi_r} sin(\theta)\,|F = 1, m_F = -1\rangle\,, \tag{49}$$

the fractional-STIRAP pulse requires $\theta = \frac{\pi}{4}$. For this particular rotation, the Stokes pulse still precedes the pump pulse, but to realise an incomplete transfer of population the following limits have to be satisfied [237]: $\lim_{t\to-\infty}\frac{\Omega_P(t)}{\Omega_S(t)} = 0$ and $\lim_{t\to\infty}\frac{\Omega_P(t)}{\Omega_S(t)} = \tan(\theta)$. Thus, for fractional-STIRAP Pump and Stokes Pulse must vanish simultaneously. This condition is satisfied by the following pulses [237]:

$$\begin{aligned}\Omega_P(t) &= \Omega_0 \sin\theta\, e^{-(t-\tau)^2/T^2},\\ \Omega_S(t) &= \Omega_0\, e^{-(t+\tau)^2/T^2} + \Omega_0 \cos\theta\, e^{-(t-\tau)^2/T^2},\end{aligned} \tag{50}$$

where substituting $\theta = \frac{\pi}{2}$ yields the required amplitude evolution for rotation 1.

Thereafter, STIRAP pulses are implemented for rotations numbered 3,5 and 7 which cycle population between different superposition states whilst maintaining a fixed phase relationship between different atom-photon states, to produce an n-photon GHZ state. These are implemented with the technique introduced in [233], as also outlined in Ch. 3. All rotations have free parameters $\Delta, \Omega_0, \vec{p}_i$ where $\vec{p}_i$ is a pulse parameter vector which corresponds to the pulse delay τ for the f-STIRAP rotation and denotes the parameters a and n for the STIRAP rotations (cf. Eq. 40 in Ch. 3). All pulse parameters are optimised numerically for maximising fidelity as shown in Tab. 5. It should be noted that the first and third rotations are less efficient in terms of fidelity, due to the reduced efficacy of f-STIRAP in a multi-level system and, because of cross-talk to the transitions which can be driven from $|F = 2, m_F = 2\rangle$.

4.2.2 *Far Off Resonant Rotations*

Far-detuned Raman rotations allow for a general implementation of any ground state rotations between state pairs which exhibit a dipole coupling to a common excited state. The scattering rate of the excited state scales as $\frac{|\Omega_1||\Omega_2|}{\Delta^2}$. Hence, a larger detuning Δ leads to reduced scattering. Implementing fast rotations is challenging, as it requires $t_\pi = \frac{2\Delta\pi}{|\Omega_1\Omega_2|}$ for square shaped $\sigma^x(\pi)$ pulses. This requires very large Rabi frequencies $|\Omega_1|, |\Omega_2| > \Delta$, which in turn increases the scattering rates, unless $\Delta^2 > |\Omega_1\Omega_2|$. This is the desirable practical regime one should be interested in. To ensure its experimental feasibility, an analysis of the scattering probability, with respect to the required laser power and chosen detuning Δ is analysed in Fig. 28. Using a detuning of 100 GHz one

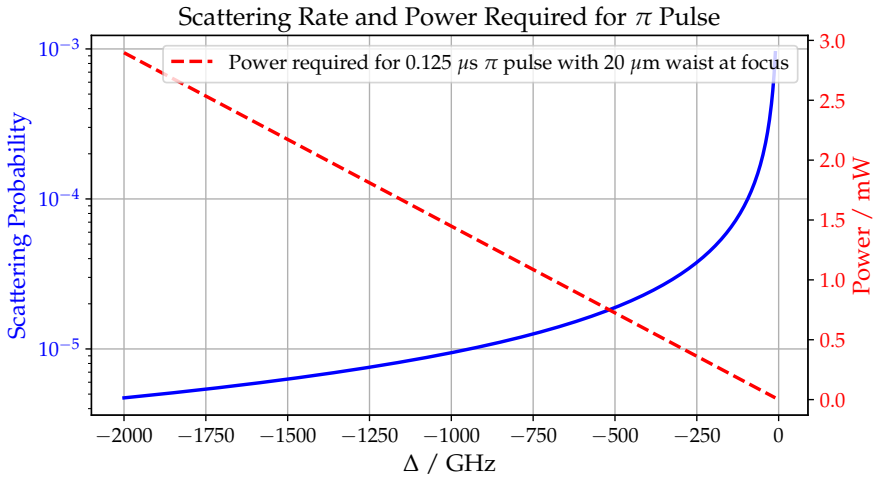

Figure 28: The total scattering probability for the pulse duration is plotted (left y axes) versus two photon detuning Δ for implementing a π pulse with a beam waist of 10μm. The required laser power in mW is shown on the right y-axis which also shows values which are experimentally feasible.

can achieve a high fidelity ground state rotations with scattering rates $< 10^{-4}$, using less than a mW of power.

Typically, smoothened square pulses are used to implement simple $\sigma^x(\theta)$ rotations. In ideal Λ systems this is a simple and effective approach for implementing high fidelity $\sigma^x(\theta = \pi, \pi/2)$ pulses as required in the cluster state scheme. However, in the presence of a small magnetic field, the ground states of ^{87}**Rb** have many spectrally close transitions. This poses a particular challenge for the rotations denoted as (3) in Fig. 26 and (C_1). in Fig. 27, because there exists an off-resonant coupling (with a detuning from two photon resonance of $\delta = 0.28$ MHz) to a transition that one does not want to address in that particular step. The effect of this is shown in Fig. 29 where it is clear that the application of standard flattop pulses [3] drives Rabi oscillations between the desired $|F = 2, m_F = 0\rangle$ and $|F = 1, m_F = 1\rangle$ states, but also simultaneously off-resonantly drives the transition between $|F = 2, m_F = -2\rangle$ and $|F = 1, m_F = -1\rangle$. This leads to detrimental fidelity losses. While using a reduced Rabi frequency helps reduce cross-talk this will also detrimentally affect expected photon generation rates.

3 The flattop Blackman pulse of duration T and rise-time t_r and fall-time t_f is defined as $B_f(t) = \begin{cases} B(t), & 0 \le t < t_r, \\ 1, & t_r \le t < T - t_f, \\ B(t - (T - t_f)), & T - t_f \le t \le T, \end{cases}$ with $B(t) = 0.42 - 0.5\cos(\frac{\pi t}{t_r}) + 0.08\cos(\frac{2\pi t}{t_r})$.

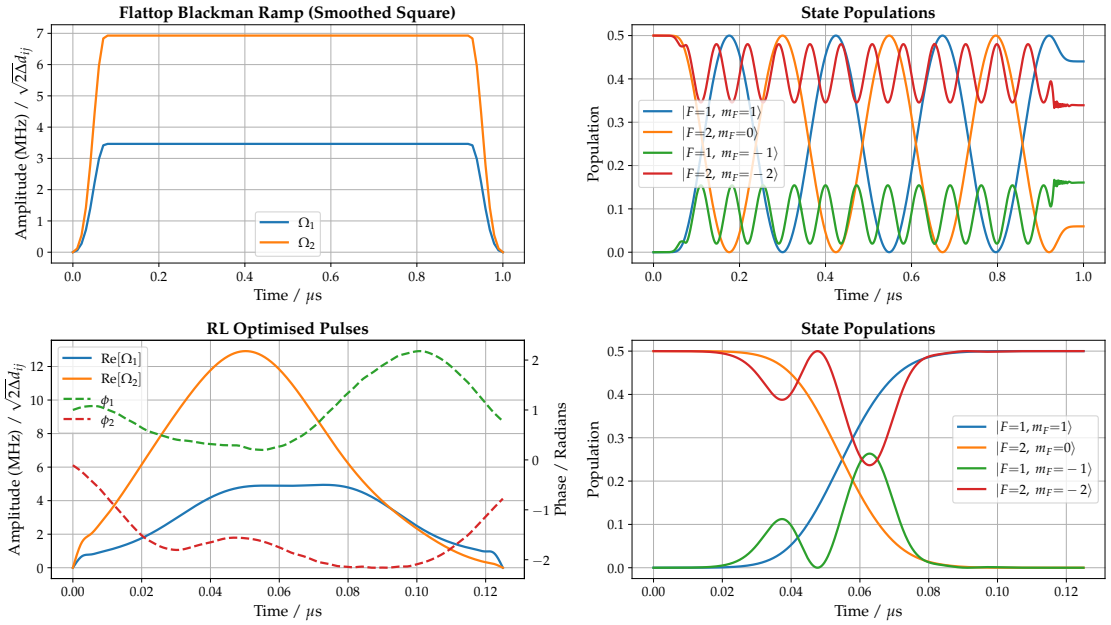

Figure 29: Explicit comparison of state population evolution for standard flattop blackman Ramp (top) versus Reinforcement Learning optimised pulse sequence (bottom) for Rotation 3 described in Fig. 26. It is clear that cross talk between the $|F = 2, m_F = -2\rangle$ and $|F = 1, m_F = -1\rangle$ states cannot be avoided with a simple flattop pulse, but optimising the pulse amplitudes and phases simultaneously [75] leads to a state evolution that returns population coherently to the desired equally weighted superposition of $|F = 2, m_F = -2\rangle$ and $|F = 1, m_F = -1\rangle$. The fidelity exceeds 0.999 and is only scattering limited.

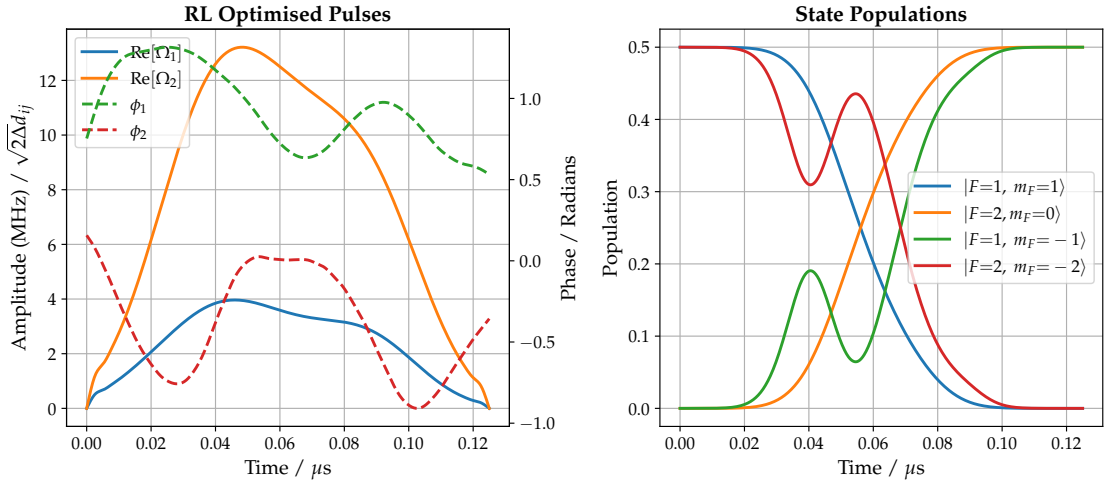

Figure 30: RL Optimised Pulses for realising simultaneous coherent population transfer between $|F = 2, m_F = -2\rangle \leftrightarrow |F = 1, m_F = -1\rangle$ and $|F = 1, m_F = 1\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ with fidelity > 0.999 (scattering limited). It is noteworthy that the smooth amplitude and phase evolution is compatible with off the shelf AWG's and AOM's.

To overcome these challenges, one can optimise the time dependent complex amplitudes or pulses with Rabi frequencies which are defined as real amplitude envelopes with time dependent phase (cf. Eq. 51), with reinforcement learning. More details on this particular technique are described in Ch. 7 and [75]. It is important to note that unlike many other optimisation techniques, emphasis is placed on the generation of experimentally realistic, smooth signals which can actually be produced with off the shelf electronics and modulators (cf. Ch. 5 for more details on a feasible experimental implementation), yet also induces the optimal time dynamics for implementing the desired rotation in just 125ns. The bottom plot of Fig.29 displays the control signals (left) and the corresponding optimised quantum state evolution (right). The system starts in the superposition state $1/\sqrt{2}(|F=2, m_F=0\rangle + |F=2, m_F=-2\rangle)$ and evolves towards $1/\sqrt{2}(|F=1, m_F=1\rangle + |F=2, m_F=-2\rangle)$. Although the state $|F=1, m_F=-1\rangle$ is populated transiently during the evolution, the population ultimately returns entirely to $|F=2, m_F=-2\rangle$. The same cross-talk problem exists for the rotation denoted as (C_1). in Fig. 27 and the same optimisation for simultaneously transferring population from $1/\sqrt{2}(|F=2, m_F=-2\rangle + |F=1, m_F=1\rangle) \rightarrow 1/\sqrt{2}(|F=1, m_F=-1\rangle + |F=2, m_F=0\rangle)$ is shown with RL optimised pulses in Fig. 30. This allows for the implementation of a $\pi/2$ pulse ((C_2) in Fig. 27) which subsequently implements the correct relative phase shift between subsequent qubits for producing a linear cluster state. The far-detuned Raman pulses are described in the following manner:

$$\begin{aligned} \Omega_S(t) &= \Omega_S^{mod}(t)\frac{\sqrt{2|\Delta|}}{d_{s\rightarrow x}}e^{-i\theta_S(t)} \\ \Omega_P(t) &= \Omega_P^{mod}(t)\frac{\sqrt{2|\Delta|}}{d_{p\rightarrow x}}e^{-i\theta_P(t)}, \end{aligned} \tag{51}$$

where $\Omega_{S/P}^{mod}(t)$ is a real time-dependent modulation of the amplitude which satisfies the boundary conditions $\Omega_{S/P}^{mod}(0) = \Omega_{S/P}^{mod}(t_{\text{pulse}}) = 0$. θ_r is defined as $\theta_P - \theta_S$. The amplitudes are scaled for order 1 effective Rabi frequencies and care is taken to adjust their relative scalings to the Clebsch-Gordan coefficients of the desired transitions. The relative phases $\theta_{S/P}$ allows for an implementation of the desired rotation direction from

a pole of the Bloch sphere for $\sigma^x(\pi/2)$ pulses and they also directly impact the relative atomic phases as described in Eq. 52 in the following section. The optimisation in time of $\theta_{S/P}(t)$ allows for the implementation of sophisticated pulses which can cancel cross talk to undesired transitions. The pulse parameters for all rotations required for the generation of linear cluster states are shown in Tab. 6.

Parameter	**Rot. 1**	**Rot. 3**	**Rot. 5**	**Rot. 7**	**Rot. C1**	**Rot. C2**
Pulse Shape	F.B.	RL	F.B.	F.B.	RL	F.B.
Ω_0^{mod}	1.656	cf. Fig. 29	1.527	1.527	cf. Fig. 30	1.656
Polarisations (Ω_S, Ω_P)	σ^+, π	σ^-, π	σ^-, π	σ^+, π	σ^-, π	σ^+, π
Duration / μs	0.125	0.125	0.25	0.25	0.125	0.125
θ_r / rad	4.687	cf. Fig. 29	4.69	4.69	cf. Fig. 30	1.596
Fidelity	> 0.9999	0.9993	> 0.9999	0.99985	0.9996	> 0.9999
Bloch Sphere Rotation	$\pi/2$	π	π	π	π	$\pi/2$

Table 6: This table shows a list of pulse parameters. F.B. denotes flattop Blackman pulse with a rise and fall time of 50 ns and RL denotes a pulse shape optimised with reinforcement learning. Ω_0^{mod} corresponds to the maximum modulation amplitude for the flattop Blackman pulses. The fidelity is calculated with respect to a feasible detuning of 100 GHz and higher fidelities can be achieved by moving further away from resonance.

The pulse parameters comprise pulse duration T, relative phases of the driving pulses θ_r, as well as the rise and fall time for flattop pulses. It shall be noted that these pulses are not robust to amplitude or timing fluctuations as will be briefly highlighted in Sec. 4.4.3.

4.3 V-STIRAP PHOTON EMISSION

This section describes the process of photon emission. The theory behind this approach is described more extensively in Sec. 3.1.3. Photons are produced coherently by a V-STIRAP process by driving either half of the atomic superposition from $|F = 1, m_F = \pm 1\rangle \longrightarrow |F = 2, m_F = \pm 2\rangle$ whilst the cavity is resonant with the transitions $|F' = 2, m_{F'} = \pm 2\rangle \longrightarrow |F = 2, m_F = \pm 2\rangle$. This particular choice of transition is made, because of favourable branching ratios as the dipole coupling for the π-polarised transition is $\sqrt{2}$ higher than that for producing a photon of orthogonal polarisation by coupling to a $\sigma_\pm$ transition,

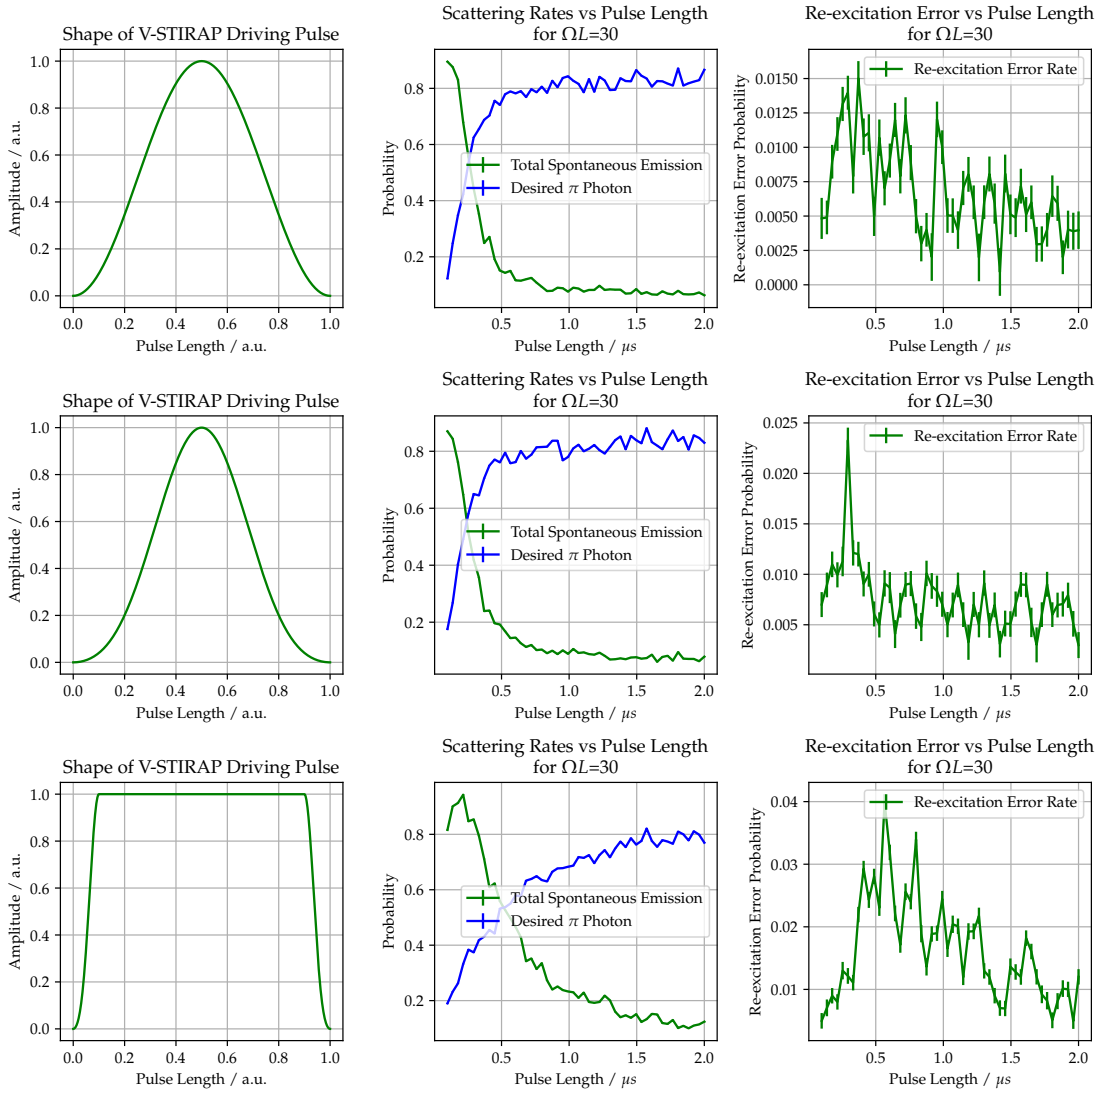

Figure 31: Comparing re-excitation errors for different pulse lengths and shapes: $\sin^2$ (top), Blackman ramp (centre) and flattop with Blackman ramp up/down (0.1 μs) (bottom) using MC analysis. The pulse area is fixed at $\overline{\Omega}L = 30$ and in general longer pulses lead to reduced spontaneous emission rates and higher photon emission rates due to an increase in adiabaticity. The re-excitation rates are calculated in accordance to Eq. 53. The shaped pulses (top and centre) generally exhibit reduced re-excitation errors, as compared to the fast rising flattop pulse, and longer pulses also generally show reduced re-excitation errors. This is in conflict with the desire for high repetition rates and will be discussed in Sec. 4.4.1. While very short photons, $< 0.3\mu$s might exhibit low re-excitation errors, they also exhibit very low emission probabilities which is undesirable.

which is not the desired photon polarisation and thus represents an error. Polarisation dependent filtering can be used to detect such errors at the cavity output and can abort a particular entanglement generation protocol. After a single cycle the ideal desired atom-photon state reads:

$$\begin{aligned} |\psi\rangle &= \frac{1}{\sqrt{2}}(|2,-2\rangle\,|e\rangle + e^{-i\phi_r}\,|2,2\rangle\,|l\rangle) \\ &\xrightarrow{\phi_r=0} \frac{1}{\sqrt{2}}(|g_e\rangle\,|e\rangle + |g_l\rangle\,|l\rangle) \end{aligned} \tag{52}$$

where ϕ_r is the relative phase which reads 0 for the ideal GHZ state. $|e\rangle, |l\rangle$ correspond to the emission time of a particular photon of fixed π-polarisation. The sequence is timed such that the emission of an $(|l\rangle, |e\rangle)$ photon corresponds precisely to the atomic states of $F = 2, m_F = \pm 2$ which are described as $|g_l\rangle$ and $|g_e\rangle$ respectively.

The fidelity of the photonic states which are emitted from the cavity can be adversely affected by re-excitation errors. These occur when there is spontaneous decay to the state $|g_i\rangle$ which is addressed with a laser near resonant with the transition $|g_i\rangle \rightarrow |x_i\rangle$. Subsequent re-excitation can then lead to the unintended emission of a photon via the cavity-coupled transition. This would be particularly problematic when the dipole matrix element for the initial transition exceeds those of the other dipole coupled transitions which represent other possible loss channels. The choice of transitions used for V-STIRAP reverses those chosen in [223]. This offers some advantages as the dipole matrix element of the cavity transition ($|F = 2', m'_F = \pm 2\rangle \leftrightarrow |F = 2, m_F = \pm 2\rangle$) is now stronger than that of the laser coupled transition ($|F = 2', m'_F = \pm 2\rangle \leftrightarrow |F = 1, m_F = \pm 1\rangle$). Therefore, the likelihood of spontaneous emission to the starting state followed by successful production of a photon is severely reduced as evident in Fig. 31. This is particularly undesirable, as it destroys coherence between the atomic states $|g_e\rangle, |g_l\rangle$ and precludes the generation of entanglement across time-bins. Spontaneous emission to another ground state is not considered to be a fidelity reducing error, because the fidelity is conditioned on the emission of a π polarised photon. The only other fidelity reducing error which could occur for the V-STIRAP photon emission is an undesired

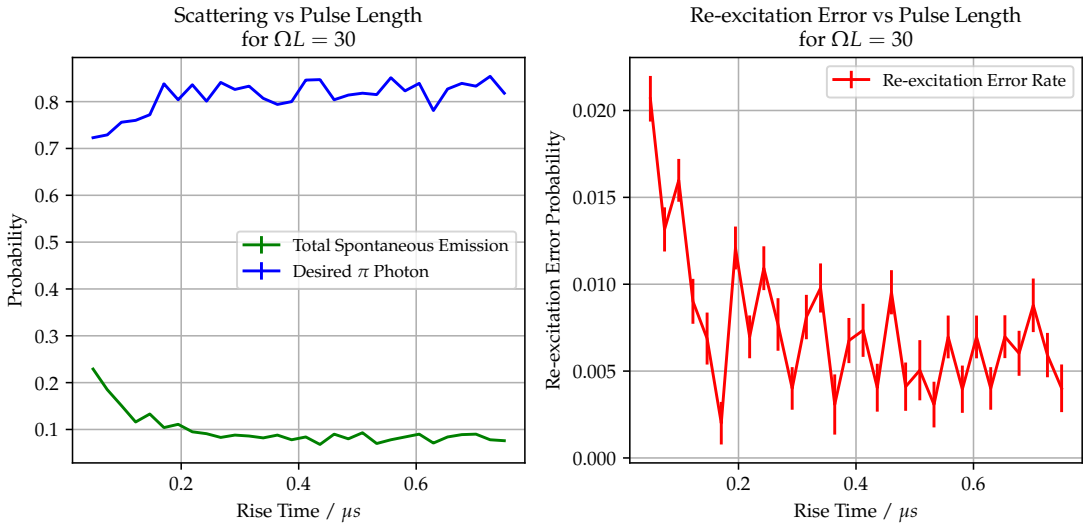

Figure 32: Analysis of different rise times of fixed length flattop Blackman pulses of total duration 1μs on re-excitation rate errors and photon emission probability. It is shown that a longer ramp up greatly reduces the spontaneous emission probability to < 0.1 and also leads to a corresponding decrease in re-excitation errors. The fall time of the pulse has close to no effect on $\mathcal{F}_{\text{rx}}$ hence it is not analysed here.

relative phase evolution away from the desired $\phi_r = 0$ phase which is shown in Eq. 52. However, this can easily be corrected by adjusting the relative phase of the Raman beams, unless it occurs randomly, due to magnetic field fluctuations or dephasing. It is shown in [225] that for an atom-cavity system with ^{87}**Rb** with active magnetic field stabilisation, coherence times on the order of hundreds of μs are achieved, so this is not a limitation for pulse durations on the order of hundreds of nanoseconds. The main fidelity reducing error which will be analysed in detail arises from re-excitation errors, which turns the coherent superposition in Eq. 52 into a mixed state $1/2\left(|g_e\rangle|e\rangle\langle e|\langle g_e| + |g_l\rangle|l\rangle\langle l|\langle g_l|\right)$.

A re-excitation, which means spontaneous emission to $|F = 1, m_F = \pm 1\rangle$ followed by coherent photon production to $|F = 2, m_F = \pm 2\rangle$ cannot be analysed in the open system evolution of the master equation (cf. Eq. 5), where photons decay out of the cavity [4]. It is simpler and more numerically tractable to consider collapse events in time, whilst maintaining an ordering of which type of collapse happened when, which

4 Whilst it is possible to treat the coupling of the photon to the outside field as coherent as in [195] this is very numerically challenging for large Hilbert spaces with several polarisation modes and atomic states.

is the framework offered by the quantum jump, or MCWF method described in Sec. 2.0.1. The fidelity reduction from re-excitation error can be approximated as:

$$1 - \mathcal{F}_{\mathrm{rx}} \approx \frac{\mathcal{N}_{rx}^{l}}{\mathcal{N}_{ems}^{l}} \cdot \frac{\mathcal{N}_{ems}^{e}}{\mathcal{N}_{t}} + \frac{\mathcal{N}_{rx}^{e}}{\mathcal{N}_{ems}^{e}} \cdot \frac{\mathcal{N}_{ems}^{l}}{\mathcal{N}_{t}}, \tag{53}$$

where $\mathcal{N}_{\mathrm{ems}}^{l,e}$ denote the number of photon emissions via the cavity-assisted transition in the late and early time-bins, respectively. $\mathcal{N}_{\mathrm{rx}}^{l,e}$ corresponds to re-excitation errors where spontaneous emission returns the system to the initial state $(F = 1, m_F = \pm 1)$ before a cavity-mediated photon is emitted. $\mathcal{N}_t = 1000$ is the total number of simulated trajectories. Note that this is not simply the sum of normalised re-excitation rates $(\frac{\mathcal{N}_{rx}^{e,l}}{\mathcal{N}_{ems}^{e,l}})$, since a re-excitation error in either the early or late time-bin is only a fidelity reducing effect if it happens in combination with an early or late photon emission. The numerical errors ϵ in this calculation are related to the total number of trajectories $\mathcal{N}_t$ and an analytic error expression in given in App. A.5.

Fig. 31 shows this effect for different pulse durations and peak Rabi frequencies, as well as for different pulse shapes. It is also worth highlighting that one can achieve very low V-STIRAP fidelity errors, at the expense of a low total emission rate $\sum_{i \in e,l} \mathcal{N}_{ems}^{i} / \mathcal{N}_{\sqcup}$, which is undesirable, as this also translates to a low experimental repetition rate of measuring the photonic states of interest. A compromise must be reached, and Fig. 31 shows that the best trade-off between high emission rates and re-excitation error rates is achieved either with sin-squared pulses or with flattop Blackman pulses[5] with long rise times.

A sweep over the shape of the Blackman pulse, by varying the rise-time exhibits some regions of optimal rate-fidelity trade-off, but the variation is somewhat erratic and it is not clear from this analysis what the true optimum pulse consists of. Looking at recent results in the literature, Fig. 32 corroborates the findings of [221] who show that Gaussian pulses do not lead to optimum rate-fidelity trade-offs for common

5 A flattop Blackman pulse consist of a variable-duration half-Blackman window for ramp-up, followed by a flattop, and end with a second half-Blackman window of variable duration for ramp-down.

coupling regimes (intermediate and strong coupling), while an asymmetric Gaussian, with a longer rise than fall time, can lead to improvements, similar to the asymmetric Blackman ramp pulse in Fig. 32. While it is clear that this particular pulse shape is better suited than a symmetric one, a detailed analysis for the true optimal pulse shape is left for future work. What can be said, is that it is advantageous to keep the leading edge of the pulse rising slowly, for maintaining adiabaticity.

4.4 RESULTS

The previous section outlined the atomic driving scheme and the time-binned emission of photons from a single atomic emitter. This section describes the detection of the resulting photonic states in a realistic experimental setup. An overview is provided of how measurements are conducted, alongside a characterisation of the expected coincidence rates and photonic state fidelities. Note, these cannot be directly inferred from the time evolution of the atom-cavity state in presence of photonic field decay.

The photon production process is fundamentally limited by the properties of the cavity, particularly its cooperativity, defined as $C = \frac{g^2}{2\kappa\gamma}$ (cf. Sec. 2.2.3 for further details). The efficiency of photon generation is constrained by the upper bound $\frac{2C}{2C+1}$ [234]. The goal of the scheme is to produce photonic entangled states with both high fidelity and high experimental coincidence rates; however, these two objectives are in conflict with one another.

The photonic state fidelity is evaluated only upon successful detection of the photon, as illustrated in Fig. 33. Photon loss following emission, or incorrect polarisation during emission, primarily affects the production rate rather than the fidelity of the detected states. Spontaneous emission occurring during the V-STIRAP process does not inherently reduce fidelity, provided the decay channel leads to a state from which photon emission is no longer possible under the driving pulse. By contrast, spontaneous emission events that return the system to a photon-producing state can lead to fidelity

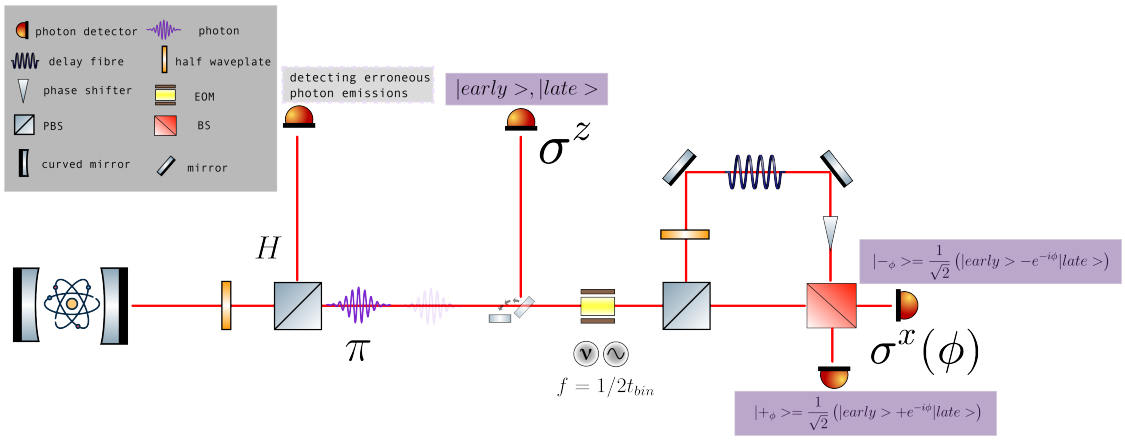

Figure 33: Shown is the detection setup required to detect a single time-bin entangled photon with our atom-cavity scheme. Emissions of the desired polarisation (π) are routed through the polarising beam splitter into a short arm for making σ^z measurements. For σ^x measurements photons are passed into the interferometer, where the early photon is routed through the long detector arm with polarisation switching realised with an electro optic modulator. Calibrating interferometer phase shift allows for picking an arbitrary measurement basis along the equator $\sigma^x(\phi)$, where the early photon interferes with the late photon on a non polarising beam splitter. It is worth highlighting that conceptually this physical detection setup for the photon coherences underpins the numerical method required to calculate off diagonal density matrix elements in the time-binned basis described in Sec. 4.4.2.

degradation, particularly if followed by re-excitation and photon emission. This effect has been analysed using Monte Carlo simulations in Figs. 31 and 32.

Additional sources of fidelity loss include asymmetries in emission efficiencies, incomplete population transfer driven by the Raman pulses, and an undesired relative phase evolution between atomic states (see θ_r in Eq. 52). Since the atom functions as an entanglement mediator, any phase error at the atomic level is directly transferred to the photonic state, thereby impacting the overall fidelity of the generated entangled multi-photon state.

4.4.1 *Expected Coincidence Rates and Measurement Setup*

Generating entangled photonic states with high expected coincidence counts is an important desideratum for various applications, such as sensing, or measurement based quantum computing. For the latter, it is directly related to the maximally achievable clock speed of any quantum operation. This reduces photon losses arising from shorter

delay lines. It is also relevant for experimental data acquisition times and important for using large resources for quantum sensing since it determines the limit of attainable sample sensing acquisition rates. A detailed analysis of different types of photon losses is provided in Tab. 7. One loss is related to inefficiencies in the manipulation of the atomic state and its interaction with the cavity mode, and other losses arise when manipulating the emitted photons in an optical setup as shown schematically in Fig. 33.

The expected photon coincidence counts are fundamentally limited by the atom-cavity interaction, bounded by the efficiency constraint $\frac{2C}{2C+1}$ [234] and we cannot get close to this limit in our scheme as some emissions correspond to an incorrect polarisation mode. Using cavity parameters from [67], losses arising due to the atom-cavity interaction per photon are approximately 17%. Furthermore, losses arise in polarisation routing, fibre coupling, delay propagation and detection which are captured in Tab. 7.

Table 7: Explicit Loss Calculation with state of the art components.

Source of Loss	**Value (%)**
Incorrect polarisation switching	2.75 [226]
Insertion loss	9.8 [240]
Photon loss in delay fibre	2.8 dB/km $\leftrightarrow$ 17.5 (for 1.25 μs photon delay) [196]
Detector loss	5 [145]
Total optical loss /photon	35.05

One of the primary motivations of the work presented in this chapter was to leverage the superior coherence properties of atomic emitters when compared to quantum dots or spontaneous parametric down-conversion (SPDC) sources, while pushing the experimentally feasible repetition rates beyond those demonstrated in [223]. Achieving this would make atomic emitters more competitive at low photon numbers. This goal was primarily pursued by significantly accelerating the Raman re-preparation pulses required between successive photon emissions. As shown in Fig. 34, the simulated ex-

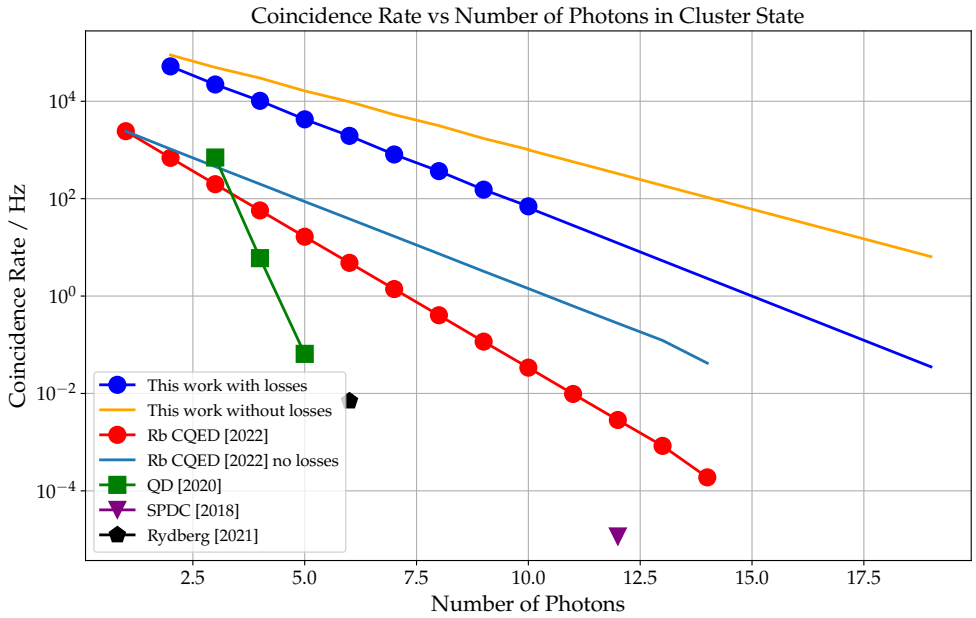

Figure 34: Photon generation rates with realistic optical losses in Hz (log scale), versus photon number for various different physical platforms. We explicitly compare Rb cQED approaches [223] and this work, with quantum dots [111], SPDC [252] and Rydberg atoms [245].

perimental coincidence rates surpass those reported in all previous works. It should be noted, however, that these results are based on simulations incorporating experimentally motivated loss parameters. Additional challenges, which were not fully accounted for in this study, may emerge in actual experimental implementations.

4.4.2 *Calculating Photonic State Fidelities*

In Sec. 4.2, a detailed discussion was provided of the optimal pulse shapes required to drive coherent atomic state transfers using Raman processes. While the final state populations after these transfers directly affect the photonic states, other physical mechanisms, such as spontaneous emission, photonic field decay and undesired phase accumulation between atomic levels, can significantly influence the temporal structure of the emitted photonic state. In particular, the production of photons via V-STIRAP sequences introduces time-dependent features that are more appropriately analysed not through time evolving density matrix elements, but through the evolution of oper-

ators. In the open system description, where photons outside the cavity are not part of the Hilbert space, all electronic coherences vanish after the first photon emission event. This does not capture the physical evolution of the atom-cavity system when no measurement of this photon is performed. An interesting alternative approach for capturing the coherences that remain even after photon emission is presented in [195] where the authors consider a continuum of environmental frequency modes, to which the cavity is coupled coherently, as opposed to treating photonic field decay phenomenologically as is usually common when using the master equation (cf. Eq. 5). However, this is very numerically challenging for the comparatively large state space considered in this chapter.

To capture the multi-photon interference and entanglement that develops during photon emission, it is conceptually advantageous to move to the *Heisenberg picture*. In this framework, the quantum state remains fixed, and the operators evolve in time. This perspective is essential for analysing temporal correlations, particularly in this context where photons are emitted at different points in time and entanglement across these emission events is to be determined.

Photon detection probabilities and correlations are extracted from the expectation values of time-evolved field operators. In this framework, diagonal density operator elements would correspond to number operators such as $\hat{n}(t) = \hat{a}^\dagger(t)\hat{a}(t)$ at specific time points. These quantities are relatively straightforward to access both experimentally[6] and in simulation.

In contrast, probing the coherence between different photonic time bins, requires evaluation of multi-time correlation functions. These involve non-trivial operator orderings of creation and annihilation operators at different points in time to capture the evolution of different photon coherences across the entire emission sequence. This squares up well with the intuition one might develop from Fig. 33, where the multi-time correlations are captured by delaying certain emitted photon components and then measuring the correlation with respect to a photon component emitted at a later

6 This would simply correspond to σ^z measurements of single photons and their emission time in Fig. 33.

time. Similarly, one applies different field operators at different times when numerically evaluating the relevant correlation functions.

To formalise this, one can define operators associated with specific time bins, labelled *early* (E) and *late* (L), as follows:

$$\hat{e} = \hat{a}(t), \quad \hat{l} = \hat{a}(t+T), \tag{54}$$

where T corresponds to the time bin width and t determines the time $[0, T](+t_0)$ where t_0 is an arbitrary offset which determines the start time of a particular time bin.

Matrices describing correlations between photon emissions in different time bins can then be constructed. These are not density matrices in the conventional sense, but rather matrices of correlation functions, with each entry corresponding to a specific temporal measurement outcome (see Fig. 33). For example, the one-photon correlation matrix takes the form:

$$\begin{bmatrix} \langle \hat{e}^\dagger \hat{e} \rangle & \langle \hat{e}^\dagger \hat{l} \rangle \\ \langle \hat{l}^\dagger \hat{e} \rangle & \langle \hat{l}^\dagger \hat{l} \rangle \end{bmatrix}, \tag{55}$$

while the two-photon case generalises to (ordering of operators correspond to photons indexed $1, 2$ as follows $a_1^\dagger a_2^\dagger a_2 a_1$):

$$\begin{bmatrix} \langle \hat{e}^\dagger \hat{e}^\dagger \hat{e} \hat{e} \rangle & \langle \hat{e}^\dagger \hat{e}^\dagger \hat{e} \hat{l} \rangle & \langle \hat{e}^\dagger \hat{e}^\dagger \hat{l} \hat{e} \rangle & \langle \hat{e}^\dagger \hat{e}^\dagger \hat{l} \hat{l} \rangle \\ \langle \hat{e}^\dagger \hat{l}^\dagger \hat{e} \hat{e} \rangle & \langle \hat{e}^\dagger \hat{l}^\dagger \hat{e} \hat{l} \rangle & \langle \hat{e}^\dagger \hat{l}^\dagger \hat{l} \hat{e} \rangle & \langle \hat{e}^\dagger \hat{l}^\dagger \hat{l} \hat{l} \rangle \\ \langle \hat{l}^\dagger \hat{e}^\dagger \hat{e} \hat{e} \rangle & \langle \hat{l}^\dagger \hat{e}^\dagger \hat{e} \hat{l} \rangle & \langle \hat{l}^\dagger \hat{e}^\dagger \hat{l} \hat{e} \rangle & \langle \hat{l}^\dagger \hat{e}^\dagger \hat{l} \hat{l} \rangle \\ \langle \hat{l}^\dagger \hat{l}^\dagger \hat{e} \hat{e} \rangle & \langle \hat{l}^\dagger \hat{l}^\dagger \hat{e} \hat{l} \rangle & \langle \hat{l}^\dagger \hat{l}^\dagger \hat{l} \hat{e} \rangle & \langle \hat{l}^\dagger \hat{l}^\dagger \hat{l} \hat{l} \rangle \end{bmatrix} \tag{56}$$

Each matrix element corresponds to a temporal correlation function, interpretable as the joint probability amplitude associated with the detection photons in specified time bins with associated phase coherences.

As a concrete example, consider the two-photon correlation corresponding to detecting one photon in the early bin and one in the late bin. The relevant correlation function is given by:

$$\mathcal{G}^{(2)}_{ELLE}(t_1, t_2) = \hat{a}^\dagger(t_1)\hat{a}^\dagger(t_2 + T)\hat{a}(t_2 + T)\hat{a}(t_1). \tag{57}$$

This expression reflects the causal and operator ordering intrinsic to the photon emission and detection process, which does not coincide with the index ordering in conventional density matrix notation. Full reconstruction of a multi-photon state is achieved by evaluating a complete set of such time-ordered correlation functions, corresponding to all relevant time-bin operator combinations.

This operator approach is consistent with recent work in quantum optics, adopting the formalism introduced in [25] for quantum dot photon sources. The present adaptation applies that framework to atom-cavity systems.

The calculation of the n photon correlation matrix elements requires the calculation of expectation values of n- order correlation functions. To evaluate the elements defined in Eq. 57, one calculates the time evolution under the master equation, accounting for the time ordering of the operators. Working in the Markovian approximation, the equations governing the evolution of non-unitary system evolution depend only on the previous time step which allows for consistent splitting of the Lindblad superoperator across different time intervals [41]. Following [25], the time evolution under the Lindbladian operators $\hat{U}^\dagger(t_i, t_j) \ldots \hat{U}(t_i, t_j)$ is expressed as $[\ldots]_{t_i \to t_j}$ for simplicity.

$$\mathcal{G}^{(2)}_{ELLE}(t_1,t_2) = Tr\left\{\left[\hat{a}\rho(t_1)\hat{a}^\dagger\right]_{t_1\to t_2+T}\hat{a}^\dagger\hat{a}\left|g_l\right\rangle\left\langle g_l\right|\right\} \tag{58}$$

$$\langle\mathcal{G}^{(2)}_{ELLE}\rangle = \frac{1}{\mathcal{N}^{(2)}}\int_0^T\int_{3T}^{4T} dt_1dt_2\mathcal{G}^{(2)}_{ELLE}(t_1,t_2), \tag{59}$$

where $\mathcal{N}^{(2)}$ is the second order normalisation term which corresponds to the sum of the time-integrals of $\mathcal{G}^{(2)}_{ijji}$:

$$\mathcal{N}^{(2)} = \int_0^T\int_{2T}^{3T} dt_1dt_2\mathcal{G}^{(2)}_{EEEE}(t_1,t_2) + \int_0^T\int_{3T}^{4T} dt_1dt_2\mathcal{G}^{(2)}_{ELLE}(t_1,t_2) + \tag{60}$$

$$\int_T^{2T}\int_{2T}^{3T} dt_1dt_2\mathcal{G}^{(2)}_{LEEL}(t_1,t_2) + \int_T^{2T}\int_{3T}^{4T} dt_1dt_2\mathcal{G}^{(2)}_{LLLL}(t_1,t_2). \tag{61}$$

The matrix element described in 58, requires the evaluation of two integrals. It shall be stressed that at the final point in time one must also insert an atomic operator which corresponds to the atomic states of the different ground states from which early $|g_e\rangle$ and late $|g_l\rangle$ emissions occur, unlike in other schemes where a single electronic ground state is used for photon emission [25]. This is particularly important for the off diagonal elements. Similarly, one can calculate a more complex off-diagonal photonic density matrix element which requires additional nested time evolution. For example, the generator for the matrix element $\mathcal{G}^{(2)}_{ELEL}$ can be expanded as:

$$\mathcal{G}^{(2)}_{ELEL}(t_1,t_2) = Tr\left\{\left[\hat{a}\left[\hat{a}\left[\rho(t_1)\hat{a}^\dagger\right]_{t_1\to t_1+T}\right]_{t_1+T\to t_2}\right]_{t_2\to t_2+T}\hat{a}^\dagger\left|g_l\right\rangle\left\langle g_e\right|\right\} \tag{62}$$

$$\langle\mathcal{G}^{(2)}_{ELEL}\rangle = \frac{1}{\mathcal{N}^{(2)}}\int_0^T\int_{2T}^{3T} dt_1dt_2\mathcal{G}^{(2)}_{ELEL}(t_1,t_2) \tag{63}$$

The remaining matrix elements described in Eq. 56 can be calculated analogously and one can extend this straightforwardly to higher photon numbers. Care must be taken when using these elements for evaluating the Pauli operators required for bounding entanglement with the stabiliser formalism (as will be explained in the following section),

since the ordering of the operators is different to the index ordering in the conventional density operator formalism.

Some more details on the numerical implementation are provided in App. A.7. What should be remarked here is that this is a very computationally expensive calculation with a density matrix of $(32 \times 4)^2$ dimensions (atomic Hilbert space, cavity modes). On a fine numerical time grid, it was infeasible to perform any matrix element calculation beyond $n = 3$ photons. Moreover, care must be taken to take small enough time-steps when approximating the integrals, (cf. Eq. 58) to avoid calculating unphysical density operators which are not positive semi-definite.

The fidelity of an n-particle quantum state with respect to a target state can, in principle, be obtained from full state tomography by reconstructing the density matrix ρ and evaluating its overlap with the ideal state, but this becomes infeasible for large systems due to exponential scaling (both in experiment and in simulation). More scalable approximations include entanglement witnesses [55], which provide lower bounds without full tomography; stabiliser-based methods [228], which exploit the structure of certain states to reduce measurement overhead; phenomenological models [228], which capture dominant noise processes in a tractable form and classical shadow techniques [107], which construct an approximate classical representation with minimal measurements. In this work, we focus on stabiliser and phenomenological approaches, as they are numerically tractable and also offer some interpretability of the dominant noise mechanisms.

4.4.2.1 *Exact Fidelity Bounds*

In the case of GHZ states, a lower fidelity bound can be expressed in terms of *stabilisers* which are Pauli operators (or products thereof) for which the ideal GHZ state (or any

other desired state) is a $+1$ eigenstate [228], such that $S_i |\psi\rangle = \psi$, for all S_i which are stabilisers of $|\psi\rangle$. For the GHZ state the stabilisers read S_1^{GHZ} and S_k^{GHZ} [228]:

$$\begin{aligned} S_1^{\text{GHZ}} &= \sigma_1^x \sigma_2^x \dots \sigma_n^x, \\ S_k^{\text{GHZ}} &= \sigma_{k-1}^z \sigma_k^z, \end{aligned} \tag{64}$$

such that the fidelity is bounded by [228]:

$$\mathcal{F}^{\text{GHZ}} \geq \frac{1}{2} \left(\langle S_1^{\text{GHZ}} \rangle + \prod_{k \geq 2} (1 + \langle S_k^{\text{GHZ}} \rangle) - 1 \right). \tag{65}$$

The fidelity of a prepared n-qubit one-dimensional cluster state ρ_C^n with respect to the ideal cluster state $|C_n\rangle$ can be similarly approximated efficiently using the same approach. This would allow for the verification of the presence of genuine multi-qubit entanglement without requiring full state tomography, which is computationally infeasible for larger photon numbers n. A cluster state $|C_n\rangle$ is uniquely defined by a set of stabiliser operators [228]:

$$S_k^C = \sigma_{k-1}^z \sigma_k^x \sigma_{k+1}^z \text{for } k = 1, \dots, n. \tag{66}$$

where σ_j^x and σ_k^z are Pauli operators acting on qubits j and k, and $\sigma_0^z = \sigma_{n+1}^z = \mathbf{1}$. To estimate the fidelity with fewer measurements, one can define a lower bound given by the sums of products of even and odd stabiliser states [228]:

$$\mathcal{F}^C \geq \prod_{k,even} \frac{1 + \langle S_k^C \rangle}{2} + \prod_{k,odd} \frac{1 + \langle S_k^C \rangle}{2} - 1 \tag{67}$$

If there exists uniformity in the stabilisers for different qubits k this further simplifies the calculation and allows for the determination of the largest possible entangled state, without explicitly determining all the required stabiliser generator values.

Example expressions of stabiliser expectation values in terms of time correlation functions discussed previously are shown App. Sec. A.6. What should be noted here is that experimentally, all stabilisers for both the GHZ and cluster state case can be measured experimentally by picking the desired $\sigma^{x,z}$ measurement bases for each photon as shown in Fig. 33. For example, for the GHZ state this requires the measurement of all photons in the time delayed interferometer setup (σ^x) and the measurement all adjacent pairs of emitted photons in the simple time arrival setup σ^z in Fig. 33.

4.4.2.2 *Phenomenological Model*

One can also estimate the infidelity induced by every part of the sequence in terms of an approximate phenomenological model. Expressing the photonic density operator as a statistical mixture of the ideal density operator corresponding to the desired photonic state and white noise allows for an approximation of the n-photon state fidelity without explicitly calculating the n-photon density operator which is computationally demanding. This approach is described in [228] in the context of entanglement witnesses and their robustness to noise. Under this phenomenological model, the density operator reads:

$$\rho_{\text{phe}} = (1 - p_{\text{noise}})\rho_{\text{ideal}} + p_{\text{noise}}\frac{\mathbf{1}}{2^n}, \tag{68}$$

where ρ_{ideal} is the ideal density matrix, $\mathbf{1}$ is the n-dimensional identity matrix, and p_{noise} represents the probability of a fidelity reducing error. Note that such errors are different from those which reduce the expected photon generation rate. The fidelity reducing errors for the GHZ state generation result from the following errors:

$$p_{\text{noise}} = 1 - (1 - p_{vs})^{2n}(1 - p_{r,1})(1 - p_{r,3})(1 - p_{r,5})^{n-1}(1 - p_{r,7})^{n-1}. \tag{69}$$

where p_{vST} describes the probability of a scattering error in the V-STIRAP photon generation process which leads to a subsequent photon emission of the correct po-

larisation. The other errors are a result of scattering or incorrect phase evolution due to the different coherent population transfer processes involved in the scheme. $p_{r,i}$ where $i \in [1,3,5,7]$ describe the STIRAP (or Raman) pulses (for further details cf. Fig. 26). Error probabilities $p_{s,i}$ for re-preparation pulses are detailed in Tabs. 5 and 6 as $p_{s,i} = 1 - \mathcal{F}_{s,i}$. It is assumed that the only contributing error during V-STIRAP photon emission occurs as a result of re-excitation, hence the error p_{VST} is set to the re-excitation error rate $\mathcal{F}_{\text{rx}}$ (cf. Eq. 53). Further discussion on this is provided in Sec. 4.3. The cluster state protocol will contain further errors given by $\prod_i (1 - p_{c,i})^{\frac{n-1}{2}}$, where p_c decomposes into $p_{c,j}$ where $j \in [1,2,3]$. Given error estimates from each individual process, one can use the phenomenological model, which captures uncorrelated errors, as an upper bound on the fidelity given by $\mathcal{F}^{GHZ} \leq \mathcal{F}(\rho_{\text{phe}}, \rho_{\text{ideal}})$. The true system evolution will not only contain these independent and uncorrelated error sources, but also contain more subtle correlated errors.

4.4.3 *Photonic State Fidelity Results*

Here, we seek a quantitative estimate of the simulated photonic state fidelities and the scaling of the fidelity with photon number n. The simulation is physically realistic with feasible experimental cavity parameters from [67] and includes the full atomic structure of the D_1 and D_2 lines of ^{87}**Rb**. In what follows, a 1μs V-STIRAP pulse of a $\sin^2$ shape is considered, with a peak Rabi frequency of 50MHz balancing the reduction of re-excitation errors and relatively high photon generation rates. An analysis of different V-STIRAP pulse shapes was conducted in Sec. 4.3 and more complex pulse shapes were not considered here, as they did not significantly alter the overall results. The fidelity of a photonic state in this section is calculated with respect to the fidelity of an ideal $n-$photon GHZ state defined in Eq. 46.

Particular fidelity reducing effects are analysed in isolation by looking at the fidelity of a single photon state of the form $|+\rangle = \frac{1}{\sqrt{2}}(|e\rangle + |l\rangle)$ using an exact method where

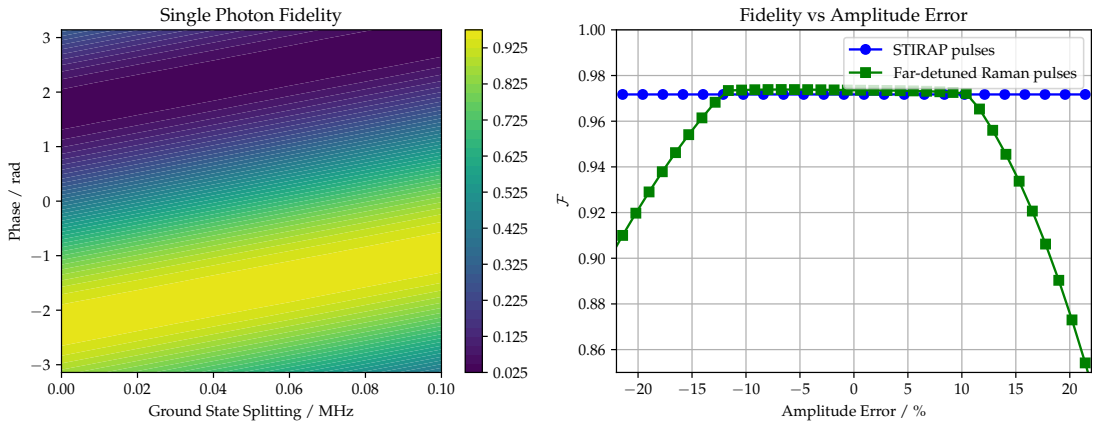

Figure 35: a) (left) Fidelity variation of single photon states $|+\rangle$ with external magnetic field inducing a ground state splitting in MHz (x), and relative laser phase (cf. θ_r in Eq. 51) (y) for far-detuned Raman rotations. b) (right) Fidelity variation of single photon states $|+\rangle$ with error (meaning a correlated error in peak Rabi frequency of both pulses Ω_S and Ω_P) in Raman pulse amplitude for STIRAP rotations (Blue) and far-detuned Raman Rotations (Green). The STIRAP rotations show very good robustness to moderate fluctuations in amplitude, whereas the Raman rotations are more susceptible.

we directly calculate the density matrix using the time-correlation functions described in Sec. 4.4.2. First we highlight the effect of correctly setting the laser phase in relation to the external magnetic field strength B as shown in Fig. 35 (left). An incorrect relative laser phase rotates the atomic states along the wrong axes of the Bloch sphere and consequentially, the relative atomic phase θ_r defined in Eq. 52 will evolve incorrectly. For an experimentally desired ground state splitting of 0.07 MHz the optimal relative phase between the re-preparation pulses ((3) in Fig. 13) is ≈ -1.8 rad.

Moreover, the robustness of two different types of Raman processes described in Sec. 4.2 to fluctuations in amplitude is compared explicitly in Fig. 35 (right). The two lines use the same parameters for V-STIRAP photon production, but the re-preparation pulses denoted 1 and 3 in Fig. 26 are implemented differently. The STIRAP rotations exhibit near perfect robustness to amplitude fluctuations, as measured by the photon fidelity, whereas the far-detuned rotations exhibit reductions in fidelity of $\approx 10\%$ for amplitude errors exceeding 15%.

For an experimental n-photon graph state to exhibit any genuine quantum correlations their fidelity with respect to the target state must exceed at least 0.5. It is instruc-

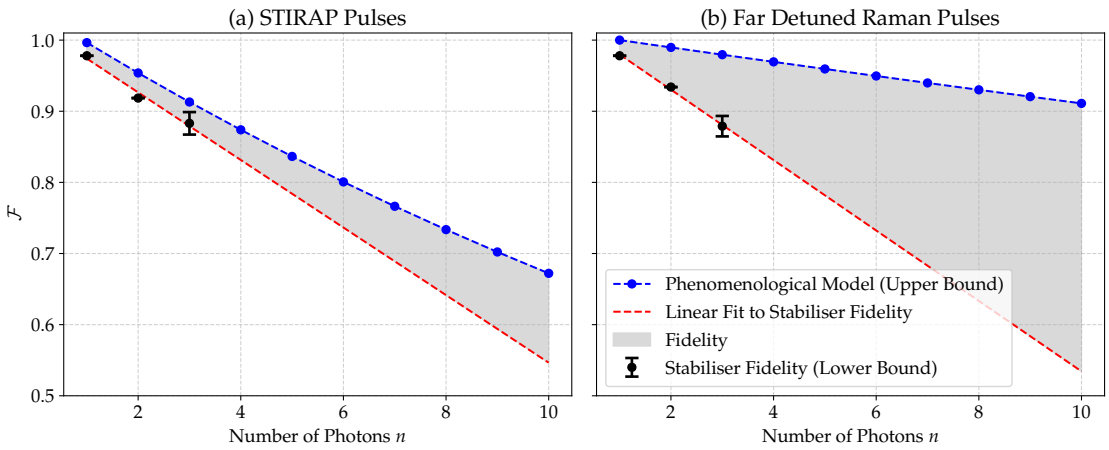

Figure 36: Fidelity of n-photon GHZ states using both phenomenological (blue) and stabiliser formalism (red) for STIRAP pulses (left) and Far-Detuned Raman Pulses (right). The estimated actual fidelity region is shaded in grey. Numerical uncertainties for the Stabiliser Calculations are detailed in App. Sec. A.7.1. It shall be highlighted that the numerical errors for the stabiliser methods become more pronounced as n increases since the numerical uncertainty scales with $\sqrt{2^{n+1}}$ where n is the photon number (see App. Sec. A.7.1). Therefore, and due to the required numerical effort, the plot is restricted to only three data-points calculated with the stabiliser formalism.

tive to analyse fidelity scaling with photon number to see what size of state can be generated and how the fidelity degrades as more photons are added. Two different methods for fidelity estimation are compared, one uses the time correlation functions and stabilisers detailed in Sec. 4.4.2.1, the other uses an oversimplified phenomenological model as detailed in Sec. 4.4.2.2. The former is proven to be a lower bound on the actual fidelity [228] and the latter serves as an upper bound, since it considers each part of the sequence in isolation and does not consider the possibility of correlated errors, as every individual process induces perfect dephasing to an even statistical mixture. This is consistent with the experimental observations in [224], where the authors note slightly lower than phenomenological fidelities in their experiment, but slightly higher experimental fidelities than those given by the stabiliser formalism.

In Fig. 36 we compare the n-photon fidelities for GHZ states produced with two different methods. One uses STIRAP pulses for re-preparation which are more amplitude robust, but have slightly higher absolute infidelities and one uses far-detuned Raman rotations which are more vulnerable to experimental imperfections. Yet, they exhibit lower absolute infidelities and are also required for the production of linear cluster

states. Although the exact fidelity is difficult to determine numerically, it is likely to be found somewhere in the shaded region. A linear trend can be observed from the stabiliser data points, which indicates that the fidelities for the $10-$ photon states are still above 0.5 for both rotation methods with comparable lower bounds in fidelities across both rotation types.

It shall be noted that for the STIRAP pulses, the stabiliser calculated fidelity initially decreases more strongly from $n = 1$ to $n = 2$ since rotations numbered 1 and 3 are less efficient and then the rate of infidelity added per photon decreases again when moving to $n = 3$ consistent with the individual process fidelities. Moreover, the gap between the phenomenological and stabiliser method in the case of the far-detuned Raman pulses is much larger than for the STIRAP pulses. This is suggestive of an error per photon not captured by individual process fidelities and which is related to photon production - as the same V-STIRAP pulses are used across both re-preparation pulses. Note that for the lower bound, it doesn't seem to matter how close the re-preparation fidelities are to 1.

While fidelities for linear cluster states were not analysed in detail in this section, earlier sections provide detailed outlines on how to generate them in this framework and how to capture their fidelity with the stabiliser formalism. Calculating the larger photon number stabilisers is computationally demanding as the protocol requires a larger number of concatenated time dependent signals. However, many of the individual steps are analogous to the GHZ state generation protocol and thus one would expect the results to be similar, particularly the scaling of fidelity with n. However the detailed calculation is left for future work.

4.5 OUTLOOK

In this chapter, time-bin encoding is investigated as a robust alternative to polarisation encoding for generating entangled photon states from atom–cavity systems, motivated

by the limitations of polarisation-based schemes such as cavity birefringence and dynamic atomic decoupling. Building on the scheme experimentally demonstrated in [223], optimal pulse sequences are proposed and simulated for the deterministic generation of Greenberger–Horne–Zeilinger (GHZ) and linear cluster states in a time-bin basis. Two Raman re-preparation protocols are compared: one based on far-detuned Raman pulses and another inspired by STIRAP, both interleaved with V-STIRAP photon emission. Benchmarking is performed using experimentally realistic cavity parameters [67]. The discussion also addresses how such states may be characterised experimentally, and the theory for determining lower bounds on photonic state fidelities is adapted for multilevel atoms following [25]. Simulation results indicate that photon state fidelities exceed 0.5 for up to 10 photons, demonstrating the feasibility of generating large entangled resource states in a time-bin basis. When realistic experimental losses are incorporated, the predicted generation rates surpass those achieved in state-of-the-art experiments. These findings suggest that time-bin encoded entangled photon generation from atom–cavity systems offers a scalable pathway toward entangled resource state generation for applications in measurement-based quantum computing and quantum sensing, while overcoming limitations in generation rates of existing cavity-based protocols. Further work is required to experimentally realise this scheme, including a detailed assessment of the impact of dipole-trapping laser light and atomic motion.

5

EXPERIMENTAL RESULTS & PROGRESS

In this chapter, the experimental progress towards the realisation of novel atom-cavity experiments is discussed. Much of the laboratory setup and equipment was inherited from my colleagues and predecessors. In particular, Tom Barrett [18] initially assembled a substantial portion of the existing apparatus that forms the foundation of the current experimental infrastructure which requires extensive control of electronic and optical components as outlined in Sec. 5.4.

The work presented here supports the experimental implementation of the schemes discussed theoretically in Chapter 3 and Chapter 4, which propose the generation of n-photon bursts for boson sampling experiments or time-bin entangled states for quantum information processing using atom-cavity platforms, by carefully calibrating and testing optimised laser pulses.

This chapter first describes two central experimental components which are required for addressing the electronic states of atoms and producing single photons, namely the Magneto Optical Trap (MOT) and the optical cavity. The MOT serves as a standard tool in atomic physics, enabling the efficient cooling and spatial localisation of neutral atoms like $^{87}\mathbf{Rb}$. The operation of an optical cavity, designed for strong atom-cavity coupling in cQED experiments is detailed alongside suggestions for future experiments.

Following this, the development of a new laser preparation system for improved Raman pulses is outlined, which includes laser beams for the MOT transitions, and several laser paths for Raman pulses and optical pumping. Much emphasis is placed on the generation of tailored laser pulse sequences, as calculated theoretically for the photon production protocols described in previous Chapters. In particular, the use of

an Acousto-Optic Modulator (AOM), which allows dynamic control of both the amplitude and frequency of the optical fields on nanosecond timescales, is discussed. AOMs, driven with distortion corrected input signals are shown to enable the implementation of complex pulse shapes.

Finally, results demonstrating enhanced population transfer with optical molasses as required by the novel single-photon generation scheme are presented.

5.1 MOT

Atoms must first be cooled and spatially localised to be able to address and manipulate their electronic states with an external laser as discussed extensively in Ch. 3 and Ch. 4. A laser can provide a directional force. Consider an atom moving at velocity v, towards a laser beam of intensity I with wave-vector k, which is near resonant with an atomic transition of line-width Γ and a saturation intensity I_{sat}.

The absorption of photons leads to an increase in momentum of the atom in the direction of beam propagation, ensuring conservation of momentum, and gives rise to the scattering force [78]:

$$F_{\text{scat}} = \hbar k \frac{\Gamma}{2} \frac{I/I_{sat}}{I/I_{sat} + 4\Delta^2/\Gamma^2}, \tag{70}$$

where the atoms experience a Doppler shift kv and the detuning from atomic resonance in the atomic frame is $\Delta = \omega - \omega_0 - kv$, where ω_0 is the atomic resonance frequency. Photon absorption is followed by spontaneous emission, which occurs at a rate set by the excited-state lifetime. The scattering force arises because absorption events impart momentum $\hbar k$ along the laser beam direction, while the subsequent spontaneous emission occurs isotropically. In the presence of detuning, atoms moving toward the beam are Doppler-shifted into resonance and preferentially absorb photons from the opposing direction, leading to a velocity-dependent damping force. For $I/I_{sat} \geq 1$, the transition becomes saturated. This implies that absorption and stimulated emission

both occur at rates much larger than the spontaneous rate, but since their momentum transfers are equal and opposite, they cancel on average and give rise to no net force. The net momentum transfer results from the imbalance introduced by spontaneous emission, which allows for further absorption cycles with a directional momentum kick. This combination of directional absorption and isotropic spontaneous emission produces the effective scattering force described in Eq. 70. A configuration of three pairs of counter-propagating beams ensures that atoms are slowed in all directions, known as an optical molasses.

To localise atoms spatially, one also requires a force with positional dependence. Using the splitting of different atomic energy eigenstates in the presence of an external magnetic field (i.e. the Zeeman effect) realises this. Introducing a magnetic quadrupole field with two coils in anti-Helmholtz configuration, which are centred on the same axes, but with opposite currents, creates a zero B-field at the MOT centre. The field gradient, in combination with particular laser polarisation induce a restoring force towards the zero B-field point. Typical operation at $I = 4$A gives a field gradient of approximately 30G / cm. The magnetic field gradient provides a spatial force due to the Zeeman shift relative to the zero field point. Counter-propagating beams will have opposite circular polarisations (in the frame of the atom) and this positional shift in resonance, will bring one atomic transition (e.g. σ^- for an atom at $z < 0$ from the centre of the MOT) closer to resonance and detune the other (σ^+) further away from resonance such that the atom experiences an effective restoring force from one of the cooling lasers. The same holds for all spatial directions and therefore, the atoms are confined to a particular position as shown in Fig. 37.

The experimental setup in the lab for ^{87}**Rb** uses cooling beams of approximately 1mW which are red-detuned 10MHz from the F=2 to $F' = 3$ transition. Sidebands are applied to the laser with an external RF drive at 6.78GHz to ensure that atoms which decay to the $F = 1$ level can be repumped to the $F = 2$ level via $F' = 2$.

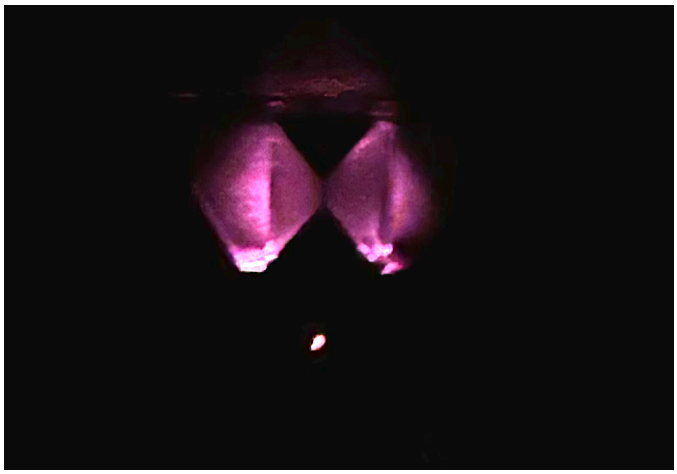

Figure 37: MOT of approximately 3000μm diameter shown 0.7 cm below the cavity.

Atoms are trapped below the cavity since this provides ample optical access for cooling beams, but for a cQED experiment they need to be transported into the cavity. This is described extensively in [18] and basically relies on detuning some MOT beams with respect to the others to launch the cloud vertically upwards.

Dipole trapping inside the cavity seems a more desirable technique for realising a deterministic atom-cavity interface. However the lack of optical access with the current experiment limits the possibility of tightly focusing a beam into the cavity. Extensive discussion on the engineering of cavities for this purpose is provided in [66] and experimental demonstrations of dipole trapping are provided by the Rempe group, cf. Refs [223] or [97], where the authors load several atoms into an optical lattice inside an optical cavity.

5.2 CAVITY

Optical cavities are central to the single-photon production schemes discussed in Chapters 3 and 4, with a more detailed theoretical treatment provided in Section 2.2.1. A key challenge in operating such cavities is the stabilisation of the cavity length, ensuring that one of the cavity's modes overlaps with the desired atomic transition frequency. This resonance must remain stable throughout an experiment, which cannot

be achieved passively due to thermal drift and acoustic noise. To maintain resonance, the cavity length is actively stabilised with piezoelectric actuators, on which the mirrors are mounted. Details of the cavity fabrication can be found in other theses from the group (see [18, 66]). A reference signal for the locking process is provided by a laser[1] that is far detuned from the atomic transition. This beam is coupled along the cavity axis and monitored in transmission or reflection using a photodiode[2].

An electro-optic modulator (EOM) adds modulation sidebands to the locking laser. The transmission peak of one cavity mode is then used for Pound-Drever-Hall (PDH) locking via a Toptica Digilock. Effective PDH locking requires coupling into the fundamental transverse mode, TEM_{00}, which provides the strongest transmissive signal. Achieving this requires precise alignment, as even small transverse or longitudinal misalignments can couple light into higher-order spatial modes.

To assist with alignment, the cavity piezos are scanned using an external DC amplifier[3] to observe the cavity's free spectral range: $\nu_{\mathrm{FSR}} = \frac{c}{2L_{\mathrm{cav}}}$ where L_{cav} is the cavity length.

Another important consideration is mode matching: a lens is used to focus the locking beam into the cavity so that its beam waist matches the cavity's mode waist. For a symmetric cavity with identical mirror curvatures, the mode waist is given by

$$w_0 = \sqrt{\frac{\lambda}{\pi}} \left(\frac{L}{2} \sqrt{\frac{R}{L} - 1} \right)^{1/2}, \tag{71}$$

where λ is the wavelength of the light and R is the radius of curvature of the mirrors. Coarse alignment is typically done using a CCD camera that views the back facet of the cavity. Once cavity modes are visible, the detection switches to a photodiode for higher sensitivity. Maximising the signal-to-noise ratio of the transmission peak is crucial, as typically only about 1% of the incident locking light is transmitted through the cavity.

1 Toptica DL Pro at 785 nm
2 Thorlabs APD110A
3 Falco Systems 20x RF Amp

Although in principle, the reflection dip has the same contrast as the transmission peak, the lower noise floor in transmission proved more conducive to locking in our setup.

Several optical cavities have been developed within our group for use in cQED experiments. One such cavity, fabricated by Tom Barrett, suffers from alignment challenges due to mirror displacement during glue curing. As a result, the fundamental TEM_{00} mode is misaligned with respect to the transverse centre of the mirrors, complicating optical coupling. This cavity has a mode waist of approximately 26 μm, necessitating the use of a 400 mm focal length lens to mode match a 3.75 mm input beam waist. Although transient locking was achieved, the cavity could not be operated reliably in a full cQED experiment.

A further limitation of this cavity is its strong birefringence, causing a hybridisation of the coupling of individual atomic polarisation modes to output modes of the cavity mirrors. This results in a time-dependent polarisation of the output photon, making it unsuitable for experiments producing polarised single photons or the production of time-bin entangled photonic states.

An alternative cavity, fabricated by Thomas Doherty and Mark IJspeert, offers improved atom-cavity coupling with a cooperativity of $\mathcal{C} \approx 10$ on the D_2 line. It also allows for greater optical access, which is advantageous for integrating dipole traps. However, its tapered mirror design rendered it highly susceptible to acoustic noise, making stable locking extremely difficult. These challenges motivated further developments in cavity design, as discussed in [212].

The most promising candidate for future cQED experiments proved to be an earlier design described in [105]. This cavity features a well-centred TEM_{00} mode, enabling easier alignment. Additionally, it possesses several beneficial properties for the protocols presented in previous chapters: a broader line-width and faster atom-cavity decay rate $\kappa \approx 12$MHz supports higher repetition rates by reducing the critical time for efficient photon emission (cf. Eq.39 in Chapter 3). While this broader line-width was a disadvantage for schemes such as that described in [18], which addressed individual

Zeeman levels with MHz-scale splittings, it poses no issue for the schemes described in this thesis that target entire hyperfine levels. The main drawback of this cavity is its elevated intra-cavity loss, which reduces the efficiency of photon extraction through the output mirror by $\approx 50\%$.

The theoretical and experimental insights gained here have important implications for future cavity design. In particular, the analysis presented in Fig. 10 (from Ch. 3) suggests that optimal photon production rates are achieved for cavities with larger line-widths and closely matched values of κ and g. This is subject to the constraint that the cavity line-width is not too large. Otherwise, sidebands for locking could not be detected by bandwidth limited photodiodes, it becomes more challenging to create a stable lock around the desired atomic transition and cross talk to other atomic transitions might become an issue. Furthermore, the reliability of mirror mounting, specifically, avoiding transverse misalignment during gluing and curing, along with pre-aligned mode-matching optics in a caged setup as in [67], enhances long-term experimental stability and eases with optical alignment.

5.3 LASER SYSTEM

5.3.1 *Laser Locking*

Laser	λ (nm)	ν (THz)	Locking scheme
Eylsa	780.246	384.228	Locked to the $F{=}2 \rightarrow C_{13}$ crossover in saturation absorption spectroscopy.
DL Pro (Stirap)	780.232	384.235	Offset locked from the Cooling laser with respect to a 6.7GHz microwave frequency reference.
DL Pro (Cavity)	785.656	381.582	Offset locked to a particular mode of the frequency comb such that it can be AOM-modulated to the right cavity mode which coincides spectrally with the desired atomic transition.
Comb	λ_n	$n\nu_r + \nu_0$	The carrier envelope offset, ν_0, is beat note locked to a Rubidium reference. The repetition rate, ν_r, is beat note locked to the Eylsa laser.

Table 8: Lasers used in the experiment and how they are locked.

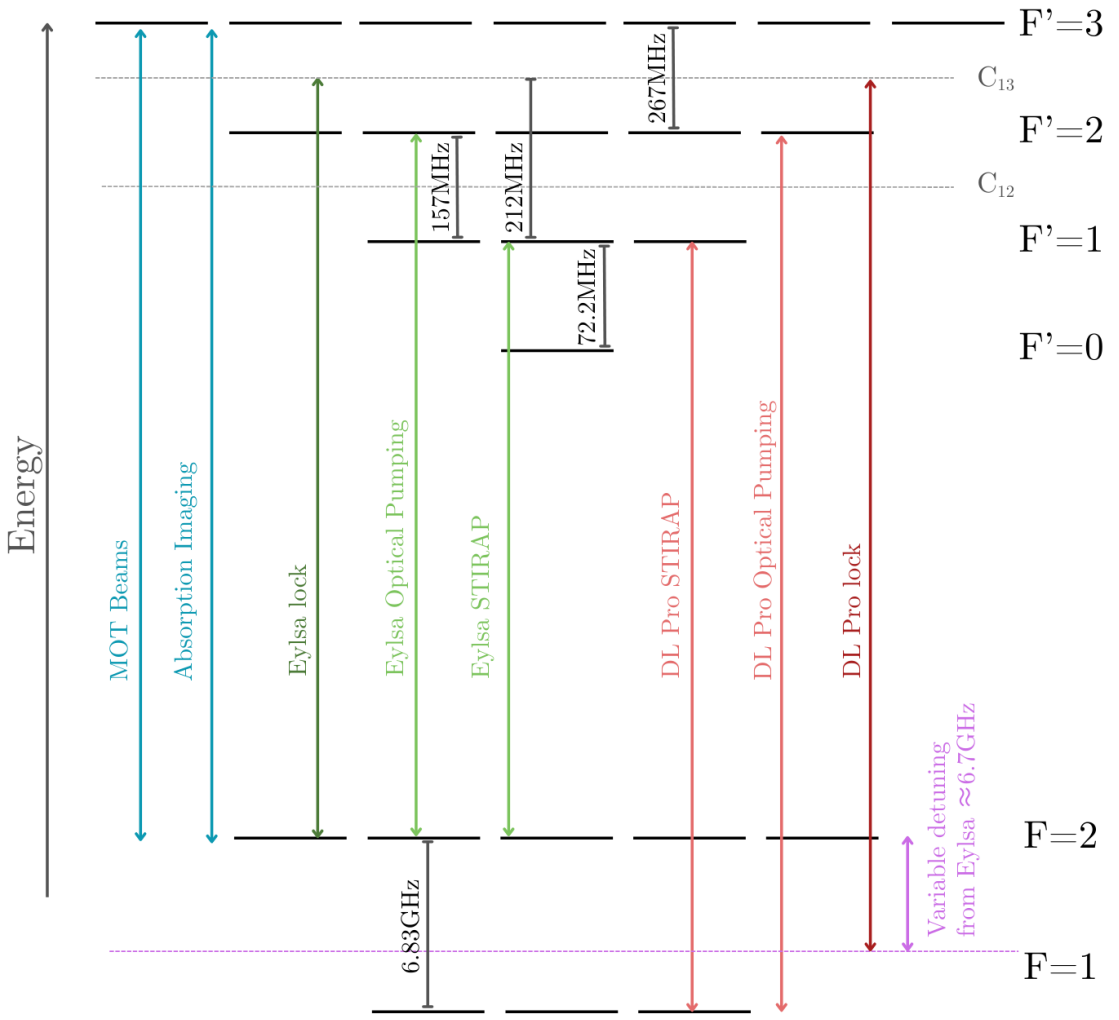

Figure 38: Diagram showing laser locking and desired transitions one wishes to address for the D_2 line of ^{87}**Rb**.

The experiments described in Ch. 3 and 4 require separate lasers which address both $F = 1 \rightarrow F' = 2$ transitions, as well as $F = 2 \rightarrow F' = 2$ transitions which are separated by 6.8GHz to be able to create Λ system type configurations in the multi-state structure of ^{87}**Rb**. It was determined theoretically that the D_1 line is favourable for any population transfer, rotation or re-preparation pulses, but stronger dipole coupling makes the D_2 line a better choice for cavity enhanced photon production. For that reason the existing lasers in our laboratories only have frequencies within the range of the D_2 transitions. A detailed figure showing how lasers are locked relative to atomic transitions is shown in Fig. 38.

The Quantel Eylsa laser is the main laser of the setup, it is used for addressing the MOT transitions, as well as for STIRAP and V-STIRAP transitions. The laser is locked to the $F = 2 \rightarrow C_{13}$ crossover line in pump-probe spectroscopy. The Pound Drever Hall (PDH) technique is used for frequency stabilisation. This sets the absolute frequency with respect to which the rest of the lasers are locked.

The DL Pro STIRAP laser is offset locked with a 6.7 GHz microwave reference to the Quantel Eylsa laser. A beat note is detected with Vescent Photonics D2-160 High-Speed Beat Note and Pulse Detector and the laser feedback is generated with a Vescent Photonics D-2 135 Offset Phase Lock Servo. The microwave reference is tunable which provides freedom in locking the laser at a specific frequency.

The frequency comb[4] must also be locked for providing a consistent spectral ruler to which one may lock the cavity. The comb frequencies are given as $\nu_n = n\nu_r + \nu_0$, where n is an integer and $\nu_{r/0}$ are the repetition rate and carrier envelope respectively. The carrier–envelope offset (CEO) frequency sets the absolute position of the comb spectrum in the optical domain. This is locked to an external Rubidium reference[5] with a 20MHz beat note offset lock[6]. The repetition rate determines the spacing between adjacent comb teeth and it is the frequency at which laser pulses are emitted from the

4 Menlo Systems FC1500-250-WG
5 Stanford Research Systems FS725
6 Beat notes are measured with Menlo Systems FPD310-FV.

mode-locked laser. This is locked via a beat note offset lock of the nearest comb mode to the Eylsa laser (60MHz).

The DL Pro Cavity laser can then be offset beat note locked (40MHz) to an arbitrary comb mode, such that the cavity mode which is one picks out for locking spectrally overlaps with the desired atomic transition for photon production. The operating frequency of this laser is far detuned from the D_1 or D_2 transitions in the ^{87}**Rb** spectrum at $\approx$ 785nm.

5.3.2 *Laser Preparation Table*

Laser	**Num**	**Type**	**$\sim$ Mod. ν (MHz)**	**Purpose**
Eylsa	3	AOM[1]	2×101	MOT cooling
Eylsa	1	EOM[2]	6.76×10^3	MOT repumping
Eylsa	1	AOM[1]	$-1 \times$ (70-130)	STIRAP initial shift
Eylsa	1	AOM[1]	$-1 \times$ (70-130)	STIRAP shaped pulse
Eylsa	1	AOM[1]	-1×106	Optical Pumping
DL Pro STIRAP	1	AOM [1]	$-1 \times$ (60-90)	STIRAP shaped pulse
DL Pro STIRAP	1	AOM [3]	$+1 \times 75.25$	Optical Pumping
DL Pro Cavity	1	AOM [3]	$2 \times$ (60-100)	Cavity locking
Cavity	1	EOM[4]	12.5	Cavity locking sidebands

Table 9: All AOMs and EOMs used for modulating lasers and their typical operational frequencies.

The experiments described in previous chapters and in what follows require a number of different laser frequencies generated with different AOM's (see App. A.8 for a detailed introduction to AOM modulation) and electro optic modulators (EOM's). EOM's use the electro-optic effect to change the refractive index of a crystal and can operate at much higher frequencies than AOM's. In Tab. 9 all the different modulators and their operational frequencies are detailed. The Eylsa laser is modulated to be

1 Crystal Technologies 3110-120
3 Crystal Technologies 3080-120
2 iXblue Photonics Photline MPZ-LN-10
4 iXblue Photonics Photline NIR-MPX800

near resonant with the cooling transitions for the MOT ($F = 2 \rightarrow F' = 3$) and sidebands generated with an EOM ensure that population can be repumped from $F = 1$. Further AOM's are then used to prepare relevant Raman and optical pumping beams. Instead of using a double pass AOM two single-pass AOM's in series are used for shaping pulses, more detail will be given in Sec. 5.5. The DL Pro STIRAP laser is also modulated for Raman transitions and optical pumping and the DL Pro cavity laser is modulated with an EOM for locking the cavity with a PDH scheme. All laser pulses are prepared on a separate optical table, fibre coupled and then sent to the optical table on which experiments are performed.

5.4 ELECTRONIC CONTROL OF LABORATORY EQUIPMENT

The control of all the required experimental equipment, with a particular focus on the required time-dependent electric fields which are eventually sent to atoms, requires the interplay of hardware and software. The Quantel Eylsa Ytterbium fibre laser is electronically controlled with a Toptica Digilock, which also performs PDH locking. Toptica DL Pro Lasers are controlled with DLC Pro Controllers, which are supplemented with Vescent Lockboxes for frequency stabilisation. The cavity is also locked with a Toptica Digilock. The lab computer runs a combination of proprietary software like the Toptica Digilock UI, but a large part of the experimental sequencing relies on a custom Python UI called Cold Control which was initially developed by Tom Barrett and has since been overhauled and modernised by myself and Matthew King [7] This handles the interplay of the ADLINK DAQ (2502) cards - which produce slower time-dependent DC voltages - as well as the AWG (Tabor Electronics WX1284C) which can send signals at much higher output sampling rates (up to 1.2GHz). The DAQ cards are used to control magnetic field coils, trigger various other instruments such as the AWG and are used to perform other experimental tasks like turning on sidebands for the Eylsa laser. Any

7 It can be found here: https://github.com/jan-o-e/cold-control

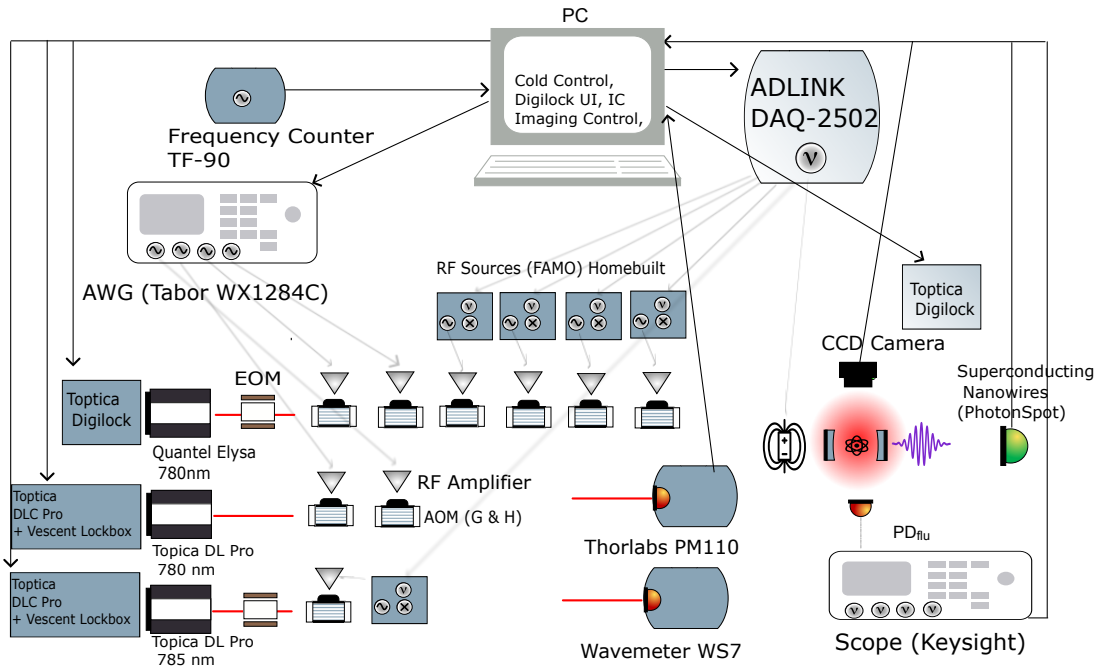

Figure 39: Shown is a simplified (and not to scale) overview of the lab setup and how every instruments are controlled and interact. A lab PC runs a custom Python UI *Cold Control*, which controls the devices which are responsible for sending voltages and signals (DAQ Cards and AWG) and also interfaces with a majority of the other equipment through proprietary software. A simplified signal chain is also shown, externally generated RF signals (either via a homebuilt RF sources or via an AWG) pass through an amplifier and are then inputted to AOM's. Homebuilt RF sources are calibrated with a frequency counter. Lasers are locked and controlled with external Toptica controllers. Laser data is acquired via Thorlabs PM, Wavemeter and experimental data is read via a combination of cameras, photodetectors and superconducting nanowires.

AOM's driven indirectly by the DAQ card pass through home-built RF sources called FAMOs (**F**requency **A**mplitude **M**onitoring **O**utput). Upon calibration [8], they provide RF signals of variable amplitude (A) at some frequency corresponding to a particular voltage input. The AWG on the other hand can directly drive AOMs with RF signals. Any instruments used for data acquisition are also controlled by the Python UI and this includes a CCD camera [9], oscilloscope [10], as well as the superconducting nanowires [11], Thorlabs Powermeter [12] and Wavemeter WS7 for monitoring laser powers.

5.5 AOM-BASED PULSE SHAPING

Producing optimal sequences of shaped pulses at short timescales of hundreds of nanoseconds as is required for the schemes described in Chs. 3 and 4 is a technical challenge. AOMs have finite rise-times (cf. 119), on the order of several nanoseconds, set by the time it takes the acoustic wave to sweep across the full optical beam, and they also have a limited instantaneous bandwidth, determined by the finite transit time of the acoustic wave across the focal spot that limits how quickly the diffracted light can follow changes in the RF drive frequency. See App. A.8 for a more detailed theoretical introduction to AOM modulation. Imperfect optical alignment or other imperfections and distortions in the signal chain consisting of driving electronics (AWG or homebuilt RF source) and an RF amplifier, will also adversely affect the shaped optical signal. Optical alignment poses a challenge for using double pass AOM's for shaping laser pulses as will be discussed in more detail. Moreover, the frequency and amplitude response of the optical signal are coupled non-linearly, as AOM efficiency is non-linear in RF amplitude and ν_{RF} (cf. App. Eq. 118). Moreover, in a single pass configuration, the diffraction angle and beam pointing is also affected by ν_{RF} which in turn affects fibre

8 Calibration of voltage to frequency is handled with a frequency counter (TF-90).
9 Imaging Source DMK 21BF04
10 Keysight InfiniiVision MSO-X-3104A
11 Photon Spot NW1FC780
12 PM-110A

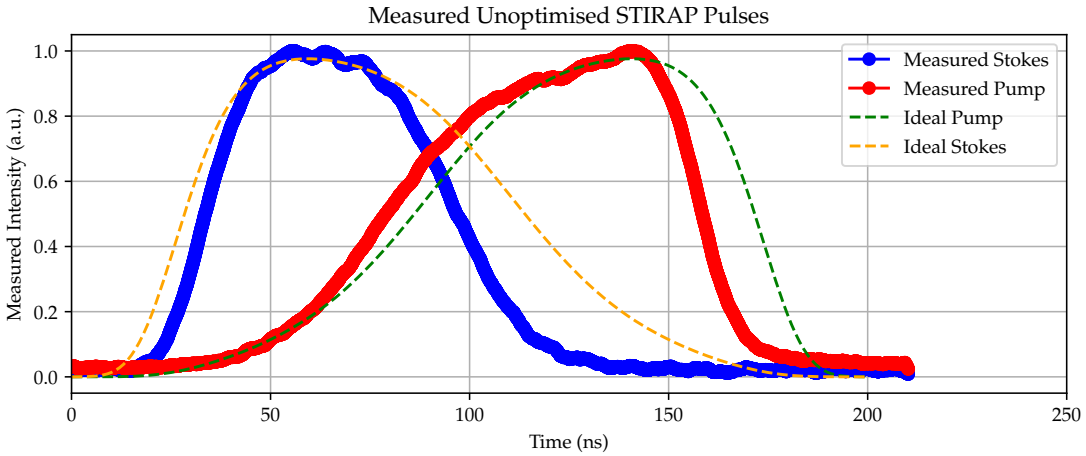

Figure 40: Measured Pump and Stokes pulses from DL Pro STIRAP and Eylsa lasers respectively with no optimised driving. It is clear that at ns timescales it is challenging to generate high quality signals and the systematic pulse shortening arising from imperfect alignment of a double pass AOM is seen for the Stokes pulse.

coupling efficiencies. This makes simultaneous amplitude and frequency modulation challenging, as frequency chirping affects the amplitude of the resultant signal.

This section focusses particularly on the experimental implementation of pulses described in Ch. 3 Fig. 13 with the analytically pre-determined and numerically optimised shapes [233] [13]. However, the methods discussed in what follows are extremely general and find application to many other types of pulses relevant for quantum information processing experiments, such as optimised pulses for Raman transitions described in Ch. 4 It is shown that with single pass AOM's and careful driving pulse calibration, any time dependent amplitude shaping and frequency chirping can be implemented, insofar as it lies within the AOM rise-time and bandwidth constraints.

5.5.1 *Time-Dependent Amplitude Shaping*

We first consider shaping the amplitude of optical signals at a fixed frequency. This requires applying a time-dependent voltage $V(t)$ consisting of an envelope modulating an RF carrier at frequency ν_{RF}:

$$V(t) = \Omega_{\mathrm{in}}(t)\sin(2\pi\nu_{\mathrm{RF}}t + \phi_0) \tag{72}$$

where $\Omega_{\mathrm{in}}(t)$ is the time-dependent Rabi frequency which satisfies $\Omega_{\mathrm{in}}(0) = \Omega_{\mathrm{in}}(t_{\mathrm{final}}) = 0$. $\Omega_{\mathrm{in}}(t)$ is given by a numerically or analytically pre-determined optimal time-dependent Rabi frequency $\Omega(t)$ which, in an ideal system, corresponds to the AOM-modulated envelope. In practice, $V(t)$ is sampled at the output rate of the AWG, fixed to 1 GHz. These sampled voltages form the discrete input vector $\mathbf{\Omega}_{\mathrm{in}}$, which is pre-loaded into AWG memory and then played to generate the desired signals.

Using the Gaussian masked pulses (see Eq. 40 in Ch. 3) this is initially tested with the DL Pro STIRAP modulated with a single pass AOM which forms the Stokes beam. The Eylsa laser, modulated with a double pass AOM forms the Pump beam and the measured envelopes are shown in Fig. 40. Intensity envelopes are measured with a photodiode [14] and one calculates the measured Rabi frequency Ω_m from the relation $I_m \propto |\Omega_m|^2$. In what follows, measuring Ω_m will imply measuring I_m and converting to an amplitude.

Pulses have a length of $150 - 200$ ns due to the desire for high repetition rates in photon production. The beam waist and acoustic velocity fundamentally limit the implementation of pulse features at such time scales. To reduce rise-time, the beam was reshaped to a smaller waist than specified in the AOM data-sheet, trading diffraction efficiency for speed. For the STIRAP AOM's where fast rise-times were desired, beams

13 This work was carried out in collaboration with Marina Llano Piñeiro who did her external Bachelors Thesis in our group. She worked on the code implementation of the filters and performed some of the measurements. This project was conceived and overseen by me, and all the optical debugging and engineering was also done by myself.

14 Thorlabs PDA 10A-EC

were focused with 150 mm lenses to a beam waist of $\approx 20\,\mu$m. Measured rise-times were 18 ns for the Eylsa STIRAP path and 25 ns for the DL Pro STIRAP path.

While minor deviations in the measured amplitude shape were observed, these could be corrected by optimising the applied drive envelope $\Omega_{\text{in}}(t)$ to compensate for experimental imperfections. A more persistent issue arose in the double-pass AOM configuration: pulses were systematically shortened relative to the expected duration of the applied envelope. This effect was absent in single-pass setups and could not be eliminated by realignment of the double-pass geometry. The likely explanation is a slight misalignment between the forward and backward acoustic passes, such that the sound waves crossed the optical waist at different spatial and therefore temporal positions. This problem is exacerbated by the small beam waist required for fast rise times. Experimentally, delays of 10–20 ns were observed, corresponding to spatial displacements of $\approx$ 40–80 μm[15]. Such effects are negligible for static or long-duration pulses, but they become significant at ns time-scales. Indeed, the literature shows that many groups working with fast signals avoid this issue by using tandem AOMs, where two single-pass devices are placed in series instead of a double-pass configuration [77, 144]. By operating in a single-pass configuration and keeping the input voltage within the RF amplifier[16] limits ($\approx$ 0.25 V, consistent with a 50 Ω input impedance and 30 dB gain, corresponding to 0.282 V), signal distortion can be avoided.

The task is to determine an optimal AWG envelope $\Omega_{\text{in}}(t) \xrightarrow{\text{sys}} \Omega_m(t)$ such that $\Omega_m(t_i) \approx \Omega(t_i)$ for all time points t_i, where $\Omega(t)$ is the desired target envelope and $\Omega_m(t)$ the measured output. The measured data form a discrete vector $\mathbf{\Omega}_m \in \mathbb{R}^N$, representing a sampled approximation of $\Omega_m(t)$. A straightforward strategy would be to model the system by an analytic transfer function and apply its inverse. However, this approach does not generalise well across different AOM drive frequencies or amplitudes and is constrained in accuracy by the chosen functional form. A better alternative is to use an adaptive filter that learns the system response directly from

15 Using an acoustic velocity of 4.2 mm/μs for a G&H AOMO 3110-120.
16 RF Bay 12-30

input–output data, a standard method in signal processing for modelling time-varying systems. Non-linear distortions affecting amplitude envelopes are treated as systematic, consistent for a given input $\mathbf{\Omega}_{\text{in}}$, at least at a given point in time. While some random, measurement-related noise is present (e.g. from cables or photodetectors), it is insignificant compared to the systematic effects. To compensate for these distortions, an adaptive Normalised Least-Mean Square (NLMS) filter [99] is used to estimate the measured signal $\mathbf{\Omega}_m$. The filter iteratively adjusts its weights to minimise the mean square error (MSE) between the measured output and the predicted output as determined by a multiplicative pass through the filter, $\widehat{\mathbf{\Omega}}_m$. These vectors contain a discrete number of N samples; in what follows, the signal processing convention is adopted and $[n]$ denotes the n-th element of a vector.

To incorporate memory of past inputs, the filter uses a window of size n_w over the most recent input vectors. The filter computes a matrix of adaptive weights $\mathbf{w} \in \mathbb{R}^{N \times n_w}$, where each row corresponds to the filter weights at time step n, and produces a predicted output $\widehat{\mathbf{\Omega}}_m \in \mathbb{R}^N$. The prediction at each time step n is given by:

$$\widehat{\mathbf{\Omega}}_m[n] = \mathcal{E}_{\text{in}}[n,:] \cdot \mathbf{w}^T[n,:] \tag{73}$$

Here, $\mathcal{E}_{\text{in}} \in \mathbb{R}^{N \times n_w}$ is the matrix of input vectors, and the dot product is computed row-wise between input vectors and the corresponding filter weights. The matrix $\mathcal{E}_{\text{in}}$ is constructed from the input vector $\mathbf{\Omega}_{\text{in}} \in \mathbb{R}^N$ by forming delayed versions of the signal over a sliding window of size n_w. Each row of $\mathcal{E}_{\text{in}}$ contains the most recent n_w input samples at time step n, with zero-padding applied for indices $k < 0$:

$$\mathcal{E}_{\text{in}}[n,:] = \begin{bmatrix} \mathbf{\Omega}_{\text{in}}[n-n_w] & \mathbf{\Omega}_{\text{in}}[n-n_w+1] & \cdots & \mathbf{\Omega}_{\text{in}}[n] \end{bmatrix} \tag{74}$$

The optimisation update rule is given in App. A.9. This method is robust for dynamic systems with moderate stochasticity, and predicting outputs with the filter is

computationally inexpensive. One important parameter which needs to be fine-tuned is the filter window size n_w which captures the memory depth of the filter [17].

Two strategies can be used for optimising amplitude envelopes with an adaptive filter; one can either sample from the physical system to establish optimal filter weights (cf. App. Eq. 120) and then aims to find an optimal input envelope $\mathbf{\Omega}'_{in}$ through the filter $\mathbf{w}$ with a gradient based optimiser such that the predicted output approaches the desired output $\mathbf{\Omega}$. Alternatively, one can invert the filter prediction equation to determine the inverted weight matrix $\mathbf{w}^{-1}$, that predicts the inputted signal $\mathbf{\Omega}'_{in}$, given a measured signal $\widehat{\mathbf{\Omega}}_m$. The advantage of this approach is that one can directly input the desired output signal and determine the necessary input voltages in time that produces such an output by passing it through the inverted weight matrix $\mathbf{w}^{-1}$.

5.5.2 *Time-Dependent Frequency Chirps*

Ultimate atomic control and the highest population transfer efficiencies or fidelities also require precise control over the time-dependent frequency of the control pulses. This is particularly relevant for quantum information processing tasks requiring high fidelities as described extensively in Ch. 4. To measure the instantaneous frequencies of the optimised Pump and Stokes pulses in time, we use a homodyne detection scheme with a reference signal of fixed amplitude and frequency (in this case the optical pumping beams described in Fig. 13 in Ch. 3 are used, as they lie within a few hundred MHz of the desired Pump and Stokes beams) with respect to a pulse signal which corresponds to a desired STIRAP signal which is amplitude (and possibly phase) modulated in time. The measurement setup consists of a fibre beam splitter where the two signals are combined. To obtain a stable beat note, the input fields must have the same polarisation and comparable optical power. The resulting beat frequency is given by

17 The implementation was carried out using the `padasip` (https://github.com/matousc89/padasip?tab=readme-ov-file) Python package.

$f_{\mathrm{BN}} = |f_{\mathrm{ref}} - f_{\mathrm{pulse}}|$ and is determined digitally after additional filtering[18]. Spectral analysis is then performed using the Fast Fourier Transform (FFT), which provides global frequency content, and the Short-Time Fourier Transform (STFT), which enables time-local frequency analysis [170]. The time-dependent phase of the beat-note cannot be extracted directly in this setup, since measurement noise obscures the absolute phase information, leaving only the digitally processed frequency interpretable.

Harmonics of the beat note frequency f_{BN}, which originate from optical detection non-linearities or higher-order AOM order interference, are removed by digital low-pass filtering. Since the sampling rate ($f_s = 10\,\mathrm{GHz}$) far exceeds the photodiode bandwidth ($f_{\mathrm{BW}} \approx 150\,\mathrm{MHz}$), the detector itself band-limits the signal and prevents aliasing. Residual high-frequency noise is further suppressed by digital filtering. For a pulse of duration 150 ns sampled at $f_s = 10\,\mathrm{GHz}$, the number of samples is $N \approx 1500$. The frequency resolution of the discrete Fourier analysis is therefore $\Delta f = \frac{f_s}{N} \approx \frac{10^{10}}{1500} \simeq 6.7\,\mathrm{MHz}$. This resolution sets the smallest distinguishable frequency difference. As a result, chirps with magnitudes below $\sim 10\,\mathrm{MHz}$ cannot be reliably measured with such short pulses.

5.5.3 *Results*

Results for two optimised pulses are presented as outlined in Sec. 5.5.1. Both rely on initially driving the system with the theoretically optimal pulse shapes $\mathbf{\Omega}$ and finding the optimal filter weights $\mathbf{w}$ (cf. Eq. 73) or $\mathbf{w^{-1}}$ which describe the action of the AOM on the input signal in time or vice versa. $\mathbf{w^{-1}}$ transforms the measured oscilloscope signal $\mathbf{\Omega}_m$ into an input matrix $\mathcal{E}'_{in}$. From this one can input the desired output shape $\mathbf{\Omega}$ and directly determine the optimal input $\mathbf{\Omega}'_{in}$. This is more computationally efficient than optimising through a filter $\mathbf{w}$ and useful for frequent re-training which is inevitable with drifts or changes in the optical setup.

18 Both Python and Matlab provide the functionality for digital signal processing, with Matlab found to be more mature.

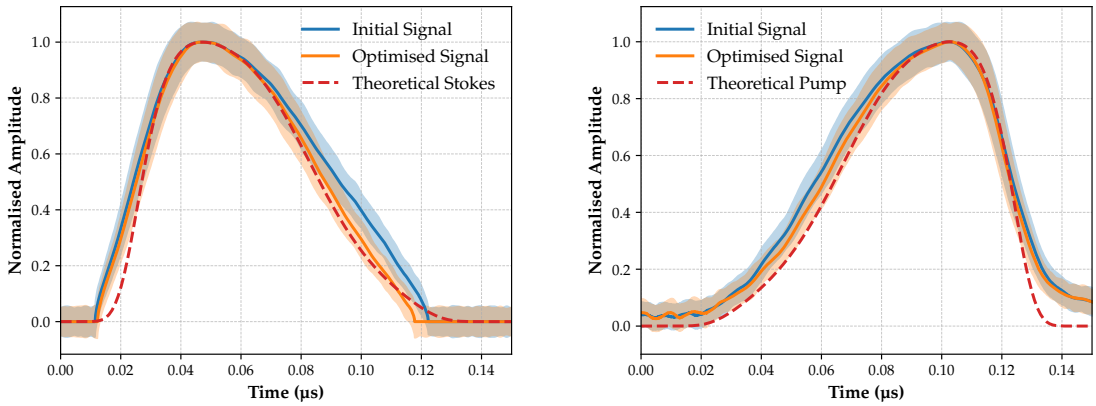

(a) Optimised Stokes Pulse

(b) Optimised Pump Pulse

Figure 41: Comparison of optimised and unoptimised driving on fixed frequency pulses with respect to the desired theoretical shapes (cf. Eq. 40).

Figure 41 shows the results for the inverse filter technique with optimised filter weight sizes of $n_w = 61, 21$ for the Eylsa and DL Pro lasers. The optimisation achieves a reduction in mean-squared error of 53% and 86% for the Pump and Stokes pulses respectively. One limitation, is that random noise on the photodiode means makes it difficult to achieve zero amplitude at the start and ends of the pulses. Moreover, measured negative amplitudes due to random electronic noise must also be corrected by a DC offset. Photo-diode bandwidth is limited to 150MHz, but data is oversampled to obtain enough points for accurate filter learning in the presence of random noise. This can be mitigated by repeating measurements several times and averaging. An example of a particular pulse is singled out here, but additional plots in App. Sec. A.10 show that this method is extremely versatile and applies to other desired pulse shape.

Spectral analysis with a digital low-pass filtering confirms agreement with the theoretical waveform: STFTs reveal the correct instantaneous frequencies, and the FFT peaks align with the expected f_{BN}. We were also able to measure larger frequency chirps of $\approx$ 20MHz. However, the limited resolution of our measurements imply that it is difficult to directly optimise the frequency evolution of the particular pulse shape in time. The experimental results, do show good agreement with theoretically predicted frequencies and thus attention is devoted to optimising time-dependent amplitudes

and ensuring that a time dependent chirp does not affect the time dependent amplitude by re-training the filter **w** with a chirped drive. A change in input RF powers or frequencies changes the imparted acoustic phase (cf. Eq. 118) and thus changes the time dynamics of the amplitude shaped pulse, requiring re-training. A plot in App. Sec. A.11 shows detailed frequency analysis of an example signal. For any novel pulse with a different frequency or maximal amplitude, it is thus necessary to re-optimise the system, which is fairly fast with the inverse filter method described previously.

5.5.4 *Power and Rabi Frequency Calibration*

The measured maximal laser power of the optimised pulses must be calibrated carefully to match the optimal peak Rabi frequency determined through simulations. For a beam of waist w_0 one can calculate the Rabi frequency Ω as a function of the laser power P as follows:

$$\frac{\Omega}{d_{F,m_F\rightarrow F',m'_F}} = \frac{d_{J,J'}}{\hbar} \cdot \sqrt{\frac{2P}{\pi\epsilon_0 c(w_0)^2}}, \tag{75}$$

where the Rabi frequency is expressed as normalised with respect to the component of the dipole matrix element pertaining to the particular hyperfine transition $d_{F,m_F\rightarrow F',m'_F}$. Given the limited operational amplitude range of the RF amplifier, of about $0.1 - 0.25$V before amplifier non-linearity becomes problematic, one can measure $P(V)$ and thus $\Omega(V)$, where V is the RF signal amplitude by the AWG using Eq. 75. This recalibration needs to be done frequently, as drifts in power and polarisation change the measured power.

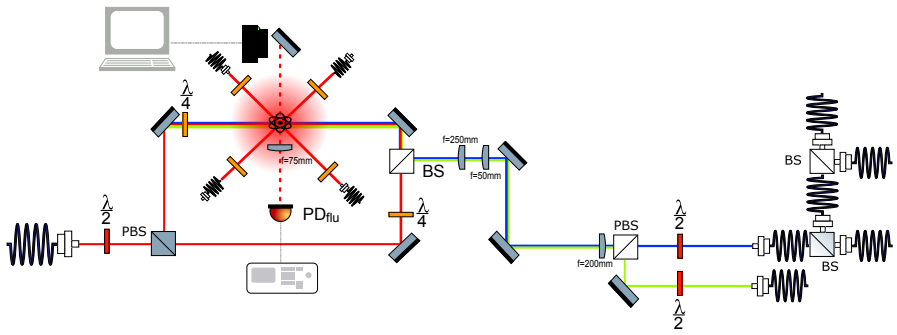

Figure 42: Shown is the experimental setup for imaging optical molasses with different STIRAP pulses to estimate state transfer efficiency through a fluorescence measurement. Atomic cloud is shown in red and six counter-propagating MOT beams are also shown in red. In green and blue are the polarisations which are along the quantisation axis and orthogonally polarised to that. These are mixed into the central MOT beams with a non-polarising beam splitter. All lasers are fibre coupled to this setup as they are prepared on a different optical table

5.6 COHERENT POPULATION TRANSFER

Isolating the population transfer process relevant to the n-photon burst scheme (as introduced in Fig. 13 of Chapter 3) enables direct measurement and optimisation of transfer efficiency [19]. In this experiment, the cavity is excluded from the setup, making it more natural to drive the Raman transition from $|F = 2, m_F = 0\rangle$ to $|F = 1, m_F = \pm 1\rangle$, since the initial state can be reliably prepared via optical pumping.

Although a previous chapter highlighted the advantages of using the D_1 transitions for Raman coupling, experimental constraints during the present work necessitated the use of the D_2 line due to frequencies of the existing lasers.

The experimental configuration is illustrated in Fig. 42, where a MOT is formed using six counter-propagating beams. Overlapping Raman beams are aligned to enable STIRAP-based coherent population transfer. Once the MOT is loaded, it is rapidly converted into an optical molasses stage. The STIRAP pulses are applied during this phase, and the final population distribution is inferred through fluorescence readout. A detailed account of the timing and control sequence is provided in Tab. 10.

The STIRAP beams are focused to a beam waist of approximately 100, μm, while the magneto-optical trap (MOT) has a slightly larger spatial extent. This configuration balances the need for sufficient optical power with maintaining good spatial overlap with the majority of molasses atoms, but crucially not with the entire cloud. Calibration of the Rabi frequency is performed as described in Sec.5.5.4, and the time-dependent pulse shaping is optimised following the procedure outlined in Sec.5.5.1. The beams are spatially overlapped on a Polarising Beam Splitter (PBS), such that they are orthogonally polarised, the Eylsa Pump beam is polarised along the quantisation axis (π in the atomic frame, H in the lab frame), while the Stokes DL Pro beam is vertically polarised in the lab frame which corresponds to a linear combination of σ_+ and σ_- in the atomic

19 Some of these experiments were done in collaboration with Marina Llano Piñeiro and Matthew King, but I led the exeprimental work.

Table 10: Control sequence for shaped pulse experiment with fluorescence readout where Channels correspond to the relevant DAQ Card Channels (Ch.) or AWG Channels (AWG Ch.) which are controlled at various points in time.

Time (µs)	Action	Channel
0	MOT coils on	Ch. 16
0	Cooling beams on	Chs. 4-6
0	Repumping EOM on	Ch. 20
14550	Start ramping imaging trigger	Ch. 23
15000	MOT coils off	Ch. 16
15000	Cooling beams off	Chs. 4-6
15100	Horizontal magnetic bias field on	Ch. 18
15100	Eylsa Repumping off	Ch. 20
16100	Trigger AWG	Ch. 22
16105	STIRAP pulses on	AWG Chs. 1 & 2
16150	Horizontal magnetic bias field off	Ch. 18
16250	Imaging beam on	Ch. 6 (or AWG Ch. 2)
16750	Imaging beam off	Ch. 6

frame. Both beams are then mixed and overlapped with the MOT beams in a Non Polarising Beam Splitter (NPBS) such that pulse can maintain the desired polarisation.

To estimate the atomic population remaining after (or before) the STIRAP pulses, fluorescence detection is performed via a cycling transition. Atoms in the $|F = 2\rangle$ state can be readout by addressing them with the σ^+-polarised Eylsa beam near resonant with the transition $F = 2 \rightarrow F' = 3$ of the D_2 line. Since this transition does not couple to the $F = 1$ ground state, it provides a reliable measure of the population that remains in $F = 2$. Furthermore, the MOT beam geometry—with counter-propagating beams along the central axis—helps to minimise Doppler shifts, ensuring a relatively high contrast optical readout.

The fluorescence readout level is calibrated by preparing atoms in the $F = 2$ state (in a random mixture of Zeeman states) and measuring the integrated fluorescence signal in the absence of any STIRAP pulses. One could use a separate imaging beam to probe population in the $F = 1$ level. However, the lack of an available cycling transition makes this too low in signal to noise to provide a high contrast readout signal.

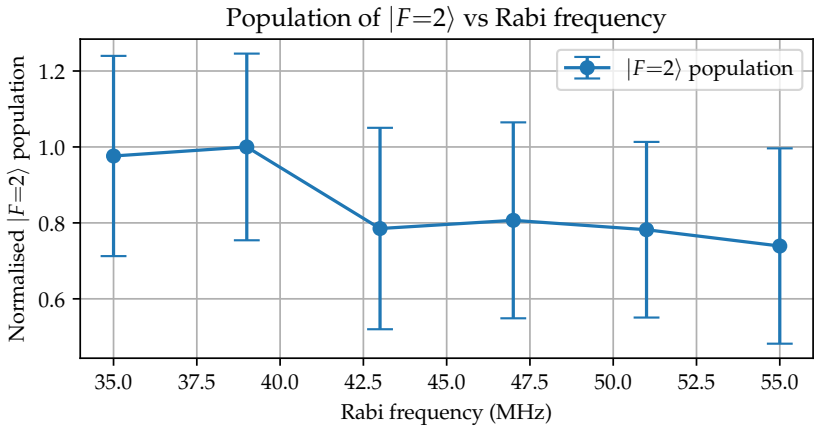

(a) Comparison of $|F = 2\rangle$ population for different Rabi frequencies with optimised shapes as outlined in Sec. 5.5. One can see that increased Rabi frequencies lead to increased depopulation (i.e. reduced population) of the initial level and 55MHz is the limit where we can go with the available laser power and beam waist used in this experiment. This is broadly consistent with the numerical simulations and optimal Rabi frequencies determined in Ch. 3

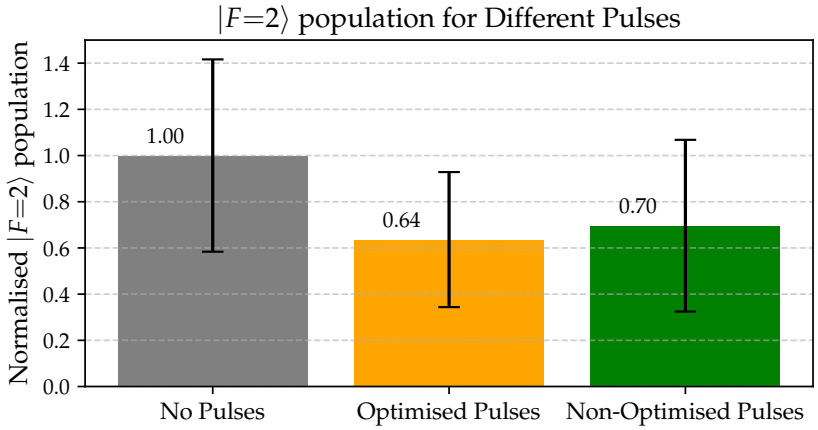

(b) Comparison of $|F = 2\rangle$ depopulation for different pulses, no STIRAP pulses (grey) is used as a reference for the maximal fluorescence readout level and then non-optimised STIRAP pulses (ignoring AOM non-linearities) and optimised pulses which are optimised with the NLMS filter method described in Sec. 5.5. All fluorescence integrals are averaged over 30 independent shots We see that the optimised shaping provides a small improvement in depopulation.

Figure 43

Experimental results shown in Fig. 43 b), based on the pulse sequence outlined in Table 10, indicate that optimised pulse shapes lead to a slight improvement in population transfer compared to non-optimised pulses. Additionally, a Rabi frequency sweep in Fig. 43 a) suggest that approaching the numerically optimal peak Rabi frequency from below enhances depopulation from the initial state. However, due to the relatively high shot noise in the fluorescence detection, reflected in the large error bars in Fig. 43, it is difficult to draw definitive quantitative conclusions from this data. Efforts are underway to reduce the shot noise by improving MOT brightness in future experiments. Moreover, it should be noted that due to short desired pulse times, they require more laser power than is available for a beam waist that completely overlaps with the molasses. Hence, even for perfect STIRAP efficiency, one would not observe complete depopulation in Fig. 43, as some atoms are not addressed by the existing STIRAP beams.

Fluorescence is detected using an avalanche photodiode [20], and the integrated signal during the readout pulse is compared to the background noise level. This yields an estimate of the population transfer efficiency. However, a limitation of this technique is that it is not state-selective, it only resolves the total population in each hyperfine ground level. Consequently, atoms that decay into the correct hyperfine level via spontaneous emission may still occupy undesired magnetic sub-states. Thus, the reported transfer efficiency represents an upper bound on the true fidelity of population transfer. The fluorescence signal was processed by integrating the detected intensity and normalising it to the background signal of the magneto-optical trap (MOT), in order to account for fluctuations in atom number caused by potential drifts in MOT beam powers, for example due to instability in fibre coupling. Additionally, background subtraction was performed using the measured noise floor in the absence of any fluorescence signal. The data processing pipeline is described in Appendix A.12.

A further experimental limitation was the insufficient optical power available for efficient optical pumping into a well-defined Zeeman sublevel. As a result, it was not

20 Thorlabs APD110A

possible to prepare atoms in a single initial state required for driving the STIRAP transitions on the timescales discussed in Section 3 (i.e., < 200ns). Instead, the atoms remained distributed over multiple Zeeman sublevels within the $F = 2$ ground-state manifold. This introduces additional complexity, as multiple dark states with varying two-photon detunings and dipole matrix elements become involved, leading to state-dependent coupling efficiencies and contributing to reduced overall transfer fidelity.

5.7 OUTLOOK

This chapter has outlined the experimental progress towards realising novel atom-cavity experiments. Building on a complex and well-established experimental infrastructure, the work presented here supports the practical implementation of schemes proposed in Chapter 3 and Chapter 4, targeting applications such as boson sampling and time-bin entanglement.

Key developments include the construction of a new laser preparation setup with various AOM's and the generation of tailored pulses using optimised acousto-optic modulation. These efforts culminated in observations of improved population transfer, validating an essential component of the single-photon production protocol.

Although a complete demonstration of deterministic single-photon generation was not achieved within the time-frame of this thesis, the insights gained highlight clear directions for future development. Some more real time optimisation of pulse sequences based on experimental data is possible as will be outlined briefly towards the end of Ch. 7. Moreover, it would be advantageous to move to the D_1 line with new lasers for coherent state re-preparation between photon emissions due to the simplified excited state coupling and slightly reduced spontaneous emission rate, as validated with extensive simulation in Chs. 3 and 4.

6

SCALING PULSE LEVEL CONTROL OF QUANTUM EXPERIMENTS

This chapter is based on the work published in [76][1].

6.1 INTRODUCTION

To realise quantum computing's potential for solving practical problems, quantum devices must scale beyond the capabilities of current systems. This motivates an investigation of the classical computational resources required to store and represent radio-frequency (RF) signals, which are essential for controlling qubits in most quantum computing architectures. Several platforms have demonstrated progress, with superconducting qubits, trapped ions, and neutral atoms among the most mature. The pulse engineering techniques discussed in this chapter are broadly applicable across different qubit technologies but are particularly well-suited for pulses that are smooth and continuous in both frequency and amplitude, lasting on the order of microseconds or longer.

In Chapter 5, progress towards the control of an atom-cavity system using frequency and amplitude optimised pulses was shown. However, the setup presented there does not readily scale to larger, more complex quantum systems. While an Arbitrary Waveform Generator can provide flexible and high-fidelity control, AWG's are expensive, limited to a small number of output channels, and thus only practical for systems com-

1 The quantisation aware fitting algorithm was developed by myself with valuable input from Jan Snoeijs and all atomic physics simulations were also carried out by myself. Jan Snoeijs focused on the Field Programmable Gate Array (FPGA) hardware implementation and the work was discussed amongst all authors.

prising a few qubits. In addition, the communication with AWG's via USB interfaces and reliance on external triggering introduces further latencies that exceed the time scales required for any real time operations in quantum systems. For example, in the experimental setup decribed in the previous chapter, loading waveforms of durations of $\approx \mu$s at a 1GHz sampling rate takes on the order of seconds which is a time overhead that makes real time wave-form reloading impossible for experiments.

The pulse shapes used to control atomic or ionic qubits - whether for gate operations or atom transport in optical tweezers - are typically smooth and continuous in amplitude and frequency, a requirement for high-fidelity operations. Gate durations vary significantly across platforms: state-of-the-art single- and multi-qubit gates typically occur on microsecond timescales [138, 197], with faster Rydberg gates reaching hundreds of nanoseconds [77, 112], and slower microwave gates extending to milliseconds [98, 185]. Atom rearrangement in optical tweezer arrays often requires pulses lasting several milliseconds to achieve defect-free configurations for analogue simulation or fault-tolerant quantum computing [17, 73, 205], and to facilitate entanglement via atom movement [32]. While the simulations in this chapter focus on neutral atoms, the method is broadly applicable to other platforms such as trapped ions, which operate on similar timescales to neutral atoms, and potentially to superconducting qubits when moving to higher output sampling rates.

Generating pulses with sufficient instantaneous bandwidth (hundreds of MHz to 1 GHz) often requires storing several to hundreds of kilobits per pulse [2] if all modulation data is held in local Field Programmable Gate Array memory. Since pulses are typically calibrated per channel, memory demands scale significantly in large quantum systems. Moreover, FPGA memory load times scale with memory size, regardless of the transfer method. Compression techniques that preserve pulse fidelity while reducing memory footprint can thus play a key role in enabling the realisation of scalable quantum control systems.

2 especially for durations exceeding a microsecond

To address these challenges and enable more scalable experimental quantum control, this chapter presents a new FPGA-based solution as presented in [76]. FPGA's allow for significantly reduced latency, high-speed parallel signal generation, and the integration of a large number of control channels. A key limitation when moving to FPGA-based control, however, is the restricted per channel memory available to store time-resolved pulse parameters, particularly when high output sampling rates are required. To overcome this constraint, an optimised pulse compression scheme designed for fixed-quantisation hardware is introduced. This approach utilises cubic spline interpolation, a well-established method for representing smooth, continuous signals [22, 31, 201], to represent control waveforms as sequences of cubic polynomials across segmented intervals. Spline fitting techniques are reviewed and it is highlighted how knot placement critically affects interpolation accuracy. It is also shown how pulse symmetries can improve data compression. A drawback of the FPGA-based spline decompression method is that the truncation of coefficients to fixed point format introduces cumulative errors, constraining segment length and compression due to increasingly inaccurate pulse representations. This is addressed with a two-stage curve fitting process, substantially reducing the memory footprint with cubic spline compression, while maintaining high interpolation accuracy between output samples due to quantisation aware fitting. Simulations of single-qubit population transfer and atom transport demonstrate that this scheme enables the realisation of high-fidelity control while dramatically lowering memory requirements as compared to a standard AWG. These advances are instrumental for scaling up the number of supported qubits and respective operations per qubit in future quantum computing architectures where memory will be a limited resource. It should be noted here, that even in industry, it is not clear that there exists a feasible blueprint for building a control system for a 1 million qubit quantum computer, even if it miraculously came into existence today.

6.2 METHODS

6.2.1 *Floating Point Cubic Spline Fitting*

Fitting smooth functions with piecewise cubic polynomials has been well studied for decades [201]. Consider a smooth function $y(t)$ (alternatively, this could also be a number of (t_i, y_i) pairs) defined over a domain $[t_i, t_f]$). A cubic spline interpolation with $k+1$ piecewise polynomials and k knots is formulated as follows:

$$y(t) := \begin{cases} P(t, \vec{p}_0), \text{ for } t_i \leq t \leq \tau_1 \\ P(t, \vec{p}_1), \text{ for } \tau_2 \leq t \leq \tau_3 \\ \vdots \\ P(t, \vec{p}_k), \text{ for } \tau_k \leq t \leq t_f \end{cases} \tag{76}$$

where $\vec{\tau} = (\tau_1, ..., \tau_k)$ is the ordered knot vector and the polynomial P is defined as $P(t, \vec{p}) = p^0 + p^1 t + p^2 t^2 + p^3 t^3$, where $\vec{p}_i$ is the polynomial coefficient vector. When performing the fit, it is possible to impose not only continuity of first- but also second- or higher-order derivatives at the knot positions τ_i [3]. Another important choice for the fitting routine is the error function. Arguably, the most common error function is the least squares error function $\epsilon = \sum_i (P(t) - y_i)^2$ for desired data points y_i and points generated by the cubic polynomial $P(t)$, but one can also consider an absolute linear error [36]. In what follows the least squares error criterion is adopted since it leads to more accurate fits for sparsely sampled higher-order polynomials and, more importantly, this allows the error to be treated as a continuous and differentiable quantity which is instrumental for using gradient-based minimisation algorithms. In general, using a larger number of knots and shorter segments improves the accuracy of the fit,

3 When moving to an FPGA, time will be quantised in terms of the sampling frequency of the output, whereas here it is treated as continuous.

but it should also be noted that the memory load of the pulse representation scales linearly with the number of knots.

The increase in memory load with the number of segments motivates an optimisation of the position of the segment boundaries, as it can noticeably improve the precision of the fit without increasing the memory load. As it turns out, the global optimisation of the knot positions is a computationally complex non-linear optimisation problem. Recently, some progress has been made in formulating it as a mixed-integer quadratically constrained problem, as well as a branch and bound algorithm [156]. Previously, only locally optimal optimisation methods have been demonstrated [36, 51] for the positioning of the segment boundaries. Even though recent results show promise, the computation times for longer pulses with a large number of points to fit to are fairly long [4] to render this solution practical on standard hardware. Additionally, the improvements in fitting accuracy for typical pulse shapes used in quantum computing are lower than that gained by exploiting pulse symmetry or fitting per the FPGA level accuracy described in Sec. 6.2.3. However, for certain critical applications where this optimisation is a one-time task, it may be worthwhile.

Typical pulse shapes used for implementing quantum gates and other quantum operations, such as atom transport, often exhibit waveforms that have some symmetry in their amplitude or phase envelope. In particular, pulse symmetry around the centre of the amplitude envelope is common for many pulses such as Gaussians (cf. [33]), Sinusoid (cf. [109]), or Blackman ramps (cf. [68]) for atom shuttling waveforms). One can utilise such symmetries to further enhance pulse compression by dividing the pulse envelope into an even number of segment with a symmetry centre. As explained in more detail in [76] this can reduce the memory load on the FPGA by a factor of two by effectively replaying the first half of the stored pulse by inverting the time direction.

4 order of hours for 6 knots with a few hundred data points on Intel Core i7-1185G7 @ 3 GHz with 8 logical cores

6.2.2 *Implementation on FPGA*

The hardware implementation uses Direct Digital Synthesis (DDS) to generate RF pulses from digital circuitry, using FPGA platforms. First a pure sinusoidal waveform is generated using a Numerically Controlled Oscillator (NCO) modulated in amplitude or phase where the modulation data is loaded and stored in memory and multiplied with the NCO output to produce a pulse at baseband frequency [38] [184] [64]. This signal is then converted to the analogue domain through a Digital to Analogue Converter (DAC). Additional electronics are used for amplification, filtering and up-conversion (if the desired frequency is outside of the baseband).

In this chapter, the focus is on the modulation of the baseband signal, which is the core part of the signal generation chain. More details on the FPGA implementation are found in [76].

The analytical function for cubic spline polynomials has a non-negligible implementation cost in digital circuitry, mainly due to the multiplication operations involved. In an integer representation, which is more resource-efficient than floating-point, the result requires twice as many bits to represent than the operands. The result is often truncated or rounded to maintain constant multiplier size, particularly in applications requiring cascaded multiplication, which introduces numerical errors when compared to a floating-point representation and requires careful consideration.

A coefficient transformation enables the computation of a digitally sampled cubic polynomial recursively as shown in [37]. The recursive expression is purely additive, resulting in lower-cost hardware implementation in terms of FPGA resource utilisation and maximum achievable clock frequency. The output voltage y at sampling time t is defined as:

$$y(t) := \begin{cases} \alpha_0 & \text{when } t = 0 \\ \alpha_{t-1} + \beta_{t-1} + \gamma_{t-1} + \delta_{t-1} & \text{when } t > 0 \end{cases} \tag{77}$$

The coefficients $(\alpha, \beta, \gamma, \delta)$ follow the following recursive definitions:

$$\begin{aligned} \gamma_t &= \gamma_{t-1} + \delta_0 \\ \beta_t &= \beta_{t-1} + \gamma_t \\ \alpha_t &= \alpha_{t-1} + \beta_t \\ \delta_t &= \delta_0 \end{aligned} \tag{78}$$

The initial coefficient definitions follow from generalising the recursive formula with respect to cubic polynomial coefficients such that $y(t) = \alpha_t = p^0 + p^1 t + p^2 t^2 + p^3 t^3$ (cf. Eq. 76):

$$\begin{aligned} \alpha_0 &= p^0 \\ \beta_0 &= p^1 - p^2 + p^3 \\ \gamma_0 &= 2p^2 - 6p^3 \\ \delta_0 &= 6p^3 \end{aligned} \tag{79}$$

The symmetry exploitation is implemented analogously for the coefficients which describe the first $N/2$ segments, for a pulse which is segmented into N total segments. Its recursive formula is very similar to that given in Eq. 77, and a detailed description is given in [76].

The recursive algorithm for cubic splines It is deployed on FPGA platforms that integrate radio frequency components and processing capabilities on a single chip. Firstly, a quantisation of the modified polynomial coefficients to a fixed-point representation is performed for which the bit-width was determined from numerical analysis. This analysis showed that regardless of the chosen bit-width, errors would accumulate inevitably due to the recursive nature of the computation. Increasing the fixed-point precision would increase the memory footprint and in turn, reduce the effective memory compression factor. Our hardware implementation uses a 36-bit fixed-point representation for the coefficients (i.e. for β) with 20 bits for the fractional part. This is

a compromise between the accumulated error and the cost for pulses in the range of up to a few to hundreds of microseconds. It is also the bit precision that was used in the latest implementation of this algorithm in the context of pulse generation for qubit control [37]. The implementation is a pipelined structure where the previous result is added to the new input in each accumulator. As observed from Eq. 77, $y(t)$ is equal to α_t at every iteration which allows for an implementation with 3 accumulators. More details on the hardware pipeline is found in [76].

The hardware implementation can be adapted to support symmetry-aware fits where, in the simplest case of a single exploited symmetry, the latter pulse half is generated from compressed data only available for the earlier pulse half. Only lightweight modifications are required, compared to the baseline implementation. They mainly consist in reloading the pipeline with the final coefficients of the first half, introducing a sign inversion, and including control logic to time the switching between computing the recursive function in forward or backward directions.

6.2.3 *Quantisation Aware Fit*

Typically, gradient-based minimisation algorithms, such as the Python Scipy CubicSpline [236] module, are used to fit functions to a representation with cubic spline segments in floating point arithmetic, i.e. operations and coefficients can be carried out and stored at a high level of accuracy (typically 64 bit). Such algorithms are very mature; they are fast, well understood, and well documented. The hardware implementation of the recursive algorithm described in Sec. 6.2.2 used to represent cubic splines is shown in Re. [76]. Note that the FPGA representation of the linear combinations of cubic polynomial coefficients α, β, γ and δ is fixed point, not floating point. This implies that there are rounding errors that arise from the quantisation, depending on the number of fractional bits used in the fixed-point representation (Figure 44 shows

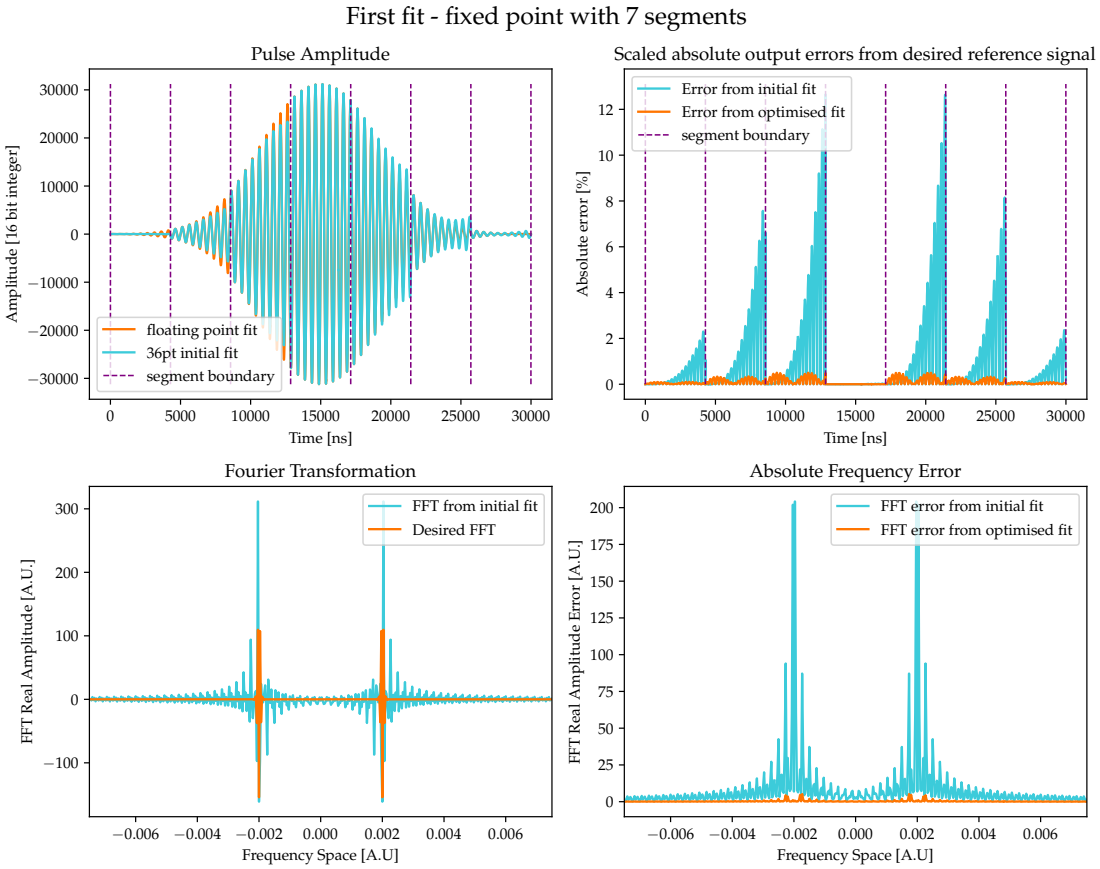

Figure 44: Shown is the effect of the quantisation-aware fitting method for a Gaussian pulse of length 30000 partitioned into 7 segments and multiplied with an NCO frequency of $f = 2$MHz for illustration purposes. Top-left: Comparison of the ideal floating point fit of the desired Gaussian amplitude envelope ($y_{ref}(t) = \sin(2\pi f t)e^{-(t-15000)^2/(1.6e^7)}$) to the quantised 36pt initial fit. Top-right: Absolute percentage error of the quantised fits with respect to the desired output. The optimised fitting method reduces the error by an order of magnitude. Bottom-left: Fourier Transformation of the desired signal and the initial fit (where x=0 corresponds to the centre frequency of 2 MHz). Bottom-right: Absolute error in Fourier Transform between quantised fits and the desired signal. The optimised fit again improves the error by an order of magnitude. The cubic polynomial coefficients for two example segments are shown for reference in Tab. 11.

	Initial Coefficients (5 s.d.)	Optimised Coefficients (5 s.d.)
α_{s2}	8821.7	8821.7
β_{s2}	3.7681	3.**3365**
γ_{s2}	0.0012048	0.00**18556**
δ_{s2}	$-6.2639e-07$	$-6.$**3287**$e-07$
α_{s3}	27812	27812
β_{s3}	3.1786	3.1**805**
γ_{s3}	-0.0014799	-0.0014**820**
δ_{s3}	$-9.3881e-10$	**2.7533e − 07**

Table 11: Example of Bowler coefficients for initial and optimised fits in Fig. 44. The Bowler coefficients are defined in Eq. 77 and the subscript denotes the segment index. It is evident that the modifications of the linear, quadratic and cubic polynomial coefficients are minor, but enough to account for the cumulative truncation errors. Here the third and fourth segments are chosen for reference. The Gaussian, which is partitioned into 7 segments, is defined over a range of 20 μs and reads $y(t) = sin(2\pi ft)e^{-(t-10000)^2/(4*10^6)}$.

a total word length of 36 bits with a fractional accuracy of 20 bits). This leads to the following error accumulation when decompressing the pulse envelope:

$$\epsilon_{seg}[n] = \epsilon_{\alpha_0} + n\epsilon_{\beta_0} + \frac{n(n+1)}{2}\epsilon_{\gamma_0} + \frac{n(n+1)(n+2)}{6}\epsilon_{\delta_0}, \tag{80}$$

where the ϵ_i denote the quantisation errors associated with the i-th coefficient described in Eq. 79, i.e. when truncating any bits that lie outside the realm of the fixed point quantisation and n denotes the number of clocks (i.e. addition) cycles per segment. The accumulated errors depend strongly on the number of sum operations n performed per segment. A larger number of additions per segment leads to better compression factors but also to an increase in the accumulated errors which are $O(n^3)$. For sufficiently large n, ϵ_δ starts dominating, and it typically becomes impossible to output non-trivial smooth pulse envelopes. For a typical example of a Gaussian pulse, with a 36-bit representation, for segment lengths exceeding 1000 rounds of addition, the accumulated errors become severe as shown in Fig. 44 and exceed the tolerances required by most quantum computers (cf. Sec. 6.3). Note that 2000 rounds of addition

corresponds to a segment length of 2000 ns at 1 GSps, or 4000 ns at 500 MSps, typical sampling rates used for pulse generation in quantum computing, which are pulse lengths realistically required for atomic qubits. These errors are particularly problematic as kinks create undesired frequency components which adversely affect the fidelity of quantum operations. One could reduce the individual segment length and use more segments to represent the pulse, but this reduces the compression factor and increases the memory load, which is undesirable in a memory-limited setting. Increasing the length of the bit representation of the coefficients also decreases the compression factor. Moreover, if the number of samples per segment is large enough, any fixed-point quantisation will introduce significant errors, particularly towards the end of an individual segment, due to the cubic dependence of the segment error function (see Eq. 80).

In [76] a strategy was developed to mitigate this, which relies on a two-stage fitting process:

1. Fit the desired function with a standard cubic spline fitting algorithm as explained in Sec. 6.2.1 which relies on floating point arithmetic.

2. Use the coefficients obtained in the first fitting round as the starting point for a second fitting algorithm which incorporates the fixed-point accuracy of the recursive algorithm described in Eq. 77.

The second step requires an evaluation of the error between a floating point implementation of Eq. 77 and a quantised implementation to construct an adequate cost function for the optimisation algorithm. One approach would be to directly implement the recursive algorithm described in Eq. 77, but this is not computationally efficient. Instead, using the error function function described in Eq. 80 is much more computationally efficient.

The constant coefficient $\alpha_0 = p^0$ is kept fixed for all segments, but the coefficients $\beta_0, \gamma_0, \delta_0$ for each segment are treated as variables which are to be optimised as $(\beta_0^*, \gamma_0^*, \delta_0^*)$.

The cost function per segment $\mathcal{E}$ for the quantisation-aware fitting method is given by:

$$\mathcal{E}_i = \sum_{n}^{N} (y_n - \left\lfloor y_n + \epsilon^*_{seg}[n] \right\rfloor)^2, \tag{81}$$

where n describes the number of additions performed per segment, y_n describes the desired pulse amplitude and ϵ^*_{seg} the quantisation error associated with the optimised coefficients $(\beta^*_0, \gamma^*_0, \delta^*_0)$ with respect to the floating point coefficients determined in the first fit $(\beta_0, \gamma_0, \delta_0)$. The least squares error can be minimised with a variety of non-gradient-descent-based optimisation methods. Gradient based methods are not efficient, because the cost function is not smoothly differentiable with respect to the input coefficients. An evolutionary algorithm was found to have a particularly good trade-off between speed and accuracy. For the examples shown in this chapter the runtime of the algorithm was on the order of minutes on a standard laptop [5]. In line with current practice in quantum computing, one would expect this optimisation routine to be carried out before quantum algorithm runtime.

6.3 BENCHMARKING

In what follows, the quantisation-aware cubic spline fitting routine is applied to simulations of a particular atomic single qubit gate as well as atom transport. An explicit comparison with the floating-point fit as well as the AWG implementation is provided, and the optimised fit is shown to exhibit a particularly favourable trade-off between the fidelity of the operation and the memory required to store the waveform.

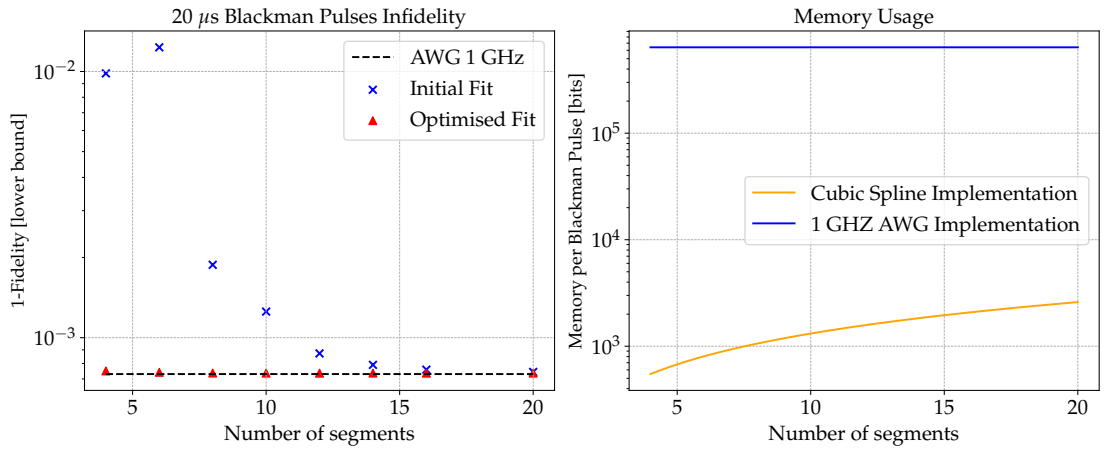

Figure 45: The STIRAP pulses are two fixed length 20 μs Blackman pulses used to implement an $X(\pi)$ gate and compare the fidelity (left) as well as the memory usage to store the pulse (right) between the optimised fit, the floating point fit and the AWG implementation. The top plot shows that as the number of segments increases, the optimised fitting fidelity no longer differs from the fidelity obtained with the initial fit. However, the point where the optimised fit reaches comparable fidelity to the AWG is already at ≈ 6 segments (compared to ≈ 20) which implies a significantly lower memory overhead. This is exhibited in the right plot, where one can see that one obtains almost three orders of magnitude of compression for 6 segments when comparing the cubic spline and AWG implementations.

6.3.1 *Coherent Quantum Population Transfer in* ^{87}**Rb** *Atom*

To benchmark the fitting methods against the state of the art, a simple model of full coherent quantum population transfer (a particular type of a single-qubit gate) is considered, and the effect of the pulse envelope shaping on the fidelity is determined. Single qubit gates in neutral atom architectures can be implemented in multiple ways. Here, STIRAP (introduced extensively in Sec. 2.1.3) is used for the coherent transfer of population between two ground states, which encode either of the qubit basis states [6]. This requires two laser pulses that have equal detunings to a common excited state to establish a coupling between the two ground states. In this section, one considers the time-dependent Rabi frequencies $\Omega_1(t)$ and $\Omega_2(t)$, as well as the detuning from the excited state Δ. In order to perform a single qubit rotation with the dark state (cf. Sec. 2.1.3), amplitude [7] modulation of the pulses $\Omega_{1/2}(t)$ is required. The optimised curve

5 Intel Core i7-1185G7 @ 3 GHz with 8 logical cores
6 While this is not a universal single qubit gate, it functions as a special exemplary case.
7 In some cases phase modulation can be used to implement amplitude modulation [139].

fitting algorithm is benchmarked on a pulse sequence used to realise qubit population transfer in ^{87}**Rb** where the two qubit states are given by the hyperfine clock states $F = 2, m_F = 0 \leftrightarrow |0\rangle$ and $F = 1, m_F = 0 \leftrightarrow |1\rangle$ which, at a particular external magnetic field, are, to first order, insensitive to fluctuations in the external magnetic field. Ω_1 and Ω_2 are both σ^- polarised. A single qubit gate implementation with two beams was experimentally demonstrated with ^{87}**Rb** atoms in [115, 247], as well as with microwaves [29, 207] and with a single laser beam which is far detuned from the transitions [139].

Following the simplified experimental model of ^{87}**Rb** introduced in Chapters 3 and 4, this analysis isolates the error induced solely by imperfections in the amplitude modulation of the pulse envelope. Other potential error sources—such as phase and frequency noise, laser power fluctuations, magnetic field instability, imperfections in the electronic signal chain, laser frequency drift, optical aberrations, and atomic motion are deliberately neglected.

This approximation is justified in the small-error regime typical of quantum computing, where errors are linearly additive to first order. Under this assumption, only scattering-induced errors are retained, while all others are excluded to focus exclusively on amplitude modulation imperfections. These imperfections are sufficiently pronounced to be visible in Fig. 45.

Due to the highly idealised nature of the model, all fidelities reported in the following results should be interpreted as strict upper bounds on what is experimentally achievable, given the inevitable presence of additional sources of error in real systems. For the purposes of this benchmark, the atom is initialised in the state $|0\rangle$ and subjected to an $X(\pi)$ rotation[8], with the goal of transferring population completely to $|1\rangle$. The state evolution is governed by the Hamiltonian defined as in Sec. 2.2.3, along with the collapse operators $\mathcal{L}(\Gamma)$ that account for spontaneous emission losses (see Chapters 2 and 3 for more details).

The fidelity is defined as the overlap between the simulated final density matrix ρ_{final} and the desired target state $\psi_{\text{des}} \leftrightarrow |1\rangle$, expressed as $\mathcal{F}(\rho_{\text{final}}, \psi_{\text{des}})$ [161]. The time-

8 This approach could be extended to other $X(\theta)$ rotations, a single $\theta = \pi$ case is considered for simplicity.

dependent amplitudes of Ω_1 and Ω_2 are both taken as Blackman envelopes[9], with a fixed temporal offset of $0.3L$, where L denotes the duration of each pulse. This configuration ensures that the root-mean-square Rabi frequency $\sqrt{\Omega_1^2(t) + \Omega_2^2(t)}$ remains approximately constant throughout the interaction. The total gate time is therefore $1.3L$, with $L = 20$ μs as shown in Fig. 45. The peak Rabi frequency is fixed at 5 MHz, and the detuning is set to a constant value of -100 MHz relative to the $F' = 1$ excited state. An external magnetic field of approximately 0.7 mT is applied to induce a ground-state Zeeman splitting of 5 MHz. All parameters used are compatible with existing experimental capabilities.

The results show that in this highly idealised simulation of a single ^{87}**Rb** atom, the quantisation-aware fitting routine has a noticeable effect on the fidelity and leads to a negligible reduction in fidelity compared to the AWG implementation, where samples representing the pulse are stored in memory at the sampling frequency, particularly for longer pulses with fewer numbers of segments. In Fig. 45 one can observe that for fixed-length pulses, the quantisation-aware fitting routine is particularly significant up to around 14 segments, but its effects trail off when individual pulse segments contain fewer than 1000 samples (when the number of segments exceeds 20) [10]. This figure shows that to achieve a maximal fidelity upper bound (given by the AWG pulse representation) one only requires 6 (with optimised fitting) compared to 20 (without optimised fitting) segments, which more than halves the required memory to store the pulse information compared to the non-optimised fitting, and is a compression ratio of ≈ 796 when comparing the AWG approach and the cubic spline with an optimised fit with 6 segments.

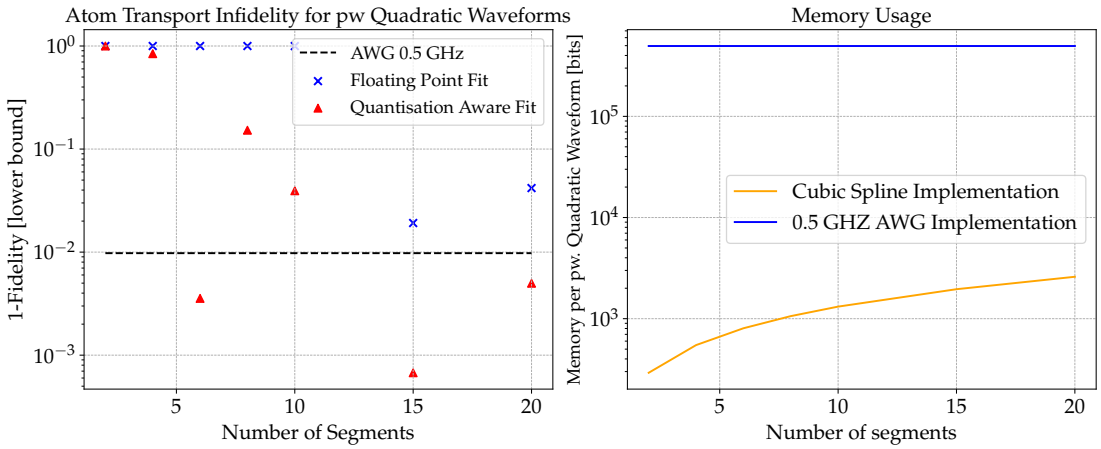

Figure 46: Comparing fidelities (left) across the floating point fit, the optimised fit, and the AWG implementation for an atom transport with a piecewise quadratic pulse (cf. Eq. 84 with a length of 31 μs and a trap depth of 20 μK) and also compare their respective memory usage (right). The left plot shows that, for a small number of segments, the cubic polynomial representations are insufficient for high-fidelity transport even with the optimised fitting because small pulse kinks lead to immediate loss of the atom from the trap. At larger segment numbers, one can simulate atom transport with extremely high fidelities and much more modest memory costs. The improvement in fidelity from floating point to quantisation-aware fit can be as large as two orders of magnitude. In fact, the fidelity can even exceed that attained with the AWG implementation, as the piecewise quadratic pulse used here is not the optimal waveform for atom transport [173]. The right plot shows similar results to those obtained in Fig. 45 and almost three orders of magnitude of compression with 6 segments when comparing the optimised fit to the AWG implementation.

6.3.2 *Atom Transport*

High-fidelity transport of neutral atoms is vital for the initialisation of stochastically loaded atomic arrays for analogue simulation and computing, as well as for the dynamic rewiring of qubit connectivity. Atoms are typically loaded into an array of static tweezers created by spatial light modulators. However, the loading probability of getting exactly one atom in an individual trap is $< 100\%$ [206], and often closer to 50%, and to form a desired topology atoms have to be moved to the desired trapping sites. This is achieved by overlaying static traps with another set of dynamic traps controlled by an Acousto Optic Deflector (AOD). Precise control of the frequency of the drives induces spatial movement of trapped atoms [17, 73, 205]. Analogously, the qubit connectivity can be modified within the algorithm runtime [32, 33]. The movement of the atoms must be carried out at high fidelities and speeds to prevent loss and to avoid introducing unnecessary latency overhead. The speed at which atoms can be displaced is bounded by the quantum speed limit, which for a distance of 3 μm is found to be around 10 μs [173]. Given the desire to move thousands of atoms with multiple tones over many trapping sites, which requires precise control over the frequency and amplitude of the transport waveforms, pulse compression is investigated to reduce the hardware overhead required for the memory-efficient representation of such signals.

The description of the transport simulation follows closely from the description in [173] and the same code is used for a numerically realistic simulation of atom transport, where the motion of a single $^{88}\mathbf{Sr}$ atom over a distance of 3 μm is considered which separates adjacent trapping sites (this is equivalent to a 1 MHz frequency modulation

9 Asymmetric time-dependent envelopes [233] can result in higher population transfer fidelities, but the symmetric Blackman ramp is used here for simplicity.

10 The cubic spline block is assumed to be implemented at an output sampling rate of 1 GHz

of the trapping waveform). Every atom is initially in a thermal state at a temperature T[173]:

$$\rho(T) = \sum_{i=1} p_i \left|\psi_i\right\rangle \left\langle\psi_i\right| \tag{82}$$

where $p_i(T) = \mathrm{e}^{-E_i/k_bT}$ are the Boltzmann weights at temperature T (for more details cf. [173] where the truncation of this sum is discussed). The optical tweezer is modelled as a Gaussian potential [173]:

$$V(t) = U_d \mathrm{e}^{\frac{-2r(t)^2}{w_0^2}}, \tag{83}$$

where U_d is the trap depth, w_0 the beam waist and r the time-dependent centre of the trap [73]. In the optical setup there is correspondence between the frequency of the driving tone $f(t)$ and the position of the tweezer $r(t)$ and the tuning of $f(t)$ creates a time-dependent motion of the trapping centre $r(t)$. Without further optimising the shape of the frequency chirp [11], an extremely simple piecewise quadratic pulse with time-dependent frequency is chosen:

$$f(t) = \begin{cases} f_0 + 2\dfrac{f_\mathrm{f} - f_0}{t_\mathrm{p}^2} t^2 & t \leq t_\mathrm{p}/2 \\ f_\mathrm{f} - 2\dfrac{f_\mathrm{f} - f_0}{t_\mathrm{p}^2}(t - t_\mathrm{p})^2 & t > t_\mathrm{p}/2 \end{cases} \tag{84}$$

where the pulse duration is fixed at $t_\mathrm{p} = 31\ \mu$s, and the frequency sweep ranges from an initial frequency f_0 to a final frequency f_f, such that $f_f - f_0 = 1$ MHz, corresponding to the spacing between adjacent trapping sites. The fidelity of atomic transport is benchmarked using different signal representations generated by an AWG, including both floating-point and quantization-aware cubic spline fits.

Following the approach described in [173], the transport fidelity is evaluated by computing the infidelity between the final density matrix $\rho_f(r)$, obtained after applying

11 For simplicity a simple pulse representation is chosen, but there are higher-order polynomials or numerical solutions determined by optimal control methods cf. ([173] for more details which lead to higher transport fidelities

the transport waveform, and the initial density matrix ρ_0, spatially shifted to the final position r_f:

$$J(r) = 1 - \mathrm{Tr}\left(\sqrt{\sqrt{\rho_\mathrm{f}(r)}\rho_0(r_\mathrm{f})\sqrt{\rho_\mathrm{f}(r)}}\right)^2 \tag{85}$$

This is averaged over 10μs due to oscillatory behaviour after the atom transport [173] which confounds the investigation of the effect of the transport waveform on the post-transport fidelity. To ensure that an atom remains trapped during transport, it is important to vary the trapping potential adiabatically and slowly. In this regime, the atom stays in thermal equilibrium with the trap and its motional state remains approximately unchanged apart from following the moving potential. The atom remains in its Boltzmann-distributed ensemble of motional states, as described in Eq. 82, determined by the local temperature and the trap parameters. However, if the trap is moved too quickly or non-adiabatically (induced by fast changes in trap position, i.e. modulation frequency cf. 84), the atom may absorb energy from the motion of the trap. This heating effect can excite the atom to higher motional states. If the energy gained exceeds the trap depth, the atom can escape the trap entirely. It shall be stressed that the effect of a memory-efficient representation of the transport waveform which has a time-dependent frequency is isolated, and does not consider any other error sources which induce noise on the trap depth, position, and waist arising from laser noise, fibre coupling, finite AOD rise time and other imperfections which are prominent in physical experiments. As in the previous section, it shall be highlighted that in the small-error regime where quantum computers typically operate, the errors are linearly additive to first order, which justifies this approximation.

The effects of the imperfect frequency modulation are also significant and visibly impact transport infidelities by orders of magnitude, as shown in Fig. 46. Considering a representation of the piecewise quadratic pulse with cubic splines using the floating point fit and quantisation-aware point fit described in Sec. 6.2.3, a comparison to AWG output at a sampling rate of 500 MHz (this is more realistic for longer shuttling wave-

forms of up to several hundreds of microseconds where higher sampling rates are not necessary) is made for a 31 μs piecewise quadratic pulse described in Eq. 84. Compared to the compressed pulse approach, the AWG approach directly stores samples in memory at the sample frequency to represent the desired output waveform. The idealised numerical simulation results for atom transport in Fig. 46 show that in the low segment limit, a cubic spline fit cannot achieve high-fidelity atom transport even with the optimised fit. However, for the case of 6 segments, the optimised fit shows over two orders of magnitude improvement in fidelity compared to the floating-point fit, and attains an even lower fidelity than the AWG representation at a memory compression factor of ≈ 620. For increasing segment numbers the results are more erratic due to the small variations in the representation of the piecewise quadratic function, but the results consistently show that the optimised fit significantly reduces atom transport infidelity, often by more than one order of magnitude. Similar results were attained with different trap depths. In addition, moving tweezers over larger distances in large atom arrays requires longer waveforms of up to hundreds of microseconds with higher memory overheads of up to order $\approx 10^6$ bits when storing the frequency envelopes directly in the memory of an AWG. This further motivates the use of the optimised compression techniques outlined in this paper.

It is worth highlighting the scalability of this approach with a simple example. Consider a control system with limited memory, capable of storing only 8 uncompressed, fully sampled waveforms (i.e. for an AWG) required for arbitrary horizontal and vertical moves in a 3x3 grid of atoms ($2 \cdot (n_{grid} - 1)!$ unique moves in each dimension). Using a constant compression factor per waveform of ≈ 625 (as shown in Fig. 46) and using the optimised cubic spline-based compression technique, the 2880 waveforms required for arbitrary moves in a 7x7 atom array can be stored within the same memory space used to store 8 waveforms directly at the output sampling rate. This effectively increases the number of qubits that this memory sparse control system can support from 9 to 49, an improvement of over 5 times in qubit capacity.

Note that in this case, the piecewise quadratic pulse represented by an AWG does not constitute an upper bound on transport fidelity. It is shown in [173] that more optimal waveforms can be determined with quantum control methods. Indeed, for a large enough number of segments the results show that waveforms represented with cubic splines in some cases achieve higher transport fidelities as shown in Fig. 46, however, this is purely by chance and not by design. The determination of an optimal transport waveform (i.e. not piecewise quadratic) which can be efficiently represented with cubic splines is left as a subject for future work.

6.4 OUTLOOK

It has been demonstrated that radio-frequency signals, particularly relevant for atomic quantum computing systems, can be efficiently compressed using cubic splines. A hardware-efficient implementation on an FPGA was presented, achieving high output sampling rates of up to 1 GHz through the use of recursive summation techniques for sample generation. Two memory compression strategies were employed: one exploiting the symmetry of pulse shapes, and another based on a fixed-point, hardware-aware fitting routine, both designed to enhance pulse representation without compromising the memory compression factor.

The effectiveness of the optimised fitting routine was evaluated through STIRAP simulations and atom transport using optical tweezers. In both cases, significant improvements in fidelity were achieved compared to the floating-point spline fit and cubic spline compression. For longer pulse durations or reduced segment counts, fidelity was maintained at levels comparable to those obtained with state-of-the-art arbitrary waveform generators (AWGs), while achieving compression factors exceeding two orders of magnitude.

Although the focus of this study is on neutral atom quantum devices, the underlying pulse compression techniques are broadly applicable to other platforms. For example,

similar approaches could be employed for ion transport in trapped-ion systems, where time-dependent electric potentials are applied to multiple trap electrodes to shuttle ions between regions. Such techniques may be particularly beneficial for implementing fast and high-fidelity transport protocols that rely on complex waveform engineering, including invariant-based inverse engineering or shortcuts to adiabaticity [183].

In contrast to atomic qubits, the pulse duration for single- and two-qubit gates in superconducting quantum computing architectures is typically on the order of tens to hundreds of nanoseconds. The bandwidth of electronic devices is generally limited to a few GHz, meaning that most signals only need hundreds of samples to accurately represent pulses within this range. Pulse shapes can vary dramatically: commonly used DRAG gates require smooth pulses [83], whereas some recent solutions leading to higher fidelity use non-smooth pulses obtained by reinforcement learning methods [26]. The time scales and the non-standard envelope shapes limit the potential of pulse compression techniques, and therefore the focus here was on applications in atomic quantum systems. However, there may be applications that require the generation of electronic signals of higher bandwidths. Furthermore, there might also be applications in other material platforms such as, but not limited to, quantum dots or colour centres, but a comprehensive study of the control requirements has not been performed here.

In general, the hope is that line of work motivates further research into resource analysis and optimisation for future-proof quantum control systems. Memory is likely a limited resource for compact integrated control systems, and this work provides some techniques which can be used to minimise memory usage without compromising pulse representation and fidelity of quantum operations.

7

REINFORCEMENT LEARNING FOR CONTROLLING QUANTUM DYNAMICS

This chapter is based on the work published in [75], carried out in collaboration with Aniket Chatterjee and Tim Franzmeyer[1].

7.1 INTRODUCTION

The optimal control of quantum systems is important for enabling the development of quantum technologies such as computing, sensing, and communication, and similarly plays an important role for quantum chemistry [45] and solid state physics. Quantum control requires the application of time-dependent signals (laser pulses, microwaves etc.) to a quantum system, to realise a desired time evolution [88, 127, 128]. Examples of such tasks include system initialisation, state preparation, gate operation, state population transfer or state measurement. Quantum control enables performing such tasks with low error rates, which is particularly important for the realisation of fault tolerant quantum computing [222].

Isolated quantum systems exhibit unitary dynamics which can be modelled with the Schrödinger equation, exhibiting linear scaling with respect to system size (cf. Eq. 3). Yet, all real quantum systems are open, subject to some interaction with the environment and require the addition of irreversible non-unitary dynamics to accurately model their evolution. The combined unitary and non-unitary evolution is typically

1 I was responsible for developing the simulations for both the multi-level atomic system and the Rydberg atom. Aniket Chatterjee focused primarily on the transmon reset problem. Tim Franzmeyer implemented the code for the multi-step problem with biased noise. The machine learning algorithms were jointly developed by the authors and I implemented the constrained formulation. All numerical experiments were conducted by myself, and the resulting data were jointly analysed and discussed.

modelled by a master equation (cf. Eq. 5), growing quadratically with the quantum system size. While stochastic methods [59] can reduce the quadratic scaling to a linear one, they require a large number of trajectories for high numerical accuracies. For large systems the only feasible option is to sample directly from a quantum device. A detailed discussion of different simulation methods is provided in Sec. 2.0.1.

Experimental imperfections and noise, arising from, e.g., signal distortion or attenuation in optical and electronic setups, or due to inherent system imperfections, pose additional challenges which are often lacking in existing approaches to quantum control. This has often been addressed with dynamical decoupling techniques, originating from techniques introduced for NMR experiments[231]. However, this comes at the expense of longer control operation durations and increased experimental complexity for implementing such pulse sequences.

In this chapter, a new approach to controlling real-world open quantum systems is presented by framing quantum control as a Reinforcement Learning problem with physical constraints. The goal is to learn control policies which produce realistic time dependent amplitudes and frequencies [2] that maximise fidelity while excluding signals that lead to overly fast quantum state dynamics. Most quantum control tasks, including those in this thesis, involve adiabatic processes — requiring *slow* evolution compared to the inverse energy gap between states ($E = \hbar\omega$). *Fast* dynamics can induce leakage errors and reduce robustness to experimental timing errors. To address this, a hard constraint on the space of possible solution dynamics is imposed to eliminate such control signals and a soft constraint is also introduced to favour smooth pulses with zero initial amplitude and finite rise time - features essential for experimental feasibility. These constraints not only enhance solution quality but also reduce computational cost by avoiding expensive simulations of fast-evolving dynamics, enabling efficient parallel hyperparameter tuning on a Graphical Processing Unit (GPU). Moreover, it is shown through various ablations with different RL algorithms, how these constraints

2 Depending on the system these signals could be either optical (typical for atoms) or microwaves for superconducting qubits (i.e. artificial atoms).

can reduce compute time and also lead to higher fidelity solutions compared to a simple fidelity maximising implementation of the same algorithm. Multi-step RL is also explored to improve performance in the case of higher noise levels, here the noise of a control signal is observed and in real-time incorporated into the determination of the optimal quantum control signal.

The approach is validated on three systems. First, a generalised electronic four level Λ system is considered, common in quantum dots, atoms, and circuit quantum electrodynamics, revisiting the STIRAP problem [238] for transferring population between states in a more complex level structure. The RL implementation successfully learns realistic control signals with state of the art fidelities, and exhibits good resilience to time-dependent noise. Then, the two-photon Rydberg gate [149] is explored, crucial for realising neutral atom quantum computers. This approach exhibits good noise robustness and achieves higher fidelities at lower pulse energy than previous works. Lastly, a superconducting transmon quantum bit (qubit) [70] is considered and qubit reset is tackled. We discover a novel physically-feasible reset waveform is discovered, which achieves an order of magnitude higher reset fidelity than any previous work. Finally, it is shown concretely how the RL approach can be adapted to real experimental implementations and closed loop experimental control.

This chapter is structured as follows: first, relevant prior work is reviewed. Next, the methods and machine learning techniques are described in detail. Finally, results from simulations on three quantum systems are presented, along with a discussion on how the approach could be adapted to real physical experiments.

7.2 RELATED WORK

Several approaches exist for optimising time-dependent control signals for quantum systems. Analytic methods like Lyapunov [106] are effective for small isolated systems but difficult to generalise to complex environments. Gradient-based methods which

typically consider the iterative evolution of a quantum system, under piece-wise controls, such as **Gr**adient **A**scent **P**ulse **E**ngineering (GRAPE) [122] or Krotov [187] (which ensures monotonic convergence) work well on simple cost landscapes, for a single objective and with good initial guesses. Since gradients need to be evaluated exactly at each time-step with these methods, they are well suited to ideal simulations with no noise, but it is challenging to adapt them to closed-loop experimental optimisation where such gradients can rarely be determined exactly and a large number of measurements are required. RL as implemented here, on the contrary, requires no intermediate quantum state information and can perform optimisation with noisy samples. Other variations of global optimal control exist which consider the optimisation of the entire time dependent control signal at once [86], which are prone to local minima. Global and local optimal control methods are explored in combination in [89]. Methods which reduce the search space by decomposing the possible signal into a basis are sensitive to basis choice [174] and evolutionary algorithms [47] lack computational scalability for larger systems or multiple objectives.

Machine learning has numerous applications in quantum science [130]. Several works successfully apply RL to the logical quantum circuit level. Examples include the determination of optimal error correction codes [169], real time error correction [211] or circuit complication techniques [230]. In contrast to these works, this chapter focusses on lower-level control of pulses.

Prior work which applies machine learning to optimising time dependent controls for quantum dynamics is reviewed, distinguishing between real device sampling and numerical simulations. In seminal work [26], the authors devise an optimal gate set on a superconducting IBM quantum device. This is the first successful demonstration of experimental sampling for a real time implementation of reinforcement learning. [189] and [179] use measurements and feedback to prepare quantum states. However, generalisation to unseen environments is difficult. A model-based Hamiltonian learning approach was applied in [121], which provided important insights into sample

efficiency and how simulation information can effectively be combined with device sampling. Although this chapter focus on realistic quantum simulations, the RL algorithm can be readily applied to physical experiments by replacing simulations with real-device sampling and some insights on this are provided in Sec. 7.5.

Several studies have explored reinforcement learning (RL) to simulated quantum systems. While RL has been applied to discrete action space control [9, 175, 251], these methods struggle in real-world settings with analogue signals exhibiting non instantaneous response times[3] and in more complex systems. This limitation when applying RL with discrete action spaces to real physical systems is noted in [251]. Some prior work focussed on controlling many-body systems [50, 154]. While these studies are extremely interesting, they lack concrete applicability to real experimental systems. This chapter extends some of this prior work by incorporating control signal noise into training as suggested in [198]. The Google Quantum Team used RL in [166] to find time-optimal gate sequences for simulated superconducting qubits using trust region policy-gradient methods, but this implementation is very computationally demanding. The methods described in this chapter enhance computational scaleability by learning entire control signals in a single step and adding physical constraints. There is some prior work on controlling typical Λ systems (STIRAP setting) with RL [48, 167]. This is extended by incorporating realistic noise models and simultaneous amplitude and frequency control to learn more optimal and realistic policies. Related work on optimising superconducting qubit gates robust to noise with RL is presented in [160], while enhanced superconducting qubit readout with RL is demonstrated in [54].

One can assess the relative strengths and weaknesses of an RL approach, as compared to other common approaches like GRAPE or Krotov, which directly differentiate with respect to the piecewise time evolution of a quantum system in Tab. 12.

3 A particular limitation is the finite rise and fall time of electronic or optical signals, which refers to the time required to transition from zero to maximum amplitude (or vice versa), typically on the order of nanoseconds or greater. Similarly the instantaneous bandwidth determines the frequency with which sinusoidal signals can change from one sample to the next without distorting the signal.

	Model Free RL (e.g. PPO)	Direct Differentiation (e.g. GRAPE)
Method	Rewards are estimated for a stochastic policy and maximised	Gradients are exactly evaluated and optimised
Flexibility	Easily handles stochasticity & multi-objective optimisation	Requires complete knowledge of time evolution of system
Efficiency	Computationally more expensive due to exploration and sampling	Faster for well-posed, deterministic problems
Robustness	Adapts to noise, parameter changes or constraints dynamically	Solution dependent on initial seed and doesn't easily adapt to complex noise

Table 12: Comparison between Model Free RL and Direct Differentiation for optimising parameters for quantum control. Note that complete knowledge of the system's time evolution is typically not easily achievable in real experiments.

7.3 BACKGROUND AND METHODS

7.3.1 *Reinforcement Learning*

Reinforcement Learning (RL) is a framework where a computer algorithm learns to make decisions by interacting with an environment to achieve a specific goal [220]. This algorithm is also often called a virtual *agent*. In quantum control, RL can be used to find the control actions $a(t)$ that steer a quantum system toward a target state ρ_{des}. The key components in this RL setup are the state s_i, which is given as the density matrix $\rho(t)$ of the quantum system, the control action a_i applied to the system, a scalar reward r_i, derived largely from the closeness of the system to its target state and the policy π that maps states to actions. The objective is to learn a policy π^* that maximises the expected cumulative expected reward, i.e. $\pi^* \in \max_\pi \mathbb{E}\left[\sum_{t=0}^{T} r_t\right]$.

BANDIT SETTING The bandit setting provides a simplified framework within reinforcement learning (RL), where the complexity of temporal dynamics and state transitions is removed. Instead of learning a policy over several quantum states, the agent focuses solely on selecting actions that yield high rewards. Each action $a \in \mathcal{A}$ corresponds to a time-dependent quantum control signal, such as a detuning $\Delta_i(t)$ or an amplitude $\Omega_i(t)$, which influences the evolution of the quantum system. The outcome of applying an action is assessed by a scalar reward, often based on the fidelity with respect to a desired target state.

In this setting, each action a is associated with a stationary, yet unknown, reward distribution with an expected value $R(a)$. The learning task reduces to identifying the optimal action a^* that maximizes the expected reward:

$$a^* = \arg\max_{a \in \mathcal{A}} \mathbb{E}[R(a)]. \tag{86}$$

This formulation serves as a foundation for more complex RL scenarios, such as those involving multi-step sequential decision processes which will also be considered in Sec. 7.4.4.

CONTINUOUS BANDIT SETTING In the *continuous bandit* setting, the problem is a single-step episode where the action space $\mathcal{A} \subseteq [-1, 1]^d$ is continuous. Here, d represents the number of discrete samples used to parametrise the control signal in time [4]. A smooth, time-continuous control pulse is then generated by interpolating between these samples[5]. The agent samples control signals from the policy π $a(t)$ and immediately observes a reward $r(a)$, typically proportional to the final state fidelity. Formally, the goal remains to identify the action $a^* \in \mathcal{A}$ that maximizes the expected reward as in Eq. 86.

Formulating a high dimensional control problem in terms of the continuous bandit setting, allows for rapid exploration of the action landscape, making it particularly suitable for settings where extensively simulating full quantum trajectories is computationally expensive. Moreover, several controls can be simulated in parallel since only final state information is required for the algorithm.

7.3.1.1 *Simulating Quantum Dynamics*

Assessing the quality of an action $a(t)$ in quantum control involves computing the fidelity (cf. Eq. 8 in Ch. 2), which measures the closeness between the evolved quantum

4 For the simulated problems large $d \geq 50$ are picked for maximal signal expressivity.
5 For instance, cubic spline interpolation can ensure physically realistic and bandwidth-limited control signals.

state and the target state. The state evolution under $a(t)$ is obtained by numerically solving the master equation (cf. Eq. 5) for $\rho(t)$, the generalised quantum state at time t. A more detailed introduction to quantum dynamics is provided in Sec. 2.0.1. To determine the time evolution of a dissipative quantum system, one requires a numerical integration method. Typically, this is done using adaptive step-size solvers that implement higher-order Runge-Kutta methods [96], which dynamically adjust internal time steps based on local error estimates. If the numerical error exceeds the desired numerical tolerance, the solver reduces its internal time step. If the error is sufficiently small, the time step is increased to enhance computational scaleability by minimising the total number of solver steps as much as possible. Therefore, control signals that lead to *slower* quantum state dynamics allow the adaptive solver to use larger time steps. Hence, *slower* quantum state dynamics require fewer solver steps and generally less computation time.

7.3.2 *Physics-Constrained Reinforcement Learning*

In practice, applying RL to efficiently find high-fidelity quantum control solutions requires restricting the space of possible solutions. The first constraints introduced are concerned with physically relevant quantities including: signal bandwidth, and signal area, reflecting the limited instantaneous bandwidth of electronics and signal components in experiments, and limitations in available signal power and duration. The constraints are implemented through the reward function cf. Eq. 88 and Sec. 7.3.3. This reminisces of the Lagrange multiplier technique introduced in [30].

Computing the reward requires a simulation of the quantum system, (cf. Sec. 7.3.1.1), which is computationally expensive. For complex quantum systems and particular actions, this simulation can require a large number of solver steps, far exceeding the computational time needed to update the RL agent. Therefore, secondly, the policies are constrained to generate actions that lead to quantum dynamics that can be simulated

within a predefined number of maximum numerical solver steps, N_{max}. This constraint incorporates priors about the physical solution time scales into the algorithm, which incentivises adiabatic quantum state dynamics. Additionally, this constraint incentivises smoothness (requiring lower bandwidth) of signals, as smooth signals generally require fewer solver steps.

RL optimisation algorithms are often sensitive to the choice of hyper-parameters [100], necessitating hyper-parameter space searching to find optimal policies. This challenge is addressed by synchronously optimising control policies for up to 1024 RL agents in parallel on a single GPU device, by implementing both the quantum solver and the RL algorithm using JAX [39], which features just-in-time compilation and automatic differentiation. This allows for the end-to-end compilation of the parallelised training and simulation loop. However, in this synchronised parallel setup, the quantum simulation time needed per RL update step is governed by the maximum quantum simulation time across all hyper-parameter configurations, as the slowest simulation among all learned policies determines the speed of the entire loop which needs to be synchronised. This bottleneck is mitigated by the constrained RL algorithm that solves the quantum control problem subject to the condition that the required number of quantum simulation steps does not exceed a chosen threshold N_{max}. Formally, this constrained RL problem, for which the quantum simulation can be executed in fewer than N_{max} steps, is defined as:

$$\pi^* \in \max_{\pi} \mathbb{E}\left[\sum_{t=0}^{T} r_t\right] \Big| \text{ for } a \in \pi \;\; N_{\text{Sim}}(a) < N_{\text{Sim}}^{\text{max}} \tag{87}$$

where π is the policy, r_t is the reward at time t, and $N_{\text{Sim}}(a)$ is the number of solver steps required for conducting the quantum simulation for an action a sampled from policy π. Implementing this constrained RL algorithm prevents bottlenecks as it ensures that all simulations within the parallelised run are completed within a fixed maximal time frame. This approach enables efficient hyper-parameter exploration with minimal computational cost.

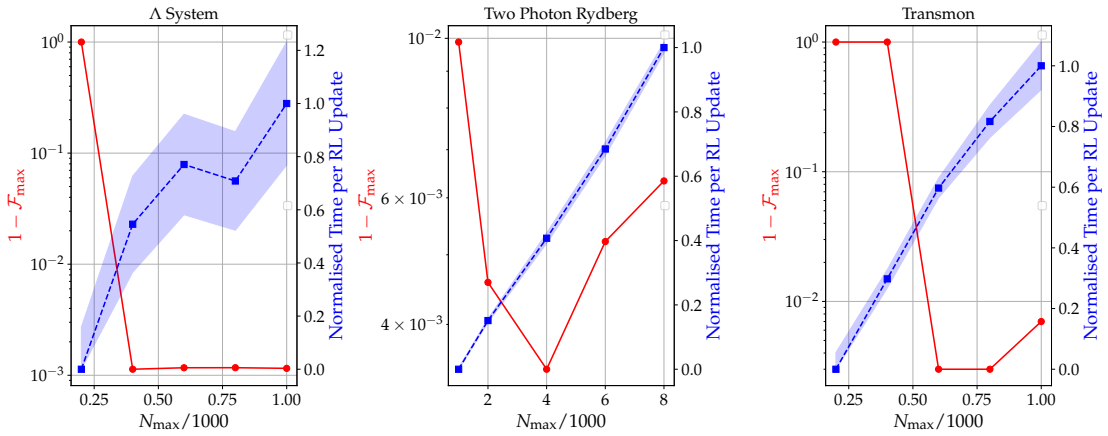

Figure 47: Minimum infidelity across hyper-parameter combinations ($1-\mathcal{F}_{\max}$, left y-axis, solid red line) and normalised GPU time per RL update (right y-axis, dotted blue line $\pm 1\sigma$) vs. permissible quantum solver steps $N_{\max}$. Limiting $N_{\max}$ - which can be understood as placing an upper bound on the rate of change of the quantum system evolution induced by the control signal - improves solution quality (lower infidelity) while also increasing computational efficiency (lower normalised time).

The constraint, though seemingly restrictive, is physically motivated for adiabatic population transfer between quantum states [129]. In adiabatic processes, the system evolves slowly relative to the inverse energy gap. The maximal effective Rabi frequency (for near resonance this is simply $\overline{\Omega} = \Omega_{\text{eff}}$ and if $|\overline{\Delta}| \gg |\overline{\Omega}|$, $\Omega_{\text{eff}} = \frac{\overline{\Omega}^2}{\overline{\Delta}}$ since the effective system dynamics are governed by this time scale instead) sets a lower bound for $N_{\max}$ via the adiabaticity condition $\Omega_{\text{eff}} \cdot \delta_t \gg 1$ [129] (where δ_t is the numerical step size). This follows as $N_{\max} \geq T/\delta_t \approx \Omega_{\text{eff}} T$ in the lower bound. In practice, $N_{\max}$ is incremented for as long a significant decrease in infidelity is observed (see Fig. 47 for infidelities at different maximum solver steps for different systems).

The constrained RL approach improves computational efficiency and promotes the selection of physically realistic control signals (cf. Fig. 49). By enforcing these constraints, one can achieve more interpretable quantum state dynamics (cf. Fig. 52) and the prioritisation of adiabaticity leads to experimentally feasible, robust and high fidelity solutions.

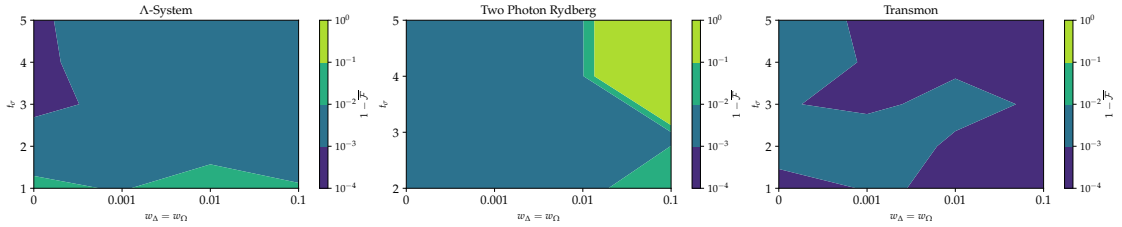

Figure 48: Ablation over smoothness penalty coefficients $w_\Delta = w_\Omega$ and filter standard deviation t_σ for three different environments. Choice of smoothing parameters is important for learning policies with low mean infidelity $1 - \overline{\mathcal{F}}$ (averaged over 64 parallel environments). For some systems, like the Λ system higher filter s.d. leads to lower infidelity, whilst for other like the Transmon higher smoothing penalties lead to lower infidelities. The relationship to experimental feasibility with limited bandwidth electronics is analysed in Fig. 49.

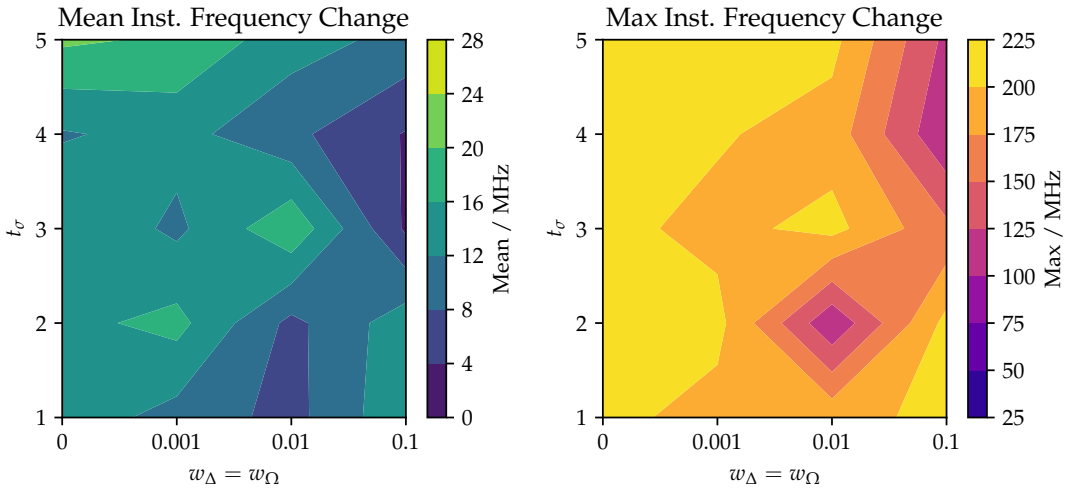

Figure 49: Using the optimised control signals for the Λ-system obtained from the smoothing hyper-parameter ablation study (Fig. 48), the instantaneous frequency change between successive output samples is analysed. The signal $S(t)$ is digitally sampled at 1GSa/s, corresponding to state-of-the-art arbitrary waveform generator (AWG) capabilities. Plotted is mean (left) and maximum (right) instantaneous frequency changes across consecutive samples. The results show that lower smoothing penalties and smaller kernel standard deviations lead to significantly larger maximum and mean instantaneous frequency variations. Instantaneous frequency changes of several hundred MHz pose serious challenges for experimental realisation. They are likely to induce severe signal distortions and fall outside the operational bandwidths of typical digital-to-analog converters (DACs) and amplifiers. In contrast, signals with maximal instantaneous frequency changes of up to $\approx$ 100MHz are far more compatible with standard experimental hardware. Thus, in addition to improving computational efficiency, the application of smoothing constraints is critical for ensuring that the optimised control signals are easily experimentally realisable. Similar results are found for the other physical systems but they are not further presented in favour of brevity.

7.3.3 *Reward Shaping*

The control signals are parametrised as a combination of time-dependent amplitudes Ω_i and time-dependent frequencies Δ_i. Smoothness constraints that facilitate efficient learning and further improve computational efficiency are implemented. Smoother waveforms are easier to implement experimentally, offer clearer interpretation of the optimal quantum state evolution, and significantly speed up simulation times by reducing the number of required solver steps. To facilitate smooth signal discovery, a Gaussian convolution filter is applied to the control signal with a standard deviation t_σ (cf. App. 127) before simulating the quantum state dynamics which improves learning dynamics by favouring slower solution dynamics (an ablation over this is found in Fig. 48). The reward function contains auxiliary smoothing penalties and is defined as:

$$\begin{aligned} L_a = & -w_{\mathcal{F}} \log_{10}(1-\mathcal{F}) - w_\Omega \mathrm{ReLu}\left(\frac{\sum S(\Omega_i)}{\sum S_{\text{base}}} - 1\right) \\ & - w_\Delta \mathrm{ReLu}\left(\frac{\sum S(\Delta_i)}{\sum S_{\text{base}}} - 1\right) - w_A \frac{\sum A(\Omega_i)}{A_{\text{base}}} \end{aligned} \tag{88}$$

The first and most important reward-function term incentives high fidelity $\mathcal{F} = \mathcal{F}(\rho_{\text{fin}}, \rho_{\text{des}})$ with respect to the desired final state ρ_{des}. This fidelity reward scales with the logarithm infidelity $\log(1-\mathcal{F})$). For example, a linear $\mathcal{F}$ reward only yields and additional reward of 0.009 when going from 0.99 to 0.999 fidelity, whereas for the fractional log, this increase in reward is 1. This is important for learning solutions with fidelity > 0.99 which is desirable for implementing error free operations with error correction. Next, smoothness penalties are defined, where $\mathrm{ReLu}(x)$ defines the ReLu function: $\mathrm{ReLu}(x) = 0$ if $x < 0$ || $\mathrm{ReLu}(x) = x$ if $x >= 0$. The smoothness S is compared to that of a reference signal S_{base} (cf. App. A.14.3 Fig. 65 for a comparison of various smoothing functions). The correct smoothness penalty coefficients w_Δ, w_Ω, balance fidelity, interpretability, and computational efficiency. A smoothing function based on the discovered control signals second derivative is found to be particularly

favourable for finding smooth high fidelity controls (cf. App. A.14.3 for more details). Fig. 48 shows an ablation over various smoothing penalties with this function. Contrary to the Λ system and Rydberg atom, the Transmon favours stronger smoothing penalties showing that this approach is adaptable to a variety of physical problem settings. Larger values of t_σ and w_Δ, w_Ω also reduce the maximum required solver steps and thereby further enhances computational scaleability. The ability to achieve high-fidelity solutions across all environments at larger convolution standard deviations, t_σ, also demonstrates that one can find optimal signals compatible with realistic electronic control systems with limited instantaneous bandwidth (cf. Fig. 49).

The final reward term penalises solutions with large pulse area (cf. Sec. 7.4.1 Fig. 54 for an ablation over different area penalties for the Λ system), for the Rydberg and Transmon problem settings area penalties are not explicitly considered, rather the signals are implemented in the fastest possible time. Low pulse areas are attractive because they implicitly reduce peak Rabi frequencies and power demands and can shorten the duration of a target operation within a fixed simulation window, thereby improving experimental clock rates and mitigating decoherence. They also connect to the broader theoretical pursuit of energy-time-optimal quantum controls.

One can introduce additional physics-informed constraints which can aid in improving the learning dynamics. One particular constraint that is explored for the Λ system, as well as the Rydberg atom, is the addition of a penalty term which penalises larger time averaged population of the excited state(s) which radiatively decay. The intuition is that actions which induce severe population of the excited state are more susceptible do spontaneous emission and should thus be avoided. This is implemented by calculating the time averaged excited state population $\langle \psi_x \rangle$ where ψ_x denotes the excited state. More details are found in the respective implementation sections for different environments.

7.3.4 *Implementation Details*

To enforce the constraint $N_{\text{Sim}}(a) < N_{\text{Sim}}^{\max}$ on actions a sampled by policies π in the bandit setting, one picks a penalty reward r_{penalty}. In the bandit setting, r_{penalty} is assigned to any policy where $N_{\text{Sim}}(a) >= N_{\text{Sim}}^{\max}$, and the value is chosen to be lower than any other possible reward in the environment, ensuring that the optimal policy cannot include states violating the constraint [6]. The final reward function is then defined as

$$L = \begin{cases} r_{\text{penalty}} & \text{if } N_{\text{Sim}}(a) >= N_{\text{Sim}}^{\max} \\ L_a & \text{else} \end{cases} \tag{89}$$

where L_a is defined in (88).

In what follows, three different model free RL algorithms are considered. Proximal Policy Optimisation (PPO) [203] is a popular algorithm in modern RL, combining the benefits of policy gradient [6] methods with stability improvements. It has been used successfully for a large number of tasks, most notably and recently, to fine-tune large language models. PPO aims to optimise the stochastic policy by ensuring that updates do not deviate too much from the previous policy. A more detailed explanation of PPO is found in App. A.13.1. Here we recap briefly how the policy and the control actions which are derived from the policy are updated over time. The learned policy π_θ is parameterised by a neural network that outputs a Gaussian action distribution in the continuous action space $\mathcal{A} \subseteq [-1, 1]^d$, which is then converted to control signals with appropriate translations and rescaling (cf. App. A.14.1). A value baseline is also predicted by the neural network, allowing the advantage $\hat{A}_t$ of each sampled action to be computed as the difference between the observed reward (Eq. 89) and the

6 Policy gradient methods are RL algorithms that directly optimise the policy by estimating the gradient of expected reward with respect to the policy parameters and updating them accordingly. Policy gradient methods differ from other types of reinforcement learning algorithms (e.g. Q-learning) in that they optimise the policy directly rather than learning a value function and deriving the policy from it.

baseline prediction, providing a measure of how much better or worse the action performed relative to expectation. During training, the policy parameters θ are updated by maximising the PPO objective $L(\theta)$ (cf. App. A.13.1) to increase the expected reward. Deep Deterministic Policy Gradient (DDPG) [142] is an actor-critic algorithm designed for continuous action spaces, where the policy is deterministic and updated through gradients from a learned critic [7]; however, in many scenarios, it suffers from overestimation bias and sensitivity to hyper-parameters. Twin Delayed Deep Deterministic Policy Gradient (TD3) [82] improves upon DDPG by introducing three key modifications: clipped learning to reduce overestimation, delayed policy updates for stability, and target policy smoothing to improve exploration of the parameter space. Further details on the implementation, including all relevant libraries, is found in App. A.14.1. All the code is fully open source at: `https://github.com/jan-o-e/RL4qcWpc`.

7.4 NUMERICAL EXPERIMENTS

This section focuses primarily on the continuous bandit reinforcement learning setting, with additional experiments showing that multi-step RL outperforms bandit methods under strong signal perturbations. Numerical experiments are conducted on three key quantum control tasks relevant to quantum information processing.

The first task involves high-fidelity population transfer in multi-level Λ systems, relevant to quantum chemistry and solid-state physics, where the effect of physics-driven constraints and the advantage of PPO over other RL baselines are demonstrated.

The second task targets optimisation of Rydberg gates in neutral atom devices, improving gate fidelity and robustness to time-dependent noise for quantum computing applications.

7 A critic is a component in RL algorithms that estimates the value or expected return of actions, guiding the actor's policy updates by providing feedback on how good the generated actions are.

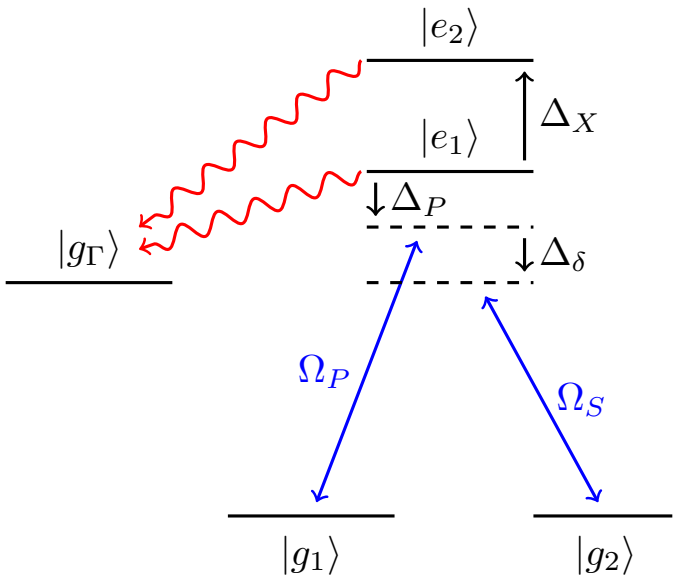

Figure 50: Energy level diagram for four level Λ system with state $|e_2\rangle$, detuned positively by Δ_X from $|e_1\rangle$ which are both coupled by the fields $\Omega_{S/P}$ at different detunings, complicating the dynamics. There is an additional state $|g_\Gamma\rangle$ which does not partake in the unitary dynamics, but to which the excited states decay (cf. red lines). This gives rise to a lower bound in attained population transfer fidelities. The laser couplings are shown in blue.

The third task focuses on developing fast and physically-feasible reset protocols for superconducting transmon qubits under bandwidth constraints, achieving an order of magnitude improvement in reset fidelity compared to prior work.

7.4.1 Λ *System*

Controlling quantum dynamics in multilevel systems is crucial for quantum information processing and relevant to solid-state physics and chemistry [238]. The system is a modified Λ system, featuring two ground and two excited states. For this task, two time-dependent control signals are considered with real Rabi frequencies Ω_S, Ω_P which couple two electronic ground states to two excited states with relative time-dependent frequency detunings to the state $|e_1\rangle$ of Δ_P (for the Pump pulse) and Δ_δ (for the two photon detuning of the Stokes pulse). The Hamiltonian is analogous to that described in Sec. 2.1.3 Eq. 19, except that there is an additional excited state $|e_2\rangle$ (in addition to $|e_1\rangle$ which replaces $|e\rangle$ in Eq. 19) which is positively detuned from $|e_1\rangle$

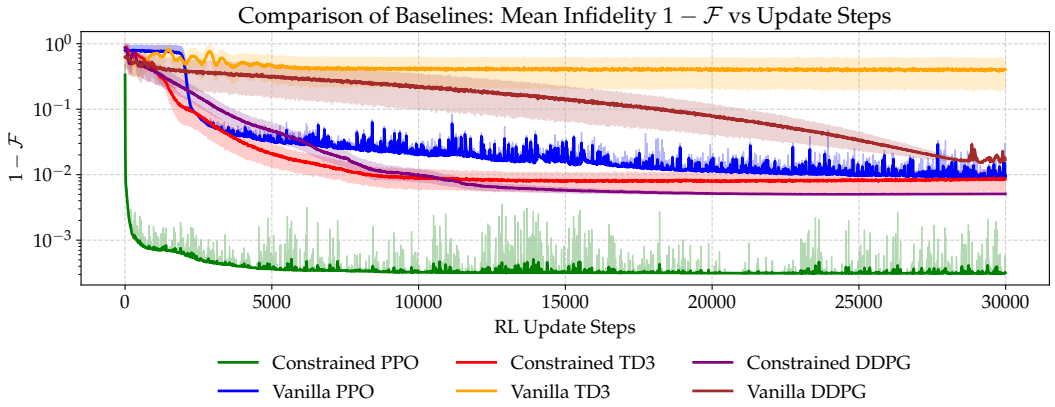

Figure 51: This figure compares three model-free RL algorithms — PPO [203], DDPG [142] and TD3 [82] in the noise-free Λ-system setting. Each algorithm is run on the same Nvidia P100 GPU, using a single set of optimised hyper-parameters per algorithm, with a common batch size and results averaged over two seeds (one s.d. shaded) for the same number of total steps. Compared are standard (vanilla) versions of the algorithms which use $L = \mathcal{F}$, to the constrained formulation, using the reward defined in Eq. 88 and incorporating step-size penalties (cf. Eq. 89). The constrained variants consistently outperform their vanilla counterparts, demonstrating the effectiveness of the constrained problem formulation and physically motivated loss function. Our implementation of PPO achieves the best performance, surpassing DDPG and TD3 by over an order of magnitude in mean fidelity. Furthermore, PPO reaches mean fidelities exceeding 0.99 up to 100× faster than the alternative RL approaches.

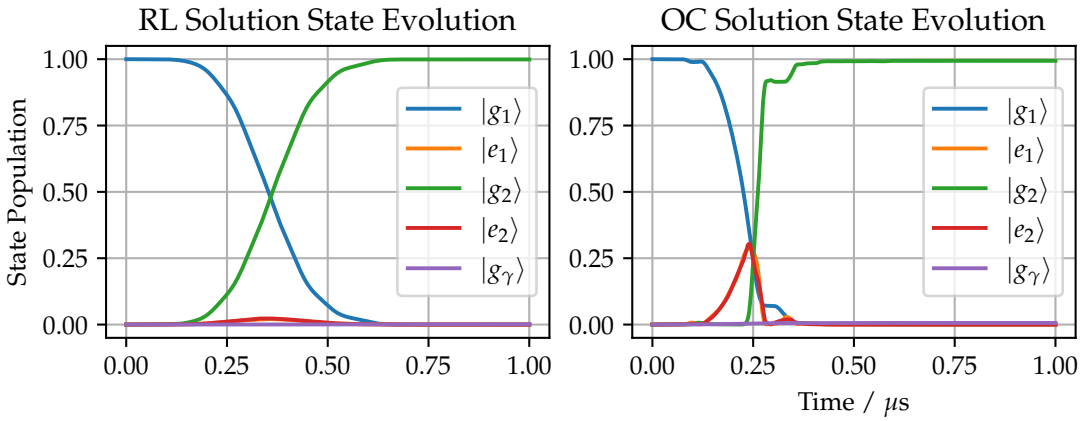

Figure 52: Comparison of quantum state dynamics for the RL solution (top) and Optimal Control solution (bottom). The RL solution exhibits a clearer interpretation of the optimal state dynamics with a smooth population of $|g_2\rangle$ and low excited state population $|e_i\rangle$. The optimal control solutions achieves high final population of the state $|g_2\rangle$, but the time evolution of the quantum states do not offer such a clear interpretation of the optimal time dynamics of the system.

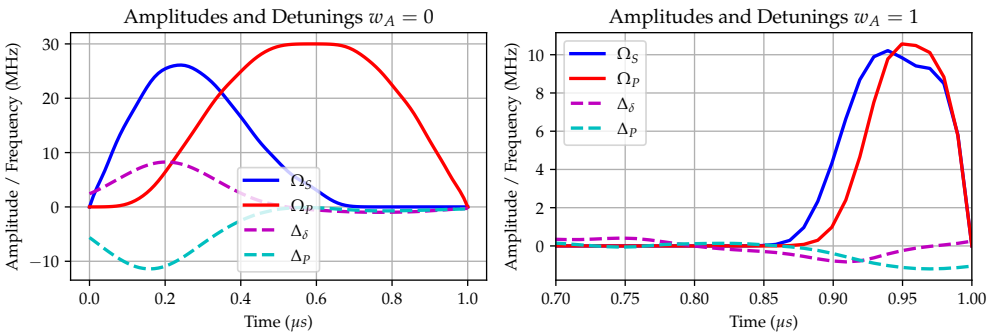

Figure 53: Shown are example control signals generated for different pulse area penalties. For $w_A = 0$ (left), the algorithm seeks to maximise Ω at all times after a fast rise and compensates cross-talk with frequency chirping. For $w_A = 1$ (right), plotted only for the time interval $[0.7, 1]$, as the pulse amplitudes are zero otherwise. Shown is the discovery of pulses which reminisce of two interleaved Gaussians, but exhibit non zero two-photon detuning $\Delta_\delta = 0$ to cancel cross-talk, which differs from the original experimental proposal for coherent population transfer [238].

by Δ_X. Moreover, the ground states $|u\rangle$ and $|g\rangle$ in Eq. 19 are renamed $|g_1\rangle$ and $|g_2\rangle$ respectively here. The energy level diagram for this system is shown in Fig. 50. Additionally, there is a sink state $|g_\Gamma\rangle$ to which spontaneous emission occurs which occurs with rate $\Gamma/2$ from either of the two excited states. This is realistic insofar as spontaneous emission can always occur outside the manifold of interest and leads to a lower bound in population transfer fidelity $\mathcal{F}$, which in this case is defined as the fidelity between the simulated final states ρ and the ground state $|g_2\rangle$ when initialising the system in $|g_1\rangle$. While we focus on open system dynamics in this chapter and generally use the master equation, this system described in this section can also be solved with an effective Schrödinger equation as is spelled out in App. A.15, this would not work if we included spontaneous emission to $|g_1\rangle, |g_2\rangle$ and re-excitation. An additional reward term is introduced an additional which reads $-w_x \cdot (\langle \overline{e_1} \rangle + \langle \overline{e_2} \rangle)$. This assigns lower rewards to dynamics which populate the excited state, since one seeks coherent population transfer. While this is implicit in the fidelity calculation, it helps to make the agents initial exploration more efficient and reduces convergence time. Other tasks like partial state population transfer can also be addressed.

While analytically optimal pulses exist for idealised three-level systems [84, 233], here, one also considers excited state dissipation and an additional excited state de-

tuned positively by Δ_X. This represents a common physical configuration, describing nitrogen vacancy centres [15], quantum dots [69], and single atoms [74]. The algorithm comparisons focus on optimising population transfer from one ground state to another, fixing $\Gamma = 1$, $\Omega_{max} = 30$ and $\Delta_X = 100$ MHz which is physically realistic. A detailed table with comparisons of different optimisation methods which each start with random guesses for time dependent Rabi frequencies and detuning is provided in [75]. The analytic approach introduced in [233], with numerically optimised shape, can yield a fidelity of 0.901 (this is also the method used in Ch. 3) with these parameters, which is a good starting point. Implementations of more conventional quantum control techniques can lead to reasonably high fidelity solutions in some cases, such as the Krotov method [91], yielding 0.99 fidelity with 10000 update steps. However, this method is orders of magnitude slower than RL due to the piecewise calculation of the quantum system dynamics [8]. Optimal Control methods [86] which directly differentiate with respect to the fidelity attained by different time dependent signal amplitudes and frequencies are also more sensitive to initial guesses and in some cases don't achieve particularly high fidelities that go much above 0.9. Moreover, the solutions determined with these methods, without any signal smoothing constraints, are much harder to implement experimentally due to higher instantaneous bandwidth requirements. As explained in Fig. 52, the RL solution leads to a much clearer interpretation of the state dynamics, with a smooth depopulation of the initial state and little excited state population throughout the process. Moreover, this implementation of PPO achieves the best fidelity > 0.999 and also achieves > 0.99 fidelity 100x times faster than alternative RL algorithms TD3 [82] or DDPG [142] as shown in Fig. 51. The fidelities achieved with the RL method are state of the art and also more robust across different random initial seeds, highlighting the superiority of RL over methods which directly differentiate the control action with respect to the fidelity. The RL learned pulses are also physically viable, while prior RL work [48, 86] found solutions, which exhibit non-zero amplitudes

8 The software implementation in [91] which is not compilable or easily GPU parallelisable also makes it significantly slower than the JAX approach in this chapter, a timing comparison is provided in [75]

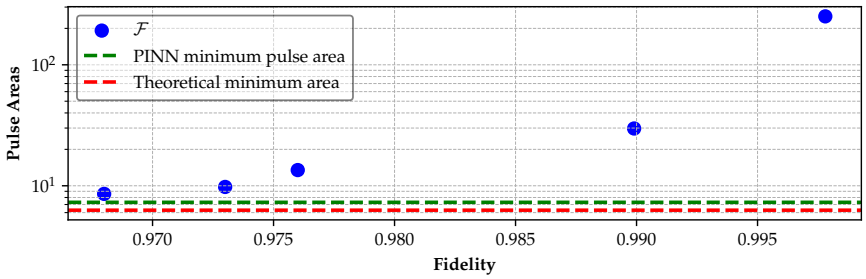

Figure 54: A sweep w_A, of the pulse area penalty weight (cf. Eq. 89) over a range of values $[0, 0.1, 0.25, 1, 2]$ is presented and the minimal pulse area achieved with a PINN is shown in [167] (green dotted line, 7.3) is approached whilst achieving significantly higher fidelities $\mathcal{F} > 0.83$ (cf. [167]). There is a clear trade off and lower pulse areas generally adversely affect fidelities. Moreover, we show the minimal pulse area achieved with a theoretically optimal solution in a three level system (2π) [146].

at the start or end or have instantaneous parameter changes which cannot be realised with bandwidth limited hardware. Moreover, the fidelities attained by using other machine learning methods, such as physically informed neural networked [167] are also lower than those achieved here. Sweeping w_A cf. Eq. 89 one can find signals which have low pulse areas which approach the lower bound quoted in [167] and [146] (cf. Fig. 54), implying minimal signal energy requirements or high energy efficiency. Example signals differ significantly for different pulse area penalties which is shown in Fig. 53. In general, Fig. 54 shows a positive correlation between pulse area and fidelity. We show that compared to the theoretical minimal pulse area 2π in an ideal three level system shown in [146] with a three-level system loss of $0.015T$ (where T is the pulse duration), we achieve a slightly higher loss of ≈ 0.03 for a realistic minimal duration of 0.15μs. However, we show that in the more challenging four-level system we can use RL techniques to approach the optimal three-level solution in terms of pulse area and fidelity and achieve much higher population transfer fidelities with higher pulse areas.

Small random fluctuations or noise of either signal $\Omega_{S/P}$ or $\Delta_{\delta/P}$ are not detrimental to the overall fidelity. To this end, training with an Ornstein–Uhlenbeck noise process is implemented for both $\Delta_{\delta/P}$ and $\Omega_{S/P}$. This noise model creates continuous noise ν_t in time with mean μ and standard deviation σ (for details cf. App. A.14.4). Such noise typically arises from a variety of imperfections in the signal chain, as well as quantum system level noise, such as magnetic field fluctuation or atomic motion. Using

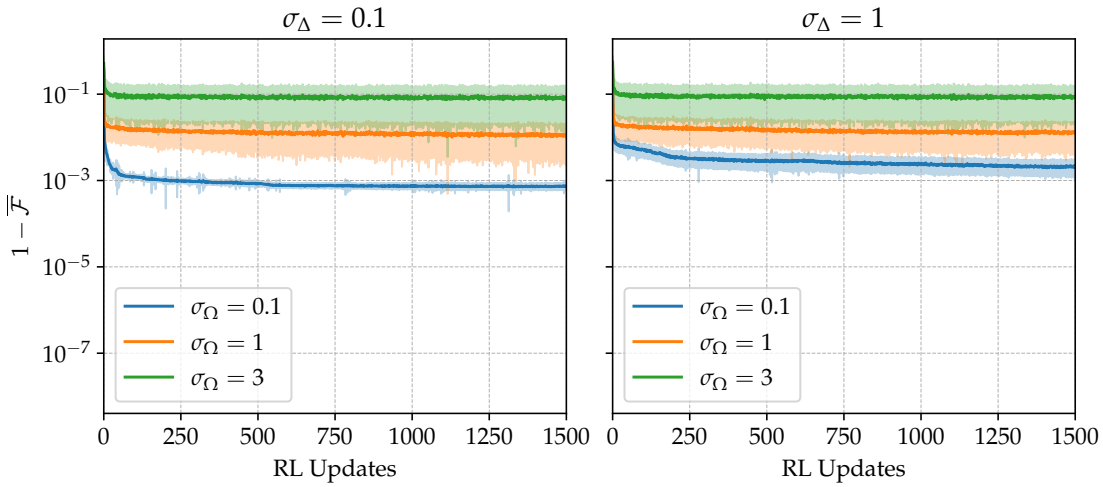

Figure 55: Robustness of PPO for Λ-system to randomly generated noise with $\mu_\Delta = 0.1$ MHz (left) and $\mu_\Delta = 1$ MHz (right) plotted on a logarithmic scale and averaged over multiple seeds. The solid lines show the average infidelity, while the shaded regions indicated one s.d. parallel environments. For small noise levels $F > 0.999$ as shown in the left plot for $\mu_\Delta = 0.1$ MHz, but as it increases fidelities drop to just below 0.97. Robustness to larger levels of noise is addressed with multi-step feedback in Sec. 7.4.4.

unbiased $(\mu = 0)$ noise with various standard deviations exemplifies good robustness to low noise levels as shown in Fig 55 where one attains > 0.99 mean fidelity for $\sigma_\Omega = \sigma_\Delta = 0.1$. Further increasing σ or adding a noise bias, leads to significantly reduced population transfer fidelities which are addressed with multi-step RL in Sec. 7.4.4.

The motivation to study a four-level Λ-type system stemmed from insights in Chapters 3 and 4, which highlighted the difficulty of designing fast, high-fidelity pulse sequences when additional excited states are involved. Compared to the full atomic structure of the D lines of ^{87}**Rb**, this setup offered a more manageable starting point as adding more states slows down simulation and training times. That said, the reinforcement learning approach can be readily extended to other systems by simply replacing the simulation and we are working to extending these ideas to a full fledged atomic model.

7.4.2 *Rydberg Atom*

Neutral atom quantum devices have shown promise for realising scalable, logical quantum computing [33]. The realisation of quantum computing requires a two-qubit gate [161] which relies on the interaction of multiple atomic qubits which are brought in relative proximity (a detailed description of the Hamiltonian is provided in App. A.16) and addressed with a (or multiple) Rydberg laser(s). In this section, the optimisation of the Rydberg gate [149] under realistic experimental conditions and signal perturbations is considered.

The most widespread implementation of a Rydberg C-Z gate is a single photon Rydberg gate [149] which relies on a single pulse of Rabi frequency Ω_P and time-dependent frequency Δ_P. This has a known time optimal solution [112]. One of the drawbacks of this implementation, is that it is not particularly robust in the face of signal imperfections and noise, such as atomic motion, laser intensity of frequency fluctuations. Some work has been done to improve its robustness to quasi-static errors [113, 155], but this often comes at the expense of the overall gate duration [81] which is undesirable. This motivates the determination of an optimal pulse sequence for the two-photon Rydberg gate which exhibits superior robustness to external perturbations without significantly increasing gate duration. A possible way to address this is to uses two time-dependent pulses with amplitudes Ω_P, Ω_S and frequencies Δ_P, Δ_S (akin to the Λ system setup). The Hamiltonian is defined in App. A.16.

Akin to the Λ system an additional reward term is introduced which reads $-w_x \cdot (\langle \overline{rr} \rangle + \langle \overline{\tilde{e}} \rangle + \langle \overline{\tilde{r}} \rangle + \langle \overline{\tilde{R}} \rangle)$, where the line denotes the time-averaged state population. This assigns lower rewards to dynamics which are susceptible to radiative decay, since one seeks coherent population transfer and speeds up the learning dynamics. Finding optimal protocols which simultaneously optimise both amplitude and frequency of Pump and Stokes beams is challenging since there exists a large number of possible control signals in a reasonably large Hilbert space. The fidelity is optimised with respect to

the Bell state fidelity, providing a realistic lower bound as in [192]. Further details are given in App. A.16.

In the optimisation the detuning of the Pump pulse is fixed to a constant value, since a time-dependent frequency chirp offers no advantages in terms of achievable maximum fidelities. Over a duration of 0.5μs one fixes Ω_S to a maximum value of 40MHz and $\Omega_P^2/(2\Delta_P)$ (the effective Rabi frequency) to a maximum value of 56.6MHz with a pump detuning of 2.5 GHz and obtains an optimal control signal which is shown in Fig. 56. It shall be noted that the signals are different from results in the literature since the amplitudes are constrained to start and end at zero. Simulated randomised benchmarking with 100 randomly generated initial states yields very similar fidelities to those obtained from the approximated Bell state fidelity given in 135 with the 0.25μs C-Z gate yielding $\mathcal{F} \approx 0.99957 \pm 0.0003$ (versus ≈ 0.99958 as obtained by taking the Bell state fidelity). Moreover, fidelities of $\overline{\mathcal{F}} > 0.999$ are simulated for time-dependent unbiased amplitude noise of $\sigma_\Omega = \pm 1$ MHz and frequency noise of $\sigma_\Delta = \pm 20$ MHz over 10 different randomly generated noise samples (cf. App. Eq. 131) for further details on the noise model.

There are existing works which addressed the two photon Rydberg gate. A similar protocol was optimised in [90], however this did not explicitly consider the decay of the Rydberg state and the optimised signals are extremely challenging to realise experimentally. Compared to [192] the solution we found with RL (cf. Fig. 56 is higher fidelity at up to $\mathcal{F} = 0.9996$ than their analytic solution $\mathcal{F} = 0.99$, as well as their numerical solution $\mathcal{F} = 0.997$. Compared to [218], this pulse sequence achieves similar fidelities but with much lower peak Rabi frequencies which implies reduced laser power requirements.

Moreover, compared to some other literature implementations, the pulse shown in Fig. 56 is a direct C-Z gate which does not require any additional single qubit rotations, which is the case for the aforementioned alternatives from the literature. When directly differentiating the input action with respect to the fidelity with a BFGS (Broy-

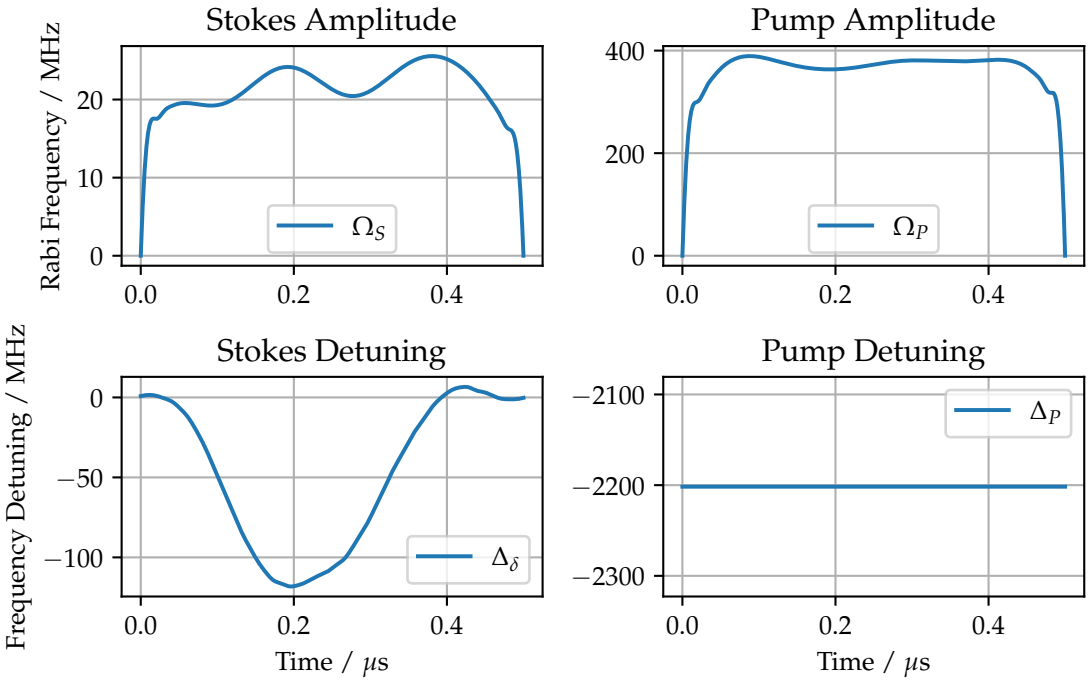

Figure 56: Shown are optimal signals for a two photon Rydberg gate directly realising a C-Z gate, with amplitudes (i.e. Rabi frequencies) for Stokes and Pump pulses in MHz shown in the top column. The effective maximum Rabi frequency of the pump pulse ($\Omega_P^2/2 * \Delta_P$) is ≈ 20 to match that of the Stokes pulse. Detunings of the Stokes and Pump pulse are shown in the bottom row. Note the symmetry of the Stokes detuning in time which shows a semblance of a reflection symmetry about its centre which ensures that a relative π phase is acquired between the basis states (cf. (135)) and their populations largely return to their initial values. This pulse yields a fidelity of 0.99917 for a 0.5μs duration and can be shortened to 0.25μs with all signals re-scaled by 2 which yields a fidelity of 0.99958 since the fidelity is mainly Rydberg level lifetime limited, with a finite blockade strength of 500MHz. This however, makes the experimental implementation challenging with high instantaneous bandwidth requirements.

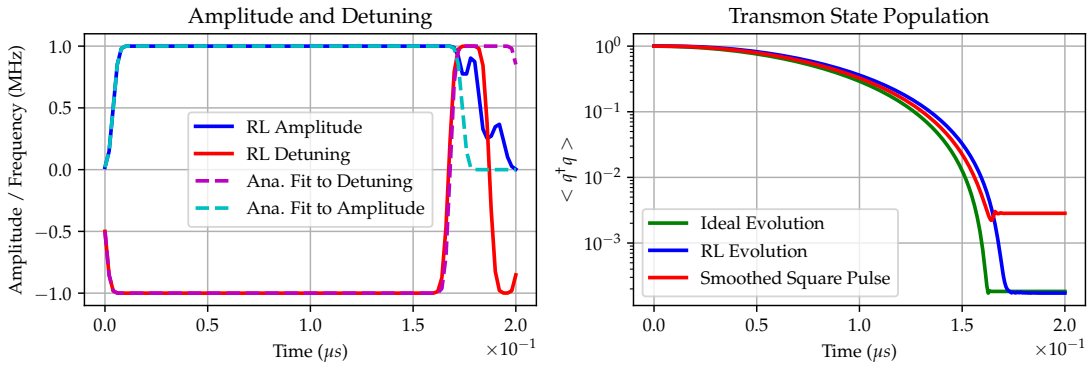

Figure 57: Optimal Waveform for Transmon Reset (left) discovered by RL and corresponding state evolution (right). The RL waveform (solid lines) amplitude evolution reminisces of a square-top Gaussian, with a smooth Heaviside-detuning that accounts for time-dependent frequency shifts. Equivalent reset performance is found by fitting a Heaviside-detuning reset and a Gaussian square amplitude waveform (dashed lines), simplifying experimental calibration. Our approach shows reset errors of $< 0.03\%$ matching the performance under an experimentally unrealistic ideal square pulse, and showing an order of magnitude improvement over a smoothed square pulse.

den–Fletcher–Goldfarb–Shanno, implemented in SciPy [236]) algorithm over 1000 iterations and for 32 random initial seeds one achieves a mean fidelity of 0.914 ± 0.0742 (one s.d.) showing the superiority of RL to reliably achieve high fidelity solutions. The enhanced computational scaleability offered by this implementation could be used to optimise multi qubit gates like a $C^k - Z$, where $k > 2$ which are also robust as in [177].

7.4.3 *Transmon Reset*

Superconducting quantum bits (qubits) have been a prominent and widely researched platform for quantum computing development, enticing extensive R & D efforts from the worlds largest technology companies like Google and IBM. Recent advances have extended transmon lifetimes beyond 0.5 ms [239], enabling longer quantum circuits and the implementation of error correction codes. To maximise circuit operations within the qubit's lifetime, transmons must be reset efficiently with high fidelity.

Two main reset techniques exist: conditional reset [190], which follows state measurement, and unconditional reset [150], which is faster and more robust. We focus on optimising waveforms for unconditional reset (cf. App. A.17 for further details).

The reset rate is proportional to drive strength, theoretically favouring high-amplitude square pulses for maximum fidelity. However, a drive-induced Stark shift alters the transmon's resonance frequencies [249]. In ideal conditions, a square pulse with a shift aware calibration frequency can be used to mitigate this. IBM demonstrated this approach experimentally, achieving 0.983 fidelity, while simulations under ideal conditions reached 0.996 fidelity [71]. This mismatch could be explained by experimentally realistic bandwidth constraints as square pulses have a finite rise and fall time, which induces a time-dependent frequency shift. While optimal control [85] has been applied to the task of reset pulse optimisation, minimal bandwidth constraints in this simulated optimisation environment, implied that no novel waveforms were found for improving the reset transition in a simple and realistic experiment. Using BFGS with direct differentiation of the input signal failed to optimise the multi-objective reward functions or satisfy realistic signal constraints. When optimising solely for fidelity, it remained slow and prone to local minima due to the large search space of non-smooth actions. Similarly, we were unable to attain any reasonable learning with alternative RL algorithms other than PPO.

We apply our RL method to optimise the transmon reset waveform under bandwidth constraints imposed by Gaussian-smoothing (for further details cf. App. Sec. A.14.3). Considering state of the art parameters, as given in the IBMQ experiment with a qubit lifetime T_1 of 500μs, we find that our RL approach achieves 0.9997 fidelity under realistic bandwidth constraints shown in Fig. 57 (cf. App. A.17 for further implementation details). This is compared with a perfect square pulse - which is not experimentally realistic - without any smoothing, and a calibrated square pulse with smoothing - which represents prior work [71]. The RL waveform matches the theoretical optimal fidelity of the perfect square pulse and improves the fidelity of waveform used in prior work [71] by an order of magnitude. In App. A.17 we explicitly compare the results with the parameters used in [71], and find that the RL-discovered reset waveform achieves the fidelity 0.997 compared to the measured fidelity of 0.983. A fitted Heaviside detuning

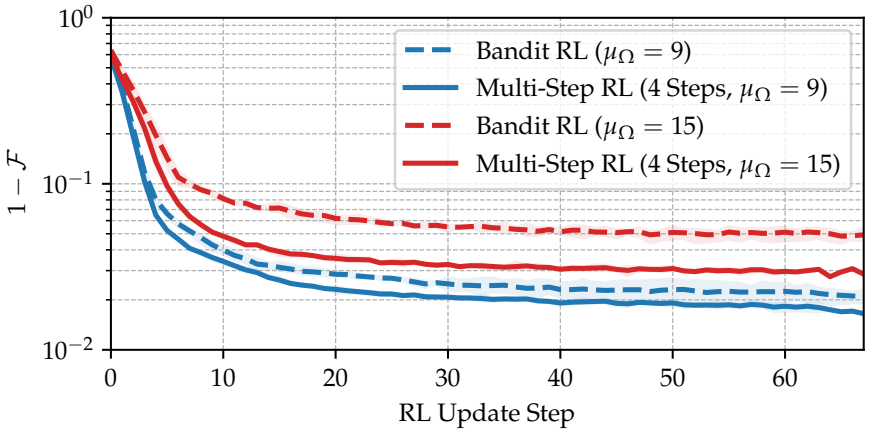

Figure 58: Comparison of mean infidelity for different values of μ_σ for bandit RL and multi-step RL. One can observe a 2x reduction in infidelity for larger noise bias μ_σ by using multi-step RL over the bandit setting.

function from the RL-discovered waveform corrects the drive-induced Stark shift, simplifying experimental calibration. The main notable feature of the waveform in Fig. 57 is a positive frequency chirp towards the ramp down of the Gaussian square pulse to account for the amplitude dependent Stark shift. Analytically, the detuning fit contains two parameters which denote the switch times t_1, t_2, as well as a maximum detuning magnitude Δ_{max} and amplitude Ω_{max} and duration t_0. These constitutes a fairly compact set of parameters that one should be able to calibrate experimentally as some signal chain or modelling imperfections will make it challenging to move directly from a simulation results to a real device.

7.4.4 *Multi-Step RL*

Here, the effectiveness of multi-step reinforcement learning strategies in achieving high fidelity control solutions under adverse noise conditions is investigated. Feedback on nanosecond timescales has been demonstrated experimentally [7], supporting the feasibility of this idea. This feedback can be realised by measuring classical signal noise without affecting quantum coherence. For example, in atomic quantum systems, laser

intensity I can be monitored separately from the quantum system, as $|\Omega|^2 \propto I$. Changes in I directly modulate $\Omega_{S/P}$ and thereby provide feedback.

In multi-step RL, the agent aims to maximise cumulative rewards over multiple steps, unlike the bandit setting where actions are independent. In this setup, at the start of each episode, a parameter μ_Ω is sampled uniformly from $[-\mu_\sigma, \mu_\sigma]$ to initialise an Ornstein–Uhlenbeck noise process (see App. A.13.2 and Eq. 131). The agent's control signal $a_t = \Omega_i(t)$ (amplitudes only) is affected by this noise, resulting in $\Omega_i' = \Omega_i + \nu_t$. In bandit RL, the agent does not observe the noise ν_t and selects the action in one step. Conversely, in multi-step RL, each episode is divided into four sections of 8 action samples corresponding to 0.25μs each. The agent initially observes $O_t = 0$ but receives the value of μ_Ω at times $t = 0.25$, 0.5, and 0.75μs (further implementation details are given in App. A.13.2), which corresponds to a measurement of the signal noise, by comparing it to a reference value. Fig. 58 illustrates that multi-step RL outperforms the bandit approach, especially as μ_σ increases beyond 10, indicating the effective use of the information encoded in the biased noise, to output improved control actions in real time.

7.5 PHYSICAL EXPERIMENTAL REALISATION

Although this whole chapter has largely focussed on simulations of physical experiments, it is instructive to see how it is physically feasible to implement the RL approach on real physical devices. There are two viable approaches to implementing our RL framework directly on physical devices, offering either more robustness or sample efficiency as will be outlined below by comparing direct on device training to device fine-tuning.

In the first approach, the RL agent is trained entirely on the physical system. The number of measurements required scales with the number of RL updates and the batch size (number of parallel simulations and measurements that were conducted on GPU

based simulations), i.e., $N_{\text{meas}} = \text{N}_{\text{updates}} \times b$ where b is the batch size. This estimation applies to systems such as the Lambda and Transmon environments, where efficient fidelity estimation is possible via single-shot measurements. In our simulations, the required number of measurements was 5100×256 and 256×128 for the Lambda and Transmon systems respectively. For the Rydberg system, fidelity estimation typically requires quantum state reconstruction (e.g., via Bell-state tomography or randomised benchmarking), introducing an additional overhead factor of at least $m = 4$. The measurement count in this case becomes: $N_{\text{meas}}^{\text{Rydberg}} = m \times \text{N}_{\text{updates}} \times b = 4 \times 6800 \times 256$.

In [26] the authors show that they generate complete gate sets on a real device with $\approx 10^6$ measurements per gate, with gate times on the order of 100ns. The reported experimental runtime is on the order of hours, with API calls dominating this time overhead, showing the feasibility of scaling to large measurement numbers in real devices. This is especially true for superconducting devices which have very fast gate times (hundreds of nanoseconds at most) and relatively fast measurement times on the order of microseconds compared to other physical platforms which can be orders of magnitude slower.

Our numerical experiments (e.g., Fig. 51) indicate that a large number of measurements may be necessary for convergence in this setting, often on the order of $\mathcal{O}(10^6)$ or higher.[9] This approach has the advantage of fully accounting for experimental imperfections or mismatches between experimental and simulation parameters. However, it may be prohibitively measurement intensive or slow, especially for systems requiring complex fidelity estimation techniques and those with slow measurement times.

An alternative and potentially more practical approach involves first training an RL policy (from which one can sample pulses) in simulation, and subsequently fine-tuning it on the experiment. This strategy exploits the numerical solution as a strong prior, allowing for correction of mismatches between the simulated and real systems

9 For instance, in the Λ system using PPO, convergence required approximately 5000×256 measurements. We note, however, that smaller batch sizes and reduced action spaces can decrease this number significantly, which we have not explored much here, albeit often at the cost of policy expressivity and robustness.

through targeted optimisation using experimental data. In [141], it was demonstrated that this hybrid approach can reduce the number of on-device steps to just a few thousand. In the context of this chapter, the numerical policies obtained through simulation can similarly serve as a valuable initialisation point for experimental demonstrations. This method significantly lowers the experimental cost and accelerates implementation, though it is unlikely to succeed when there is a big mismatch between simulation and experimental parameters.

Table 13: Example timing breakdown for experiment described in Sec. 5.6. Note that one may take many shots per RL policy update.

Experimental Step	$\approx$ **Time (s)**	**Note**
RL policy update (classical)	0.0001	
AWG pulse loading	2.5	
MOT loading	2.00	n_{shots} per update
State preparation	0.00005	n_{shots} per update
Quantum control pulse	0.0000001	n_{shots} per update
Fluorescence Readout pulse	0.0005	n_{shots} per update
Data processing	0.001	

It is important to identify and mitigate various timing bottlenecks when moving from simulation to experiment. The time it takes to play a particular pulse is negligible compared to the other time scales. The time scales for the MOT fluorescence experiment in Ch. 5 are shown in Tab. 13. The biggest bottleneck is the loading of pulses onto the AWG and reloading the MOT after every shot. The latter dominates if several shots are taken per RL Update, as is desirable for reducing variance and shot noise on the fidelity estimate. While these timings are not preclusive for performing closed-loop RL, they do show how it makes it difficult to scale to a large number of shots and updates in reasonable times. For example running 100 shots per update for 1000 updates would take over 50 hours. It might therefore be more promising to start with a simulation based policy and to fine-tune this with experimental sampling to reduce the total number of measurements.

Furthermore, in an experiment, measurement outcomes are subject to noise that is unrelated to the optimality of the pulse, primarily due to imperfect state preparation, and shot noise or systematic readout errors. Readout shot noise is fairly significant in the setup described in Sec. 5.6. This shot noise can significantly affect the learning dynamics and stability of RL-based optimisation. More samples can generally reduce shot noise, but this comes at the expense of longer time overhead per RL update. While some control signal noise was considered in this chapter, no readout noise was explicitly incorporated into the learning process. It is an interesting extension to try and train an algorithm which can learn with those limitations in mind.

7.6 OUTLOOK

A physics-constrained reinforcement learning framework was developed for quantum control, enabling high-fidelity operations in open quantum systems by incorporating physical constraints and priors about the time dynamics of optimal solutions. Fidelities exceeding 0.999 were achieved across three different tasks, while improved robustness to time-dependent noise was also demonstrated. In principle, this technique can be applied to any quantum system which can be accurately simulated and where a cost function can be defined. Limitations of this approach will materialise for systems with highly non-adiabatic dynamics, where the constraints imposed no longer represent physically desired solutions.

For the Transmon reset problem, smooth waveforms were discovered that are compatible with standard hardware and can be described using compact functional parametrisations which enables straightforward experimental calibration. Experimental implementation of the simulation learned pulses would greatly underscore the effectiveness of the RL approach.

Although the methods were conducted in simulation, the generated waveforms are physically feasible and suitable for implementation on experimental hardware. This

might prove difficult for systems where it is difficult to establish accurate models, especially if they have complex noise models or are too large to simulate. One can mitigate this by directly sampling from a device. In this practical experimental setting, timing bottlenecks, such as slow AWG waveform loading and MOT reloading, must be considered, as they limit the scalability of a closed-loop RL approach. Despite these constraints, there is hope that on-device RL remains viable when using simulated policies as starting points to reduce the device time and the number of samples for convergence.

Looking ahead, extending this approach to more complex pulse sequences and tasks such as multi-qubit gates would be an interesting extension.

8

CONCLUSIONS

Quantum technology is a rapidly evolving field, with researchers in both academia and industry continuously advancing a variety of hardware platforms. A central challenge remains the time-dependent control of quantum systems, an essential requirement for realising the full theoretical potential of quantum computing.

A scheme for efficiently generating high-coincidence-rate n-photon bursts was presented in Chapter 3, with extensive numerical simulations demonstrating its promise. It was compared to alternative single-photon sources in terms of expected brightness. Our results show that achieving quantum advantage demonstrations in boson sampling is possible, but not at high frequencies. Atomic sources remain valuable for a range of other applications. In particular, their long coherence times and the ability to precisely shape photons enable multi-photon interference experiments, where photons are routed through optical networks or MMIs with measurement-induced feedback for controlled spatial routing, as demonstrated in [7].

Measurement-based quantum computing relies on the generation of entangled resource states, upon which adaptive measurements are performed to implement arbitrary computations. Generating such states with high fidelity and at high rates remains a major challenge. The scheme outlined in Chapter 4 builds on the seminal work of [223], which demonstrated that atomic cQED platforms outperform all other known physical systems for generating linear cluster and GHZ states with more than four photons. This advantage persists despite quantum dots typically offering repetition rates an order of magnitude higher than those of atomic sources [227], due to the superior coherence properties of an atomic source.

We proposed several key modifications to the scheme described in [223], including a switch from polarisation to time-bin encoding and the use of tailored fast Raman pulses to significantly improve expected coincidence rates. Additionally, we employed Heisenberg picture evolution to calculate photon correlations across different emission windows, a technique not previously discussed for atomic cQED systems. This provides new insight into their temporal dynamics. Future work could explore the decoherence mechanisms that were numerically bounded in this work in greater depth to uncover more general system behaviours. Much work also remains to make such a scheme experimentally viable, especially since measurement-based quantum computing typically requires the generation of multi-dimensional, not merely linear, cluster states. Encouraging progress in this direction has been made experimentally (e.g., [224]), while theoretical proposals such as that in [208] point toward further promising avenues for generating 2-d cluster states.

Implementing optimised nanosecond-scale pulses in real-world experimental setups is challenging due to numerous experimental imperfections. Chapter 5 describes the development of a new laser system and an optical pulse calibration technique that optimises the electronic driving signals applied to an AOM-based pulse shaping setup. This leads to significantly improved output pulse envelopes. These pulses were successfully tested in a coherent re-preparation task using an optical molasses to provide cold atoms, demonstrating improvements over non-distortion corrected pulses. While the realisation of a fully deterministic atom-cavity photon source capable of supporting the schemes described in earlier chapters remains a work in progress, valuable insights into cavity stability and control are also presented. Although atom-cavity systems face scalability challenges - due to the overhead of vacuum and optical components per source - they continue to offer unmatched emitter controllability, which is critical for advanced cQED experiments.

Scaling pulse level quantum control to larger systems introduces substantial challenges. As qubit counts grow, so does the complexity of generating and managing

thousands of parallel, high-fidelity control waveforms, imposing high demands on hardware, particularly in memory, latency, bandwidth, and timing precision. Chapter 6 shows that considerable memory compression and high-fidelity waveform generation can be achieved using a tailored cubic spline fitting and hardware-aware optimisation routine. Simulations demonstrate that fidelities are comparable to standard AWG output, indicating strong potential for real-time control with FPGA platforms. This is particularly advantageous for atomic systems, which require long (μs-scale) pulses.

This work lays the foundation for future investigations into hardware resource requirements for scalable control, particularly with respect to memory. Further exploration is needed to extend these techniques to non-atomic platforms. One promising direction is integrating such compression routines into closed-loop experimental control systems, or using them to design optimal atom transport waveforms (or other waveforms required for different control tasks) in a memory-efficient spline format.

Quantum systems featuring large Hilbert spaces, multiple dissipative channels, and real-world imperfections remain notoriously difficult to control. Chapter 7 introduces a reinforcement learning (RL) framework for quantum control. This approach enables high-fidelity operations that are robust to time-dependent noise in open quantum systems. By incorporating physical constraints and priors into the RL process, smooth and hardware-compatible waveforms are learned. The RL framework also eliminates the need for costly step-by-step simulations, treating system evolution as a black box. GPU parallelisation provides the necessary computational resources to make this approach practical for more complex systems. While limitations persist for systems with highly non-adiabatic dynamics or uncharacterised noise, the method is applicable to a wide range of systems that one can simulate and have well-defined cost functions. Future extensions could target more complex pulse sequences or multi-qubit gates, as well as theoretical optimisation of re-preparation pulses for the complex D_2 and D_1 line structure in Chapter 3, or coherence-optimised photon emission relevant to the time-bin scheme of Chapter 4. On-device sampling also offers a compelling path forward,

helping to overcome mismatches between simulation and experiment, or to account for hidden noise sources. Some initial steps in this direction are discussed in Chapter 5, where experimental latencies currently remain a challenge. Real time integrations of GPUs into quantum computing experiments is also an interesting direction worth exploring.

No one can predict where the field will be in ten years, but it is likely that the design and implementation of fast, high-fidelity control mechanisms will remain a central challenge in the development of quantum technologies and associated experiments.

A

APPENDIX

A.1 ^{87}Rb TRANSITION LINE DATA

The ^{87}**Rb** D_2 Transition line data is shown in Fig. 59.

A.2 BOSON SAMPLING

A.2.1 *Connection Between Boson Sampling and Matrix Permanent Calculations*

The interference of multiple bosons in a multimode interferometer is related to the calculation of the matrix permanent, which is known to be #P-hard [1]. Boson sampling involves sending photons through a linear optical interferometer and measuring the output distribution of these photons. This involves an n-photon input state and an m-mode interferometer characterised by a unitary matrix U. The matrix permanent is a function similar to the determinant but without alternating signs [199].

Experimentally one starts with a particular input distribution and measures an output distribution, subject to losses. The probability of detecting a particular output con-

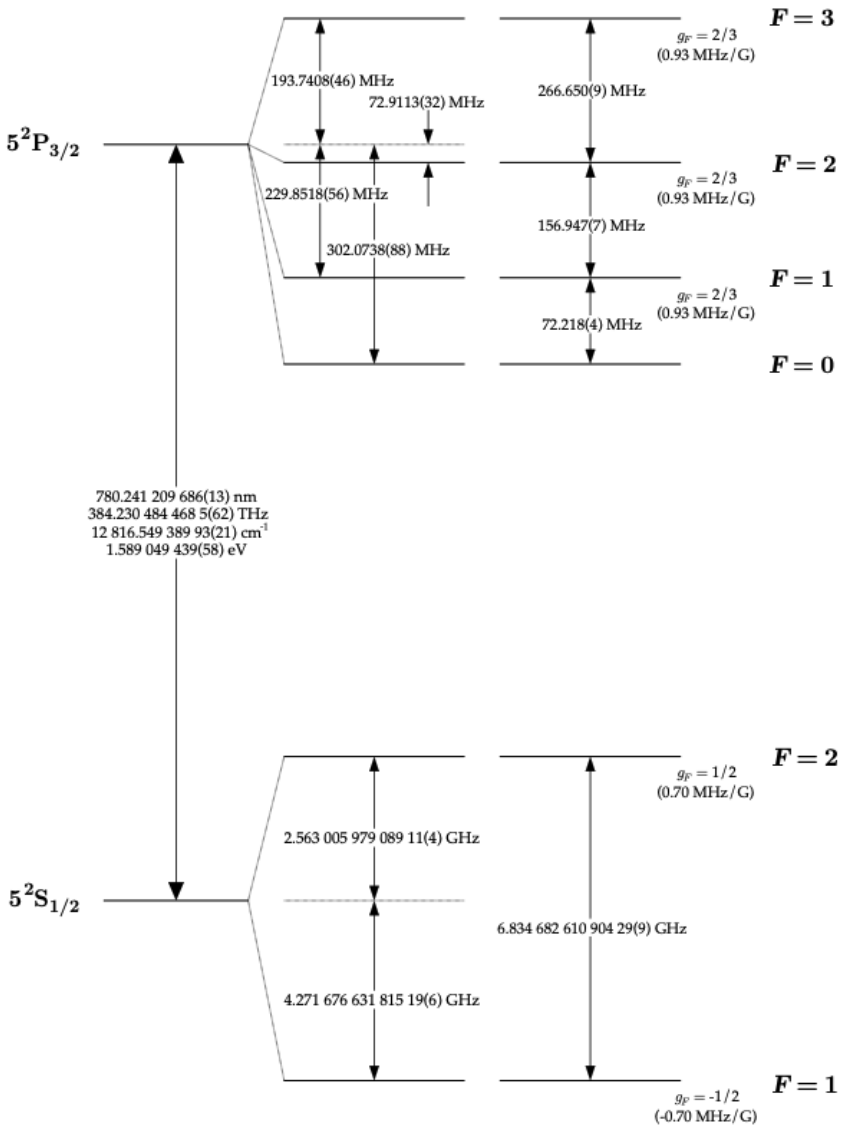

Figure 59: The D_2 transition between the $5^2S_{1/2}$ level and the $5^2P_{3/2}$ levels are shown with hyperfine structure. Diagram taken from [216].

figuration of photons is related to the absolute value squared of the permanent of a sub-matrix of U which reads [199]:

$$U_{S,T} = \begin{pmatrix} u_{s_1,t_1} & u_{s_1,t_2} & \cdots & u_{s_1,t_n} \\ u_{s_2,t_1} & u_{s_2,t_2} & \cdots & u_{s_2,t_n} \\ \vdots & \vdots & \ddots & \vdots \\ u_{s_n,t_1} & u_{s_n,t_2} & \cdots & u_{s_n,t_n} \end{pmatrix} \tag{90}$$

Let $S = \{s_1, s_2, \ldots, s_n\}$ and $T = \{t_1, t_2, \ldots, t_{n'}\}$ be two subsets of size n, n' representing the input and output modes, respectively (for post-selecting with losses in accordance to [3] $n' < n$, but for simplicity in the following $n = n'$). The sub matrix $U_{S,T}$ is formed by taking the rows indexed by S and the columns indexed by T. The probability $P(T|S)$ of detecting the output configuration T given the input configuration S reads [199]:

$$P(T|S) = \frac{|\text{perm}(U_{S,T})|^2}{\prod_{i=1}^{n} s_i! \prod_{j=1}^{n} t_j!} \tag{91}$$

This shows that the experimental determination of $P(T|S)$ is, up to combinatorial factors, equivalent to the absolute square of the the permanent of the sub-matrix $U_{S,T}$.

A.2.2 *Classical Simulation of Boson Sampling*

The algorithm in [188] can produce an approximation of the output distribution measured from a noisy boson sampling setup for a permanent of order k (this also corresponds to the number of indistinguishable bosons), for a photon indistinguishability

x^2 and a total photon number n with l lost photons after post selection the error is given by E:

$$\frac{x^{2(k+1)}\left(\frac{n-l}{n}\right)^{k+1}}{1-\left(x^2\frac{n-l}{n}\right)} \geq E^2 \tag{92}$$

This bound can be used when determining the total photon number n, as well as the number of lost photons l for a source with a given indistinguishability and serves as a lower bound on the classic simulation error for n photons interfering on a multi-mode interferometer, such that an experimental measurement of a distribution of k photons constitutes a genuine demonstration of quantum advantage.

A.3 DYNAMICAL STARK SHIFT CALCULATION

An optical field $\mathbf{E}$ with angular frequency ω can be defined as a complex vector $\mathbf{E} = 2\mathrm{Re}(\mathbf{E}_0^+ e^{i\omega t})$, such that $I = 2|\mathbf{E}_0^+|^2\epsilon_o c$ and its three field components $\{-1, z, 1\}$ denote the three polarisation components $\{\sigma_-, \pi, \sigma_+\}$. This field induces a hyperfine state dependent Stark shift which reads [215]:

$$\begin{aligned}\Delta E(F, m_F; \omega) = &-\alpha^{(0)}(F;\omega)\left|\mathbf{E}_0^{(+)}\right|^2 - \alpha^{(1)}(F;\omega)\left(|\mathbf{E}_{0,-1}^{(+)}|^2 - |\mathbf{E}_{0,1}^{(+)}|^2\right)_z \frac{m_F}{F} \\ &- \alpha^{(2)}(F;\omega)\frac{\left(3\left|E_{0z}^{(+)}\right|^2 - \left|\mathbf{E}_0^{(+)}\right|^2\right)}{2}\left(\frac{3m_F^2 - F(F+1)}{F(2F-1)}\right).\end{aligned} \tag{93}$$

The scalar, vector and tensor polarisabilities are defined as [215]:

$$
\begin{aligned}
\alpha^{(0)}(F;\omega) &= \sum_{F'} \frac{2\omega_{F'F}\,|\langle F\|\mathbf{d}\|F'\rangle|^2}{3\hbar\left(\omega_{F'F}^2-\omega^2\right)} \\
\alpha^{(1)}(F;\omega) &= \sum_{F'}(-1)^{F+F'+1}\sqrt{\frac{6F(2F+1)}{F+1}}\begin{Bmatrix} 1 & 1 & 1 \\ F & F & F' \end{Bmatrix}\frac{\omega\,|\langle F\|\mathbf{d}\|F'\rangle|^2}{\hbar\left(\omega_{F'F}^2-\omega^2\right)} \\
\alpha^{(2)}(F;\omega) &= \sum_{F'}(-1)^{F+F'}\sqrt{\frac{40F(2F+1)(2F-1)}{3(F+1)(2F+3)}}\begin{Bmatrix} 1 & 1 & 2 \\ F & F & F' \end{Bmatrix}\frac{\omega_{F'F}\,|\langle F\|\mathbf{d}\|F'\rangle|^2}{\hbar\left(\omega_{F'F}^2-\omega^2\right)},
\end{aligned}
\tag{94}
$$

where the transition frequencies between different lines are denoted $\omega_{F'F}^2$. Using time dependent fields $\mathbf{E}(t)$ one can derive the state dependent Stark shifts as a sum of the different shifts arising from different transitions the state is coupled to, for given laser polarisations and frequencies. More details and a complete derivation is found in [215].

A.4 FAR DETUNED RAMAN PULSE HAMILTONIAN DERIVATION

Without loss of generality, the derivation for an effective far off resonant coupling focuses on a single $\Lambda-$system comprised of two ground states $|g_1\rangle$ and $|g_2\rangle$, and a single excited state $|e\rangle$. One laser couples $|g_1\rangle \leftrightarrow |e\rangle$ with Rabi frequency Ω_1 and detuning Δ_1, and another laser couples $|g_2\rangle \leftrightarrow |e\rangle$ at Rabi frequency Ω_2 and detuning Δ_2. $\Delta = (\Delta_1+\Delta_2)/2$ and $|\Delta_1-\Delta_2| = \delta$. The derivation of the effective Hamiltonian is presented in the limit of large detuning $\Delta \gg \Omega_1, \Omega_2$, where the excited state can be adiabatically eliminated. It is convenient to move to the interaction picture with a frequency shift $\delta/2$, such that $H = H_0 - \delta/2$. The interaction-picture Hamiltonian for

this system in the rotating-wave approximation can be written in matrix form using the basis $\{|g_1\rangle, |g_2\rangle, |e\rangle\}$:

$$\hat{H} = \frac{\hbar}{2}\begin{bmatrix} -\delta & 0 & \Omega_1^* \\ 0 & \delta & \Omega_2^* \\ \Omega_1 & \Omega_2 & 2\Delta \end{bmatrix}. \tag{95}$$

In the regime $\Delta \gg \Omega_1, \Omega_2$, and $\Delta \gg \delta$ the excited state $|e\rangle$ can be adiabatically eliminated. The key idea is that the coupling through the far-detuned excited state introduces an effective interaction between the ground states.

The effective Hamiltonian can be derived using the projection operator method. Consider a projection operator P onto the ground-state subspace $|g_1\rangle\langle g_1| + |g_2\rangle\langle g_2|$ to express the effective Hamiltonian as:

$$H_{\text{eff}} = -\frac{\hbar}{2}\begin{bmatrix} \delta + \frac{|\Omega_1|^2}{2\Delta} & \frac{\Omega_1\Omega_2^*}{2\Delta} \\ \frac{\Omega_2\Omega_1^*}{2\Delta} & -\delta + \frac{|\Omega_2|^2}{2\Delta} \end{bmatrix}. \tag{96}$$

This method can be extended to an atom with n ground states and m excited states by considering a projection into subspaces defined by every possible Λ system (as permitted by the dipole selection rules) formed by laser beam pairs of particular polarisations.

This convenient representation does have some limitations, which are discussed extensively in [176]. In particular, it doesn't yield good approximations for $>$ two photon transitions, small Δ and large Ω. Additionally, it does not provide an accurate description of excited state populations and assumes the that the derivative of $|e(t)\rangle$ vanishes, this can lead to issues with normalisation. Some more rigorous alternatives are offered in [176] using the first order Markov approximation but they are not required for the

analysis of Raman transitions in a multilevel atom where traditional adiabatic elimination provides sufficient accuracy.

To analyse the physics of the Raman process accurately and realistically one must consider the effects of spontaneous emission when the excited state is eliminated. Spontaneous emission induces incoherent decay from $|e\rangle$ to $|g_i\rangle$ with rates proportional to $|C_{e\to g_i}|^2\Gamma$, where Γ is the total decay rate of the excited state. After adiabatic elimination of $|e\rangle$, the effective Lindblad operators for the decay between ground states are:

$$\hat{L}_{ij} = \sqrt{\Gamma_{ij}}|g_i\rangle\langle g_j|, \tag{97}$$

where the effective decay rate is determined by the subspace projection described for the determination of the effective Hamiltonian. One can approximate $\langle e|\,\rho\,|g_i\rangle \approx \Omega_i/(2\Delta)$ which leads to an effective Lindblad operator:

$$\Gamma_{ij} \approx \Gamma\frac{\Omega_i\Omega_j^*}{4\Delta^2}C_{e\to g_i}C^*_{e\to g_j}. \tag{98}$$

Here the Clebsch-Gordan coefficients capture the branching ratios for decay into different ground states. It is evident from the Δ^{-2} dependence in Eq. 98 that a large detuning can significantly suppress spontaneous emission.

A.5 TIME-BIN V-STIRAP MCWF ERROR PROPAGATION

Beginning with the expression for the V-STIRAP re-excitation error:

$$f = \frac{1}{\mathcal{N}_t}\left(\frac{\mathcal{N}^l_{\text{rx}}}{\mathcal{N}^l_{\text{ems}}}\cdot\mathcal{N}^e_{\text{ems}} + \frac{\mathcal{N}^e_{\text{rx}}}{\mathcal{N}^e_{\text{ems}}}\cdot\mathcal{N}^l_{\text{ems}}\right)$$

Define the two terms:

$$f_1 = \frac{\mathcal{N}^l_{\text{rx}}}{\mathcal{N}^l_{\text{ems}}\mathcal{N}_t} \cdot \mathcal{N}^e_{\text{ems}}, \quad f_2 = \frac{\mathcal{N}^e_{\text{rx}}}{\mathcal{N}^e_{\text{ems}}\mathcal{N}_t} \cdot \mathcal{N}^l_{\text{ems}}, \quad \text{such that } f = (f_1 + f_2) \tag{99}$$

From the MCWF approach (cf. Sec. 2.0.1), all fractional errors $\delta\mathcal{N}_i/\mathcal{N}_i$ carry the same relative error, denoted as: $\varepsilon_r = \frac{1}{\mathcal{N}_t}$. Then, one applies standard error propagation rules for products and quotients. Each term in f depends on three independent variables, so:

$$\delta f_1 = f_1 \cdot \sqrt{\varepsilon_r^2\,(1+1+1)} = f_1 \cdot \sqrt{3} \cdot \epsilon_r$$

$$\delta f_2 = f_2 \cdot \sqrt{\varepsilon_r^2\,(1+1+1)} = f_2 \cdot \sqrt{3} \cdot \varepsilon_r$$

The total uncertainty in f is then:

$$\delta f = \sqrt{(\delta f_1)^2 + (\delta f_2)^2} = \sqrt{f_1^2 + f_2^2} \cdot \sqrt{3} \cdot \epsilon_r. \tag{100}$$

Therefore, the relative error in f is given by:

$$\frac{\delta f}{f} = \sqrt{3} \cdot \epsilon_r \cdot \frac{\sqrt{f_1^2 + f_2^2}}{f_1 + f_2}$$

Substituting the definitions of f_1, f_2 and ϵ_r, one arrives at the final symbolic expression for the relative uncertainty:

$$\varepsilon_f = \frac{\sqrt{3}}{\mathcal{N}_t} \cdot \frac{\sqrt{\left(\frac{\mathcal{N}^l_{\text{rx}}}{\mathcal{N}^l_{\text{ems}}} \cdot \mathcal{N}^e_{\text{ems}}\right)^2 + \left(\frac{\mathcal{N}^e_{\text{rx}}}{\mathcal{N}^e_{\text{ems}}} \cdot \mathcal{N}^l_{\text{ems}}\right)^2}}{\frac{\mathcal{N}^l_{\text{rx}}}{\mathcal{N}^l_{\text{ems}}} \cdot \mathcal{N}^e_{\text{ems}} + \frac{\mathcal{N}^e_{\text{rx}}}{\mathcal{N}^e_{\text{ems}}} \cdot \mathcal{N}^l_{\text{ems}}} \tag{101}$$

A.6 STABILISER EXPECTATION VALUES AS OPERATOR EXPECTATION VALUES

For the $n = 2$ photon case we can write example expressions for stabilisers relevant for the calculation of GHZ states in terms of two point correlation functions. For the S_k^{GHZ} stabiliser this can be expressed using the following elements elements of the two point correlation matrix (cf. 56):

$$\langle S_k^{\text{GHZ}} \rangle = \frac{1}{\mathcal{N}^{(2)}} \langle \mathcal{G}^{(2)}_{EELL}(t_{k-1}, t_k) + \mathcal{G}^{(2)}_{LLEE}(t_{k-1}, t_k) + \mathcal{G}^{(2)}_{ELLE}(t_{k-1}, t_k) + \mathcal{G}^{(2)}_{LEEL}(t_{k-1}, t_k) \rangle. \tag{102}$$

Similarly, for the cluster state one can express the stabiliser in terms of three photon correlation matrix elements:

$$\langle S_j^{\text{cluster}} \rangle = \frac{1}{\mathcal{N}^{(3)}} \langle \mathcal{G}^{(3)}_{Z_{j-1}\text{-}X_j\text{-}Z_{j+1}}(t_{j-1}, t_j, t_{j+1}) \rangle, \tag{103}$$

where $\mathcal{N}^{(3)}$ is the normalisation for the three photon correlation matrix and:

$$\mathcal{G}^{(3)}_{Z_{j-1}\text{-}X_j\text{-}Z_{j+1}}(t_{j-1}, t_j, t_{j+1}) = \langle (\hat{e}^\dagger_{j-1}\hat{e}_{j-1} - \hat{l}^\dagger_{j-1}\hat{l}_{j-1})(\hat{e}^\dagger_j \hat{l}_j + \hat{l}^\dagger_j \hat{e}_j)(\hat{e}^\dagger_{j+1}\hat{e}_{j+1} - \hat{l}^\dagger_{j+1}\hat{l}_{j+1}) \rangle. \tag{104}$$

A.7 NUMERICAL INTEGRATION OF PHOTONIC DENSITY MATRIX ELEMENTS

Consider the double integral

$$\langle \mathcal{G}^{(2)}_{ELEL} \rangle = \int_0^T \int_{2T}^{3T} dt_1 \, dt_2 \, \mathcal{G}^{(2)}_{ELEL}(t_1, t_2), \tag{105}$$

where $\mathcal{G}^{(2)}_{ELEL}(t_1, t_2)$ is a two-time correlation function evaluated over the domains $t_1 \in [0, T]$ and $t_2 \in [2T, 3T]$.

To approximate this integral numerically, each time interval is discretised into n_t uniformly spaced grid points:

$$t_1^{(i)} = i\frac{T}{n_t - 1}, \quad i = 0, 1, \ldots, n_t - 1, \tag{106}$$

$$t_2^{(j)} = 2T + j\frac{T}{n_t - 1}, \quad j = 0, 1, \ldots, n_t - 1. \tag{107}$$

The function $\mathcal{G}^{(2)}_{ELEL}(t_1^{(i)}, t_2^{(j)})$ is evaluated at each grid point $(t_1^{(i)}, t_2^{(j)})$.

A cubic interpolation $\mathcal{G}^{(2)}_{\text{interp}}(t_1, t_2)$ is then constructed from the discrete data to ensure smoothness across the domain, guaranteeing continuity of the first derivatives with respect to both t_1 and t_2.

The integral is approximated by applying a composite trapezoidal rule to the interpolated surface:

$$\langle \mathcal{G}^{(2)}_{ELEL} \rangle \approx \sum_{i=0}^{n_t-1} \sum_{j=0}^{n_t-1} \mathcal{G}^{(2)}_{\text{interp}}(t_1^{(i)}, t_2^{(j)})\, \Delta t_1 \Delta t_2, \tag{108}$$

where the now, a smaller grid spacing ($n_s > n_t$) is used:

$$\Delta t_1 = \frac{T}{n_s - 1}, \qquad \Delta t_2 = \frac{T}{n_s - 1}. \tag{109}$$

This method straightforwardly applies to other photonic density matrix elements or dimensions. The rationale behind this method is that the quantum correlations are structurally similar in time to the emitted photon shapes which are smooth in time due to the exponential decay out of the cavity and the smooth shape of the driving pulses. Linearly interpolating between discrete points or calculating trapezoidal integrals without smoothing often leads to an overestimation of off-diagonal matrix elements (which are the most computationally demanding, making the evaluation of finely sampled grids impractical) such as $\langle G^{(2)}_{ELEL} \rangle$ such that one obtains unphysical correlations or fidelity bounds > 1.

A.7.1 *Uncertainty Propagation for GHZ Pauli Expectation Values*

Uncertainties are propagated for the stabilisers S_1 and S_k, assuming a fixed uncertainty δ for each real or imaginary part of the measured integrals. While this is not exactly the same for all the different matrix elements which are comprised of discretised time integrals and interpolations between time grid points, the order of magnitude in absolute numerical accuracy will be the same. The total sum of diagonal terms is defined as:

$$\mathcal{N} = \sum_{\tilde{i}} |\tilde{i}\rangle \langle\tilde{i}| \tag{110}$$

where $\tilde{i}$ are super-indices corresponding to the are the real-valued diagonal density matrix elements of an $n-$photon density operator. The associated uncertainty is

$$\Delta(\mathcal{N}) = \sqrt{2^n}\,\delta \tag{111}$$

with 2^n the number of diagonal elements corresponding to the photon number n. The numerator of S_1 is the sum over the off-diagonal complex elements which are not yet normalised by $\mathcal{N}$, which one can denote $\mathcal{X} = \mathcal{N} S_1^{GHZ}$. Thus,

$$\Delta(\mathcal{X}) = \sqrt{2^{n+1}}\delta, \tag{112}$$

the uncertainty includes an additional factor of 2 due to the numerical real and imaginary uncertainties. The expectation value S_1^{GHZ} has an uncertainty which propagates from the uncertainty in the numerator and the normalisation term:

$$\Delta S_1^{GHZ} = |S_1^{GHZ}| \sqrt{\left(\frac{\Delta(\mathcal{X})}{\mathcal{X}}\right)^2 + \left(\frac{\Delta(\mathcal{N})}{\mathcal{N}}\right)^2}. \tag{113}$$

Similarly, the S_k^{GHZ} stabilisers (associated with σ_z-like operators) are linear combinations of diagonal density operator terms: $\mathcal{Z}_k = \mathcal{N}\sigma_{k-1}^z\sigma_k^z$. Each $\mathcal{Z}_k$ has an uncertainty:

$$\Delta(\mathcal{Z}) = \sqrt{2}\delta \tag{114}$$

and the propagated uncertainty on S_k is

$$\Delta S_k^{GHZ} = \left|S_k^{GHZ}\right| \sqrt{\left(\frac{\Delta(\mathcal{Z}_k)}{(\mathcal{Z}_k)}\right)^2 + \left(\frac{\Delta(\mathcal{N})}{\mathcal{N}}\right)^2}. \tag{115}$$

These uncertainties of individual stabiliser values can then be combined in accordance with Eq. 65.

A.8 AOM BACKGROUND

Acousto-optic modulators are widely used devices for controlling the amplitude and frequency of laser beams. Their operation is based on the acousto-optic effect, where an acoustic wave propagating through an optical medium induces a modulation of the refractive index. This modulation acts as a diffraction grating for the incident optical beam.

When an AOM is driven at a radio-frequency (RF) ν_{RF}, an acoustic wave with wavelength $\Lambda = \frac{v}{\nu_{\mathrm{RF}}}$ where v is the acoustic velocity, is generated inside the medium. The incident laser light of wavelength λ can be diffracted according to the Bragg condition:

$$2\Lambda \sin\theta_B = m\lambda, \tag{116}$$

where Λ is the acoustic wave-vector, m the diffraction order and θ_B is the Bragg angle, defined as the angle between the incident light and the acoustic wavefronts.

The first-order diffracted beam experiences a frequency shift equal to the acoustic frequency ν_{RF}. If the input optical frequency is ν_0, then the diffracted beam has frequency

$$\nu_1 = \nu_0 \pm \nu_{RF}. \tag{117}$$

The diffraction efficiency η, defined as the fraction of incident optical power diffracted into the first order can be approximated by [137]:

$$\begin{aligned} \eta &= \sin^2\left(\frac{\Delta\phi}{2}\right), \\ \Delta\phi &= \frac{\pi c}{\nu}\sqrt{2 r_A M_2 P_{RF}} \end{aligned} \tag{118}$$

where r_A is the aspect ratio (length to height) of the sound field and M_2 is defined as $n^6 p^2/v^2\rho$ where n is the refractive index of the medium, p is the effective photo-elastic constant, ρ is the density and v is the acoustic velocity. The acoustic velocity is given by the product of the acoustic wavelength Λ and the external RF frequency ν_{RF}. Thus, the input power P_{RF} should be chosen to induce a desired phase $\pi/2$ phase in Eq. 118, maximising the power diffracted into the first order.

The response time of an AOM determines how rapidly the device can respond to changes in the external RF signal, and determines how fast the optical beam can be modulated. The primary parameters influencing this are the acoustic velocity, the interaction length, and the size of the laser beam within the crystal. The response time τ of an AOM is fundamentally limited by the time required for the acoustic wave to traverse the optical interaction region and can be approximated as:

$$\tau \approx \frac{d}{v_{\mathrm{acoustic}}}, \tag{119}$$

where d denotes the optical beam diameter at the interaction point inside the crystal, and v_{acoustic} is the speed of sound in the AOM crystal. Minimising the optical beam

diameter and utilising materials with higher acoustic velocities both help in achieving shorter response times.

There exist inherent design trade-offs in optimising the operation of an AOM. A smaller optical beam increases the achievable bandwidth but may reduce diffraction efficiency if the beam under-fills the acoustic interaction region.

In a single-pass setup, the laser beam enters the AOM, interacts with the acoustic wave, and the first-order diffracted beam is used. The frequency of the output beam is shifted by $\pm\nu_{RF}$, and the angular deviation is determined by the Bragg condition. An illustration of this configuration is shown in Fig. 60.

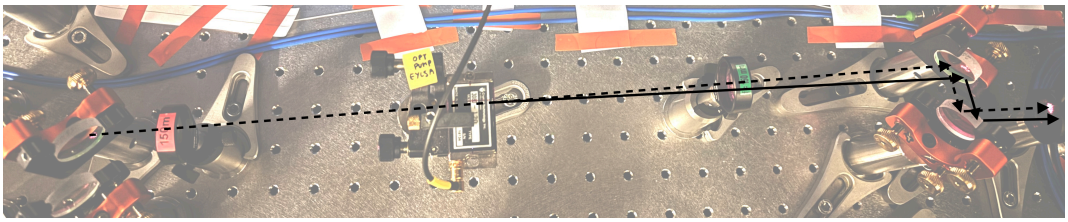

Figure 60: Single-pass configuration of an AOM. The laser beam is shifted by $-\nu_{RF}$ after a single diffraction. The zeroth order is shown as a black dotted line whereas the -1st order is shown in a solid black line.

In a double-pass setup, the first-order diffracted beam is retroreflected back into the AOM, so that it undergoes a second diffraction. This second pass results in a further frequency shift of $\pm\omega_{RF}$, leading to a total frequency shift of $\pm 2\omega_{RF}$.

A.9 NLMS FILTER UPDATE RULE

The NLMS update rule is [99]:

$$\Delta \mathbf{w}[n] = \frac{\mu \cdot \epsilon[n] \cdot \mathcal{E}_{\text{in}}[n,:]}{\alpha + \|\mathcal{E}_{\text{in}}[n,:]\|^2}, \tag{120}$$

where μ is the learning rate, α is a small regularisation term used to avoid instability when $\|\mathcal{E}_{\text{in}}[n:]\| \approx 0$. The prediction error is given by $\epsilon[n] = \Omega_m[n] - \widehat{\Omega}_m[n]$ which one

is minimising. The vector $\mathbf{w}[n]$ denotes the adaptive filter weights at sample n, with dimension n_w and the updates would look as follows:

$$\mathbf{w}[n+1,:] = \mathbf{w}[n,:] + \Delta\mathbf{w}[n]. \tag{121}$$

A.10 ADDITIONAL AMPLITUDE OPTIMISED PULSES

Fig. 61 shows an additional amplitude optimisation for the chirped pulsed determined numerically in Ch. 4.

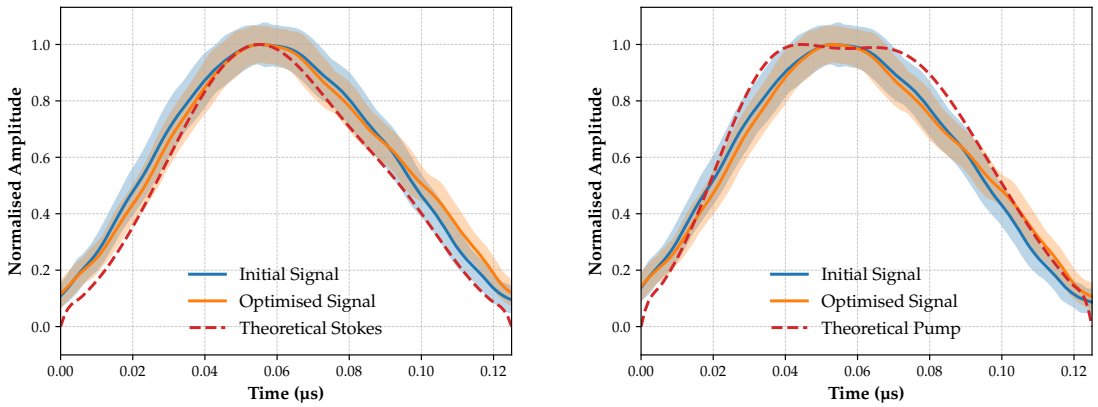

(a) Optimised Stokes Pulse (b) Optimised Pump Pulse

Figure 61: Comparison of optimised and unoptimised driving on fixed frequency pulses with respect to the desired theoretical shapes that were determined with RL in Fig. 29. It is clear that given the short duration and frequency chirps, it is difficult to get the exact desired pulse shape, especially the features near the ends, but the optimised fitting shows some improvements over the initial driving.

A.11 FREQUENCY ANALYSIS OF PULSES

Fig. 62 shows the spectral analysis of the optimised Stokes pulse with a reference beam to extract a beat note frequency upon applying some digital filtering.

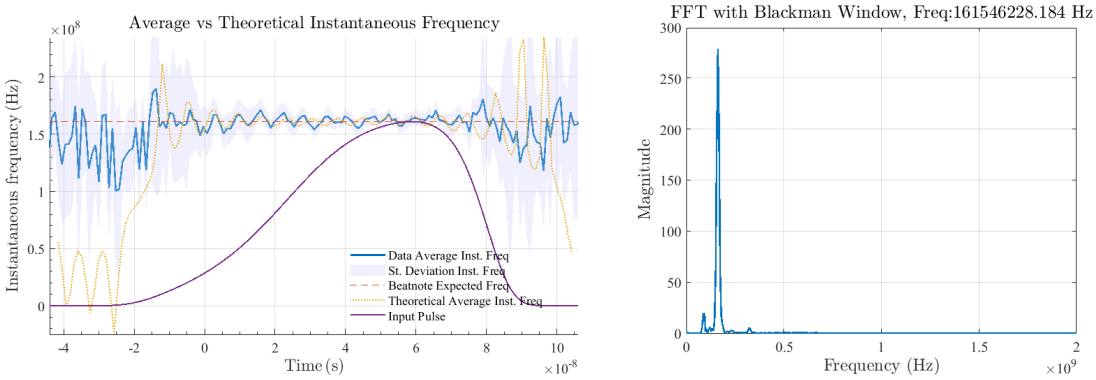

Figure 62: Left: Instantaneous frequency extracted from the homodyne detection scheme, averaged over 100 measurements. Outside the beat note duration and at low pulse amplitudes, the signal is dominated by noise and contains no useful spectral information. For reference, the Pump pulse envelope is overlaid to indicate the time window in which the beat note is active. Right: Fast Fourier Transform (FFT) of the beat note signal, filtered to retain frequency components between 20 and 180MHz. The dominant frequency peak is automatically detected and highlighted.

A.12 DATA ANALYSIS FOR EXPERIMENTAL RESULTS

Fig. 63 shows a representative trace from the fluorescence-based readout. The initial trace up to $\approx$ 0.6 ms corresponds to the fluorescence from the MOT, which can vary from shot to shot due to fluctuations in atom number, polarisation or laser power. Immediately before the readout pulse, the signal reflects the optical and electronic background noise of the photodiode. Analysis of multiple traces confirms that this noise is random and unbiased.

To reduce high-frequency noise, the entire oscilloscope trace is smoothed using a rolling average filter with a window size of 64 samples. The mean signal level in the 50μs window immediately preceding the readout pulse is taken as the background reference. The MOT fluorescence level is defined as the average signal over the entire visible MOT window.

The fluorescence signal is then integrated over a 500μs window during which the readout beam is on. This integral is background-subtracted and normalised to the MOT fluorescence level by dividing by the average MOT fluorescence level. This yields the fluorescence signal as a percentage of the maximum attainable level, enabling a shot-to-shot comparison that is robust against MOT fluctuations.

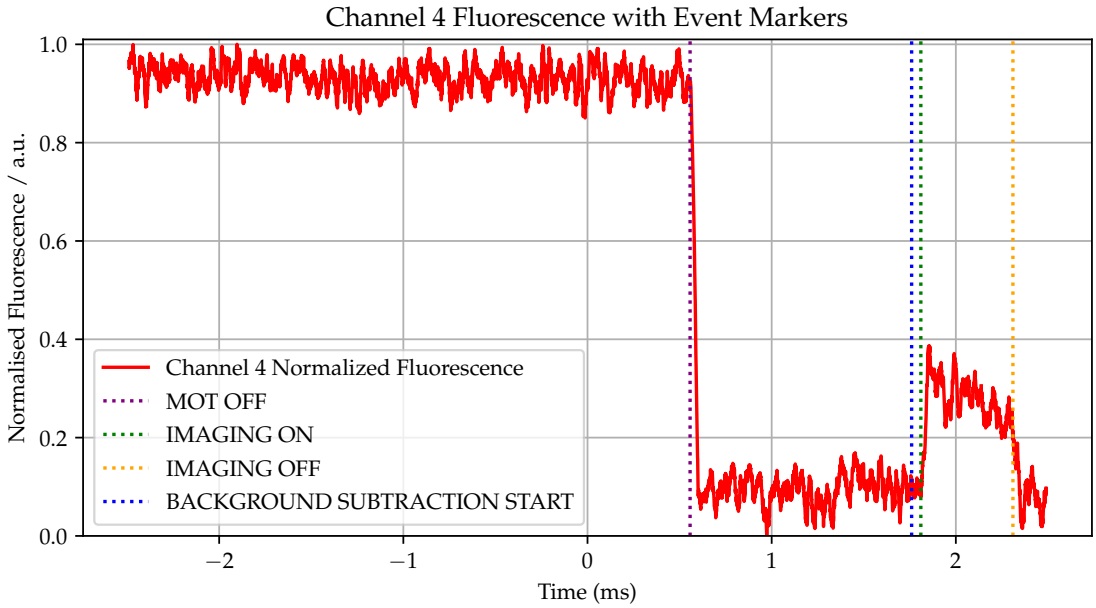

Figure 63: Normalised fluorescence for example state readout plotted against time. Event markers when the MOT is turned off and the imaging starts, as well as when one averages for background subtraction.

A.13 RL BACKGROUND

A.13.1 *Proximal Policy Optimisation (PPO)*

The PPO clipped objective is defined as [203]:

$$L^{\text{CLIP}}(\theta) = \mathbb{E}_t \left[\min \left(r_t(\theta)\hat{A}_t,\ \text{clip}(r_t(\theta), 1 \pm \epsilon)\hat{A}_t \right) \right], \tag{122}$$

with:

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}, \quad \hat{A}_t = \text{advantage estimate}, \quad \epsilon = \text{clipping parameter}.$$

In the continuous bandit setting, the advantage is defined as the difference between the observed reward r and a learned baseline:

$$\hat{A}_t = r(a_t) - V_\phi(s_t), \tag{123}$$

where $r(a_t)$ is the reward obtained after applying action a_t, and $V_\phi(s_t)$ is the baseline value estimate for the current state s_t. This measures how much better or worse the action performed relative to expectation. Clipping prevents large policy updates, stabilising training. An entropy bonus is added to encourage exploration, giving the overall objective:

$$L(\theta) = \mathbb{E}_t \left[L^{\text{CLIP}}(\theta) + c_1 \hat{A}_t + c_2 E[\pi_\theta](s_t) \right], \tag{124}$$

where c_1, c_2 are coefficients and $E[\pi_\theta](s_t)$ is the policy entropy at state s_t.

A.13.2 *Extended Time Horizon in Multi-step RL*

For multi-step RL, an extended time horizon is considered. In contrast to the bandit setting, each episode is divided into four sections, each of $1/4$ the length of the total number of action samples comprising the control signal. The agent does not observe any information about the noise at time step $t = 0$, with the observation $O_t = 0$. However, at time steps $t = 0.25$, 0.5, and 0.75, the agent receives the value of mean noise μ_Ω sampled at the beginning of the episode. Formally, the observation function $\mathcal{O}_t$ is defined as:

$$\mathcal{O}_t = \begin{cases} 0 & \text{if } t = 0 \\ \mu_\Omega & \text{if } t = 0.25k \text{ for } k = 1, 2, 3 \end{cases} \tag{125}$$

The agent's policy π then uses this observation to decide the action at each time step, where $\tilde{S}_t$ is the union of the state in the bandit setting s_t and the observation $\mathcal{O}_t$ which defines an action a_t through a conditional probability distribution $\mathcal{P}$:

$$\pi(\tilde{s}_t, a_t) = \mathcal{P}[a_t | \tilde{s}_t, \mathcal{O}_t] \tag{126}$$

In general extended time horizon RL, the agent must consider the long-term consequences of its actions. This is formalised through the discount factor γ, which ensures that future rewards are appropriately weighted. Given that there are a fixed number of four steps the discount factor is set to zero.

A.14 RL IMPLEMENTATION DETAILS

A.14.1 *Training Implementation Details*

The Qiskit-Dynamics Solver interface [182] is used [1] for constructing both Hamiltonians and collapse operators, enabling the simulation of open quantum systems through the dissipative master equation. The Diffrax ODE solver [123] is used for quantum system simulation, which utilise adaptive step-sizing techniques to efficiently integrate the first-order linear differential equations, PureJAX-RL is used for implementing PPO algorithms [148] and CleanRL [108] for TD3 and DDPG.

A.14.2 *Benchmarking of Simulation Speed*

Fig. 64 highlights the advantages of GPU parallelisation for quantum simulations. One can observe up to a two-order-of-magnitude improvement in speed per environment step, showcasing the performance benefits of running parallelised quantum simulations on GPUs.

1 Similar functionality is available in Dynamiqs with comparable performance characteristics [95].

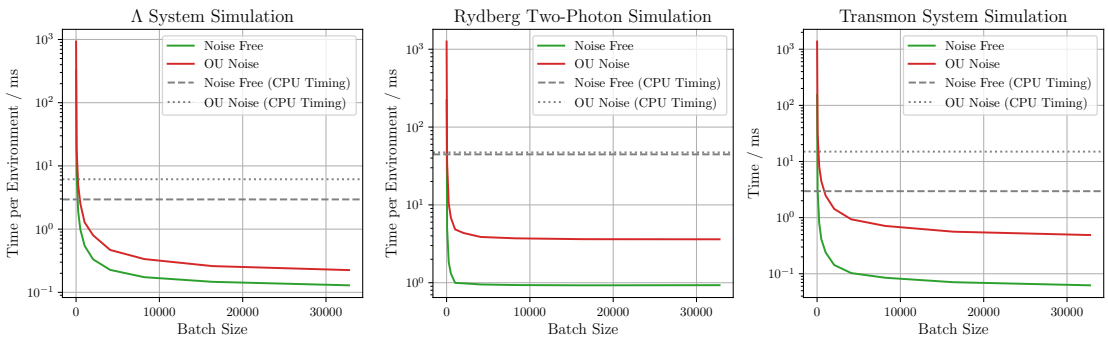

Figure 64: The time per environment step (log scale) for Qiskit Dynamics simulation compared across multiple environments (Λ system, two photon Rydberg gate and Transmon) under noise-free and Orstein Uhlenbeck noise conditions. The left panel shows the Λ system simulation timings, while the right panel illustrates the Rydberg two-photon simulation timings on a V-100 Nvidia GPU where one can parallelise the simulation of several environments with different random actions. The solid lines represent the simulation times obtained with a GPU, while the dashed and dotted horizontal lines indicate the corresponding CPU timings (Apple Silicon M1). In the best case one can get up to about two orders of magnitude improvement in simulation time per environment in a larger batch by moving to a GPU.

A.14.3 *Signal Processing & Analysis*

The actions a_i are smoothed with a Gaussian convolution $(a * G)(t) = \int_{-\infty}^{\infty} \mathcal{A}(\tau) G(t - \tau)\, \mathrm{d}\tau$, where the Gaussian function $G(t)$ is defined as:

$$G(t) = N(t_\sigma) \exp\left(-\frac{t^2}{2 \cdot t_\sigma^2}\right), \tag{127}$$

where t_σ defines the standard deviation and its value corresponds to the strength of the convolution filter. An ablation over this is provided in Fig. 48. This ensures that the generated time-dependent control signals are smooth and give rise to dynamics which can be solved in fixed number of time-steps. Pulse amplitude ends are always fixed at zero to ensure experimental viability. Additionally, cubic spline interpolation between action samples is used which is efficient for use with adaptive step size solvers.

The pulse smoothness is defined in terms of different pulse smoothness functions. The first smoothing function is constructed by calculating the second derivative of $\mathcal{A}(t)$:

$$S_{der}(a(t)) = \int_0^{t_f} \left(\frac{d^2\mathcal{A}}{dt^2} \right)^2 dt. \tag{128}$$

An alternative smoothing function is defined in terms of the difference in output to that generated by a low pass Butterworth filter. This requires an expression of the filtered action which is the convolution of $\mathcal{A}(t)$ with the impulse response $h(t)$ of the Butterworth filter:

$$a_{\text{filter}}(t) = (h * A)(t) = \int_0^{t_f} h(t-\tau)A(\tau)\,\mathrm{d}\tau \tag{129}$$

Calculating the difference with respect to the unfiltered signal, one gets an expression for the low-pass smoothness with respect to a cut-off frequency ω_{max} and the filter order n_{order}:

$$S_{lp}(a(t), n_{order}, \omega_{max}) = \int_0^{t_f} |a_{\text{filter}}(t) - a(t)|\,\mathrm{d}t. \tag{130}$$

It shall be noted that since all signals are discretised, the integrals decompose into discrete sums. The reference smoothness for an action is given by $S(B(t))$, where $B(t)$ is the Blackman window comprised of n samples where n also defines the number of signal samples corresponding to Ω_i or Δ_i. This minimises spectral leakage. Penalising pulse smoothness is required because even after applying a convolution filter, the agent does not output signals which exhibit low enough smoothness (see Sec. 7.3.2 for more discussion).

Choosing the right smoothness penalty in the construction of the reward function is important as it can determine the learning speed and the extent to which realistic and interpretable controls are generated. A low-pass filter approach with the right cut-off frequency generally works well and provides the fastest learning of *smooth* signals as

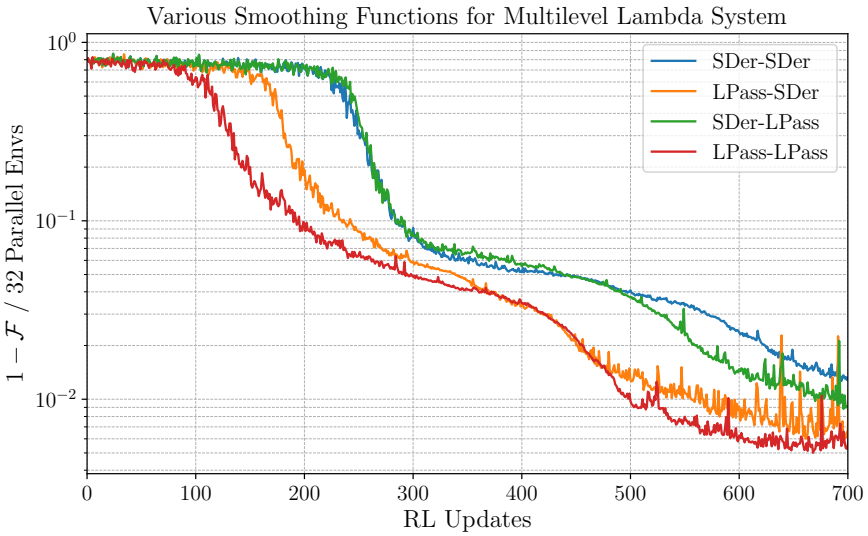

Figure 65: Comparison of different smoothing functions for mean infidelity (1- $\mathcal{F}$) across 32 parallel environments with different seeds plotted against the number of RL Updates for multi-level Lambda system. The legend corresponds to the type of smoothness penalty used where the ordering of the labels describes the amplitude and detuning smoothness functions respectively. One can observe that particularly for the amplitudes Ω_i, using a low pass filter (LPass cf. (130)) instead of a second derivative penalty (SDer cf. (128)) allows for significantly sped up learning and also higher mean fidelity. For this ablation, the smoothness penalties $w_\Omega = w_\Delta$ are fixed to 0.001.

shown in Fig. 65. Other simpler smoothness functions such as the L_1 or L_2 norm are not considered because they were found to be less effective.

Picking the right hyper-parameters for the Gaussian convolution filter standard deviation t_σ defined in 127, as well as the right smoothing penalties w_Δ and w_Ω (cf. (89)) is crucial to ensure the optimal trade-off between smooth signal discovery to facilitate parallel optimisation, improved interpretability and discovery of high fidelity solutions. Overly strong signal smoothing or smoothing penalties result in the optimiser focussing largely on signal smoothness over fidelity of the quantum control task which is the primary objective. This is shown clearly in Fig. 48, where the Λ system benefits from higher strict smoothing in form of a larger Gaussian kernel and higher weak smoothing in form of a larger pulse smoothness penalty, compared to the two photon Rydberg gate.

A final objective which competes with the fidelity, are the pulse areas $A(\Omega_i)$ and implicitly the pulse duration. Ω_{max} is limited physically by laser, RF or microwave power.

Additionally, minimising pulse area is important for reducing the pulse energy and in turn the amount of heat introduced into the system, particularly for those quantum systems operating at cryogenic temperatures. Generally faster pulse sequences increase the clock cycles of a particular quantum operation which is desirable, but secondary to their fidelity, so implementing optimal control for some maximal amplitude Ω_{max} but with a minimal pulse area is considered in the example of a Λ system. The baseline pulse area (cf. (89)), which is particularly relevant for the results shown in Fig. 54 and Fig. 53 is computed by comparing the generated pulse area $A = \int_{t=0}^{t=t_f} \Omega(t)dt$ to the area of a Blackman window A_B defined over the same timescale.

A.14.4 *Noise Model*

The model used for injecting noise into various environments is an Ornstein-Uhlenbeck noise model defined with standard deviation σ and mean μ which defines time-dependent noise in time t:

$$\nu_t = \nu_{t-1}(1 - \alpha^2) + \sqrt{2}\sigma X(t)\alpha + \sigma^2\mu, \tag{131}$$

where α defines the characteristic time scale of the noise fluctuations and $X(t)$ is random Gaussian noise at time t with a standard deviation of 1 and a mean of 0.

A.15 LAMBDA SYSTEM

We include an effective Hamiltonian for the Λ system here used in the main text. The Hamiltonian in the rotating wave approximation in the basis $\{|g_1\rangle, |g_2\rangle, |e_1\rangle, |e_2\rangle\}$ reads:

$$\frac{H_\Lambda^{\text{eff}}}{\hbar} = \begin{pmatrix} 0 & 0 & \frac{\Omega_P}{2} & \frac{\Omega_P}{2} \\ 0 & \Delta_P - \Delta_\delta & \frac{\Omega_S}{2} & -\frac{\Omega_S}{2} \\ \frac{\Omega_P}{2} & \frac{\Omega_S}{2} & \Delta_P - i\frac{\Gamma}{2} & 0 \\ \frac{\Omega_P}{2} & -\frac{\Omega_S}{2} & 0 & \Delta_P + \Delta_X - i\frac{\Gamma}{2} \end{pmatrix}. \tag{132}$$

Here Ω_P and Ω_S denote the (real-valued) Rabi frequencies of the pump and Stokes laser fields, respectively. The detuning Δ_P is the one-photon detuning of the pump laser from the $|g_1\rangle \leftrightarrow |e_1\rangle$ transition, while Δ_X denotes the energy splitting between the two excited states $|e_1\rangle$ and $|e_2\rangle$. The quantity $\Delta_\delta \equiv \Delta_P - \Delta_S$ is the two-photon detuning between the ground states $|g_1\rangle$ and $|g_2\rangle$, with Δ_S the one-photon detuning of the Stokes laser. Spontaneous emission from the excited states is incorporated phenomenologically via the non-Hermitian terms $-i\Gamma/2$ on the excited-state diagonal elements, where Γ is the total excited-state decay rate. As is common in multilevel atomic systems we pick a an excited state $|e_2\rangle$ which has a Clebsch-Gordan coefficient of different polarity on the Stokes transition to $|g_2\rangle$ such that couplings of Stokes laser to the different excited states interfere with one another.

A.16 RYDBERG ATOMS

Following the Hamiltonian definition $H_{r_2} = H_{0,2} + H_{\text{int},2}$ for a two-photon Rydberg gate given in [218] (where h.c. denotes the hermitian conjugate):

$$\begin{aligned}\frac{H_{0,2}}{\hbar} = &\frac{\Omega_P(t)}{2}\left|10\right\rangle\left\langle e0\right| + \frac{\Omega_S(t)}{2}\left|e0\right\rangle\left\langle r0\right| + \text{ h.c.} \\ &+ \Delta_P(t)\left|e0\right\rangle\left\langle e0\right| + \Delta_S(t)\left|r0\right\rangle\left\langle r0\right|\end{aligned} \tag{133}$$

with time-dependent Rabi frequencies $\Omega_P(t), \Omega_S(t)$, and values for the one photon detuning Δ_P and two-photon detuning Δ_S. The Hamiltonian terms for $|01\rangle$ follow analogously from symmetry considerations by swapping all qubits in their respective state in $H_{0,2}$. The interaction Hamiltonian $H_{int,2}$ for the state $|1,1\rangle$ consists of the atom light interaction as well as the dipole-dipole interaction. A basis transformation simplifies the Hamiltonian, the new basis states read $|\tilde{e}\rangle = (|e1\rangle + |1e\rangle)/\sqrt{2}, |\tilde{r}\rangle = (|r1\rangle + |1r\rangle)/\sqrt{2}$ and $\left|\tilde{R}\right\rangle = (|re\rangle + |er\rangle)/\sqrt{2}$, after the rotating wave approximation, and effectively neglecting $|ee\rangle$, as in this regime of interest $\Delta_P >> \Delta_S$, $H_{int,2}/\hbar$ can be expressed as:

$$\begin{aligned}\frac{H_{int,2}}{\hbar} = &\frac{\Omega_P(t)}{\sqrt{2}}\left|11\right\rangle\left\langle\tilde{e}\right| + \frac{\Omega_S(t)}{2}\left|\tilde{e}\right\rangle\left\langle\tilde{r}\right| + \\ &\frac{\Omega_P(t)}{2}\left|\tilde{r}\right\rangle\left\langle\tilde{R}\right| + \frac{\Omega_S(t)}{\sqrt{2}}\left|\tilde{R}\right\rangle\left\langle rr\right| + \text{ h.c. } + \\ &\Delta_P(t)\left|\tilde{e}\right\rangle\left\langle\tilde{e}\right| + (2\Delta_P(t) + B)\left|rr\right\rangle\left\langle rr\right| + \\ &(\Delta_P(t) + \Delta_S(t))\left|\tilde{R}\right\rangle\left\langle\tilde{R}\right| + \Delta_S(t)\left|\tilde{r}\right\rangle\left\langle\tilde{r}\right|\end{aligned} \tag{134}$$

Parameters $\Omega_{S/P}, \Delta_{S/P}, B$ are the Rabi frequencies, detunings and the blockade strength. The Lindbladian decay terms read; $\sum_i \Gamma_r(|g_\Gamma\rangle)(\langle r,i| + \langle i,r|) + \Gamma_r(|g_\Gamma\rangle\langle r,r|) + \sum_i \Gamma_e(|g_\Gamma\rangle(\langle e,i| + \langle i,e|) + \Gamma_e(|g_\Gamma\rangle\langle e,e|)$, where Γ_r describes the decay rate of the Rydberg level and Γ_e the decay of the excited level $|e\rangle$ where for typical atoms $\Gamma_e >> \Gamma_r$. Realistic values for ^{87}**Rb** [218] are picked.

A comparison of different RL algorithms for the two photon Rydberg gate is shown in Ref. [75].

The fidelity $\mathcal{F}_R$ is defined by the Bell state fidelity as is common in optimisation protocols of the Rydberg gate [81, 192]:

$$\mathcal{F}_R = \frac{1}{16}|1 + \sum_{\psi_q^0 \in 10,01,11} e^{-i\theta_q} \left\langle q|\psi_q^0 \right\rangle|^2, \tag{135}$$

where q is the simulated computational basis state. Without loss of generality, the optimisation here focusses on the case where $\theta_q = 0$, except $\theta_{1,1} = \pi$, this is particularly useful insofar as it does not require additional single qubit rotations (in comparison to a general CZ gate) and does not introduce any further complexity associated with additional rotations. In practice, the computational complexity of simulating the Rydberg gate can be reduced by focusing on initialising the system in the coherent superposition state $|+,+\rangle$. The optimisation then targets both the populations and the relative phases of the resulting superposed states, ensuring they match the desired angles θ_i and maintain equal amplitudes. This approach effectively reproduces the transformation expected from an ideal logical C-Z gate: $\frac{1}{\sqrt{4}}(|00\rangle + |10\rangle + |10\rangle - |11\rangle)$.

A.17 TRANSMON RESET

Unconditional Transmon qubit reset with fixed-frequency devices involves using the coupling of a Transmon to a decaying resonator. One particular protocol is based on a cavity-assisted Raman transition utilising the drive-induced coupling between $|f,0\rangle$ and $|g,1\rangle$, where $|s,n\rangle$ denotes the tensor product of a Transmon in $|s\rangle$ (where $s = 0,1,2,.. \leftrightarrow g,e,f,..$) and a readout resonator mode in the Fock state $|n\rangle$. By driving the Transmon simultaneously at the $|e,0\rangle \leftrightarrow |f,0\rangle$ transition and the $|f,0\rangle \leftrightarrow |g,1\rangle$ transition, one can realise a Λ system in the Jaynes-Cummings Model (cf. Eq. 25) which

can be used to reset the Transmon through the emission of a photon via the resonator. The Transmon reset Hamiltonian reads:

$$\begin{aligned}\frac{H}{\hbar} =&\chi a^\dagger a q^\dagger q + \frac{g\alpha}{\sqrt{2}\delta(\delta+\alpha)}\Omega(t)(q^\dagger q^\dagger a + \text{h.c. }) + \\ &(\Delta(t) + \delta_S(t))q^\dagger q\end{aligned} \tag{136}$$

where $a(a^\dagger)$ are the resonator lowering (raising) operators, $q(q^\dagger)$ are the Transmon lowering (raising) operators, χ the Transmon-resonator dispersive shift, α the Transmon anharmonicity, g the Transmon-resonator coupling rate, δ the difference in the Transmon and resonator resonant frequencies, $\Omega(t)$ the Transmon drive amplitude, $\Delta(t)$ the Transmon drive detuning, and $\delta_S(t)$ the drive-induced stark shift. As determined in [249], this Stark Shift is to first order quadratic in the drive amplitude, $\delta_S(t) \propto \Omega^2(t)$. For the Transmon mode we consider three levels $|g\rangle, |e\rangle, |f\rangle$ coupled with a two level resonator. We neglect self-Kerr terms in the resonator mode as we target single photon populations where such non-linearities are not significant.

The dissipative master equation terms for the Transmon reset simulation read: $\kappa\mathcal{D}_\kappa[\rho] + \Gamma\mathcal{D}_\Gamma[\rho]$, with κ describing the resonator decay rate, and Γ the Transmon decay rate.

We construct the Transmon reset environment to match the physical parameters in [71], with maximum drive amplitudes of 330 MHz, however with an additional small detuning control of up to $\pm$100 kHz for frequency corrections. To represent bandwidth constraints, we add a Gaussian convolution of duration 14ns to the amplitude and detuning defined in (127). The maximal solver step limit is set to 900. Reset is optimised for a qubit lifetime of $T_1 = 500\mu$s, which represents the Transmon lifetimes currently attainable in state of the art experiments. Optimal waveforms and corresponding Transmon populations are shown in Fig. 57, where the RL pulse can achieve fidelities of 0.9997 even with realistic bandwidth constraints. Other parameter regimes are also explored in the Appendix of [75].

BIBLIOGRAPHY

[1] Scott Aaronson. "A linear-optical proof that the permanent is # P -hard." In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 467.2136 (July 2011), pp. 3393–3405. ISSN: 1471-2946. DOI: 10.1098/rspa.2011.0232.

[2] Scott Aaronson and Alex Arkhipov. "The Computational Complexity of Linear Optics." In: *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*. STOC '11. New York, NY, USA: ACM, 2011, pp. 333–342. ISBN: 978-1-4503-0691-1. DOI: 10.1145/1993636.1993682.

[3] Scott Aaronson and Daniel J. Brod. "BosonSampling with lost photons." In: *Phys. Rev. A* 93 (1 Jan. 2016), p. 012335. DOI: 10.1103/PhysRevA.93.012335.

[4] Rajeev Acharya et al. "Suppressing quantum errors by scaling a surface code logical qubit." In: *Nature* 614.7949 (Feb. 2023), pp. 676–681. ISSN: 1476-4687. DOI: 10.1038/s41586-022-05434-1.

[5] Dorit Aharonov and Michael Ben-Or. "Fault-tolerant quantum computation with constant error rate." In: *SIAM Journal on Computing* 38.4 (1999), pp. 1207–1282.

[6] Eitan Altman. *Constrained Markov decision processes*. Routledge, 2021.

[7] Juan-Rafael Alvarez et al. "How to administer an antidote to Schroedinger's cat." In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 55.5 (Mar. 2022), p. 054001. DOI: 10.1088/1361-6455/ac5674.

[8] Kyungwon An. "Controlled rapid generation of single photons in an atom-cavity quantum composite." In: *2011 13th International Conference on Transparent Optical Networks*. 2011, pp. 1–2. DOI: 10.1109/ICTON.2011.5970816.

[9] Zheng An, Hai-Jing Song, Qi-Kai He, and D. L. Zhou. "Quantum optimal control of multilevel dissipative quantum systems with reinforcement learning." In: *Phys. Rev. A* 103 (1 Jan. 2021), p. 012404. DOI: 10.1103/PhysRevA.103.012404.

[10] Martin Hayhurst Appel et al. "Entangling a hole spin with a time-bin photon: a waveguide approach for quantum dot sources of multiphoton entanglement." In: *Physical review letters* 128.23 (2022), p. 233602.

[11] Alex Arkhipov and Greg Kuperberg. "The bosonic birthday paradox." In: *Geometry and Topology Monographs* (June 2011). DOI: 10.2140/gtm.2012.18.1.

[12] Juan Miguel Arrazola and Thomas R. Bromley. "Using Gaussian Boson Sampling to Find Dense Subgraphs." en. In: *Physical Review Letters* 121.3 (July 2018), p. 030503. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.030503. (Visited on 06/01/2024).

[13] Frank Arute et al. "Quantum supremacy using a programmable superconducting processor." In: *Nature* 574.7779 (2019), pp. 505–510. DOI: 10.1038/s41586-019-1666-5.

[14] Alain Aspect, Philippe Grangier, and Gerard Roger. "Experimental Tests of Realistic Local Theories via Bell's Theorem." In: *Phys. Rev. Lett.* 47 (7 Aug. 1981), pp. 460–463. DOI: 10.1103/PhysRevLett.47.460.

[15] Gopalakrishnan Balasubramanian et al. "Ultralong spin coherence time in isotopically engineered diamond." In: *Nature Materials* 8.5 (2009), pp. 383–387. DOI: 10.1038/nmat2420.

[16] Leonardo Banchi, Mark Fingerhuth, Tomas Babej, Christopher Ing, and Juan Miguel Arrazola. "Molecular docking with Gaussian Boson Sampling." In: *Science Advances* 6.23 (2020), eaax1950. DOI: 10.1126/sciadv.aax1950.

[17] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys. "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays." In: *Science* 354.6315 (2016), pp. 1021–1023. DOI: 10.1126/science.aah3778.

[18] Thomas Barrett. "Polarised Single Photons from a Cavity-Enhanced Atom-Light Interface in Photonic Quantum Networks." PhD thesis. University of Oxford, 2018.

[19] Thomas D Barrett, Thomas H Doherty, and Axel Kuhn. "Pushing Purcell enhancement beyond its limits." In: *New Journal of Physics* 22.6 (June 2020), p. 063013. ISSN: 1367-2630. DOI: 10.1088/1367-2630/ab8ab0.

[20] Thomas D Barrett, Dustin Stuart, Oliver Barter, and Axel Kuhn. "Nonlinear Zeeman Effects in the Cavity-Enhanced Emission of Polarised Photons." In: *New Journal of Physics* 20.7 (July 2018), p. 073030. ISSN: 1367-2630. DOI: 10.1088/1367-2630/aad14e.

[21] H G Barros et al. "Deterministic single-photon source from a single ion." In: *New Journal of Physics* 11.10 (Oct. 2009), p. 103004. DOI: 10.1088/1367-2630/11/10/103004.

[22] Richard H Bartels, John C Beatty, and Brian A Barsky. *An introduction to splines for use in computer graphics and geometric modeling*. Morgan Kaufmann, 1995.

[23] Oliver Barter. "Deterministic Quantum Feedback Control in Probabilistic Atom-Photon Entanglement." PhD thesis. University of Oxford, 2016.

[24] Sara Bartolucci et al. "Fusion-based quantum computation." In: *Nature Communications* 14.1 (Feb. 2023). ISSN: 2041-1723. DOI: 10.1038/s41467-023-36493-1.

[25] David Bauch, Nikolas Köcher, Nils Heinisch, and Stefan Schumacher. "Time-bin entanglement in the deterministic generation of linear photonic cluster states." In: *APL Quantum* 1.3 (July 2024), p. 036110. ISSN: 2835-0103. DOI: 10.1063/5.0214197.

[26] Yuval Baum et al. "Experimental Deep Reinforcement Learning for Error-Robust Gate-Set Design on a Superconducting Quantum Computer." In: *PRX Quantum* 2.4 (Dec. 2021). arXiv: 2105.01079 Publisher: American Physical Society. ISSN: 26913399. DOI: 10.1103/PRXQuantum.2.040324.

[27] Charles H Bennett and Gilles Brassard. "Quantum cryptography: Public key distribution and coin tossing." In: *Theor. Comput. Sci.* 560 (1984), pp. 7–11.

[28] Klaas Bergmann, Hanns-Christoph Nägerl, Cristian Panda, et al. "Roadmap on STIRAP applications." In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 52.20 (Sept. 2019), p. 202001. DOI: 10.1088/1361-6455/ab3995.

[29] I.I. Beterov et al. "Implementation of one-qubit quantum gates with individual addressing of two rubidium atoms in two optical dipole traps." In: *Quantum Electronics* 51.6 (June 2021), p. 464. DOI: 10.1070/QEL17583.

[30] Shalabh Bhatnagar and K Lakshmanan. "An online actor–critic algorithm with function approximation for constrained markov decision processes." In: *Journal of Optimization Theory and Applications* 153 (2012), pp. 688–708.

[31] Garrett Birkhoff and Carl R De Boor. "Piecewise polynomial interpolation and approximation." In: *Approximation of functions* (1965), pp. 164–190.

[32] D. Bluvstein et al. "A quantum processor based on coherent transport of entangled atom arrays." In: *Nature* 604.7906 (Apr. 2022), pp. 451–456. DOI: 10.1038/s41586-022-04592-6.

[33] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays." In: *Nature* 626.7997 (Dec. 2023), pp. 58–65. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06927-3.

[34] J. Bochmann et al. "Fast Excitation and Photon Emission of a Single-Atom-Cavity System." In: *Phys. Rev. Lett.* 101 (22 Nov. 2008), p. 223601. DOI: 10.1103/PhysRevLett.101.223601.

[35] Yuval Boger. *Field-Programmable Qubit Arrays: The Quantum Analog of FPGAs — eetimes.eu*. https://www.eetimes.eu/field-programmable-qubit-arrays-the-quantum-analog-of-fpgas/.

[36] C. de Boor. "Good Approximation by Splines with Variable Knots." In: *Spline Functions and Approximation Theory: Proceedings of the Symposium held at the University of Alberta, Edmonton May 29 to June 1, 1972*. Ed. by A. Meir and A. Sharma. Basel: Birkhäuser Basel, 1973, pp. 57–72. ISBN: 978-3-0348-5979-0. DOI: 10.1007/978-3-0348-5979-0_3.

[37] R. Bowler. "Coherent Ion Transport in a Multi-electrode Trap Array." In: *CU Scholar University Library* (2015), pp. 93–94. URL: https://www.nist.gov/system/files/documents/2017/05/09/Bowler15_thesis-2.pdf.

[38] R. Bowler, U. Warring, J. W. Britton, B. C. Sawyer, and J. Amini. "Arbitrary waveform generator for quantum information processing with trapped ions." In: *Rev. Sci. Instruments* 84 (3 Feb. 2013), p. 033108. DOI: https://doi.org/10.1063/1.4795552.

[39] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018. URL: http://github.com/google/jax.

[40] Samuel L Braunstein and Peter van Loock. "Quantum information with continuous variables." In: *Reviews of Modern Physics* 77.2 (2005), p. 513. DOI: 10.1103/RevModPhys.77.513.

[41] H. P. Breuer and F. Petruccione. *The theory of open quantum systems*. Great Clarendon Street: Oxford University Press, 2002.

[42] H. J. Briegel, D. E. Browne, W. Duer, R. Raussendorf, and M. Van den Nest. "Measurement-based quantum computation." In: *Nature Physics* 5.1 (Jan. 2009), pp. 19–26. DOI: 10.1038/nphys1157.

[43] Hans J Briegel, Wolfgang Dür, J Ignacio Cirac, and Peter Zoller. "Quantum repeaters: The role of imperfect local operations in quantum communication." In: *Physical Review Letters* 81.26 (1998), p. 5932.

[44] Hans J Briegel and Robert Raussendorf. "Persistent entanglement in arrays of interacting particles." In: *Physical Review Letters* 86.5 (2001), p. 910. DOI: 10.1103/PhysRevLett.86.910.

[45] Constantin Brif, Raj Chakrabarti, and Herschel Rabitz. "Control of quantum phenomena: past, present and future." In: *New Journal of Physics* 12.7 (July 2010), p. 075008. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/7/075008.

[46] Daniel J Brod et al. "Photonic implementation of boson sampling: a review." In: *Advanced Photonics* 1.3 (2019), pp. 034001–034001.

[47] Jonathon Brown, Mauro Paternostro, and Alessandro Ferraro. "Optimal quantum control via genetic algorithms for quantum state engineering in driven-resonator mediated networks." In: *Quantum Science and Technology* 8.2 (Apr. 2023), p. 025004. ISSN: 2058-9565. DOI: 10.1088/2058-9565/acb2f2. (Visited on 03/24/2024).

[48] Jonathon Brown et al. "Reinforcement learning-enhanced protocols for coherent population-transfer in three-level quantum systems." In: *New Journal of Physics* 23.9 (Sept. 2021), p. 093035. ISSN: 1367-2630. DOI: 10.1088/1367-2630/ac2393.

[49] Kenneth R Brown, Aram W Harrow, and Isaac L Chuang. "Arbitrarily accurate composite pulse sequences." In: *Physical Review A—Atomic, Molecular, and Optical Physics* 70.5 (2004), p. 052318.

[50] Marin Bukov et al. "Reinforcement Learning in Different Phases of Quantum Control." In: *Physical Review X* 8.3 (Sept. 2018). arXiv: 1705.00565 Publisher: American Physical Society. ISSN: 21603308. DOI: 10.1103/PhysRevX.8.031086.

[51] H. G. Burchard. "Splines (with optimal knots) are better." In: *Applicable Analysis* 3.4 (1974), pp. 309–319. DOI: 10.1080/00036817408839073.

[52] H. J. Carmichael. *An Open Systems Approach to Quantum Optics*. Vol. 18. Lecture Notes in Physics. Springer, 1993. DOI: 10.1007/3-540-47620-7.

[53] Esteban Castro-Ruiz and Eduardo Nahmad-Achar. *On the Relativistic Invariance of Entanglement*. 2011. arXiv: 1109.5747 [quant-ph]. URL: https://arxiv.org/abs/1109.5747.

[54] Aniket Chatterjee, Jonathan Schwinger, and Yvonne Y. Gao. "Enhanced qubit readout via reinforcement learning." In: *Physical Review Applied* 23.5 (May 2025). ISSN: 2331-7019. DOI: 10.1103/physrevapplied.23.054057.

[55] Dariusz Chruściński and Gniewomir Sarbicki. "Entanglement witnesses: construction, analysis and classification." In: *Journal of Physics A: Mathematical and Theoretical* 47.48 (Nov. 2014), p. 483001. ISSN: 1751-8121. DOI: 10.1088/1751-8113/47/48/483001.

[56] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. "Quantum State Transfer and Entanglement Distribution among Distant Nodes in a Quantum Network." In: *Phys. Rev. Lett.* 78.16 (Apr. 1997), pp. 3221–3224. DOI: 10.1103/PhysRevLett.78.3221.

[57] William R. Clements, Peter C. Humphreys, Benjamin J. Metcalf, W. Steven Kolthammer, and Ian A. Walmsley. "Optimal design for universal multiport interferometers." In: *Optica* 3.12 (Dec. 2016), pp. 1460–1465. DOI: 10.1364/OPTICA.3.001460.

[58] Peter Clifford and Raphaël Clifford. *The Classical Complexity of Boson Sampling*. 2017. arXiv: 1706.01260 [cs.DS]. URL: https://arxiv.org/abs/1706.01260.

[59] Jean Dalibard, Yvan Castin, and Klaus Mølmer. "Wave-function approach to dissipative processes in quantum optics." In: *Physical Review Letters* 68.5 (1992), pp. 580–583. DOI: 10.1103/PhysRevLett.68.580.

[60] E. B. Davies. "Markovian master equations." In: *Communications in Mathematical Physics* 39.2 (1974), pp. 91–110. DOI: 10.1007/BF01608389.

[61] Christian L Degen, Friedemann Reinhard, and Paola Cappellaro. "Quantum sensing." In: *Reviews of Modern Physics* 89.3 (2017), p. 035002.

[62] Emma Deist et al. "Mid-Circuit Cavity Measurement in a Neutral Atom Array." In: *Phys. Rev. Lett.* 129 (20 Nov. 2022), p. 203602. DOI: 10.1103/PhysRevLett.129.203602.

[63] David Deutsch. "Quantum theory, the Church–Turing principle and the universal quantum computer." In: *Proceedings of the Royal Society of London A* 400.1818 (1985), pp. 97–117.

[64] J. P. G. van Dijk et al. "Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control." In: *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS* 67 (12 Dec. 2020), pp. 5380–5393. URL: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9189938.

[65] Jerome Dilley, Peter Nisbet-Jones, Bruce W. Shore, and Axel Kuhn. "Single-Photon Absorption in Coupled Atom-Cavity Systems." In: *Phys. Rev. A* 85.2 (Feb. 2012), p. 023834. DOI: 10.1103/PhysRevA.85.023834.

[66] T Doherty. "Optical cavities for quantum information processing with trapped neutral atoms." PhD thesis. University of Oxford, 2022.

[67] Thomas H. Doherty, Axel Kuhn, and Ezra Kassa. "Multi-resonant open-access microcavity arrays for light matter interaction." In: *Opt. Express* 31.4 (Feb. 2023), pp. 6342–6355. DOI: 10.1364/OE.475921.

[68] S. Ebadi et al. "Quantum phases of matter on a 256-atom programmable quantum simulator." In: *Nature* 595.7866 (July 2021), pp. 227–232. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03582-4.

[69] Sophia E. Economou et al. "Scalable qubit architecture based on holes in quantum dot molecules." In: *Phys. Rev. B* 86 (8 Aug. 2012), p. 085319. DOI: 10.1103/PhysRevB.86.085319.

[70] D.J. Egger et al. "Pulsed Reset Protocol for Fixed-Frequency Superconducting Qubits." In: *Phys. Rev. Appl.* 10 (4 Oct. 2018), p. 044030. DOI: 10.1103/PhysRevApplied.10.044030.

[71] D.J. Egger et al. "Pulsed Reset Protocol for Fixed-Frequency Superconducting Qubits." In: *Phys. Rev. Appl.* 10 (4 Oct. 2018), p. 044030. DOI: 10.1103/PhysRevApplied.10.044030.

[72] A. Einstein, M. Born, and H. Born. *The Born-Einstein Letters: Correspondence Between Albert Einstein and Max and Hedwig Born from 1916-1955, with Commentaries by Max Born*. Macmillan, 1971. ISBN: 9780333112670.

[73] M. Endres et al. "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays." In: *Science* 354.6315 (2016), pp. 1024–1027. DOI: 10.1126/science.aah3752.

[74] Jan Ole Ernst, Juan-Rafael Alvarez, Thomas D Barrett, and Axel Kuhn. "Bursts of polarised single photons from atom-cavity sources." In: *Journal of Physics B: Atomic, Molecular and Optical Physics* (2023). URL: http://iopscience.iop.org/article/10.1088/1361-6455/acf9d2.

[75] Jan Ole Ernst, Aniket Chatterjee, Tim Franzmeyer, and Axel Kuhn. "Reinforcement Learning for Quantum Control with Physical Constraints." In: *Proceedings of the 42nd International Conference on Machine Learning (ICML)*. Vancouver, Canada: PMLR, 2025. URL: https://arxiv.org/abs/2501.14372.

[76] Jan Ole Ernst, Jan Snoeijs, Mitchell Peaks, and Jochen Wolf. "Memory-Optimized Cubic Splines for High-Fidelity Quantum Operations." In: *IEEE Transactions on Quantum Engineering* (2025). DOI: 10.1109/TQE.2025.3574463.

[77] Simon J. Evered et al. "High-fidelity parallel entangling gates on a neutral-atom quantum computer." In: *Nature* 622.7982 (Oct. 2023), pp. 268–272. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06481-y.

[78] C. J. Foot. *Atomic Physics*. Oxford Master Series in Physics 7. Atomic, Optical, and laser physics. Oxford ; New York: Oxford University Press, 2005. ISBN: 978-0-19-850695-9 978-0-19-850696-6.

[79] Austin G Fowler and et al. "Surface codes: Towards practical large-scale quantum computation." In: *Physical Review A* 86.3 (2012), p. 032324.

[80] Mark Fox. *Quantum Optics: An Introduction*. Oxford Master Series in Physics 15. Oxford ; New York: Oxford University Press, 2006. ISBN: 978-0-19-856672-4 0-19-856672-7 978-0-19-856673-1 0-19-856673-5.

[81] Charles Fromonteil, Dolev Bluvstein, and Hannes Pichler. "Protocols for Rydberg Entangling Gates Featuring Robustness against Quasistatic Errors." In: *PRX Quantum* 4.2 (June 2023). ISSN: 2691-3399. DOI: 10.1103/prxquantum.4.020335.

[82] Scott Fujimoto, Herke Hoof, and David Meger. "Addressing Function Approximation Error in Actor-Critic Methods." In: *Proceedings of the 35th International Conference on Machine Learning (ICML)*. PMLR, 2018, pp. 1587–1596. URL: https://proceedings.mlr.press/v80/fujimoto18a.html.

[83] J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm. "Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator." In: *Phys. Rev. A* 83 (1 Jan. 2011), p. 012308. DOI: 10.1103/PhysRevA.83.012308.

[84] U Gaubatz, P Rudecki, S Schiemann, and K Bergmann. "Population transfer between molecular vibrational levels by stimulated Raman scattering with partially overlapping laser fields. A new concept and experimental results." In: *The Journal of Chemical Physics* 92.9 (1990), pp. 5363–5376.

[85] Ronan Gautier, Élie Genois, and Alexandre Blais. "Optimal Control in Large Open Quantum Systems: The Case of Transmon Readout and Reset." In: *Phys. Rev. Lett.* 134 (7 Feb. 2025), p. 070802. DOI: 10.1103/PhysRevLett.134.070802.

[86] Luigi Giannelli et al. "A tutorial on optimal control and reinforcement learning methods for quantum technologies." In: *Physics Letters A* 434 (2022), p. 128054. ISSN: 0375-9601. DOI: https://doi.org/10.1016/j.physleta.2022.128054.

[87] Mercedes Gimeno-Segovia, Pete Shadbolt, Dan E. Browne, and Terry Rudolph. "From Three-Photon Greenberger-Horne-Zeilinger States to Ballistic Universal Quantum Computation." In: *Phys. Rev. Lett.* 115 (2 July 2015), p. 020502. DOI: 10.1103/PhysRevLett.115.020502.

[88] Steffen J. Glaser et al. "Training Schrödinger's cat: quantum optimal control: Strategic report on current status, visions and goals for research in Europe." In: *The European Physical Journal D* 69.12 (Dec. 2015), p. 279. ISSN: 1434-6060, 1434-6079. DOI: 10.1140/epjd/e2015-60464-1. (Visited on 04/30/2024).

[89] Michael H Goerz, K Birgitta Whaley, and Christiane P Koch. "Hybrid optimization schemes for quantum control." In: *EPJ Quantum Technology* 2.1 (Sept. 2015). ISSN: 2196-0763. DOI: 10.1140/epjqt/s40507-015-0034-0.

[90] Michael H. Goerz, Eli J. Halperin, Jon M. Aytac, Christiane P. Koch, and K. Birgitta Whaley. "Robustness of high-fidelity Rydberg gates with single-site addressability." In: *Phys. Rev. A* 90 (3 Sept. 2014), p. 032329. DOI: 10.1103/PhysRevA.90.032329.

[91] Michael H. Goerz et al. "Krotov: A Python implementation of Krotov's method for quantum optimal control." In: *SciPost Phys.* 7 (2019), p. 80. DOI: 10.21468/SciPostPhys.7.6.080.

[92] T. M. Graham et al. "Midcircuit Measurements on a Single-Species Neutral Alkali Atom Quantum Processor." In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041051. DOI: 10.1103/PhysRevX.13.041051.

[93] Daniel M Greenberger, Michael A Horne, Abner Shimony, and Anton Zeilinger. "Going beyond Bell's theorem." In: *American Journal of Physics* 58.12 (1990), pp. 1131–1143. DOI: 10.1119/1.16243.

[94] Lov K Grover. "A fast quantum mechanical algorithm for database search." In: *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (1996), pp. 212–219.

[95] Pierre Guilmin, Ronan Gautier, Adrien Bocquet, and Élie Genois. *dynamiqs: an open-source Python library for GPU-accelerated and differentiable simulation of quantum systems*. 2024. URL: https://github.com/dynamiqs/dynamiqs.

[96] Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Vol. 8. Springer, 1993.

[97] Lukas Hartung, Matthias Seubert, Stephan Welte, Emanuele Distante, and Gerhard Rempe. "A quantum-network register assembled with optical tweezers in an optical cavity." In: *Science* 385.6705 (2024), pp. 179–183. DOI: 10.1126/science.ado6471.

[98] T. P. Harty et al. "High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves." In: *Phys. Rev. Lett.* 117 (14 Sept. 2016), p. 140501. DOI: 10.1103/PhysRevLett.117.140501.

[99] Simon Haykin. *Adaptive Filter Theory*. 4th. Prentice Hall, 2002.

[100] Peter Henderson et al. "Deep Reinforcement Learning that Matters." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. 1. 2018. URL: https://arxiv.org/abs/1709.06560.

[101] M. Hennrich, T. Legero, A. Kuhn, and G. Rempe. "Vacuum-Stimulated Raman Scattering Based on Adiabatic Passage in a High-Finesse Optical Cavity." In: *Phys. Rev. Lett.* 85 (23 Dec. 2000), pp. 4872–4875. DOI: 10.1103/PhysRevLett.85.4872.

[102] Loic Henriet et al. "Quantum computing with neutral atoms." In: *Quantum* 4 (Sept. 2020), p. 327. ISSN: 2521-327X. DOI: 10.22331/q-2020-09-21-327.

[103] B. Hensen et al. "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres." In: *Nature* 526.7575 (Oct. 2015), pp. 682–686. ISSN: 1476-4687. DOI: 10.1038/nature15759.

[104] Mark Hillery, Vladimír Bužek, and André Berthiaume. "Quantum secret sharing." In: *Physical Review A* 59.3 (1999), p. 1829. DOI: 10.1103/PhysRevA.59.1829.

[105] Annemarie Holleczek et al. "Quantum Logic with Cavity Photons From Single Atoms." In: *Phys. Rev. Lett.* 117.2 (July 2016), p. 023602. DOI: 10.1103/PhysRevLett.117.023602.

[106] S. C. Hou, M. A. Khan, X. X. Yi, Daoyi Dong, and Ian R. Petersen. "Optimal Lyapunov-based quantum control for quantum systems." In: *Phys. Rev. A* 86 (2 Aug. 2012), p. 022321. DOI: 10.1103/PhysRevA.86.022321.

[107] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements." In: *Nature Physics* 16.10 (June 2020), pp. 1050–1057. ISSN: 1745-2481. DOI: 10.1038/s41567-020-0932-7.

[108] Shengyi Huang et al. "CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms." In: *Journal of Machine Learning Research* 23.274 (2022), pp. 1–18. URL: http://jmlr.org/papers/v23/21-1342.html.

[109] A. C. Hughes et al. "Benchmarking a High-Fidelity Mixed-Species Entangling Gate." In: *Phys. Rev. Lett.* 125 (8 Aug. 2020), p. 080504. DOI: 10.1103/PhysRevLett.125.080504.

[110] D. Hunger et al. "A Fiber Fabry-Perot Cavity with High Finesse." In: *New J. Phys.* 12.6 (June 2010), p. 065038. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/6/065038.

[111] D. Istrati et al. "Sequential generation of linear cluster states from a single photon emitter." In: *Nature Communications* 11.1 (Oct. 2020).

[112] S. Jandura and G. Pupillo. "Time-Optimal Two- and Three-Qubit Gates for Rydberg Atoms." In: *Quantum* 6 (May 2022), p. 712. ISSN: 2521-327X. DOI: 10.22331/q-2022-05-13-712.

[113] Sven Jandura, Jeff D. Thompson, and Guido Pupillo. "Optimizing Rydberg Gates for Logical-Qubit Performance." In: *PRX Quantum* 4 (2 June 2023), p. 020336. DOI: 10.1103/PRXQuantum.4.020336.

[114] J.R. Johansson, P.D. Nation, and Franco Nori. "QuTiP: An open-source Python framework for the dynamics of open quantum systems." In: *Computer Physics Communications* 183.8 (2012), pp. 1760–1772. ISSN: 0010-4655. DOI: https://doi.org/10.1016/j.cpc.2012.02.021.

[115] M. P. A. Jones et al. "Fast quantum state control of a single trapped neutral atom." In: *Phys. Rev. A* 75 (4 Apr. 2007), p. 040301. DOI: 10.1103/PhysRevA.75.040301.

[116] Richard Jozsa. "Fidelity for Mixed Quantum States." In: *Journal of Modern Optics* 41.12 (1994), pp. 2315–2323. DOI: 10.1080/09500349414552171.

[117] F. Kaneda and P. G. Kwiat. "High-efficiency single-photon generation via large-scale active time multiplexing." In: *Science Advances* 5.10 (2019), eaaw8586. DOI: 10.1126/sciadv.aaw8586.

[118] Mark Kasevich and Steven Chu. "Atomic interferometry using stimulated Raman transitions." In: *Physical Review Letters* 67.2 (1991), pp. 181–184. DOI: 10.1103/PhysRevLett.67.181.

[119] E Kassa, W Hughes, S Gao, and J F Goodwin. "Effects of cavity birefringence in polarisation-encoded quantum networks." In: *New Journal of Physics* 25.1 (Jan. 2023), p. 013004. DOI: 10.1088/1367-2630/acabe0.

[120] Matthias Keller, Birgit Lange, Kazuhiro Hayasaka, Wolfgang Lange, and Herbert Walther. "Continuous generation of single photons with controlled waveform in an ion-trap cavity system." In: *Nature* 431.7012 (Oct. 2004), pp. 1075–1078. DOI: 10.1038/nature02961.

[121] Irtaza Khalid, Carrie A. Weidner, Edmond A. Jonckheere, Sophie G. Schirmer, and Frank C. Langbein. "Sample-efficient model-based reinforcement learning for quantum control." In: *Physical Review Research* 5.4 (Oct. 2023), p. 043002. ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.5.043002. (Visited on 01/26/2024).

[122] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. "Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms." In: *Journal of Magnetic Resonance* 172.2 (2005), pp. 296–305. ISSN: 1090-7807. DOI: https://doi.org/10.1016/j.jmr.2004.11.004.

[123] Patrick Kidger. *On Neural Differential Equations*. 2022. arXiv: 2202.02435 [cs.LG]. URL: https://arxiv.org/abs/2202.02435.

[124] Z. KIS and S. STENHOLM. "Optimal control approach for a degenerate STIRAP." In: *Journal of Modern Optics* 49.1-2 (Jan. 2002), pp. 111–124. DOI: 10.1080/09500340110074736.

[125] Morten Kjaergaard and et al. "Superconducting qubits: Current state of play." In: *Annual Review of Condensed Matter Physics* 11 (2020), pp. 369–395.

[126] E. Knill, R. Laflamme, and G. J. Milburn. "A Scheme for Efficient Quantum Computation with Linear Optics." In: *Nature* 409.6816 (Jan. 2001), pp. 46–52. ISSN: 0028-0836. DOI: 10.1038/35051009.

[127] Christiane P Koch. "Controlling open quantum systems: tools, achievements, and limitations." In: *Journal of Physics: Condensed Matter* 28.21 (May 2016), p. 213001. ISSN: 1361-648X. DOI: 10.1088/0953-8984/28/21/213001.

[128] Christiane P. Koch et al. "Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe." In: *EPJ Quantum Technology* 9.1 (Dec. 2022), p. 19. ISSN: 2662-4400, 2196-0763. DOI: 10.1140/epjqt/s40507-022-00138-x. (Visited on 04/30/2024).

[129] Petr Král, Ioannis Thanopulos, and Moshe Shapiro. "Colloquium: Coherently controlled adiabatic passage." In: *Rev. Mod. Phys.* 79 (1 Jan. 2007), pp. 53–77. DOI: 10.1103/RevModPhys.79.53.

[130] Mario Krenn, Jonas Landgraf, Thomas Foesel, and Florian Marquardt. "Artificial intelligence and machine learning for quantum technologies." In: *Phys. Rev. A* 107 (1 Jan. 2023), p. 010101. DOI: 10.1103/PhysRevA.107.010101.

[131] Axel Kuhn, Markus Hennrich, and Gerhard Rempe. "Deterministic Single-Photon Source for Distributed Quantum Networking." In: *Phys. Rev. Lett.* 89 (6 July 2002), p. 067901. DOI: 10.1103/PhysRevLett.89.067901.

[132] JR Kuklinski, U Gaubatz, and K Bergmann. "Adiabatic population transfer in a three-level system driven by delayed laser pulses." In: *Physical Review A* 40.11 (1989), p. 6741.

[133] Xavier Laforgue, Ghassen Dridi, and Stephane Guerin. "Optimal robust stimulated Raman exact passage by inverse optimization." In: *Phys. Rev. A* 105 (3 Mar. 2022), p. 032807. DOI: 10.1103/PhysRevA.105.032807.

[134] S. Langenfeld, P. Thomas, O. Morin, and G. Rempe. "Quantum Repeater Node Demonstrating Unconditionally Secure Key Distribution." In: *Phys. Rev. Lett.* 126 (23 June 2021), p. 230506. DOI: 10.1103/PhysRevLett.126.230506.

[135] J P Lee et al. "A quantum dot as a source of time-bin entangled multi-photon states." In: *Quantum Science and Technology* 4.2 (Mar. 2019), p. 025011. DOI: 10.1088/2058-9565/ab0a9b.

[136] Thomas Legero, Tatjana Wilk, Axel Kuhn, and Gerhard Rempe. "Characterization of Single Photons Using Two-Photon Interference." In: *Advances In Atomic, Molecular, and Optical Physics* 53 (2006), pp. 253–289.

[137] John Lekavich. "Basics of Acousto-Optic Devices." In: *Lasers & Applications* (Apr. 1986).

[138] H. Levine et al. "Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms." In: *Phys. Rev. Lett.* 123 (17 Oct. 2019), p. 170503. DOI: 10.1103/PhysRevLett.123.170503.

[139] Harry Levine et al. "Dispersive optical systems for scalable Raman driving of hyperfine qubits." In: *Phys. Rev. A* 105 (3 Mar. 2022), p. 032618. DOI: 10.1103/PhysRevA.105.032618.

[140] Jin-Peng Li et al. "Multiphoton Graph States from a Solid-State Single-Photon Source." In: *ACS Photonics* 7.7 (June 2020), pp. 1603–1610. DOI: 10.1021/acsphotonics.0c00192.

[141] Shengyong Li et al. "Robust quantum control using reinforcement learning from demonstration." In: *npj Quantum Information* 11.1 (July 2025). ISSN: 2056-6387. DOI: 10.1038/s41534-025-01065-2.

[142] Timothy P Lillicrap et al. "Continuous control with deep reinforcement learning." In: *International Conference on Learning Representations (ICLR)*. 2016. URL: https://arxiv.org/abs/1509.02971.

[143] Netanel H. Lindner and Terry Rudolph. "Proposal for Pulsed On-Demand Sources of Photonic Cluster State Strings." In: *Physical Review Letters* 103.11 (Sept. 2009). ISSN: 1079-7114. DOI: 10.1103/physrevlett.103.113602.

[144] Joanna W Lis et al. "Midcircuit operations using the omg architecture in neutral atom arrays." In: *Physical Review X* 13.4 (2023), p. 041035.

[145] Adriana E. Lita, Aaron J. Miller, and Sae Woo Nam. "Counting near-infrared single-photons with 95% efficiency." In: *Opt. Express* 16.5 (Mar. 2008), pp. 3032–3040. DOI: 10.1364/OE.16.003032.

[146] Kaipeng Liu, Dominique Sugny, Xi Chen, and Stéphane Guérin. "Optimal Pulse Design for Dissipative-Stimulated Raman Exact Passage." In: *Entropy* 25.5 (May 2023), p. 790. ISSN: 1099-4300. DOI: 10.3390/e25050790.

[147] Chao-Yang Lu and Jian-Wei Pan. "Quantum-dot single-photon sources for the quantum internet." In: *Nature Nanotechnology* 16.12 (Dec. 2021), pp. 1294–1296. DOI: 10.1038/s41565-021-01033-9.

[148] Chris Lu et al. "Discovered policy optimisation." In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 16455–16468.

[149] M. D. Lukin et al. "Dipole blockade and quantum information processing in mesoscopic atomic ensembles." In: *Physical Review Letters* 87.3 (2001), p. 037901. DOI: 10.1103/PhysRevLett.87.037901.

[150] P. Magnard et al. "Fast and Unconditional All-Microwave Reset of a Superconducting Qubit." In: *Phys. Rev. Lett.* 121 (6 Aug. 2018), p. 060502. DOI: 10.1103/PhysRevLett.121.060502.

[151] Hannah J. Manetsch et al. *A tweezer array with 6100 highly coherent atomic qubits*. 2024. arXiv: 2403.12021 [quant-ph]. URL: https://arxiv.org/abs/2403.12021.

[152] J. McKeever et al. "Deterministic Generation of Single Photons from One Atom Trapped in a Cavity." In: *Science* 303.5666 (2004), pp. 1992–1994.

[153] M. Meraner et al. "Indistinguishable photons from a trapped-ion quantum network node." In: *Phys. Rev. A* 102 (5 Nov. 2020), p. 052614. DOI: 10.1103/PhysRevA.102.052614.

[154] Friederike Metz and Marin Bukov. "Self-correcting quantum many-body control using reinforcement learning with tensor networks." In: *Nature Machine Intelligence* 5.7 (July 2023). ISSN: 2522-5839. DOI: 10.1038/s42256-023-00687-5.

[155] Madhav Mohan, Robert de Keijzer, and Servaas Kokkelmans. "Robust control and optimal Rydberg states for neutral atom two-qubit gates." In: *Physical Review Research* 5.3 (July 2023). ISSN: 2643-1564. DOI: 10.1103/physrevresearch.5.033052.

[156] R. Mohr, M. Coblenz, and P. Kirst. "Globally optimal univariate spline approximations." In: *Computational Optimization and Applications* 85.2 (Feb. 2023), pp. 409–439. ISSN: 1573-2894. DOI: 10.1007/s10589-023-00462-7.

[157] Christopher Monroe and Jungsang Kim. "Scaling the ion trap quantum processor." In: *Science* 339.6124 (2013), pp. 1164–1169.

[158] Julian Münzberg, Franz Draxl, Saimon Filipe Covre da Silva, et al. "Fast and efficient demultiplexing of single photons from a quantum dot with resonantly enhanced electro-optic modulators." In: *APL Photonics* 7.7 (July 2022), p. 070802. DOI: 10.1063/5.0091867.

[159] Murat Muradoglu et al. *Quantum-assured magnetic navigation achieves positioning accuracy better than a strategic-grade INS in airborne and ground-based field trials*. 2025. arXiv: 2504.08167 [quant-ph]. URL: https://arxiv.org/abs/2504.08167.

[160] Ho Nam Nguyen et al. "Reinforcement learning pulses for transmon qubit entangling gates." In: *Machine Learning: Science and Technology* 5.2 (June 2024), p. 025066. ISSN: 2632-2153. DOI: 10.1088/2632-2153/ad4f4d.

[161] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[162] Georgios M. Nikolopoulos. "Cryptographic one-way function based on boson sampling." In: *Quantum Information Processing* 18.8 (2019). arXiv: 1907.01788 Publisher: Springer US ISBN: 1112801923729, pp. 1–25. ISSN: 15700755. DOI: 10.1007/s11128-019-2372-9.

[163] Georgios M. Nikolopoulos and Thomas Brougham. "Decision and function problems based on boson sampling." In: *Physical Review A* 94.1 (2016). arXiv: 1607.02987, pp. 1–13. ISSN: 24699934. DOI: 10.1103/PhysRevA.94.012315.

[164] Peter B R Nisbet-Jones, Jerome Dilley, Annemarie Holleczek, Oliver Barter, and Axel Kuhn. "Photonic qubits, qutrits and ququads accurately prepared and delivered on demand." In: *New Journal of Physics* 15.5 (May 2013), p. 053007. ISSN: 1367-2630. DOI: 10.1088/1367-2630/15/5/053007.

[165] Peter B R Nisbet-Jones, Jerome Dilley, Daniel Ljunggren, and Axel Kuhn. "Highly Efficient Source for Indistinguishable Single Photons of Controlled Shape." In: *New Journal of Physics* 13.10 (Oct. 2011), p. 103036. ISSN: 1367-2630. DOI: 10.1088/1367-2630/13/10/103036.

[166] Murphy Yuezhen Niu, Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven. "Universal quantum control through deep reinforcement learning." In: *npj Quantum Information* 5.1 (Dec. 2019). Publisher: Nature Partner Journals. ISSN: 20566387. DOI: 10.1038/s41534-019-0141-3.

[167] Ariel Norambuena, Marios Mattheakis, Francisco J González, and Raúl Coto. "Physics-informed neural networks for quantum control." In: *Physical Review Letters* 132.1 (2024), p. 010801.

[168] M. A. Norcia et al. "Midcircuit Qubit Measurement and Rearrangement in a ^{171}Yb Atomic Array." In: *Phys. Rev. X* 13 (4 Nov. 2023), p. 041034. DOI: 10.1103/PhysRevX.13.041034.

[169] Jan Olle, Remmy Zen, Matteo Puviani, and Florian Marquardt. "Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent." In: *npj Quantum Information* 10.1 (Dec. 2024). ISSN: 2056-6387. DOI: 10.1038/s41534-024-00920-y.

[170] Alan V. Oppenheim and Ronald W. Schafer. *Discrete-Time Signal Processing*. 2nd. Prentice Hall, 1999. ISBN: 9780137549207.

[171] Román Orús. "A practical introduction to tensor networks: Matrix product states and projected entangled pair states." In: *Annals of Physics* 349 (2014), pp. 117–158. DOI: 10.1016/j.aop.2014.06.013.

[172] Mark Oxborrow and Alastair G Sinclair. "Single-photon sources." In: *Contemporary Physics* 46.3 (2005), pp. 173–206.

[173] Alice Pagano, Daniel Jaschke, Werner Weiss, and Simone Montangero. *Optimal control transport of neutral atoms in optical tweezers at finite temperature*. Sept. 2024. DOI: 10.1103/PhysRevResearch.6.033282.

[174] Alice Pagano, Matthias M. Müller, Tommaso Calarco, Simone Montangero, and Phila Rembold. "Role of bases in quantum optimal control." In: *Phys. Rev. A* 110 (6 Dec. 2024), p. 062608. DOI: 10.1103/PhysRevA.110.062608.

[175] Iris Paparelle, Lorenzo Moro, and Enrico Prati. "Digitally stimulated Raman passage by deep reinforcement learning." In: *Physics Letters A* 384.14 (May 2020), p. 126266. ISSN: 03759601. DOI: 10.1016/j.physleta.2020.126266. (Visited on 01/26/2024).

[176] Vanessa Paulisch, Han Rui, Hui Khoon Ng, and Berthold-Georg Englert. "Beyond adiabatic elimination: A hierarchy of approximations for multi-photon processes." In: *The European Physical Journal Plus* 129 (2014), pp. 1–14.

[177] G Pelegrí, A J Daley, and J D Pritchard. "High-fidelity multiqubit Rydberg gates via two-photon adiabatic rapid passage." In: *Quantum Science and Technology* 7.4 (Aug. 2022), p. 045020. DOI: 10.1088/2058-9565/ac823a.

[178] Asher Peres and Daniel R Terno. "Quantum information and relativity theory." In: *Reviews of Modern Physics* 76.1 (2004), p. 93.

[179] Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt. "Deep Reinforcement Learning for Quantum State Preparation with Weak Nonlinear Measurements." In: *Quantum* 6 (June 2022), p. 747. ISSN: 2521-327X. DOI: 10.22331/q-2022-06-28-747.

[180] PsiQuantum Team. "A manufacturable platform for photonic quantum computing." en. In: *Nature* (Feb. 2025).

[181] E. M. Purcell. "Spontaneous emission probabilities at radio frequencies." In: *Physical Review* 69 (1946), p. 681.

[182] Daniel Puzzuoli, Christopher J. Wood, Daniel J. Egger, Benjamin Rosand, and Kento Ueda. "Qiskit Dynamics: A Python package for simulating the time dynamics of quantum systems." In: *Journal of Open Source Software* 8.90 (2023), p. 5853. DOI: 10.21105/joss.05853.

[183] L. Qi, J. Chiaverini, H. Espinós, M. Palmero, and J. G. Muga. "Fast and robust particle shuttling for quantum science and technology." In: *Europhysics Letters* 134.2 (Apr. 2021), p. 23001. ISSN: 1286-4854. DOI: 10.1209/0295-5075/134/23001.

[184] J. Raftery et al. *Direct digital synthesis of microwave waveforms for quantum computing*. 2017. arXiv: 1703.00942v1 [quant-ph].

[185] J. Randall et al. "Efficient preparation and detection of microwave dressed-state qubits and qutrits with trapped ions." In: *Phys. Rev. A* 91 (1 Jan. 2015), p. 012322. DOI: 10.1103/PhysRevA.91.012322.

[186] Robert Raussendorf and Hans J. Briegel. "A One-Way Quantum Computer." In: *Phys. Rev. Lett.* 86 (22 May 2001), pp. 5188–5191. DOI: 10.1103/PhysRevLett.86.5188.

[187] Daniel M. Reich, Mamadou Ndong, and Christiane P. Koch. "Monotonically convergent optimization in quantum control using Krotov's method." In: *The Journal of Chemical Physics* 136.10 (Mar. 2012), p. 104103. ISSN: 0021-9606. DOI: 10.1063/1.3691827.

[188] Jelmer Renema, Valery Shchesnovich, and Raul Garcia-Patron. *Classical simulability of noisy boson sampling*. 2019. arXiv: 1809.01953 [quant-ph].

[189] Kevin Reuer et al. "Realizing a deep reinforcement learning agent for real-time quantum feedback." In: *Nature Communications* 14.1 (Nov. 2023). ISSN: 2041-1723. DOI: 10.1038/s41467-023-42901-3.

[190] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo. "Feedback Control of a Solid-State Qubit Using High-Fidelity Projective Measurement." In: *Phys. Rev. Lett.* 109 (24 Dec. 2012), p. 240502. DOI: 10.1103/PhysRevLett.109.240502.

[191] Stephan Ritter et al. "An Elementary Quantum Network of Single Atoms in Optical Cavities." en. In: *Nature* 484.7393 (Apr. 2012), pp. 195–200. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature11023.

[192] M. Saffman, I. I. Beterov, A. Dalal, E. J. Páez, and B. C. Sanders. "Symmetric Rydberg controlled-Z gates with adiabatic pulses." In: *Phys. Rev. A* 101 (6 June 2020), p. 062309. DOI: 10.1103/PhysRevA.101.062309.

[193] Mark Saffman. "Quantum computing with neutral atoms." In: *National Science Review* 6.1 (Sept. 2018), pp. 24–25. ISSN: 2095-5138. DOI: 10.1093/nsr/nwy088.

[194] Mark Saffman, Thad G Walker, and Klaus Mølmer. "Quantum information with Rydberg atoms." In: *Reviews of Modern Physics* 82.3 (2010), p. 2313.

[195] A. Saharyan, B. Rousseaux, Z. Kis, S. Stryzhenko, and S. Guérin. "Propagating single photons from an open cavity: Description from universal quantization." In: *Phys. Rev. Res.* 5 (3 July 2023), p. 033056. DOI: 10.1103/PhysRevResearch.5.033056.

[196] Hesham Sakr et al. "Hollow core optical fibres with comparable attenuation to silica fibres between 600 and 1100nm." In: *Nature Communications* 11.1 (Nov. 2020). ISSN: 2041-1723. DOI: 10.1038/s41467-020-19910-7.

[197] S. Saner et al. "Breaking the Entangling Gate Speed Limit for Trapped-Ion Qubits Using a Phase-Stable Standing Wave." In: *Phys. Rev. Lett.* 131 (22 Dec. 2023), p. 220601. DOI: 10.1103/PhysRevLett.131.220601.

[198] Frank Schaefer, Michal Kloc, Christoph Bruder, and Niels Larch. "A differentiable programming method for quantum control." In: *Machine Learning: Science and Technology* 1.3 (Sept. 2020), p. 035009. ISSN: 2632-2153. DOI: 10.1088/2632-2153/ab9802. (Visited on 04/30/2024).

[199] Stefan Scheel. *Permanents in linear optical networks*. 2004. arXiv: quant-ph/0406127 [quant-ph].

[200] Maximilian Schlosshauer. "Decoherence, the measurement problem, and interpretations of quantum mechanics." In: *Rev. Mod. Phys.* 76 (4 Feb. 2005), pp. 1267–1305. DOI: 10.1103/RevModPhys.76.1267.

[201] I. J. Schoenberg. "Contributions to the problem of approximation of equidistant data by analytic functions. Part A. On the problem of smoothing or graduation. A first class of analytic approximation formulae." In: *Quarterly of Applied Math-*

ematics 4 (1946), pp. 45–99. URL: https://api.semanticscholar.org/CorpusID:125923957.

[202] Erwin Schrödinger. "Über das Verhältnis der Heisenberg-Born-Jordanschen Quantenmechanik zu der meinen." In: *Annalen der Physik* 384.8 (1926), pp. 734–756. DOI: 10.1002/andp.19263840802.

[203] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. *Proximal Policy Optimization Algorithms*. 2017. arXiv: 1707.06347 [cs.LG].

[204] I. Schwartz et al. "Deterministic generation of a cluster state of entangled photons." In: *Science* 354.6311 (2016), pp. 434–437.

[205] K. Schymik et al. "Enhanced atom-by-atom assembly of arbitrary tweezer arrays." In: *Phys. Rev. A* 102 (6 Dec. 2020), p. 063107. DOI: 10.1103/PhysRevA.102.063107.

[206] Kai-Niklas Schymik et al. "In situ equalization of single-atom loading in large-scale optical tweezer arrays." In: *Phys. Rev. A* 106 (2 Aug. 2022), p. 022611. DOI: 10.1103/PhysRevA.106.022611.

[207] C. Sheng et al. "High-Fidelity Single-Qubit Gates on Neutral Atoms in a Two-Dimensional Magic-Intensity Optical Dipole Trap Array." In: *Phys. Rev. Lett.* 121 (24 Dec. 2018), p. 240501. DOI: 10.1103/PhysRevLett.121.240501.

[208] Yu Shi and Edo Waks. "Deterministic generation of multidimensional photonic cluster states using time-delay feedback." In: *Phys. Rev. A* 104 (1 July 2021), p. 013703. DOI: 10.1103/PhysRevA.104.013703.

[209] Peter W Shor. "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer." In: *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509.

[210] Deepesh Singh et al. "Proof-of-work consensus by quantum sampling." In: *Quantum Science and Technology* 10.2 (Feb. 2025), p. 025020. ISSN: 2058-9565. DOI: 10.1088/2058-9565/adae2b.

[211] V. V. Sivak et al. "Real-time quantum error correction beyond break-even." In: *Nature* 616.7955 (Mar. 2023), pp. 50–55. ISSN: 1476-4687. DOI: 10.1038/s41586-023-05782-6.

[212] Chloe So, Thomas Doherty, Mark IJspeert, Jan Ole Ernst, and Axel Kuhn. "Fabrication, Characterisation and Prospects of a High-Finesse Microcavity Created From Tapered Fused Silica Substrates." In: *Quantum 2.0 Conference and Exhibition*. Optica Publishing Group, 2024, QW3A.46. DOI: 10.1364/QUANTUM.2024.QW3A.46.

[213] N. Somaschi et al. "Near-optimal single-photon sources in the solid state." In: *Nature Photonics* 10.5 (Mar. 2016), pp. 340–345. DOI: 10.1038/nphoton.2016.23.

[214] R Srinivas et al. "An elementary quantum network of entangled optical atomic clocks." In: *Nature* 609.7928 (2022), pp. 689–694.

[215] Daniel A. Steck. *Quantum and Atom Optics*. http://steck.us/teaching. Version Revision 0.16.4, 7 May 2025. 2025.

[216] Daniel A. Steck. *Rubidium 87 D Line Data*. http://steck.us/alkalidata/rubidium87numbers.pdf. 2001, Revised 2021.

[217] LJ Stephenson et al. "High-rate, high-fidelity entanglement of qubits across an elementary quantum network." In: *Physical review letters* 124.11 (2020), p. 110501.

[218] Yuan Sun. "Off-resonant modulated driving gate protocols for two-photon ground-Rydberg transition and finite Rydberg blockade strength." In: *Opt. Express* 31.2 (Jan. 2023), pp. 3114–3121. DOI: 10.1364/OE.480513.

[219] Patrik I. Sund, Ravitej Uppu, Stefano Paesani, and Peter Lodahl. *Hardware requirements for realizing a quantum advantage with deterministic single-photon sources*. Apr. 2024. DOI: 10.1103/PhysRevA.109.042613.

[220] Richard S Sutton and Andrew G Barto. "Reinforcement learning: An introduction." In: *Robotica* 17.2 (1999), pp. 229–235.

[221] Kazufumi Tanji, Hiroki Takahashi, Wojciech Roga, and Masahiro Takeoka. "Rate-fidelity tradeoff in cavity-based remote entanglement generation." In: *Phys. Rev. A* 110 (4 Oct. 2024), p. 042405. DOI: 10.1103/PhysRevA.110.042405.

[222] Barbara M. Terhal. "Quantum error correction for quantum memories." In: *Reviews of Modern Physics* 87.2 (Apr. 2015), pp. 307–346. ISSN: 1539-0756. DOI: 10.1103/revmodphys.87.307.

[223] Philip Thomas, Leonardo Ruscio, Olivier Morin, and Gerhard Rempe. "Efficient generation of entangled multiphoton graph states from a single atom." In: *Nature* 608.7924 (Aug. 2022), pp. 677–681. DOI: 10.1038/s41586-022-04987-5.

[224] Philip Thomas, Leonardo Ruscio, Olivier Morin, and Gerhard Rempe. "Fusion of deterministically generated photonic graph states." In: *Nature* 629.8012 (May 2024), pp. 567–572. ISSN: 1476-4687. DOI: 10.1038/s41586-024-07357-5.

[225] René Philip Thomas. "Photonic graph state generation from individual atoms in an optical cavity." en. PhD thesis. Technische Universität München, 2024, p. 101. URL: https://mediatum.ub.tum.de/1747824.

[226] Inc. Thorlabs. *PBS25-780 Laser Line Polarizing Beamsplitter Cube*. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=7490&pn=PBS25-780. Accessed: 2024-11-21. 2024.

[227] Natasha Tomm et al. "A bright and fast source of coherent single photons." In: *Nature Nanotechnology* 16.4 (Jan. 2021), pp. 399–403. DOI: 10.1038/s41565-020-00831-x.

[228] Géza Tóth and Otfried Gühne. "Detecting Genuine Multipartite Entanglement with Two Local Measurements." In: *Phys. Rev. Lett.* 94 (6 Feb. 2005), p. 060501. DOI: 10.1103/PhysRevLett.94.060501.

[229] Takeru Utsugi, Akihisa Goban, Yuuki Tokunaga, Hayato Goto, and Takao Aoki. "Gaussian-wave-packet model for single-photon generation based on cavity quan-

tum electrodynamics under adiabatic and nonadiabatic conditions." In: *Phys. Rev. A* 106 (2 Aug. 2022), p. 023712. DOI: 10.1103/PhysRevA.106.023712.

[230] Stan Van Der Linde, Willem De Kok, Tariq Bontekoe, and Sebastian Feld. "qgym: A gym for training and benchmarking rl-based quantum compilation." In: *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 2. IEEE. 2023, pp. 26–30.

[231] L. M. K. Vandersypen and I. L. Chuang. "NMR techniques for quantum control and computation." In: *Rev. Mod. Phys.* 76 (4 Jan. 2005), pp. 1037–1069. DOI: 10.1103/RevModPhys.76.1037.

[232] Luca Vannucci and Niels Gregersen. "Highly efficient and indistinguishable single-photon sources via phonon-decoupled two-color excitation." In: *Phys. Rev. B* 107 (19 May 2023), p. 195306. DOI: 10.1103/PhysRevB.107.195306.

[233] G. S. Vasilev, A. Kuhn, and N. V. Vitanov. "Optimum Pulse Shapes for Stimulated Raman Adiabatic Passage." In: *Phys. Rev. A* 80.1 (July 2009), p. 013417. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.80.013417.

[234] Genko S. Vasilev, Daniel Ljunggren, and Axel Kuhn. "Single Photons Made-to-Measure." In: *New J. Phys.* 12.6 (June 2010), p. 063024. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/6/063024.

[235] Arian Vezvaee, Evangelia Takou, Paul Hilaire, Matthew F. Doty, and Sophia E. Economou. "Avoiding Leakage and Errors Caused by Unwanted Transitions in Lambda Systems." In: *PRX Quantum* 4 (3 July 2023), p. 030312. DOI: 10.1103/PRXQuantum.4.030312.

[236] P. Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[237] N V Vitanov, K-A Suominen, and B W Shore. "Creation of coherent atomic superpositions by fractional stimulated Raman adiabatic passage." In: *Journal of*

Physics B: Atomic, Molecular and Optical Physics 32.18 (Sept. 1999), p. 4535. DOI: 10.1088/0953-4075/32/18/312.

[238] Nikolay V. Vitanov, Andon A. Rangelov, Bruce W. Shore, and Klaas Bergmann. "Stimulated Raman adiabatic passage in physics, chemistry, and beyond." In: *Rev. Mod. Phys.* 89 (1 Mar. 2017), p. 015006. DOI: 10.1103/RevModPhys.89.015006.

[239] Chenlu Wang et al. "Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds." In: *npj Quantum Information* 8.1 (Jan. 2022). ISSN: 2056-6387. DOI: 10.1038/s41534-021-00510-2.

[240] Hui Wang et al. "Boson Sampling with 20 Input Photons and a 60-Mode Interferometer in a 10^{14}-Dimensional Hilbert Space." In: *Phys. Rev. Lett.* 123 (25 Dec. 2019), p. 250503. DOI: 10.1103/PhysRevLett.123.250503.

[241] Stephanie Wehner, David Elkouss, and Ronald Hanson. "Quantum internet: A vision for the road ahead." In: *Science* 362.6412 (2018), eaam9288.

[242] T. Wilk, H. P. Specht, S. C. Webster, G. Rempe, and A. Kuhn. "Scheme for generating a sequence of single photons of alternating polarization." In: *Journal of Modern Optics* 54.11 (July 2007), pp. 1569–1580. DOI: 10.1080/09500340600736777.

[243] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. "Single-Atom Single-Photon Quantum Interface." In: *Science* 317.5837 (July 2007), pp. 488–490. DOI: 10.1126/science.1143835.

[244] T. Wilk, S. C. Webster, H. P. Specht, G. Rempe, and A. Kuhn. "Polarization-Controlled Single Photons." In: *Physical Review Letters* 98.6 (Feb. 2007).

[245] Chao-Wei Yang et al. "Sequential generation of multiphoton entanglement with a Rydberg superatom." In: *Nature Photonics* 16.9 (Aug. 2022), pp. 658–661. ISSN: 1749-4893. DOI: 10.1038/s41566-022-01054-3.

[246] Xing-Can Yao et al. "Experimental demonstration of topological error correction." In: *Nature* 482.7386 (Feb. 2012), pp. 489–494. DOI: 10.1038/nature10770.

[247] D. D. Yavuz et al. "Fast Ground State Manipulation of Neutral Atoms in Microscopic Optical Traps." In: *Phys. Rev. Lett.* 96 (6 Feb. 2006), p. 063001. DOI: 10.1103/PhysRevLett.96.063001.

[248] Shang Yu et al. "A universal programmable Gaussian boson sampler for drug discovery." en. In: *Nature Computational Science* 3.10 (Oct. 2023), pp. 839–848. ISSN: 2662-8457. DOI: 10.1038/s43588-023-00526-y. (Visited on 06/01/2024).

[249] S. Zeytino ğlu et al. "Microwave-induced amplitude- and phase-tunable qubit-resonator coupling in circuit quantum electrodynamics." In: *Phys. Rev. A* 91 (4 Apr. 2015), p. 043846. DOI: 10.1103/PhysRevA.91.043846.

[250] Chao Zhang, Yun-Feng Huang, Bi-Heng Liu, Chuan-Feng Li, and Guang-Can Guo. "Spontaneous Parametric Down-Conversion Sources for Multiphoton Experiments." In: *Advanced Quantum Technologies* 4.5 (Mar. 2021), p. 2000132. DOI: 10.1002/qute.202000132.

[251] Xiao-Ming Zhang, Zezhu Wei, Raza Asad, Xu-Chen Yang, and Xin Wang. "When does reinforcement learning stand out in quantum control? A comparative study on state preparation." In: *npj Quantum Information* 5.1 (Oct. 2019), p. 85. ISSN: 2056-6387. DOI: 10.1038/s41534-019-0201-8. (Visited on 04/26/2024).

[252] Han-Sen Zhong et al. "12-Photon Entanglement and Scalable Scattershot Boson Sampling with Optimal Entangled-Photon Pairs from Parametric Down-Conversion." In: *Phys. Rev. Lett.* 121 (25 Dec. 2018), p. 250505. DOI: 10.1103/PhysRevLett.121.250505.

[253] Han-Sen Zhong et al. "Quantum computational advantage using photons." In: *Science* 370.6523 (2020), pp. 1460–1463.

[254] Marek Zukowski, Anton Zeilinger, Michael A Horne, and Artur K Ekert. ""Event-ready-detectors" Bell experiment via entanglement swapping." In: *Physical Review Letters* 71.26 (1993), p. 4287.

DECLARATION

I declare, that unless otherwise indicated, all of the work is my own.

Oxford, December 2025

Jan Ole Ernst