

ABSTRACT

Title of Dissertation: RYDBERG ENSEMBLES FOR
QUANTUM NETWORKING

Alexander Nicholas Craddock
Doctor of Philosophy, 2020

Dissertation Directed by: Professors Steve Rolston and Trey Porto
Joint Quantum Institute,
National Institute of Standards and Technology
and
Department of Physics, University of Maryland

Rydberg ensembles, atomic clouds with one or more atoms excited to a Rydberg state, have proven to be a good platform for the study of photon-photon interactions. This is due to the nonlinearities they exhibit at the single photon level arising from Rydberg-Rydberg interactions. As a result, they have shown promise for use in a multitude of applications, among them quantum networking.

In this thesis I describe the construction and operation of an apparatus for the purpose of cooling, trapping and probing Rydberg ensemble physics in a cloud of ^{87}Rb atoms. In addition, I describe a pair of projects undertaken with the apparatus. In the first, I report our demonstration of a Rydberg ensemble based on-demand single photon source. Here, we make use of Rydberg blockade to allow us to prepare a single collective Rydberg excitation in the cloud. The spin wave excitation is then retrieved by coherently mapping it onto a propagating photon. Our source is highly pure and efficient, while producing narrow bandwidth and indistinguishable pho-

tons. Such sources are important devices for the purposes of quantum networking, computation and metrology. Following from this, I describe a collaborative project where we show time resolved Hong-Ou-Mandel interference between photons produced by our Rydberg ensemble source, and a collaborators source based on a single trapped barium ion. This demonstration is a critical step in the entanglement, and hybrid quantum networking, of these two disparate systems.

RYDBERG ENSEMBLES FOR QUANTUM NETWORKING

by

Alexander Nicholas Craddock

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:

Professor Steve Rolston, Chair/Advisor

Professor Trey Porto, Co-Advisor

Professor Alexey Gorshkov

Professor Mohammad Hafezi (Dean's Representative)

Professor Norbert Linke

© Copyright by
Alexander Nicholas Craddock
2020

Dedication

For my wife, Clarissa, who always believed in me, even when I didn't.

Acknowledgments

When I arrived at UMD the rubidium Rydberg lab was just an vacant room. It has been a long road since then, with plenty of frustration along the way. However, the glimmers of success are always incredibly satisfying and I am truly grateful for the experience I have had. While this thesis formally bears my name, there are a number of people without which I would not be writing it.

Firstly, I have to thank both of my advisors, Trey Porto and Steve Rolston who have patiently endured my ignorant questions (because there is no such thing as a stupid question, only ignorant ones), Britishisms, and a verbose and profanity-laden lab book. Their guidance has been invaluable and without it I would not be the physicist I am today, and for that I am truly gratefully.

Throughout the course of my time on the experiment there have been a number of other members of the rubidium Rydberg team who deserve a shout out. Dalia Ornelas Huerta has been on this project since day one. Her contributions to the experiment are far too numerous to reasonably list here, and I am truly in awe of patience and dedication she has shown throughout the years. During the early years of my grad school career our post-doc Mary Lyon served as an additional source of guidance, helping me develop my skills as an experimental physicist as well as teaching me many important laws of experimental physics (e.g the conservation of

effort - the effort required in finishing a task is approximately constant regardless of the method taken to complete it). Although their tenures in the lab were brief, the experiment still bears contributions from both Nathan Fredman and AJ Hachtel, to which I am grateful. I have to also thank recent additions to the team Patrick Banner and Deniz Kurdak for providing me with an opportunity to recall and explain the (sometimes wildly unintuitive) design and operational practices of the experiment. We are leaving the experiment in good hands, and I look forward to seeing what direction you go with it.

Outside of my lab, there are a number of collaborators to whom I am immensely grateful. Elizabeth Goldschmidt has probably forgotten more about quantum optics than I think I will ever know, and I have nothing but thanks for all that she has taught me on the subject. A major part of this thesis concerns a collaboration between the rubidium Rydberg group and the barium trapped ion group run by Qudsia Quirashi, James Siverns, and John Hannegan. Getting one experiment running on its own is tricky enough, and its a testament to them at how seamless it was to perform the joint experiment between our labs. I hope our two labs can continue to work together to do some incredible things. Thanks also need to made to our theory collaborators, Alexey Gorshkov, Przemek Bienias, Yidan Wang, Michael Gullans, and Jeremy Young who have always showed immense excitement and support for our work, and been willing to discuss potential ideas for projects (and debate sign and factor of two conventions). Although not directly affiliated with the work, I would be remiss if I did not thank members of the wider JQI and NIST community, who have provided invaluable advice, help and equipment

over the years, in no particular order: Gretchen Campbell, Paul Lett, Alan Migdall, Alessandro Restelli, Sergey Polyakov, Zachary Levine, Dan Barker, Ben Reschovsky, Neal Piseni, Sarthak Subhankar, Carlos Bracamontes, James Maslek, Tsz-Chun Tsui, Varun Vaidya, Thomas Boulier, Eric Magnan, Matt Reed, Zach Smith, Joe Tiamsuphat, Madison Anderson, Swarnav Banik, Peter Elgee, Monica Gutierrez Galan, Hirokazu Miyake, Ananya Sitaram, Hector Sosa Martínez, Hyok Sang Han, and many others. I also owe a huge thank you to the JQI and IREAP staff: Melissa Britton, Kelly Hickman, Taylor Prendergast, Nancy Boone, Jessica Crosby, Paulina Alejandro, Lorraine DeSalvo, Josiland Chambers, Ruth Heffes, Pauline Rirksoa. You all helped make our experiment a success, and to make my time at UMD as enjoyable as it was.

It is potentially a dangerous statement to make in a thesis read by your advisors, but sometimes you need to take a break from science. Through UMD I am privileged to have got to know Jon Vannucci, Rodney Snyder, Molly Snyder, Jimmy Juno, Anna Wright, Zach Eldredge, Katie Goff, Julia Sell, and Sarthak Subhankar. Slightly further afield, I have a great appreciation for the members of the bocce team, the CUA group, and my old school mates from Cowbridge who tolerate my ramblings about physics far longer than is socially necessary. Sanity is a precious thing and you all have some hand in helping me keep mine.

I owe a great deal of thanks to my family, my mum and dad, siblings Rosie and Jos, and their partners Tom and Gwennan. I would not be where I am today without your love, support, and the whole keeping me alive for the first eighteen years of my life thing. And thanks should also go to my in-laws Carleen, Ron and

Eric, who have always shown an interest in my work. I am truly lucky to have a family, and place to call home, on both sides of the Atlantic.

Most of all I have to say thank you to my wife, Clarissa, to whom this thesis is dedicated. Thank you for moving across an ocean so I could pursue my PhD. Thank you for always believing in me, and helping me believe in myself. Thank you for listening to me babble about physics when I was frustrated, and for celebrating with me in my successes. My years as a student are coming to a close, and I cannot wait to start the next chapter of my life with you.

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	vii
List of Tables	x
List of Figures	xi
List of Abbreviations	xv
1 Introduction	1
1.1 Quantum Networks	2
1.2 Thesis Outline and Statement on Contributions	11
2 Rydberg Atoms	14
2.1 Lifetime	14
2.2 Electric Field Sensitivity	18
2.3 Interactions	20
3 Rydberg Ensembles	26
3.1 Electromagnetically Induced Transparency	26
3.2 Collective Excitations	34
3.3 Rydberg Electromagnetically Induced Transparency	37
4 Experimental Setup and Techniques	40
4.1 Science Chamber	40
4.1.1 Chamber Design	40
4.1.2 In-Vacuum Components	44
4.1.3 Initial Vacuum Pump-Down Procedure	53
4.1.4 Magnetic Field Generation	55
4.1.5 Rubidium Oven	58
4.2 Laser Systems and Beam Paths	61
4.2.1 Magneto-Optical Trap and Optical Molasses	61
4.2.2 Optical Pumping	65
4.2.3 Optical Dipole Trap	68

4.2.4	Rydberg Excitation	74
4.2.5	General Laser System Information	84
4.3	Experimental Sequence	84
4.3.1	Magneto-Optical Trap	86
4.3.2	Optical Molasses	87
4.3.3	Optical Pumping	88
4.4	Computer Control	89
4.5	Diagnostic Tools	92
4.5.1	Spectroscopy	93
4.5.2	Probe Transmission Imaging	95
4.5.3	Absorption Imaging	96
4.6	Photonic Data Collection	97
4.7	Experimental Techniques	101
4.7.1	Spatial Alignment	101
4.7.2	Polarization Alignment	105
4.7.3	Temporal Synchronization	107
5	Measurement and Calculation of Photonic Correlation Functions	109
5.1	First Order Coherence	109
5.2	Second Order Coherence	110
5.2.1	Formal Definition	110
5.2.2	Classical Limit	110
5.2.3	Quantum Mechanically	111
5.2.4	Probabilistic definition	112
5.2.5	Measurement	113
5.3	Higher Order Coherences	115
5.4	Note on Losses	116
5.5	Note on Normalization	119
5.6	Calculation	122
5.6.1	Problem to be Solved	122
5.6.2	Methods	123
5.6.3	Computational implementation	129
6	On-demand, Pure and Indistinguishable Single Photon Source Based on a Rydberg Ensemble	140
6.1	General Method	141
6.2	Source Optimizations	144
6.2.1	Single Photon Purity	144
6.2.2	Source Efficiency	146
6.3	Source Characterization	151
6.3.1	Efficiency and Rate	151
6.3.2	Purity	154
6.3.3	Indistinguishability	157
6.4	Contaminants	159
6.5	Issues, Improvements and Outlook	167

7	Quantum Interference Between Photons from an Atomic Ensemble and a Remote Atomic Ion	172
7.1	Experimental Configuration	173
7.1.1	Barium Trapped Ion Details	174
7.1.2	Atomic Ensemble Details	177
7.1.3	HOM Interferometer	178
7.2	Interference for Stochastically Produced Photons	179
7.2.1	Trapped Ion Stochastic Photon Production	179
7.2.2	Atomic Ensemble Stochastic Photon Production	180
7.2.3	Stochastic Hong-Ou-Mandel Results	181
7.3	Interference for On-Demand Produced Photons	187
7.3.1	Trapped Ion On-Demand Photon Production	188
7.3.2	Atomic-Ensemble On-Demand Photon Production	189
7.3.3	Synchronization and Measurement Scheme	190
7.3.4	On-Demand Hong-Ou-Mandel Results	192
7.4	Outlook	204
8	Conclusion and Outlook	210
A	Detailed C_6 Coefficient Calculations	212
B	Calculation of Expected Coincidences for Self-Hong-Ou-Mandel Measurement	219
	Bibliography	224
	Bibliography	224

List of Tables

4.1	Low voltage electrode connections	47
4.2	Time tagger connections	98
4.3	Anatomy of time tagger high words	99
4.4	Anatomy of time tagger low words	99
5.1	Specifications for computer used for correlations code benchmarking .	130
6.1	List of the efficiencies along the post-cloud probe path.	152
6.2	Decay probability of the $139S_{1/2}$ state due to blackbody induced transitions	162
B.1	Outcomes for two photons in the same pulse	221
B.2	Outcomes for two photons in the adjacent pulses	221
B.3	Outcomes for two photons in the adjacent but one pulses	221

List of Figures

1	Legend for optics components	xvii
1.1	Example silica fiber attenuation curve	5
1.2	Two node DLCZ setup	7
1.3	Four node DLCZ setup	8
2.1	Zero temperature rubidium $nS_{1/2}$ state lifetimes	16
2.2	Calculated lifetime of rubidium $nS_{1/2}$ states at $T = 300$ K	17
2.3	Calculated polarizability of rubidium $nS_{1/2}$ states	19
2.4	Calculated C_6 coefficients of rubidium $nS_{1/2}nS_{1/2}$ states	23
2.5	Cartoon of Rydberg blockade effect for a pair of atoms	24
3.1	Energy level diagram for EIT in three-level ladder configuration	27
3.2	Theoretical EIT transmission curves	30
3.3	Theoretical EIT susceptibility curves	31
3.4	Spin-wave emission angular dependence for forward retrieval	35
3.5	Spin-wave emission angular dependence for backward retrieval	36
4.1	Render of science vacuum chamber	41
4.2	Render of cage washer for 2.75" viewports	43
4.3	Render of cage washers for 8" viewports	43
4.4	Image of the interior of the central chamber of the science vacuum apparatus	44
4.5	Annotated render of pinout for high voltage feedthrough	49
4.6	Image of one of the crosses during construction showing "false ceiling"	53
4.7	Cross-sectional render of 8" viewport with quadrupole coil mounted	56
4.8	Annotated image of chamber with quadrupole and bias coils mounted.	57
4.9	Atomic level diagram, with relevant experimental couplings	60
4.10	Beam schematic for MOT and molasses	62
4.11	Schematic of laser setup for MOT and optical molasses	63
4.12	Atomic level diagram showing laser scheme for the MOT and optical molasses	64
4.13	Beam schematic for optical pumping	66
4.14	Atomic level diagram showing laser scheme for optical pumping	67
4.15	Schematic of setup for D1-pump laser	68
4.16	Beam schematic for dipole trap	69

4.17	Atomic level diagram showing laser scheme for dipole trapping	71
4.18	Schematic of laser setup for dipole trapping	73
4.19	Beam schematic for Rydberg excitation	75
4.20	Image of ULE cavity vacuum chamber	77
4.21	Schematic of probe laser setup	78
4.22	Servo bump seen in probe lock error signal	79
4.23	Atomic level diagram showing laser scheme for Rydberg excitation . .	80
4.24	Schematic of control laser setup	82
4.25	Waveforms for various parameters during the experimental sequence .	85
4.26	Example of EIT transmission spectra	93
4.27	Schematic for generalized photonic data collection	97
4.28	Fraction of tags from time tagger with a temporal payload bit value equal to one	100
4.29	Images of atoms trapped in optical dipole and probe trap	104
4.30	Coincidence measurement used for temporal synchronization of SPADs	107
5.1	Hanbury Brown and Twiss setup	114
5.2	Generalized Hanbury Brown and Twiss setup	115
5.3	Example $g^{(2)}$ for coherent light calculated using different normaliza- tion methods	120
5.4	Benchmarks for the compute time as a function of threads for a multi-threaded implementation of the reduced double histogram coin- cidence calculation method	132
5.5	Benchmarks for the CPU and GPU based implementations of the reduced double histogram method	134
5.6	Benchmarks for the compute time as a function of threads for a multi- threaded implementation of the reduced per-photon histogram coin- cidence calculation method	136
5.7	Benchmarks for the Docker and native based implementations of the reduced per-photon histogram method	137
5.8	Comparison of the benchmarks for the optimized implementations of the two methods	138
6.1	Sequence for on-demand single photon generation	141
6.2	Level scheme for production of on-demand photons	142
6.3	Experimental setup for production and characterization of single pho- tons	143
6.4	$g^{(2)}$ dependence on cloud length	145
6.5	Normalized photon production efficiency as a function of the control field detuning during retrieval	150
6.6	Normalized photon production efficiency as a function of MOT load time/optical depth	151
6.7	Effect of duty cycle on photon production	152
6.8	Effect of production cycle time on photon production	153
6.9	On-demand single photon source purity measurement scheme	154

6.10	Gating window for on-demand single photon measurements	155
6.11	On-demand single photon source purity measurement	156
6.12	On-demand single photon source indistinguishability measurement scheme	157
6.13	On-demand single photon source indistinguishability measurement . .	158
6.14	Decay behavior in $g^{(2)}$ and the photon production efficiency in a pulse train	160
6.15	Normalized retrieval efficiency as a function of microwave Rabi flop- ping time	163
6.16	Blockade effect on further production as a function of microwave Rabi flopping time	164
6.17	Lifetime measurement of artificially created contaminant	165
6.18	Contaminant production probability as a function of the peak atomic density of the cloud	166
6.19	Contaminant probability as a function of the spin wave storage time .	167
6.20	Photon production probability as a function of the spin wave storage time	169
7.1	Joint experimental layout	174
7.2	Experimental layout and energy level diagram for barium ion system	175
7.3	Experimental layout and energy level diagram for atomic ensemble system	177
7.4	Experimental layout for HOM interferometer	178
7.5	$g^{(2)}(\tau)$ for stochastically produced photons from the ion source	179
7.6	$g^{(2)}(\tau)$ for stochastically produced photons from the atomic ensemble source	180
7.7	HOM measurement results for stochastically produced photons	181
7.8	Beamsplitter input and output operators	182
7.9	Photon profile and area-integrated $g^{(2)}(\tau)$ for on-demand photons pro- duced by the trapped ion source	188
7.10	Photon profile and area-integrated $g^{(2)}(\tau)$ for on-demand photons pro- duced by the atomic-ensemble source	189
7.11	Photon profiles for the two sources for one measurement period	192
7.12	Photon profiles and three-fold coincidences for on-demand photons . .	193
7.13	Method for determining temporally overlapped and non-overlapped coincidences	194
7.14	HOM measurement results for on-demand photons	195
7.15	Comparison to data of time resolved HOM theory, including account- ing for non-transform limited barium photon	197
7.16	Comparison to data of time resolved HOM theory, when accounting for constant center frequency difference between the photons from the two sources	200
7.17	Barium Photon Frequencies Including Zeeman Splittings.	201
7.18	Comparison to data of time resolved HOM theory, when accounting for the probabilistic spectrum of the ion-produced photon	202

7.19	Comparison to data of time resolved HOM theory, when accounting for relative drift of the center frequencies of the two photons	203
7.20	Comparison to data of time resolved HOM theory, when accounting for the combination of experimental imperfections	203
7.21	Setup to herald entanglement between distant matter qubits	206
7.22	Inferred entanglement fidelity and generation rate as a function of the coincidence bin width	207
7.23	Theoretical entanglement fidelity and rate speedup as a function of bin width assuming imperfection free interference	208
A.1	Angular dependence of C_6 coefficient for different $ r, r\rangle$ pair states	217
B.1	Schematic for self-HOM measurement	219

List of Abbreviations

ABS	Acrylonitrile butadiene styrene
AOM	Acousto-optic modulator
API	Application programming interface
AR	Anti-reflection
AWG	American wire gauge
BB84	Bennet and Brassard 1984
BBM92	Bennet, Brassard and Mermin 1992
BS	Beamsplitter
c.c.	Complex conjugate
CCD	Charge-coupled device
CMOT	Compressed magneto-optical trap
CPU	Central processing unit
CF	Conflat
CUDA	Compute unified device architecture
DC	Direct current
DDS	Direct digital synthesizer
DFG	Difference frequency generation
DLCZ	Duan, Lukin, Cirac and Zoller
DLL	Dynamic-link library
E91	Eckert 1991
EIT	Electromagnetically induced transparency
EMCCD	Electron multiplying charge-coupled device
EOM	Electro-optic modulator
FIFO	First in first out
FSR	Free spectral range
GPU	Graphical processing unit
GUI	Graphical user interface
HDF5	Hierarchical data format version 5
HOM	Hong-Ou-Mandel
HWP	Half-wave plate
JQI	Joint quantum institute
LVTTL	Low-voltage transistor-transistor logic
MCP	Multi-channel plate detector
MOT	Magneto-optical trap

NA	Numerical aperture
OD	Optical depth
PBS	Polarizing beam splitter
PDH	Pound-Drever-Hall
PEEK	Polyether ether ketone
PID	Proportional, integral, derivative
PLL	Phase locked loop
PMF	Polarization maintaining fiber
PPLN	Periodically poled lithium niobate
QUESS	Quantum experiments at Space Scale
RAM	Random access memory
RF	Radio frequency
RGA	Residual gas analyzer
RMS	Root mean square
SHG	Second harmonic generation
SMF	Single mode fiber
SNSPD	Superconducting nanowire single photon detector
SPAD	Single photon avalanche detector
SPD	Single photon detector
TA	Tapered amplifier
TTL	Transistor-transistor logic
ULE	Ultra-low expansion
ULEC	Ultra-low expansion cavity
VBG	Volume Bragg grating




Figure 1: Legend for optics components. Some of these symbols are taken or adapted from Alexander Franzen's ComponentLibrary (version 3, available at <http://www.gwoptics.org/ComponentLibrary/>), which is licensed under the Creative Commons Attribution-NonCommercial 3.0 Unported License (<https://creativecommons.org/licenses/by-nc/3.0/>). The remainder were designed by Dalia Ornelas-Huerta.

Chapter 1: Introduction

Since its discovery over a century ago, quantum mechanics has become a mainstay of physics research. However, the emergence of the field of quantum information, which seeks to marry concepts from physics and computer science, is relatively recent. During the 1980s the concept of quantum computers, that is computation performed with quantum information, started to be discussed by academics such as Richard Feynman [1] and David Deutsch [2]. Then, in the 1990s, there was an explosion of interest in the field following the discovery of algorithms such as that of Deutsch-Jozsa [3], Shor [4] and Grover [5], as well as error correction schemes [6, 7] that suggested the potential for enormous error-free quantum speedups relative to classical computation. Parallel to this came the subfield of quantum communication, concerned with the distribution of quantum information. Here, protocols such as BB84 [8], E91 [9] and BBM92 [10] promised communications secured by the laws of nature, while others, like superdense coding [11] and quantum teleportation [12], found novel ways to use quantum entanglement as a resource. Despite this interest, practical physical realizations of both quantum computation and communication are still lacking. For computation, the group at Google has, controversially, claimed a demonstration of quantum supremacy [13], but large scale universal quan-

tum computers remain elusive. At the forefront of quantum communications, the QUESS experiment has managed to demonstrate entanglement over an enormous distance [14]. However, as a point of comparison, secure key distribution rates based on this technology still pale in comparison to their classical counterparts [15].

1.1 Quantum Networks

Having run through a brief history of the field, and before we get too ahead of ourselves, it is perhaps a good idea to define what we mean when we say ‘quantum information’. Much like the ‘bit’ is the fundamental unit of classical information, the basic unit of quantum information is the ‘qubit’. While the bit can take either the value of 0 or 1, the qubit we instead describe by a quantum mechanical wavefunction

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \tag{1.1}$$

where the coefficients α and β are complex numbers that satisfy $|\alpha|^2 + |\beta|^2 = 1$. If we measure our qubit, we will find it to be in state $|0\rangle$ with a probability $P_{|0\rangle} = |\alpha|^2$, and state $|1\rangle$ with a probability $P_{|1\rangle} = |\beta|^2$. Given that, one might be forgiven for thinking that qubits are simply probabilistic classical bits. However, there is a major difference due to the complex nature of the coefficients, which can lead to interferometric phenomenon that would not be seen with probabilistic classical bits.

While the qubit may seem like an abstract concept, it actually has a plethora of natural physical realizations [16]. For example, a spin 1/2 particle, such as an electron, has a pair of potential outcomes when the spin is projected onto the z-

axis, either spin-up or spin-down, which can be thought of as our $|0\rangle$ and $|1\rangle$ states. Another physical realization that shall be of interest in this thesis is that of the photon, where we can use any of a number of the photon properties to encode the qubit state, such as polarization, temporal mode, frequency, transverse spatial mode etc. Given their speed of propagation, and their robustness against decoherence due to weak environmental interactions, flying photons are a natural qubit for quantum information transfer [17].

While the qubit is the fundamental unit of quantum information, it is perhaps also important to touch on the idea of entanglement, a fundamental resource within quantum information. We consider a pair of two or more qubits to be entangled when we are unable to write their wavefunctions as a tensor product of their individual wavefunctions. At first glance this results in behavior that may resemble something like a classical correlation, however, it is important to note that entanglement is a purely quantum mechanical phenomenon. A classic example (if you will excuse the pun) of entanglement is that of the so-called Bell states [16]

$$\begin{aligned} |\Phi_{\pm}\rangle &= \frac{|0\rangle \otimes |0\rangle \pm |1\rangle \otimes |1\rangle}{\sqrt{2}} \\ |\Psi_{\pm}\rangle &= \frac{|0\rangle \otimes |1\rangle \pm |1\rangle \otimes |0\rangle}{\sqrt{2}}, \end{aligned} \tag{1.2}$$

which are considered ‘maximally entangled’, and are critical in many quantum networking applications.

Having defined what we mean about quantum information we can now start to talk about quantum networks. Like its counterpart, a classical network, which

allows the transfer of classical bits between network nodes, a quantum network facilitates the transfer of qubits and entanglement between the quantum nodes of its network. Such a network has a number of critical uses. Most obvious of these would be for quantum communication, such as the quantum key distribution protocols like BB84 [8], E91 [9] or BBM92 [10], which require the distribution of qubits or entangled pairs of qubits between remote nodes that wish to share a one-time pad. Additionally, quantum networks are also likely to have an important role to play as quantum computers becomes more mature and need to start communicating with one another. Indeed, some proposals [18,19] envision large scale quantum computing achieved using a modular approach where small quantum computer nodes are networked together, similar to classical cluster computing. A further use for quantum networks relates to quantum sensing. If entanglement can be generated between a number of network nodes, we then have the ability to perform measurements with a quantum-enhanced sensitivity [20–22].

The construction of large scale quantum networks is, however, challenging, primarily due to the fact that quantum information is inherently fragile. The standard terrestrial method of transporting photonic qubits is along optical fibers. For the typical silica fibers in current fiber networks, the minimum absorption is seen in the telecom bands, as we can see in the example¹ in figure 1.1, which run from around 1250-nm to 1700-nm. However, even in the telecom bands the attenuation is typically² on the order of 0.2 dB/km. This means that a distance of only 15 km

¹taken from https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=351

²The example fiber in the figure was chosen as it is a broadband multimode fiber. Lower attenuation can be achieved with single mode fiber designed at a specific wavelength, see https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=1362 as an example.




Figure 1.1: Example silica fiber attenuation curve.

is enough to bring the probability of successful transmission of a photonic qubits to $1/2$. Furthermore, despite photons being relatively robust to decoherence, they are not immune from it with in-fiber effects, such as birefringence, potentially causing issues [23, 24]. Similarly, satellite based approaches to quantum network construction are not free from these issues, with turbulence and absorption in the lower atmosphere leading to both decoherence and loss of flying qubits [14, 15, 25].

Classically, the problems of networking are solved with the idea of amplification and repetition. Light pulses encoding the bits are amplified up and propagated some distance, where they are again amplified, and so on across the network. However, this is not possible for qubits in quantum networks due to the no-cloning theorem [26], a no-go theorem of quantum mechanics that prohibits us from making copies of an arbitrary quantum state.

But, where quantum mechanics closes the door of cloning, another one opens in the form of quantum teleportation [12]. Here, we assume that our two communicating parties, Alice and Bob, have a shared Bell state pair, with each possessing one of the qubits. If Alice wants to send a particular qubit to Bob she can perform a measurement in the Bell-state basis on her two qubits, both the one that is to be sent, and the one that formed part of the entangled pair. The outcome of this measurement projects Bob's qubit into some state. But Bob is able to recover the qubit that Alice wanted to send to him by performing a local operation that depends on the outcome of Alice's measurement, which can be communicated classically. So provided we can generate this entanglement between the pair, we can transmit quantum information using only classical communication. What we have essentially done is converted the problem of sending qubits down lossy and/or noisy channels to a problem of distributing entanglement across the channels. Now this may not sound like a particularly helpful thing to do. After all to have Alice and Bob share some entanglement we will certainly have to send some qubits down the channels. However, the key here is that as we cannot clone the qubit, we essentially get one shot at the transmission from Alice to Bob. But for the teleportation method we can keep trying over and over to generate the entanglement until we are successful, and only then will we attempt the teleportation protocol. This can have a dramatic increase in the transfer fidelity over the brute force method of simply trying to send the qubit through a lossy and/or noisy channel..

This, then, brings us on to discussing how to generate entanglement between two remote and potentially distant parties. Quantum repeaters are an idea that shot




Figure 1.2: Two node DLCZ setup. Relevant atomic level structure shown on left, and experimental system on right.

to prominence with the discovery of what is now known as the DLCZ protocol [27]. The idea behind the DLCZ protocol is that Alice and Bob each have an atomic ensemble, where the atoms are assumed to have a three-level λ like structure, like that shown in figure 1.2. For a single ensemble, weak excitation of one of the transitions, which we shall call the $|1\rangle$ to $|2\rangle$ transition, has the potential to lead to spontaneous decay from the $|2\rangle$ to $|3\rangle$ state emitting a photon. On the detection of a spontaneously scattered photon, as it is unknown which atom in the ensemble decayed, the ensemble is projected into a collective state known as a spin-wave, a symmetric superposition with one atom in the $|3\rangle$ state and the rest in the $|1\rangle$ state. Now, to entangle two separate ensembles we take the spontaneously decayed photons from both systems and direct them towards a beamsplitter, with a pair of single photon detectors at the two beamsplitter outputs, as shown in figure 1.2. The detection of a single photon projects the two ensembles into an entangled state where one has an excitation and the other does not and vice versa, as we are unable to determine which ensemble the photon was emitted from.

To entangle distant parties, the DLCZ protocol takes things a step further by




Figure 1.3: Four node DLCZ setup. Here, black lines joining ensembles indicate entanglement, and the black arrowed lines indicate the entanglement swapping procedure.

envisioning a series of nodes between Alice and Bob, with adjacent nodes located within some characteristic attenuation length of each other. Each node is assumed to have an ensemble and performs the entanglement generation protocol described above, so that pairs of nodes along the chain share entanglement. Next, entanglement swapping is performed. This works by taking a pair of adjacent nodes, which are not presently entangled with one another, and retrieving their stored excitations as photons, which are directed to another beamsplitter-single photon detector setup. The detection of a single photon here heralds the ‘swapping’ of the entanglement on to the nodes which are next-next-nearest-neighbors, as again we are unable to determine from which ensemble the photon originated. This entanglement swapping procedure can then be repeated to entangle more and more distant nodes, until finally Alice and Bob have a entangled pair of ensembles. A four node example of this procedure is shown in figure 1.3.

This method of distributing entanglement has an advantage, compared with the brute force method of locally generating an entangled pair of photonic qubits and directly transmitting one of them, as photons are only ever required to travel the distance between nearest-neighbor nodes. However, there are also overheads associated with the protocol. Using some reasonable estimates, it can be shown [28] that the DLCZ repeater approach becomes beneficial even for relatively modest distances, ≈ 500 km, and far outperforms it for greater length scales.

While the DLCZ protocol for quantum repeaters has regimes where it has considerable gains over direct transmission, it is not without its drawbacks [28]. For one, the ensemble excitation must be performed weakly, so that the probability

of multiple photon emissions are suppressed, which inevitably means working in a regime where a single photon emission event is unlikely. This in turn means that the per-attempt likelihood of generating entanglement between a pair of ensembles is also small, which will limit how quickly we can generate entanglement between Alice and Bob. Another, potentially more problematic issue, is the requirement on stability between the channels linking the nodes to the beamsplitter-photodetection setups. A differential phase accumulated along the photon paths from the nodes to the beamsplitter will change the entangled state created upon photon detection. The differential phase accumulation, therefore, needs to be very well controlled, which usually means controlling or compensating for path lengths to much better than an optical wavelength, which is extremely challenging in practice.

As a result there exist a number of subsequent quantum repeater schemes that have built upon the DLCZ protocol, resolving some of its issues. For the differential phase issue, the use of coincidence detection of photon pairs from the two nodes [29, 30], rather than single photon detection, relaxes the experimental stability requirements. The problem of low probability photon emission is addressed by using (semi) deterministic single-photon production methods [31–33], which leads to reductions in entanglement generation times.

We turn now to discussing the practical side of constructing quantum networks. As so often happens in physics, the implementation has somewhat lagged behind theory. At the forefront of the field, in terms of generating entanglement over long distances, is the satellite based QUESS project [14], mentioned at the outset. They have been able to demonstrate reasonable fidelity entanglement, ≈ 0.87 , over

a distance just over 1200 km, using direct transmission of entangled photon pairs to two ground stations from a spontaneous-parametric-downconversion source. However, the entanglement generation rate of $\approx 1.1 \text{ s}^{-1}$ is likely too low for practical use. Terrestrially, much higher entanglement generation rates, $> 100 \text{ s}^{-1}$, have been demonstrated between pairs of trapped ions [34], and similarly with quantum dots [35], albeit at significantly smaller distances of a few meters. Entanglement generation over large distances and at reasonable rates, therefore, remains an ongoing experimental goal.

Before concluding our discussion on the present state of quantum networking, it is important to point out that much of the work in the field has been focused on the entangling of homogeneous systems, that is systems that inherently possess near-identical properties. However, given the plethora of different quantum platforms in the broad field of quantum information, it is likely that future quantum networks will need to have an architecture featuring heterogeneous nodes [36–38]. Therefore, it is important that we not only continue to improve the practical implementation of homogeneous quantum networks, but start to develop hybrid quantum networks which can leverage the benefits of the differing constituent nodes within the network.

1.2 Thesis Outline and Statement on Contributions

In this thesis I describe the construction of an apparatus for the trapping and probing of a Rydberg ensemble, and a pair of projects that relate to its use for quantum networking. In this chapter I have provided a brief introduction to the field of

quantum information, and quantum networking. In chapters 2 and 3 I derive some experimentally relevant properties of Rydberg atoms and Rydberg ensembles respectively and discuss their implications for designing a Rydberg ensemble experiment. Chapter 4 contains the description of our experimental apparatus, its construction, operation, and a few techniques developed to aid in its modification and maintenance. This is followed by an, admittedly verbose, but hopefully useful, discussion of photonic correlations and the software I developed for their calculation from time tagged data in chapter 5. In chapter 6, I discuss our development of a Rydberg ensemble based high-efficiency single photon source, work which was recently published [39]. Finally in chapter 7, based on the published work in reference [40], I will discuss a project where we, along with our collaborators at Army Research Lab, demonstrated interference between photons generated by our Rydberg ensemble, and their remotely located trapped ion.

Experimental physics is rarely a solo effort, and the projects I have worked on have been no exception. Therefore, I want to acknowledge the people who made the various chapters possible, and try to highlight some specific things that I have worked on. Starting with the construction of the apparatus, discussed in chapter 4, this endeavor was primarily undertaken by myself, Dalia Ornelas-Huerta, and Mary Lyon, with honorable mention to Nathan Fredman. Here, I focused on the in-vacuum electronics, the magnetic field generation, the current rubidium oven and optical pumping scheme, janky microwave excitation electronics, computer control, and data collection.

The correlation software we use, which is the main topic of chapter 5, has

been a labor of love of mine for over three years now. However, it would not be the monstrosity it is today without discussions with Dalia, Mary, Steve, and Trey from my group in addition to Sergey Polyakov, Zachary Levine, Elizabeth Goldschmidt and James Juno.

The experimental work presented on the on-demand single-photon source in chapter 6 was undertaken by Dalia, AJ Hachtel and myself. Theory support was provided by Yidan Wang, Przemek Bienias and Alexey Gorshkov. Here, my efforts were primarily focused in the identification and (unsuccessful) elimination of the contaminant states, in addition to some of the early theory work in understanding our source efficiencies.

As regards the collaborative work discussed in chapter 7, the trapped ion apparatus, which included the all-important frequency conversion setup, is entirely the work of John Hannegan and James Siverns. The Rydberg ensemble side of things was again operated by Dalia, AJ and myself. For this work I was primarily responsible for the analysis of the data, including the background subtraction, and the theory for both the stochastic and on-demand cases. A special mention should go to John Hannegan for his work on determining the messy photon spectrum of the barium ion, and for his calculations on projected entanglement fidelities and rates of our hybrid setup.

Chapter 2: Rydberg Atoms

The term ‘Rydberg atom’ refers to an atom in which one or more electrons has been promoted to a high principal quantum number, n . In these highly excited states the electron(s) lie far from the atomic nucleus, which causes the Rydberg atoms to exhibit various interesting properties. This has led to an explosion in recent years in the study of Rydberg physics owing to its wide range of applicability in emerging fields such as quantum optics [41–45], quantum simulation [46–49], digital quantum computation [50–54] and quantum-enhanced sensing [55, 56]. As a result, there presently exists a large body of literature, including multiple review articles [48, 57, 58] and books [59], that discuss at length the properties of Rydberg atoms. Therefore, rather than a comprehensive review, in what follows I shall discuss a subsection of those properties which are relevant to the experiments contained within this thesis.

2.1 Lifetime

The zero-temperature radiative lifetime, τ , of a state, $|i\rangle$, can be calculated

$$\tau_i = \left(\sum_f \Gamma_{f,i} \right)^{-1}, \quad (2.1)$$

where $\Gamma_{f,i}$ is the spontaneous emission rate from state $|i\rangle$ to $|f\rangle$, and the sum is taken over all states to which $|i\rangle$ can spontaneously decay. We can write the emission rate as [60]

$$\Gamma_{f,i} = \frac{\omega_{f,i}^3}{3\pi\epsilon_0\hbar c^3} \left| \langle f | \hat{d} | i \rangle \right|^2 \quad (2.2)$$

where $\omega_{f,i} = \omega_f - \omega_i$ is the transition frequency, and \hat{d} is the dipole operator. Although the dipole matrix elements, $\left| \langle f | \hat{d} | i \rangle \right|$, are significantly larger between states where $n_f \approx n_i$, $\omega_{f,i}$ is significantly smaller for these transitions. As a result, for the low angular momentum states of alkali atoms, the sum in equation 2.1 is dominated by transitions to states with low principal quantum number, where $\omega_{f,i}$ is large. For an alkali atom we can use the Rydberg formula for the binding energy to find

$$\omega_{f,i} = \frac{R_y}{\hbar} \left(\frac{1}{(n_f^*)^2} - \frac{1}{(n_i^*)^2} \right), \quad (2.3)$$

where $n_{(i,f)}^* = n_{(i,f)} - \delta_{(i,f)}$ is the quantum-defect-modified principal quantum number for the state $|i\rangle$ or $|f\rangle$, and R_y is the mass-corrected Rydberg constant. For transitions where $n_i \gg n_f$ we can make the approximation that $\omega_{f,i}$ is roughly independent of n_i , while the matrix elements exhibit a scaling [59]

$$\left| \langle f | \hat{d} | i \rangle \right| \propto (n_i^*)^{-3/2}. \quad (2.4)$$

This leads to the following scaling relationship for the lifetime of state $|i\rangle$,

$$\tau_i \propto (n_i^*)^3. \quad (2.5)$$




Figure 2.1: Calculated [61] zero-temperature lifetime of rubidium $nS_{1/2}$ states. Orange line is cubic fit of form $\tau = a(n^*)^3$, with a fitted value of $a = 1.257$ ns.

This scaling means that the Rydberg state lifetime increases rapidly as a function of the principal quantum number. We can see this in figure 2.1, which shows the calculated values [61] for the zero-temperature radiative lifetime of the $nS_{1/2}$ states of rubidium.

Until now we have been ignoring any effect of finite temperature on the Rydberg lifetime. However, this is not a luxury we can experimentally afford due to the interaction of blackbody photons with the atoms. To take account of this interaction we must modify equation 2.1

$$\tau_i = \left(\sum_f \Gamma_{f,i} + \Gamma_{f,i}^{BB} \right)^{-1} = \left(\frac{1}{\tau_0} + \frac{1}{\tau_{BB}} \right)^{-1}, \quad (2.6)$$

where $\Gamma_{f,i}^{BB}$ quantifies the transition rate from $|i\rangle$ to $|f\rangle$ due stimulated emission or




Figure 2.2: Calculated [61] lifetime of rubidium $nS_{1/2}$ states at $T = 300$ K. Orange line is fit of form $\tau = \left(\frac{1}{a(n^*)^3} + \frac{1}{b(n^*)^2} \right)^{-1}$, with fitted values of $a = 1.257$ ns and $b = 53.92$ ns.

absorption of blackbody radiation, τ_0 is the zero-temperature lifetime, and τ_{BB} is the blackbody induced lifetime. It can be shown [59] that, for a near-room temperature environment, the blackbody decay rate scales as $\Gamma_{f,i}^{BB} \propto (n_i^*)^{-2}$, hence the blackbody induced lifetime scales as $\tau_{BB} \propto (n_i^*)^2$. The effect of blackbody radiation can be significant in shortening the Rydberg state lifetime, as can be seen in figure 2.2. Nevertheless, even with this reduction, these states are often long lived enough that they can be considered metastable for the purposes of applications.

2.2 Electric Field Sensitivity

An atom in an electric field, \vec{F} , is subject to the Hamiltonian

$$\hat{H} = -\hat{d} \cdot \vec{F}, \quad (2.7)$$

where \hat{d} is again the dipole operator. Treating this perturbatively for a state $|n, l, j, m_j\rangle$, the lowest order correction to the energy is the second order, or quadratic Stark, term

$$\Delta E_{n,l,j,m_j}^{(2)} = -\frac{1}{2}\alpha|F|^2 = |F|^2 \sum_{n',l',j',m'_j} \frac{|\langle n', l', j', m'_j | \hat{d} | n, l, j, m_j \rangle|^2}{E_{n',l',j',m'_j}^{(0)} - E_{n,l,j,m_j}^{(0)}}, \quad (2.8)$$

where we have defined the polarizability, α , to describe the state's electric field sensitivity. The scaling of the terms of the fraction means that the dominant contribution to the sum comes from nearby states. The denominator of the fraction

$$E_{n',l',j',m'_j}^{(0)} - E_{n,l,j,m_j}^{(0)} = R_y \left(\frac{1}{(n^*)^2} - \frac{1}{(n'^*)^2} \right) \quad (2.9)$$

can be expanded in n'^* about n^*

$$E_{n',l',j',m'_j}^{(0)} - E_{n,l,j,m_j}^{(0)} \approx \frac{2R_y}{(n^*)^3} (n'^* - n^*). \quad (2.10)$$

Given that we are only interested in terms where $n'^* \approx n^*$, then $n'^* - n^*$ is going




Figure 2.3: Calculated [61] polarizability of rubidium $nS_{1/2}$ states. Orange line is fit of form $\alpha = a(n^*)^7$, with fitted value $a = 72.91 \mu\text{Hz (V/cm)}^{-2}$.

to be roughly independent of n . Hence,

$$E_{n',l',j',m'_j}^{(0)} - E_{n,l,j,m_j}^{(0)} \propto \frac{1}{(n^*)^3}. \quad (2.11)$$

The matrix element between two nearby Rydberg states scales as $(n^*)^2$ [59]. We can therefore see that the Rydberg state polarizability is going to scale as

$$\alpha \propto (n^*)^7. \quad (2.12)$$

This rapid scaling of the polarizability in principal quantum number means that high-lying Rydberg states are extremely sensitive to electric fields. Figure 2.3 shows the calculated [61] polarizability of the $nS_{1/2}$ Rydberg states for rubidium. As can

be seen, for $n > 100$, the polarizability becomes such that electric fields of only a few mV/cm can cause shifts comparable to the inverse lifetime of the Rydberg state. Therefore, fine control of the electric field is usually required to prevent unwanted shifts in the Rydberg energy.

2.3 Interactions

In the absence of an electric field Rydberg states do not possess a permanent electric dipole. However, dipole-dipole interactions between a pair of Rydberg atoms still occur due to the transition dipole moments of the atomic states. An intuitive way to understand this phenomenon is to consider a pair of atoms. The first atom, with an oscillating transition dipole moment, radiates an electric dipole field at the transition frequency, ω . This electric field can then interact with the second atom. This results in a potential of the form¹

$$\hat{V}_{dd} = \frac{1}{4\pi\epsilon_0 R^3} \left[\hat{d}_1 \cdot \hat{d}_2 - 3(\hat{d}_1 \cdot \vec{n})(\hat{d}_2 \cdot \vec{n}) \right], \quad (2.13)$$

where \hat{d}_i is the dipole operator for atom i , R is the inter-atomic distance, and \vec{n} is the unit vector that points along the inter-atomic axis. The potential couples pair states and leads to an effective interaction between the two atoms. This is most straightforwardly seen by first considering a Hilbert space consisting of two pair states, $|r_1, r_2\rangle$ and $|r'_1, r'_2\rangle$, where the subscript labels the atom in the pair. The

¹Note that this is strictly speaking valid only when the inter-atomic distance is much larger than the spatial extent of the electronic wavefunctions of the two atoms

Hamiltonian for the system, in matrix form in the pair states basis, is

$$\hat{H} = \begin{pmatrix} 0 & \frac{C_3}{R^3} \\ \frac{C_3}{R^3} & -\Delta \end{pmatrix}, \quad (2.14)$$

where $C_3/R^3 = \langle r_1, r_2 | \hat{V}_{dd} | r'_1, r'_2 \rangle$ and $\Delta = E_{r_1} + E_{r_2} - E_{r'_1} - E_{r'_2}$. The eigenenergies of the Hamiltonian are

$$E_{\pm} = -\frac{\Delta}{2} \pm \frac{1}{2} \sqrt{4 \left(\frac{C_3}{R^3} \right)^2 + \Delta^2}. \quad (2.15)$$

In the limit that $R \gg \sqrt[3]{C_3/\Delta}$ the eigenstates have barely any admixture, $|\psi_+\rangle \approx |r_1, r_2\rangle$ and $|\psi_-\rangle \approx |r'_1, r'_2\rangle$, and the eigenenergies have the form

$$E_+ \approx \frac{C_3^2}{R^6 \Delta} = \frac{C_6}{R^6} \quad (2.16)$$

$$E_- \approx -\Delta - \frac{C_3^2}{R^6 \Delta} = -\Delta - \frac{C_6}{R^6}, \quad (2.17)$$

with the atoms exhibiting van der Waals interactions with an interaction strength defined by $C_6 = C_3^2/\Delta$. In the opposite limit where $R \ll \sqrt[3]{C_3/\Delta}$, the states heavily admix and the eigenenergies take the form

$$E_{\pm} \approx \pm \frac{C_3}{R^3}, \quad (2.18)$$

i.e. of resonant dipole-dipole form.

There are in general a multitude of pair states for which coupling is non-

negligible and an exact calculation can become impractical. Instead, to calculate the interaction strength for a pair of atoms in Rydberg states, we tend to treat the interaction potential perturbatively. For this thesis we shall be concerned with interactions between atoms in the same Rydberg state, so let us calculate the lowest order energy shift to a state $|r, r\rangle$, which is given by

$$\Delta E_{r,r}^{(2)} = \sum_{|r',r''\rangle \neq |r,r\rangle} \frac{|\langle r, r | V_{dd} | r', r'' \rangle|^2}{2E_r^{(0)} - E_{r'}^{(0)} - E_{r''}^{(0)}} = \frac{C_6}{R^6}. \quad (2.19)$$

While the sum is taken to be over infinitely many pair states, it is in general dominated by only a few states for which both the energy defect, $2E_r^{(0)} - E_{r'}^{(0)} - E_{r''}^{(0)}$, is small and the magnitude of the dipole-dipole matrix element, $|\langle r, r | V_{dd} | r', r'' \rangle|$, is large. This occurs for states $n \approx n', n''$, where $n^{('')}$ is the principal number for state $r^{('')}$. In this limit the energy defect, following from equation 2.10, scales as $2E_r^{(0)} - E_{r'}^{(0)} - E_{r''}^{(0)} \propto (n^*)^{-3}$. While the exact form of the dipole-dipole matrix elements is somewhat complicated, and is covered in more detail in appendix A, from inspection we can see that it is going to have a dependence like

$$\langle r, r | V_{dd} | r', r'' \rangle \sim \frac{\langle n, l, j, m_j | \hat{d}_1 | n', l', j', m'_j \rangle \langle n, l, j, m_j | \hat{d}_2 | n'', l'', j'', m''_j \rangle}{4\pi\epsilon_0 R^3}, \quad (2.20)$$

where we have again used the primes to indicate the principal quantum, orbital angular, total angular and magnetic quantum numbers for the corresponding r 's. Again considering only low angular momentum states, for which the dipole matrix elements scale as $(n^*)^2$ [59], then we expect the dipole-dipole matrix element to scale




Figure 2.4: Calculated [61] C_6 coefficients of rubidium $nS_{1/2}nS_{1/2}$ states. Orange line is fit of form $\alpha = a(n^*)^{11}$, with fitted value $a = 8.58 \text{ nHz } \mu\text{m}^6$.

as $(n^*)^4$. The scaling of the van der Waals coefficient is thus

$$C_6 \propto (n^*)^{11}. \quad (2.21)$$

This scaling can be seen in figure 2.4, which shows the calculated [61] values for the rubidium $|nS_{1/2}, nS_{1/2}\rangle$ C_6 coefficients. As demonstrated, the rapid scaling of the C_6 coefficient with the principal quantum number means that large energy shifts can be seen over macroscopic distances. For example a pair of $|100S_{1/2}\rangle$ atoms separated by $10 \mu\text{m}$ experience an energy shift on the order of 50 MHz. A point of note here is that even though the scaling of the C_6 coefficient with n^* is much stronger than that for the resonant dipole-dipole interactions, $C_3 \propto (n^*)^4$, it is still a second-order effect. As with the case where we only considered two pair states,




Figure 2.5: Cartoon of Rydberg blockade effect for a pair of atoms, alternately excited from the ground to the Rydberg state.

when the two atoms are sufficiently close together the interactions become resonant dipole-dipole like with a $1/R^3$ form.

It is predominantly these large, long-ranged interactions that make Rydberg atoms well suited for quantum applications. An example of its usefulness can be seen by considering what happens when we attempt to excite a pair of atoms to a Rydberg state. For ease of discussion let us assume that we can individually address the two atoms with two light fields, denoted by Rabi frequencies $\Omega_{1,2}$, that couple some ground state, $|g\rangle$, to a Rydberg state, $|r\rangle$. We can take our first atom and excite it to the Rydberg state without issue. With our first atom in the excited state, we then attempt to excite the second atom to the excited state. If the atoms are far from each other, no interactions occur and the excitation of the second atom to the Rydberg state proceeds without issue. However, if the atoms are closer than what is known as the blockade radius, r_B , the interactions change the $|r, r\rangle$ pair state energy

such that the excitation field is no longer resonant with the $|r, g\rangle$ to $|r, r\rangle$ transition, as shown in figure 2.5. This is the principle of so-called Rydberg blockade, which is the underlying principle behind a vast majority of the applications that involve Rydberg atoms [41–54]. It is important to realize that the concept does not just hold simply for pairs of atoms, but we could also envisage an ensemble of atoms all within a blockade radius of one another. Here, when we try and excite the atoms in the ensemble to a Rydberg state, we will only be able to get a single Rydberg excitation within the entire ensemble. This will be key when we discuss the optical non-linearities associated with a Rydberg ensemble in the rest of this thesis.

Chapter 3: Rydberg Ensembles

Having discussed Rydberg atoms, we can now talk about Rydberg ensembles. Quite simply, these are atomic ensembles where we have promoted one or more atoms to a Rydberg state. Here, the ensemble nature of the system provides us with good light-matter coupling [62], while the Rydberg component provides us with interactions, discussed in the previous chapter. Such systems have already been shown to be able to produce novel photon-photon interactions [41–43,63], in addition to being useful tools for the creation of non-classical light [41,64]. Meanwhile, theory proposals have been put forward for a number of further applications, including Rydberg-ensemble based quantum repeaters [31,32] for use in quantum networks.

There is a body of literature on the topic of Rydberg ensembles, with some comprehensive review articles [57,65]. Therefore, here I shall discuss only a few of the properties of Rydberg ensembles that inform the experimental design, and experiments undertaken in the latter chapters of the thesis.

3.1 Electromagnetically Induced Transparency

Let us start by considering a three-level atom with a ladder configuration, like the one shown in figure 3.1. A weak probe field couples the ground, $|g\rangle$, and




Figure 3.1: Energy level diagram for EIT in three-level ladder configuration.

intermediate excited state, $|e\rangle$, while a strong control field couples $|e\rangle$ and a Rydberg state, $|r\rangle$. For now, we will not worry too much about the Rydberg nature of state $|r\rangle$, other than acknowledging that it is likely to have a relatively long lifetime compared with the intermediate state, given what we saw in chapter 2. With the driving fields as shown in the figure, under the rotating wave approximation we have a Hamiltonian

$$\hat{H} = \hbar \begin{pmatrix} 0 & \frac{\Omega_p^*}{2} & 0 \\ \frac{\Omega_p}{2} & -\Delta_p & \frac{\Omega_c^*}{2} \\ 0 & \frac{\Omega_c}{2} & -(\Delta_p + \Delta_c) \end{pmatrix}. \quad (3.1)$$

We will assume that both the intermediate excited state and the Rydberg state have finite lifetimes, given by $1/\Gamma$ and $1/\gamma$ respectively, and we shall take $|g\rangle$ to be infinitely long lived. The evolution of the density matrix of the system, $\hat{\rho}$, is given

by the Linblad equation

$$\dot{\hat{\rho}} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}] + \sum_i \hat{L}_i \hat{\rho} \hat{L}_i^\dagger - \frac{1}{2} \{ \hat{L}_i^\dagger \hat{L}_i, \hat{\rho} \}, \quad (3.2)$$

where the $[\]$ brackets are taken to be the commutator, the $\{ \}$ brackets are taken to be the anti-commutator, and the \hat{L}_i terms are the so-called ‘jump’ operators. For this system there are two jump operators we need to be concerned with, that corresponding to decay from $|e\rangle$ to $|g\rangle$, given by $\hat{L}_1 = \sqrt{\Gamma} |g\rangle \langle e|$, and that corresponding to decay from $|r\rangle$ to $|g\rangle$, which is captured by $\hat{L}_2 = \sqrt{\gamma} |g\rangle \langle r|$.

From the Linblad equation we find

$$\dot{\rho}_{gg} = \Gamma \rho_{ee} + \gamma \rho_{rr} - \frac{1}{2} (i\Omega_p^* \rho_{eg} + c.c.) \quad (3.3)$$

$$\dot{\rho}_{ee} = -\Gamma \rho_{ee} - \frac{1}{2} (i\Omega_c^* \rho_{re} - i\Omega_p^* \rho_{eg} + c.c.) \quad (3.4)$$

$$\dot{\rho}_{rr} = -\gamma \rho_{rr} + \frac{1}{2} (i\Omega_c^* \rho_{re} + c.c.) \quad (3.5)$$

$$\dot{\rho}_{eg} = -\left(\frac{\Gamma}{2} - i\Delta_p\right) \rho_{eg} - \frac{i}{2} (\Omega_c^* \rho_{rg} - \Omega_p \rho_{ee} + \Omega_p \rho_{gg}) \quad (3.6)$$

$$\dot{\rho}_{rg} = -\left(\frac{\gamma}{2} - i(\Delta_p + \Delta_c)\right) \rho_{rg} + \frac{i}{2} (\Omega_p \rho_{re} - \Omega_c \rho_{eg}) \quad (3.7)$$

$$\dot{\rho}_{re} = -\left(\frac{\Gamma + \gamma}{2} - i\Delta_c\right) \rho_{re} + \frac{i}{2} (\Omega_p^* \rho_{rg} + \Omega_c (\rho_{rr} - \rho_{ee})). \quad (3.8)$$

In general, the system is dynamic with each of the density matrix elements evolving in time. A solution to the above equations is therefore generally non-trivial. However, practically speaking, as can be verified numerically for the experimental parameters we discuss later, the density matrix rapidly approaches a steady state. Therefore, it is sufficient to look at the steady state solution of the above equations.

Additionally, we can make another simplification here by assuming we are in the weak probe limit, i.e. that $\Omega_p \ll \Omega_c$. In this regime the atoms spend almost all their time in the ground state, and we can make the approximation $\rho_{gg} \approx 1$, $\rho_{ee} \approx 0$ and $\rho_{rr} \approx 0$. This simplifies things further and we find

$$\rho_{eg} \approx \frac{\Omega_p}{2\Delta_p + i\Gamma - \frac{|\Omega_c|^2}{2(\Delta_p + \Delta_c) + i\gamma}}, \quad (3.9)$$

where we have ignored terms higher than linear order in Ω_p .

A medium composed of an ensemble of three-level atoms has a probe susceptibility given by [60]

$$\chi = -\frac{\mathcal{N}c\sigma_0\Gamma}{\omega_0\Omega_p}\rho_{eg}, \quad (3.10)$$

where \mathcal{N} is the atomic density, ω_0 is the ground to excited state resonant frequency, and $\sigma_0 = \frac{2\omega_0}{c\epsilon_0\Gamma\hbar}|d_{eg}|^2$ is the resonant cross-section, with $d_{eg} = \langle e|\hat{d}|g\rangle$ as the dipole matrix element between the excited and ground state. From equation 3.9 we find the probe susceptibility to be

$$\chi = \frac{\mathcal{N}c\sigma_0}{\omega_0} \frac{\Gamma}{\frac{|\Omega_c|^2}{2(\Delta_p + \Delta_c) + i\gamma} - 2\Delta_p - i\Gamma}. \quad (3.11)$$

It is often helpful to separate out the susceptibility into its real and imaginary parts, corresponding to the dispersive and absorptive properties, respectively, of the




Figure 3.2: Theoretical EIT transmission curves for a medium with $\text{OD} = 20$, $\Gamma = 2\pi \times 6$ MHz, $\gamma = 2\pi \times 100$ kHz and $\Delta_c = 0$.

medium

$$\text{Re}[\chi] = \frac{\mathcal{N}c\sigma_0}{\omega_0} \frac{2\Gamma|\Omega_c|^2\delta - 2\Delta_p\Gamma(4\delta^2 + \gamma^2)}{(4\delta^2 + \gamma^2)(4\Delta_p^2 + \Gamma^2) + 2|\Omega_c|^2(\gamma\Gamma - 4\Delta_p\delta) + |\Omega_c|^4} \quad (3.12)$$

$$\text{Im}[\chi] = \frac{\mathcal{N}c\sigma_0}{\omega_0} \frac{\Gamma^2(4\delta^2 + \gamma^2) + \Gamma\gamma|\Omega_c|^2}{(4\delta^2 + \gamma^2)(4\Delta_p^2 + \Gamma^2) + 2|\Omega_c|^2(\gamma\Gamma - 4\Delta_p\delta) + |\Omega_c|^4}, \quad (3.13)$$

where we have defined the two-photon detuning $\delta = \Delta_p + \Delta_c$ for brevity.

In most cases χ is sufficiently small that we can approximate the complex refractive index as $n \approx 1 + \chi/2$. From this we find that the transmission, T , of the probe through medium with a resonant optical depth, $\text{OD} = \sigma_0 \int dz \mathcal{N}(z)$, is given by

$$T = \exp \left[-\text{OD} \frac{\Gamma^2(4\delta^2 + \gamma^2) + \Gamma\gamma|\Omega_c|^2}{(4\delta^2 + \gamma^2)(4\Delta_p^2 + \Gamma^2) + 2|\Omega_c|^2(\gamma\Gamma - 4\Delta_p\delta) + |\Omega_c|^4} \right]. \quad (3.14)$$




Figure 3.3: Theoretical $\text{Re}[\chi]$ curves for a medium with $\Gamma = 2\pi \times 6$ MHz, $\gamma = 2\pi \times 100$ kHz and $\Delta_c = 0$.

Looking at the spectrum for $\Omega_c = 0$ and for a finite control field¹, as seen in figure 3.2, we can see the presence of the control field opens up a narrow transparency window for the probe around $\Delta_p = 0$, hence the name ‘electromagnetically induced transparency’ (EIT).

Another interesting property of EIT comes from the real part of the susceptibility. The group velocity of the probe light in the medium is given by [60] $v_g = c/n_g$, where $n_g = \text{Re}[n] + \omega_p \frac{d\text{Re}[n]}{d\omega_p}$ is the group index. From figure 3.3 we can see that, when the control field is present, the real part of χ , and thus the real part of the index of refraction, changes rapidly as a function of the probe frequency inside the transparency window seen in 3.2. In this region the derivative of $\text{Re}[n]$ is large, and therefore n_g can be significantly greater than the phase index of the medium. To

¹One might expect things to be nonsensical as $\Omega_c \rightarrow 0$ as we used the fact that $\Omega_c \gg \Omega_p$ in deriving the above expressions. However, it is easy to verify that in the limit of $\Omega_c = 0$ the transmission for the two-level system is recovered.

get a sense of this let us evaluate the group index for the doubly resonant condition,

$$\Delta_p = \Delta_c = 0$$

$$n_g = 1 + \omega_p \frac{\mathcal{N}c\sigma_0}{\omega_0} \frac{\Gamma (|\Omega_c|^2 - \gamma^2)}{(\gamma\Gamma + |\Omega_c|^2)^2} \approx \mathcal{N}c\sigma_0 \frac{\Gamma}{|\Omega_c|^2}, \quad (3.15)$$

where we have used the fact that in most cases $|\Omega_c| \gg \gamma$. From this we see

$$v_g \approx \frac{|\Omega_c|^2}{\mathcal{N}\sigma_0\Gamma}. \quad (3.16)$$

Plugging in some numbers for the rubidium stretched D2 transition [66], $\sigma_0 \approx 2.9 \times 10^{-9} \text{ cm}^2$, $\Gamma \approx 2\pi \times 6 \text{ MHz}$, and assuming some reasonable experimental density, $\mathcal{N} \approx 3 \times 10^{11} \text{ cm}^{-3}$, and control field Rabi frequency, $\Omega_c \approx 2\pi \times 20 \text{ MHz}$, we find $v_g \approx 5 \text{ km/s}$. This phenomenon, where the group velocity of the light moving through the EIT medium is much less than c , is where the term ‘slow-light’ comes from.

Thus far we have been treating the probe field classically. However, for some applications this is no longer a correct approach as the flux of probe photons is sufficiently small. Here, we need to treat the probe as a quantum field. In doing so it becomes natural to talk about ‘dark-state polaritons’ [67], which are a quasiparticle superposition of an excitation in the electromagnetic mode of the probe (a probe photon) and an atomic excitation, in our case to the Rydberg state. Under the doubly resonant condition discussed above, the polariton field creation operator has the form [68]

$$\hat{\psi}^\dagger(z) = \frac{\Omega_c \hat{\mathcal{E}}^\dagger(z) - g\sqrt{\mathcal{N}} \hat{\mathcal{R}}^\dagger(z)}{\sqrt{|\Omega_c|^2 + g^2\mathcal{N}}}, \quad (3.17)$$

where g describes the single atom-photon coupling, \mathcal{N} is again the atomic density, $\hat{\mathcal{E}}^\dagger$ is the creation operator for the probe field, and $\hat{\mathcal{R}}^\dagger$ is the Rydberg excitation creation operator. Equation 3.17 allows us to gain some interesting insights into the system. Firstly we note that the admixture of the photonic and atomic components depends on the relative coupling strengths Ω_c and $g\sqrt{\mathcal{N}}$ (where the $\sqrt{\mathcal{N}}$ enhancement arises from the fact that the states being coupled to are collective excitations which will be discussed in further detail in the next section). It is this admixture of the atomic component that is responsible for the reduced group velocity of light traveling through an EIT medium, seen in equation 3.16. Intuitively, as we reduce Ω_c the polariton becomes more Rydberg-like and the group velocity decreases, and as we increase Ω_c the polariton becomes more photon-like and the group velocity increases.

Taking the idea of slow-light to its extreme, we can think about a probe pulse which is sent into an EIT medium resulting in a propagating polariton. If we then turn the control field off, the polariton is completely mapped or stored as a Rydberg excitation. We can then reverse this process by turning the control field back on, retrieving the pure Rydberg excitation back as a propagating polariton in the probe mode. This idea of coherently converting a propagating polariton into a pure atomic excitation and then back again is the underlying principle behind EIT-based quantum memories, which have been shown to have the high storage efficiencies [69] and long lifetimes [70] critical for the purposes of quantum networking.

3.2 Collective Excitations

From equation 3.17, it mathematically makes sense that the turning on and off of the control field can map from a propagating polariton, to a pure Rydberg excitation, and back again. However, intuitively it might sound peculiar that the retrieved polariton would propagate in the same mode as before the storage. If we think about an ensemble of two-level atoms and direct a resonant photon at the ensemble, there is no preference for the scattered photon to resume propagation in the input photon mode, if scattering occurs. The difference in the EIT case is that the excitation is a collective, or spin-wave excitation. When the control field is turned off the resulting state we have is a single (assuming a single input photon) symmetric, Dicke-like [67, 71], excited state

$$|R\rangle = \frac{1}{\sqrt{N}} \sum_j e^{i(\vec{k}_p + \vec{k}_c) \cdot \vec{r}_j} |g_1 g_2 \dots r_j \dots g_N\rangle, \quad (3.18)$$

because each of the N atoms is equally likely to be excited. To see why this state results in directed emissions, it is perhaps easiest to consider an approach similar to reference [72]. Here, we imagine using a second control field, still resonant with the $|e\rangle$ to $|r\rangle$ transition, but with a different wavevector $k_{c'}$, to perform a fast π rotation, resulting in the state

$$|E\rangle = \frac{1}{\sqrt{N}} \sum_j e^{i(\vec{k}_p + \vec{k}_c - \vec{k}_{c'}) \cdot \vec{r}_j} |g_1 g_2 \dots e_j \dots g_N\rangle. \quad (3.19)$$




Figure 3.4: Spin-wave emission probability as a function of the angle between the retrieved and input probe mode for $\vec{k}_c = \vec{k}_{c'}$. P_r here has been normalized such that the phase matched probability is equal to unity. For this plot the atom positions were randomly generated with a Gaussian distribution equivalent to a cloud with an RMS radius of $20 \mu\text{m}$ in every direction. The wavelengths for the fields used are those for Rubidium D2 Rydberg excitations, $\lambda_{c,c'} = 480 \text{ nm}$ and $\lambda_{p,r} = 780 \text{ nm}$.

The emission probability, P_r , into a mode with polarization $\hat{\mathcal{E}}_r$, wavevector \vec{k}_r , and corresponding raising operator \hat{a}_r^\dagger is given by

$$P_r \propto \left| \langle G, 1_r | \left(\hat{\mathcal{E}}_r \cdot \hat{d} \right) \hat{a}_r^\dagger e^{-i\vec{k}_r \cdot \vec{r}} | E, 0_r \rangle \right|^2, \quad (3.20)$$

where we have used the notation $|E, 0_r\rangle$ with the first element denoting the atomic ensemble state, and the second the photon occupation number of the retrieved field, and where $|G\rangle$ has all atoms in the ground state $|g\rangle$. Under the assumption that the matrix elements $\langle g_j | \left(\hat{\mathcal{E}}_r \cdot \hat{d} \right) | e_j \rangle$ are identical for all atoms, then we have

$$P_r \propto \frac{1}{N} \left| \sum_j e^{i(\vec{k}_p + \vec{k}_c - \vec{k}_{c'} - \vec{k}_r) \cdot \vec{r}_j} \right|^2. \quad (3.21)$$




Figure 3.5: Spin-wave emission probability as a function of the angle between the retrieved and input probe mode for $\vec{k}_c, \vec{k}_{c'}$ at an angle $\approx 128^\circ$ and $\approx 52^\circ$ respectively to the input probe field. P_r here has been normalized such that the phase matched probability is equal to unity. For this plot the atom positions were randomly generated with a Gaussian distribution equivalent to a cloud with an rms radius of $20 \mu\text{m}$ in every direction. The wavelengths for the fields used are those for Rubidium D2 Rydberg excitations, $\lambda_{c,c'} = 480 \text{ nm}$ and $\lambda_{p,r} = 780 \text{ nm}$.

If $\vec{k}_p + \vec{k}_c - \vec{k}_{c'} - \vec{k}_r$ is non-zero, as we perform the sum we are going to get some partial cancellations due to the atoms all being located at slightly different positions. But, when the phase matching condition is met, i.e. when $\vec{k}_p + \vec{k}_c - \vec{k}_{c'} - \vec{k}_r = 0$, all the phase factors add constructively and we see an enhancement in P_r . Under the condition that $\vec{k}_c = \vec{k}_{c'}$ the phase matching condition is clearly satisfied when $\vec{k}_r = \vec{k}_p$, i.e. the retrieved photon is in the same mode as the input probe field, as seen in figure 3.4. It is this phase matching condition that was implicitly met when we previously discussed the EIT storage and retrieval. However, one can also retrieve into a different mode with appropriate choice of the two control fields, as seen in figure 3.5, where the two control field wavevectors have been chosen so

that the phase matching condition preferentially retrieves the photon into a mode counter-propagating that of the input probe field. A further point to make here is that the enhancement of the phase-matched emission improves with increasing atom number, and the magnitude of the dipole matrix element of the $|g\rangle$ to $|e\rangle$ transition. This means that to achieve high directed retrieval efficiency we want a system which has a high optical depth [73].

It should be pointed out that we have ignored the effect of any dephasing mechanisms, such as motional dephasing, or dephasing due to inhomogenous differential energy shifts, that may occur between the creation of the initial $|R\rangle$ spin-wave and the retrieval of the photon. Such dephasing leads to an overall reduction in the retrieval probability that tends to worsen as we increase the storage time of the spin-wave. In practice one has to try to mitigate these dephasing mechanisms, which usually requires working with a cold atomic sample, and which may place some more stringent restrictions on the choice of the wavevectors used, over and above the phase matching condition.

3.3 Rydberg Electromagnetically Induced Transparency

In the previous two sections we have not really troubled ourselves with the Rydberg nature of state $|r\rangle$. From equation 3.17 we know that polaritons propagating in the EIT medium have some admixture of the Rydberg state, and from chapter 2 we know that pairs of Rydberg atoms will interact with each other, meaning that propagating polaritons in a Rydberg-EIT medium will also interact with one an-

other. This polariton-polariton interaction allows us to engineer effective optical non-linearities that are strong even at the single photon level. A prime example of this is the experiment described in reference [41]. Here, the researchers took a doubly resonant, $\Delta_p = \Delta_c = 0$, Rydberg-EIT system utilizing a high lying Rydberg state, $n > 77$, and high optical depth, $OD > 10$. They observed that, for a coherent state input probe, the light exiting the cloud exhibited strong anti-bunching statistics. This behavior is understood fairly intuitively by considering a pair of photons sent into the medium, one after the other. The first propagates, with a reduced velocity, through the medium as a polariton. When the second photon enters the medium, if it is sufficiently close to the first, then the Rydberg-Rydberg interactions mean that the control field is no longer resonant with the $|e\rangle$ to $|r\rangle$ transition, and the photon sees an ensemble of two-level atoms with which it is resonant, and it has a high likelihood of being scattered away. Therefore, the probability of multiple photons simultaneously exiting the medium is suppressed.

In addition to that experiment showing dissipative interactions, a wealth of work has been performed in the dispersive regime. Here, the control and probe are detuned off single-photon resonance, usually with $\Delta_c, \Delta_p \gg \Gamma$, but near two-photon resonance $\Delta_c \approx -\Delta_p$. In this regime the Rydberg-Rydberg interactions results in a change in the real part of χ when the two photons are near, with the imaginary part of χ , related to absorption, being less significant. This provides more flexibility, allowing one to realize attractive polariton-polariton potentials [42, 43], leading to photon bunching, non-trivial three-photon interactions [63], and, with some additional couplings, molecular-like [45] and non-lossy repulsive photonic

interactions [74].

Outside of its use as a diagnostic tool and briefly in chapter 7, we will not be specifically discussing Rydberg EIT much further in this thesis, although it should be pointed out that the initial design of the apparatus was made with Rydberg EIT experiments in mind. We will, however, be making use of a number of the related phenomena discussed in this chapter, such as collective excitations and the non-linearities associated with a Rydberg ensemble.

Chapter 4: Experimental Setup and Techniques

Owing to their low temperatures, and potentially high atomic densities, ultracold atomic systems serve as a good test bed for exploring Rydberg ensemble physics. In this chapter I describe the experimental apparatus and procedures we have developed for cooling, trapping and probing a ^{87}Rb Rydberg ensemble. Additionally, I wish to share some tricks and techniques that we have discovered in the hope that they may benefit future graduate students.

Unfortunately, a graduate student's career is not full of unmitigated successes. However, our failures, while often unreported, are as important, if not more so, to pass on to future generations of graduate students as our successes are. Therefore, throughout this section I have endeavored to also include details of some of the apparatus and techniques that were unsuccessful.

4.1 Science Chamber

4.1.1 Chamber Design

Ultracold atomic physics in general requires high vacuum in order to reduce the collision rate of the atoms of interest with background particles. Thus, we




Figure 4.1: Render of science vacuum chamber. Note the co-ordinate system here as it is referenced throughout this chapter.

perform all experiments within a vacuum apparatus, the exterior design of which is shown in figure 4.1. The central region of this apparatus is our main vacuum chamber (Kimball Physics 8" spherical octagon chamber)¹, which is where atoms are trapped, cooled and probed. The chamber has ten ports, eight 2.75" conflat (CF) and two 8" CF flanges, giving us a potential for high optical access. To all the flanges, except those along the y axis we attach viewports (Kurt Lesker VPZL-275 for 2.75" flanges and Kurt Lesker VPZL-800 to the 8" flanges), with the 2.75" viewports off the x axis anti-reflection (AR) coated for 780 nm for the magneto-optical trapping light. The two viewports along the x-axis are coated for 479 nm, 780 nm and 1004 nm, to be compatible with our 780-nm probe and 479-nm control Rydberg excitation lasers, as well as our 1004-nm optical dipole trapping light.

¹throughout this section I have, where known, tried to include part numbers

To the main chamber we connect a pair of four-way crosses (Kurt Lesker C-0275). We attach a further 780-nm AR coated viewport (Kurt Lesker VPZL-275) to the y-axis flange of both crosses. In addition, an ion pump (Gamma Vacuum 45s TiTan) and D-sub 9 feedthrough (Accu-Glass 9D-275) is attached to one of the crosses, while the other connects to a bellows and a flange multiplexer (Kimball Physics MCF275-FlgMplx-Cr1A5), which houses four MHV and one BNC (Accu-Glass BNC-GS-133) 1.33" feedthrough. The ion pump, used for maintaining vacuum in the system, has an additional port which we connect to an all-metal valve. In normal operation the valve is sealed shut and a blank (Kurt Lesker F0133X000N) is attached. A breakable glass ampoule (Alfa Aesar 10315-03), containing rubidium in its natural abundance, is placed inside the bellows.

For all the flange-flange connections we use copper gaskets to create the seal. This has the potential for causing problems with time dependent magnetic fields, due to eddy currents. However, we require sufficiently high vacuum, and desired to bake at sufficiently high temperature, that Viton gaskets were deemed to be inadequate.

All the viewports are bolted to the chamber along with either one or two washers, shown in figures 4.2 and 4.3, that have various tapped holes for attaching hardware directly to the chamber. The 2.75" viewports have a single washer attached that has a set of through holes, which allow the washer to be bolted to the chamber along with the viewports, and 4-40 tapped holes, at spacings compatible with the 30mm Thorlabs cage system. The 8" viewports have a pair of washers that are stacked on top of each other. The inner washer is bolted directly to the




Figure 4.2: Render of cage washer for 2.75" viewports.




Figure 4.3: Render of cage washers for 8" viewports. Outer (inner) washer shown on the left (right).




Figure 4.4: Image of the interior of the central chamber of the science vacuum apparatus. Note the co-ordinate system here as it is referenced throughout this chapter.

chamber, along with the viewport, using through holes, similar to the 2.75" washer. It has a set of 1/4"-20 tapped holes, which line up with a matching set on the outer washer allowing the two to be attached together. The outer washer has a number of additional 1/4"-20 tapped holes that allow equipment to be attached to the chamber.

4.1.2 In-Vacuum Components

The main chamber houses a number of in-vacuum components, that can be seen in figure 4.4. These are essential to the experimental operation and each are discussed in detail below. To hold the components in place they are screwed into either one or both of a pair of custom designed stainless steel plates. The plates themselves are fixed to the upper and lower lips of the main chamber using a set

of groove grabbers (Kimball Physics MCF800-GrvGrb-C01). The groove grabbers are electrically conductive meaning that the plates are at the same potential as the chamber, which itself is grounded using a thick metal braid, helping to eliminate stray electric fields that might cause Stark shifts to the Rydberg states. The plates both have a set of slits, which sit at 45° to the x and y-axes, designed to reduce undesirable eddy currents. In addition, the plates have an $\approx 1.1''$ by $1.1''$ square hole in their center to allow passage of the z-axis magneto-optical trap (MOT) light.

4.1.2.1 Lenses

Two 1" diameter aspheric lenses (Asphericon AFL25-40) with a numerical aperture of ≈ 0.3 lie along the chambers x-axis. These are AR coated at 780 nm, 479 nm and 1004 nm and are present to allow tight focusing of the probe and control Rydberg excitation light, in addition to that of the optical dipole trapping light. The lenses sit in custom design stainless steel mounts, and are held in place using a custom PEEK washer and a metal ring, which screws into the mount. The lens mounts each have a pair of tapped screw holes which line up with through holes on the lower plate, allowing us to directly attach the two.

A small amount of play exists when screwing the mounts into the plate, allowing for small rotations and displacements of the lenses. In order to ensure correct positioning of the lenses a custom cuboid aluminum block was machined whose long side matched the desired distance between the two lens mounts. The lens mounts were set to be flush with the block while they were being initially screwed into the

plate. Further fine alignment was performed before sealing the chamber by observing the back-reflections of visible red light from a fiber pen off the lens pair. Slight adjustments to the orientation of the lens mounts were made to ensure all the reflections off the lenses were concentric. To achieve this, some small vertical correction was required and so a small shim, made out of a folded piece of aluminum foil, was placed under the negative-x lens mount. During the alignment process a screw became cross-threaded with the negative-y screw-hole of the positive-x lens mount. Nevertheless, the mount was able to be aligned and robustly secured to the plate.

In addition to the necessity for the lenses to be parallel, it is also desirable for the lenses to *not* be parallel with the x-axis windows, in order to prevent any etaloning effect. Therefore, the support plates are turned slightly such that the windows and lenses make a slight angle with each other.

4.1.2.2 Low Voltage Electrodes

Given the sensitivity of Rydberg states to electric fields that we saw in chapter 2, we require precise control of the local electric field felt by the atoms. For this, we have a set of eight low-voltage electrodes, four attached to each plate in a square, side length ≈ 1.1 ", around the MOT light hole. Each electrode has a pair of tapped holes which, along with corresponding through holes, allows it to be attached to the support plates. For each screw we use a pair of ceramic "hat" washers (McAllister Alumina "hat" washer) on either side of the plate that electrically isolate the screw and electrode from the plate itself.

D-Sub 9 Pin Number	Electrode
1	X+Z+
2	X-Z+
3	X+Z-
4	X-Z-
5	Y-Z-
6	Y+Z+
7	Y-Z+
8	Y+Z-

Table 4.1: Low voltage electrode connection matrix. Electrodes are labelled by their position relative to the previously defined experimental axes.

The electrodes are attached to the D-Sub 9 feedthrough with a Kapton-coated wire ribbon cable (Accuglass KAP-R9). On the feedthrough side the ribbon cable simply plugs into the D-Sub 9 connector. For each electrode the Kapton was stripped slightly from the 28AWG wire and the bared wire wrapped directly around one of the electrode screws. A pinout for this wiring can be found in table 4.1. During chamber construction each electrode was checked to ensure that it was not in electrical contact with the plate. However, after pumping down and sealing the chamber it was discovered that the X-Z- electrode had become shorted to the chamber. While undesirable, full control of the electric field is possible even with one of the electrodes shorted. Therefore, the problem has not been rectified due to desire to avoid breaking vacuum.

The two sets of electrodes have an ≈ 34 mm spacing, and when trapped the atoms sit roughly equidistant between the two. Experimentally, we generate electric fields by applying a differential potential to electrodes along the desired axis, e.g. to apply field along the x-direction all the electrodes in the positive-x direction have some potential applied to them, while those in the negative-x direction

have the opposite applied. This procedure is always performed ensuring that the shorted electrode is at ground. We have spectroscopically determined the electric field at the atoms per volt of differential potential applied to be $\approx 0.14(\text{V}/\text{cm})/\text{V}$, $\approx 0.10(\text{V}/\text{cm})/\text{V}$ and $\approx 0.08(\text{V}/\text{cm})/\text{V}$ for the z, x and y axes respectively.

4.1.2.3 High Voltage Electrodes and Ion Detection Electronics

It can often be desirable to identify exactly what Rydberg states are present in a system. One can use state selective field ionization [59] for this purpose, using a large electric field to selectively ionize certain Rydberg states, with the resulting ions then detected. To that end, we have a pair of high voltage electrodes in the vacuum chamber, along with a multi-channel plate detector (MCP) (Photonis Microtron 2 APD 5.5/32/25/8). These all lie along a line which makes an angle $\approx 17^\circ$ to the y-axis.

One of the electrodes is a solid stainless steel block which sits on the positive-y side of the chamber. The other electrode is of a similar design but has a circular hole covered in a wire mesh, with the mesh fixed in place by a circular ring bolted to the electrode block. The MCP sits in the steel block immediately behind the mesh covered electrode, with a small set screw securing it in place. All the blocks are bolted to both the upper and lower plates, with ceramic hat washers again used for electrical isolation from the plate.

The electrodes and MCP are connected to the various MHV and BNC feedthroughs using high voltage Kapton coated wire and coaxial cable respectively (Accu-glass




Figure 4.5: Annotated render of pinout for high voltage feedthrough.

100670 and 100705). The wire is attached to the feedthroughs, the MCP anode and MCP output tabs by stripping a small amount of insulation from the wire tip and securing the wire to feedthrough/tab using a barrel connector. Similar to the low voltage case, the connection to the two electrodes is made by stripping the insulation from the wire and winding the bared conductor around one of the screws which secure the electrode block to the supporting plate, taking care to ensure there are no shorts to the plate itself. The outer casing of the MCP acts as its electrode, which is in electrical contact with the block it sits in, therefore, this is wired up by also winding a stripped wire around one of the screws securing the MCP block. A diagram of the pinout for these connections is shown in figure 4.5.

The solid and meshed electrode are separated by a distance ≈ 91 mm, with the

atoms trapped in a region approximately in the middle of the two. We have spectroscopically determined the electric field generated due to a differential potential between the electrodes to be $\approx 0.03(\text{V}/\text{cm})/\text{V}$.

While the original intention for MCP and high-voltage electrodes was for the diagnostic purpose of state-selectively field ionizing Rydberg states, several issues have prevented their use for this purpose. Most fundamentally there seems to be an issue with using the MCP in the presence of the 479 nm Rydberg excitation light. When monitoring the MCP we have observed detection events with the 479 nm light present, even with the absence of trapped atoms. Given that the 479 nm photons are sufficiently energetic to induce desorption of rubidium atoms from surfaces [75], our hypothesis is that some fraction of the blue light injected into the chamber is scattered and proceeds to reflect off the various shiny in-vacuum surfaces. The blue light rattling around the chamber desorbs rubidium atoms from the interior surfaces, a process which has the potential to produce ions [76], which are then detected by the MCP. This issue makes it impossible to perform Rydberg ion detection while the blue light is present, and even for a time (of order 10s of microseconds) after it has been extinguished. Therefore, at present, we have been unable to use the MCP. It should be pointed that there are other groups with similar experimental setups who do not experience this problem [77], which we believe is due to better shielding of the MCP from blue light and/or the desorbed ions.

In addition to issues with the MCP, we have a related problem with our high voltage electrodes when used for field ionization. Large electric fields, of the magnitude required for field ionization, seem to cause some charging effect within the

chamber when the 479-nm light is present. Similar to with the MCP, we have also observed this charging when the electric field and blue are not on simultaneously, but the field is turned on some 10s of microseconds after the light has been extinguished. We believe that this is also related to desorption of rubidium ions from chamber surfaces, which are jettisoned away from the surface by the electric field, leaving the surface charged. Due to the lack of electrical conductivity and the large potential for interaction with the blue light, the in-vacuum lenses are likely candidates for the surfaces becoming charged. As a result of the sensitivity of the Rydberg states to electric fields, the charging has been observed to produce measurable shifts, e.g. several MHz for $n=96$. Therefore, at present, we have also been unable to utilize the electrodes for their originally intended purpose.

In spite of these issues, the in-vacuum high-voltage electrodes have been given a new lease of life as microwave antennas, for the purpose of driving microwave transitions. Given the enormous electric dipole matrix elements between nearby Rydberg states [65], a relatively small amount of microwave power is required to produce large Rabi frequencies. For example, for the stretched $139S_{1/2}$ to $139P_{3/2}$ transition a microwave intensity of ≈ 5 pW/cm² is sufficient to produce a Rabi frequency of $\approx 2\pi \times 1$ MHz. Even with the sub-optimal design, a modest amount of RF power, of order a few milliwatts, connected to the solid electrode feedthrough has been found to be sufficient to produce microwave intensities at the atoms of a few hundred nW/cm².

In addition to driving electric dipole transitions between adjacent Rydberg states, microwaves are also of use for driving magnetic dipole transitions between

the different sub-levels of rubidium’s $5S_{1/2}$ ground state manifold. While the coupling is typically much smaller for these type of transitions, we have observed Rabi frequencies on the order of $\approx 2\pi \times 1$ kHz by applying ≈ 2 W of RF power to the solid electrode feedthrough. This is sufficiently high for a number of purposes, such as preparing the cloud in a Zeeman sublevel for which optical pumping is not straightforward.

Due to the fact that the in-vacuum high-voltage electrodes were not designed as microwave antennas, they are not natively well impedance matched to any RF drive. To overcome this, we use a triple stub tuner (Maury microwave 1819B) which allows us to suppress unwanted reflections that might damage RF generating equipment.

4.1.2.4 Wiring

The way the vacuum apparatus is configured, wires run from the feedthroughs, which reside on the two crosses, to various components in the main chamber. This causes a potential problem as wires dangling in the central region of the cross will obscure optical access along the y-axis. To eliminate this issue a “false ceiling” was created using coils of stainless steel wire which run along the inner part of the cross, as seen in figure 4.6. These are set up so that any wires are pinned between the coil and the inner wall of the cross, keeping them out of the central region. Additionally, within the main chamber the wires are strategically routed behind the plates, with most of the connections to the in-vacuum electronics made to bolts on the outer




Figure 4.6: Image of one of the crosses during construction showing “false ceiling” facing part of the plates.

4.1.3 Initial Vacuum Pump-Down Procedure

Several measures were taken during and after construction of the science chamber to ensure that we could reach the high vacuum required for an ultracold atomic physics experiment. The metal in-vacuum components, described in the previous section (electrodes, lens mounts, plates etc.), were electropolished after machining. Several components were salvaged from old experiments, rather than bought new, or were deemed to be sufficiently dirty prior to construction. This included the crosses, bellows, and groove grabbers. These, along with the Allen keys used to tighten the in-chamber screws, were all cleaned using a standard procedure detailed on the JQI

wiki². The chamber was then constructed on an optical table covered in vacuum compatible aluminum foil, using nitrile gloved hands to ensure no contamination of the chamber.

Upon completion of the construction of the vacuum apparatus, a roughing turbopump was connected to the, normally blanked, port of the valve and the valve was opened. An initial leak test was performed by running the roughing pump for ≈ 17 hours, during which time we observed an exponential decrease in the pressure. A further test was performed using a residual gas analyzer (RGA), while pumping, to monitor the helium pressure in the chamber while spraying helium around the various potential entryways, with no signs of a leak being observed.

The entire vacuum apparatus was then transferred to an oven with an integrated roughing (used for the initial pumping) and ion pump (used for the majority of the bake duration) and RGA. The system was baked at $\approx 150^\circ$ C for ≈ 1 week, while constantly pumping. While a higher temperature would have resulted in a shorter bake time, the custom AR coatings for the in-vacuum lenses and x-axis viewports were not recommended to exceed 150° C. To prevent cracking of the viewports due to temperature gradients, the oven temperature was ramped up and down from room temperature to the baking temperature over a period of ≈ 10 hours. Although the temperature was well below the maximum rated temperature of the on-chamber ion pump magnets, $\approx 250^\circ$ C, out of an abundance of caution we removed them during the bake.

Post-bake, with the chamber at room-temperature, the magnets were re-

²https://jqi-wiki.physics.umd.edu/d/wiki/lab_tips/vacuum_preparation

attached to the ion pump and it was turned on. An initial pressure increase was observed, believed to be due to some arcing within the ion pump causing material to be ejected into the vacuum system. However, with the oven's ion pump operating in combination with the on-chamber pump the pressure was seen to drop back down relatively quickly. The valve was then sealed shut and the chamber moved to its present position on the optical table. After sealing the valve the pressure was seen to decrease about an order of magnitude over a few hours. Since sealing the chamber, the pressure has been observed to fluctuate within the range $5 - 20 \times 10^{-10}$ mBar, as measured from the current draw of the ion pump. However, these fluctuations do not seem to be obviously correlated with anything.

4.1.4 Magnetic Field Generation

A pair of coils in anti-Helmholtz configuration generates a quadrupole magnetic field used for the MOT. The coils were salvaged from an old experiment and each consist of hollow, square-profile Kapton coated tubing wound with seven turns and seven layers with an inner diameter of ≈ 7 cm and outer diameter of ≈ 13 cm. The two ends of both coils are stripped of their Kapton, for electrical connection, and have Swagelok connectors attached to the tips, to allow for water cooling. We connect the coils such that both current and water are run through the pair in series. Each coil is epoxied (3M Scotch-Weld 3501) to a 3D printed ABS plastic mount, which is bolted to the outer washers of the 8" viewports, as shown in figure 4.7. When mounted to the chamber the coils are separated by ≈ 10 cm, from which we




Figure 4.7: Cross-sectional render of 8" viewport with quadrupole coil mounted.

calculate the magnetic field gradient in the center of the chamber, along z-axis to be $\approx 0.9 \text{ G/cm/A}$.

In addition to the quadrupole coils, a set of bias coils are also attached to the chamber which allow us to produce near-uniform fields across the trapped atomic cloud in all three axis. These are essential for shimming the magnetic zero of the MOT, zeroing out stray magnetic fields, defining quantization axes, and lifting Zeeman degeneracies. For the z-axis bias field, the coils were formed by winding 5 turns and 2 layers of insulated 22 AWG wire directly around the exterior of the chambers 8" viewports, as seen in figure 4.8. The wire is clamped in place by a series of 3D ABS printed clips which bolt to the 1/4"-20 screw holes of the 8" viewport outer washer. Current is run through the coil pair in series in Helmholtz configuration. The coils have a separation of $\approx 7 \text{ cm}$ and an inner diameter of $\approx 20 \text{ cm}$, from which we expect a field of $\approx 1 \text{ G/A}$ at the atoms, which has been verified spectroscopically.

For the x-y plane a set of four coils are mounted to the chamber on the 2.75"




Figure 4.8: Annotated image of chamber with quadrupole and bias coils mounted. viewports that lie along the axes $\pm 45^\circ$ to the x-axis. They were formed by winding four layers by four turns of insulated 22 AWG wire around a 3D ABS printed mount. Both the coils and mount can be seen in figure 4.8. As the coils were wound, epoxy (3M Scotch-Weld 3501) was applied to cement the wire in place. The mount has through holes compatible with Thorlabs 60 mm cage system. We use a 30 mm to 60 mm cage adapter, along with cage rods, to attach the coils to the washers of the 2.75" viewports. Given the ≈ 21 cm separation of the coils pairs and their ≈ 10 cm

inner diameter we would expect fields on the order of ≈ 0.3 G/A at the atoms. However, spectroscopically we determine the value to be about half of this. Upon discovering this suppression, the coil pair was removed and tested off the chamber, with the field magnitude matching the expected value. In addition, we have observed hysteresis that we believe to be related to the field suppression issue. This hysteresis has not been quantified spectroscopically, but measurements performed using a single coil mounted on one of the x-y viewports, with a magnetometer placed near the opposite viewport, indicated a difference of ≈ 0.1 G in the remnant, following the application of a large field with the coil. A similar effect has been observed in another, more recently constructed, set of coils. While there appears to be something within the chamber itself causing the issue, we have, as of today, been unable to determine its exact origin.

4.1.5 Rubidium Oven

As previously mentioned, the bellows attached to the chamber houses a glass ampoule containing rubidium. After baking out and sealing the chamber this ampoule was cracked by bending the bellows. The exposure of the rubidium within the ampoule results in a vapor within the chamber, from which we load our MOT. We control the rubidium vapor pressure within the chamber by adjusting the temperature of the bellows. Temperature control was originally implemented using a Peltier heat pump, which had one side in thermal contact with the bellows and the other with a water cooling block. This allowed us to heat up the bellows during the day,

and cool it at night. However, several catastrophic failures with the water cooling, including a rather messy leak, convinced us to transition to the current oven design.

The present oven design has the bellows enclosed by set of three metal blocks, with copper wool wedged between the bellows and blocks to ensure good thermal contact. The blocks are bolted to a washer, similar to those around the viewports, fixing them in place around the bellows. These are wrapped in heater tape, which is then wrapped in aluminum foil to reduce heat loss. We drive the heater tape with a variac, using a solid state relay to allow power to the tape to be turned on and off. A thermistor wedged in between the copper wool allows us to monitor the oven temperature, which we feed back to a standalone computer (Raspberry Pi) which controls the solid state relay. We implement a rudimentary feedback control system in software to regulate the temperature to a set point to better than $\pm 5^\circ\text{C}$. A web-GUI interface allows altering of the set temperature as well as the defining of set temperatures to be implemented at fixed times. Both pieces of software can be found in https://github.com/acraddoc91/RbRy_Oven_Control. For a typical day the oven is heated to 60°C an hour prior to experimental use. We have observed that lower oven temperatures than this can significantly reduce the final atom density of the cloud, while much higher temperatures significantly reduce the lifetime of the clouds. Although an hour seems to be sufficient time for the vapor pressure in the main chamber to reach acceptable levels for loading the MOT, it should be noted that it took significantly longer, multiple days, for this to be the case the first time the oven was turned on.

Rubidium has a tendency to stick to and kill ion pumps [78, 79]. To try and




Figure 4.9: Atomic level diagram, with relevant experimental couplings. For the couplings shown to be off-resonant, the detuning (black arrows) is intended to indicate the sign of the detuning from resonance.

mitigate this as an issue we also constantly heat our ion pump to $\approx 60^\circ\text{C}$. The idea here is that a temperature gradient between the pump and the rest of the chamber will cause rubidium to diffuse away from the pump.

4.2 Laser Systems and Beam Paths

We perform all experiments with a cloud of atoms in an optical dipole trap. Several processes are necessary for producing a high density, low temperature, and high state purity collection of atoms in the trap suitable for the purpose of probing Rydberg ensemble physics. We first load atoms from the background rubidium vapor into a MOT, followed by optical molasses to cool the atoms into the optical dipole trap. The atoms are then optically pumped into a well-defined Zeeman sublevel prior to probing. Here, I detail the optical schemes for these stages, as well as the laser systems used in the experiment. Figure 4.9 shows the atomic level system and indicates the transitions addressed by the various beams/light fields.

4.2.1 Magneto-Optical Trap and Optical Molasses

Three circularly polarized beams, e^{-2} diameter ≈ 25 mm, are injected into the chamber, two in the x-y plane and a third along the z-axis, as shown in figure 4.10. At the exit of the chamber all the beams pass through a quarter waveplate before being retro-reflected back into the chamber. The two beams in the x-y plane contain both cooling light, addressing the $F = 2$ to $F' = 3$ transition on the D2 line, and repump light, addressing the $F = 1$ to $F' = 2$ transition of the D2 line. The beam




Figure 4.10: Beam schematic for MOT and molasses. Due to the 2D nature of the diagram, the third MOT beam, which propagates along the z-axis, is difficult to see. All three beams consist of light with a 780-nm wavelength. Beams are roughly to scale.

in the z-axis consists only of cooling light.

The cooling and D2-repump light for the MOT and molasses stages originate from two laser systems, a Toptica TA-pro and DL-pro respectively. A schematic of the setup for the two laser systems can be seen in figure 4.11. While kHz-narrow laser linewidths are not required here, some level of frequency stability is. To that end, light from the D2-repump laser is picked off for the purpose of a saturated absorption lock. The error signal from the lock is used to feed back to the laser piezo which controls the grating of the ECDL. As seen in figure 4.12, we lock the light coming out of the laser to the $F = 1$ to $F' = 1/0$ crossover peak of the D2 line.

To stabilize the light from the cooling laser, we use a further pick off to take




Figure 4.11: Schematic of laser setup for MOT and optical molasses. Wavelength for all beams shown is 780 nm.




Figure 4.12: Atomic level diagram showing laser scheme for the MOT and optical molasses.

additional light from the D2-repump laser, which we mix using a 50:50 fiber combiner with light taken from the cooling laser. The mixed light is sent to a high bandwidth detector (Vescent D2-160), which detects the resulting optical beatnote. We send the electronic beatnote signal to a phase-locked loop (PLL), the design of which is similar to that in reference [80], along with a reference tone. We use the output of the PLL to feed back to the cooling laser current (Mod-DC input). This locking scheme is beneficial as it allows us to make large, 100s of MHz, on-the-fly changes to the cooling laser frequency by altering that of the reference³. Additionally, the PLL operates by trying to match the phase of the beatnote with that of the reference signal which results in cooling and D2-repump light being phase locked. This actually ends up

³Note that the cooling laser frequency does not instantaneously change when the reference is altered, with some time taken for the laser to re-lock that is variable, depending on the frequency jump. For jumps of order tens to hundreds of megahertz this re-lock time is on the order of hundreds of microseconds to milliseconds.

being critical for the gray molasses [81] that we will discuss later. In order to prevent long term drifts from causing mode-hops, we utilize a slower lock which feeds back to the grating piezo of the cooling seed ECDL, that drives the current feedback to zero.

For the D2-repump light sent to the chamber, we use a double pass AOM (Gooch & Housego 3110-197), which allows for rapid adjustment of the D2-repump frequency, as well as fast shuttering of the light. While for the cooling light, we use a single pass AOM (Gooch & Housego 3110-197), driven at ≈ 80 MHz, for fast shuttering. After the AOMs, we combine light from both the cooling and D2-repump lasers using 50:50 fiber combiner, whose output arms are sent to the science chamber as the two x-y plane MOT beams. Additional post-AOM cooling light is coupled into another fiber and sent to the chamber for the z-axis MOT beam.

4.2.2 Optical Pumping

Initialization of all the atoms into a single state is desirable for several reasons. Firstly, it presents a cleaner system to work with than one in which two or more Zeeman sublevels are populated. Additionally, as we saw in chapter 3, and as will be important in the experiments later, having a large OD is advantageous when dealing with light-matter interactions. We, therefore, want to work in a Zeeman sublevel for which the probe coupling will have a large Clebsch-Gordan coefficient, as this maximizes the OD per atom. For the D2 line of rubidium, the maximum transition strength is on the stretched $|F = 2, m_f = \pm 2\rangle$ to $|F' = 3, m'_f = \pm 3\rangle$ transition. As




Figure 4.13: Beam schematic for optical pumping. Beam consists of D2-repump light with a 780-nm wavelength, and D1-pump light with a 795-nm wavelength. Beam is roughly to scale.

such, we optically pump our atoms into the $|F = 2, m_f = 2\rangle$ state. Experimentally, this is achieved using a single circularly-polarized beam, e^{-2} diameter ≈ 8 mm, injected into the chamber along the y-axis, as shown in figure 4.13. The beam consists of pumping light, which addresses the $F = 2$ to $F' = 2$ transition on the D1 line, and a repump addressing the $F = 1$ to $F' = 2$ transition of the D2 line. This scheme pumps atoms into the $|F = 2, m_f = 2\rangle$ state, which is dark for the pure σ_+ polarized beam.

Previously, we had attempted to pump using the same transition on the D2 line. In this scheme off resonant scatter of the $|F = 2, m_f = 2\rangle$ state off of $F' = 3$ can occur, which can lead to depolarization of the cloud. Thus, for the D2 pumping scheme, there is a delicate balance between pumping fidelity and heating the cloud,




Figure 4.14: Atomic level diagram showing laser scheme for optical pumping.

with both depending on the pumping duration. In the D1 scheme the scattering of the dark state only depends on the polarization purity of the beam.

In addition to pumping on the D2 line, historically attempts were made to perform pumping along the x-axis, which is typically the major axis of the optical dipole trapped cloud of atoms. The high optical density of the cloud in that direction exacerbated the problems of depolarization of the D2 pumping scheme. For this reason, pumping along the y-axis, which is typically the minor axis of the cloud, was empirically found to be a better approach. However, in order to correctly define the quantization axis for both probing and pumping, this approach does require the B-field to be adiabatically swept through ninety degrees between pumping and probing.




Figure 4.15: Schematic of setup for D1-pump laser (795-nm wavelength).

The optical pumping beam, shown in figure 4.13, is formed of light from the D2-repump laser, which we have already discussed, and that of another Toptica DL-Pro which produces the 795-nm D1 pump. A schematic of the D1-pump laser setup is shown in figure 4.15. We use a saturation absorption setup to frequency lock the D1-pump laser, using the grating piezo of the ECDL for feedback, to the D1 $F = 1$ to $F' = 2$ transition, as shown in figure 4.14. An AOM, driven at ≈ 80 MHz, is used for fast shuttering of the pump light. The D1-pump and D2-repump light is combined with a 50:50 fiber combiner, with one of the output arms sent to the chamber.

4.2.3 Optical Dipole Trap

The optical dipole trap is comprised of three separate beams which intersect inside the chamber, two ‘arms’ that pass through the in-vacuum lens and cross at




Figure 4.16: Beam schematic for dipole trap. All beams composed of light with a 1004-nm wavelength. Beams *not* to scale.

a shallow angle, and a ‘dimple’ which travels along the y-axis near-perpendicular to the other arms, shown in figure 4.16. This approach allows for a high degree of configurability of the shape of the trap, and therefore atomic cloud, by varying the relative power between the arms and dimple, and/or varying the crossing angle of the arms.

As seen in figure 4.16, we use a system for the arms where we have a first beam that propagates through the chamber, which is then recycled into a ‘retro’ beam, counter-propagating the first. This is used to maximize the trap depth for the optical power available. The arms are coupled into the chamber by a pair of dichroic mirrors (Chroma T800 DCSPXR), which are shortpass mirrors with a cutoff wavelength of ≈ 800 nm, allowing us to spatially overlap the Rydberg excitation and optical dipole trap light. The first arm enters the chamber linearly polarized, with the polarization along the z-axis. It is aligned off-axis of the first in-vacuum lens it interacts with, which causes the beam to focus down to a e^{-2} radius of $\approx 30 \mu\text{m}$ at a half-angle to the x-axis of⁴ $\approx 12^\circ$. After passing through the second in-vacuum lens and exiting the chamber, the beam is roughly collimated. The light is then directed to a polarizing beam splitter (PBS), telescope and a retarder, to ensure polarization purity, adjust the longitudinal alignment and polarization of the retro arm respectively. The retro beam is directed back into the chamber, roughly mirroring the first beam about the x-axis, so that it focuses down to a similar waist and subtends a similar half-angle. For the work presented in this thesis we used a half-wave plate (HWP) for the retarder so that first and retro arms had

⁴this can be varied by changing the off-axis displacement of the beam on the lens




Figure 4.17: Atomic level diagram showing laser scheme for dipole trapping.

orthogonal polarization, preventing the formation of an optical lattice. However, we have recently replaced this with a liquid-crystal retarder (Thorlabs LCC1411-B) to allow us to dynamically rotate between a standard optical dipole trap and optical lattice. Due to losses through the chamber, and in the post chamber reshaping and polarization optics, the retro arm power is $\approx 80\%$ of that in the first arm. The dimple beam is elliptically shaped, and when focused down has e^{-2} radii of $\approx 50\ \mu\text{m}$ and $\approx 30\ \mu\text{m}$ along the x and z axes respectively. It too is linearly polarized, with its polarization along the x-axis.

For the work shown in this thesis we use ≈ 1004 nm light for our three trapping beams. This has the advantage that, in theory, we can tune the trapping light to the ‘magic’ wavelength [82]. Here, the light is near-blue detuned from the $6P_{1/2}$ to $nS_{1/2}$ transition, as shown in figure 4.17, where the Rydberg and ground state have the same scalar polarizability, which allows us to eliminate inhomogeneous broadening of the Rydberg-ground state transition due to the trap. However, at the magic wavelength the vector polarizability of the Rydberg and ground states can differ significantly. Due to the polarization configuration of the dipole trap arms, for the work shown in this thesis, we have an inhomogeneous differential vector shift that makes working at the magic wavelength unfavorable. Therefore, we work at an empirically determined somewhat magic wavelength, which minimizes the the inhomogeneous broadening due to the combination of the differential scalar and vector light shifts. For the Rydberg states we typically work with, $n > 80$, the somewhat magic detuning tends to be of order a few hundred megahertz to gigahertz blue detuned from the $6P_{1/2}$ to $nS_{1/2}$ transition. Going forward, operating the arms in a lattice configuration, with the polarization of the outgoing and retro arm being parallel, the vector light shift for both states will disappear, which should permit working at the actual magic wavelength.

The optical dipole trapping light originates from a 1004-nm Toptica TA-pro, containing a seed ECDL and a TA, and a standalone Moglabs MOA TA, both shown in figure 4.18. Some light from the seed ECDL is picked off and sent to a wavemeter (HighFinesse Angstrom WS/U-2). We use the locking feature of the wavemeter to stabilize the frequency of light by feeding back to the grating piezo of the seed




Figure 4.18: Schematic of laser setup for dipole trapping. Wavelength of all beams shown is 1004 nm.

ECDL. The remaining light from the ECDL seeds a first TA, integrated into the TA-pro itself, which produces ≈ 1.5 W power. A fraction of this, ≈ 75 mW, is picked off and sent to seed the Moglabs TA which amplifies the power up to ≈ 1.5 W. The remaining light from the TA-pro is sent to the chamber to become the dimple beam, while that from the Moglabs TA forms our dipole arms. Prior to entering the fiber we send the light from each TA through an AOM (Gooch & Housego 3080-122) for shuttering. In order to prevent interference between the arms and dimple, the drive frequency of these two AOMs is offset by 10 MHz. Due to losses from the AOMs, optics and coupling into the large mode area fibers (NKT Photonics LMA-PM-15) we are able to get ≈ 750 mW of power delivered to the chamber for both the dimple and first arm.

4.2.4 Rydberg Excitation

For coupling the ground and Rydberg state we use a two-photon excitation, shown in figure 4.9, with a ≈ 780 -nm ‘probe’, addressing the D2 line of rubidium, and a ≈ 479 -nm ‘control’, addressing the $5P_{3/2}$ to $nS_{1/2}$ transition. These both pass on-axis through the in-vacuum lenses, as shown in figure 4.19. In order to minimize the Doppler broadening of the ground-Rydberg transition we counter propagate the two beams. The in-vacuum lens focuses the beams down to e^{-2} radii of $\approx 3.3 \mu\text{m}$ and $\approx 18 \mu\text{m}$, for the probe and control respectively. We chose this probe waist in order to reduce the transverse dimension of the probe to be smaller than a blockade radius⁵, effectively making all the Rydberg interactions one-dimensional along the axis of the probe. The choice of control waist is a compromise, on the one hand we want the beam to be small to ensure a large intensity and thus Rabi frequency of the control field. However, we also want the beam to be large for two reasons. Firstly, we want the control Rabi frequency to be homogeneous across the entirety of the probe. Additionally, as the control field acts as a anti-trap for the rubidium $5S_{1/2}$ ground state, the beam also needs to be large enough that it does not significantly reduce the atom density along the probe path. Outside of the chamber the two colors are combined and separated on a pair of dichroic mirrors (Semrock FF757-D101), while also passing through a further pair of dichroic mirrors (Chroma T800 DCSPXR), which allow the spatial overlap of the probe and control with the dipole trapping light, as previously discussed.

⁵This waist was chosen along with picking an appropriate Rydberg level where this condition is met.




Figure 4.19: Beam schematic for Rydberg excitation. Probe beam (780-nm wavelength) path shown in red, and control beam (479-nm wavelength) path shown in blue. Beams *not* to scale.

As mentioned, the maximal electric dipole matrix element for the D2 line occurs for the $|F = 2, m_f = \pm 2\rangle$ to $|F' = 3, m'_f = \pm 3\rangle$ transition. Therefore, to achieve the maximum optical density per atom, the probe is circularly polarized to address the stretched D2 transition. For this configuration the control field only couples the $5P_{3/2}$ to $nS_{1/2}$ states if it has the opposite circular polarization. Therefore, it too is circularly polarized to maximize the Rabi frequency of the control field.

The laser systems that produce the probe and control light are arguably the most important in the experiment, warranting their slightly lengthier discussion. We generate the light for the probe using a Toptica DL Pro ECDL. While the control field light comes from a Toptica TA-SHG Pro, where a 958-nm seed ECDL is amplified using a TA, then frequency doubled to 479 nm in a SHG ring cavity.

As we saw in section 2.1, Rydberg states can be extremely long lived. For the states we shall be experimentally dealing with in this thesis, the state lifetimes are on the order of 100s of microseconds. As a result, the natural linewidth of the ground-Rydberg transition is orders of magnitude smaller than that of the free-running linewidths of the Rydberg excitation lasers, hundreds of kilohertz. We, therefore, want to spectrally narrow both the probe and control light sent to the chamber. This is achieved by Pound-Drever-Hall (PDH) locking [83] both the 780-nm probe and the 958-nm control seed ECDL to a high-finesse ultra-low expansion (ULE) cavity (custom design Advanced Thin Films coated ULE cavity), which has a linewidth < 10 kHz.

In order to ensure stability of the lasers locked to the ULE cavity, several




Figure 4.20: Image of ULE cavity vacuum chamber, used for the two Rydberg excitation laser PDH locks.

measures are taken to stabilize the cavity itself. Firstly, the cavity is housed inside its own vacuum chamber, separate from the main science chamber, which is primarily comprised of a 6" diameter tee (Kurt Lesker T-0600-275), pumped down to of order 10^{-8} mBar by a 2 l/s ion pump (Modion C-1765M). Within the chamber the cavity sits on a custom designed cradle, via a set of Viton pads, with the cradle resting on the inside wall of the tee on a set of Viton rods. An image of cavity housed within the tee can be seen in figure 4.20. Furthermore, the exterior of the tee is stabilized to the ULE cavities zero crossing temperature⁶, $\approx 33^{\circ}\text{C}$. More extensive information

⁶This is the temperature at which the derivative of the cavity length as a function of temperature goes to zero.




Figure 4.21: Schematic of probe laser (780-nm wavelength) setup

on the design of the ULE cavity system can be found in reference [84], upon which our system is based.

For the probe ECDL, shown in figure 4.21, the PDH lock is implemented by taking light and passing it through a fiber coupled electro-optic modulator (EOM) (Photline NIR-MPX800-LN-0.1), driven at ≈ 18 MHz. This light, $\approx 200 \mu\text{W}$, is then directed to the cavity chamber off a longpass dichroic mirror, with the reflected power collected on a fast photodiode (Newport New-Focus 1801). The signal from the photodiode is demodulated and input into a fast analog servo (Vescent D2-125) which feeds back to the current (Mod-DC input) of the probe laser. As with the cooling laser, we want to prevent long term laser drifts from causing mode-hopping. Therefore, we also perform a slower lock with the grating piezo of the ECDL to drive the current feedback to zero.

We optimize the servo in order to increase the lock bandwidth, as measured by




Figure 4.22: Power spectrum of in-loop error signal for the probe lock, showing the servo bump. Black dashed line shows the approximate cavity linewidth.

looking at the Fourier transform of the in-loop error signal. With the servo optimized and the probe laser locked to the cavity we have verified the laser linewidth, using a self-heterodyne measurement, to be consistent with less than that of the cavity linewidth. However, a precise measurement could not be performed due to the requirement for a longer delay line that we had at hand.

With the probe lock optimized, we observe, in the Fourier transform of the error in-loop error signal, that noise at frequencies less than ≈ 500 kHz are suppressed, with a ‘servo bump’ at this frequency, shown in figure 4.22. These servo bumps can be thought of as phase noise on the laser and have been shown to cause additional dephasing within Rydberg systems [85, 86] when performing two-photon excitations. Given that the ULE cavity has a linewidth much narrower than that of the lock bandwidth, this phase noise is not present on the light exiting the cavity. Therefore, it is the probe laser light from the cavity that we send to the main science




Figure 4.23: Atomic level diagram showing laser scheme for Rydberg excitation.

chamber, as seen in figure 4.21.

Unfortunately, the cavity resonances do not necessarily align with the atomic resonances. Therefore, post-cavity we use an AOM (Brimrose GPF-1000-500-.780), driven at ≈ 1 GHz, either in single or double pass, followed by a double passed AOM (Gooch & Housego 3080-122), driven at ≈ 80 MHz, to set the exact frequency of the probe light arriving at the chamber, as shown in figure 4.23. Typically, the frequency and RF power of the ≈ 1 GHz AOM is not changed, while the lower frequency AOM is used for shuttering and adjustment of the probe light detuning. Given the ≈ 1.5 GHz free spectral range (FSR) of the cavity at ≈ 780 nm, by adjusting the diffraction order taken from the GHz AOM, and its single/double

pass configuration, we can arbitrarily tune the probe frequency.

A similar locking scheme is used for the control light, albeit with some small differences. A small amount of light, ≈ 3 mW, is picked off from the seed diode and passed through a ≈ 1 GHz AOM (Brimrose GPF-1000-500-970), with the positive first order sent to the ULE cavity, around 50 to 100 μ W power. In contrast to the probe setup, the seed diode current itself is modulated at ≈ 20 MHz to produce the required sidebands for the PDH lock. At the cavity chamber, the 958-nm light passes through the longpass dichroic, with the light reflected from the cavity directed to another fast photodiode (Newport New-Focus 1801). We use the demodulated photodiode signal as an input for another fast analog servo (Toptica FALC 110), which feeds back to the current of the control laser (Mod-DC input). The lock bandwidth for the control is of a similar order of magnitude to the probe. However, we have seen some variation as the frequency of the ≈ 1 GHz AOM is adjusted, changing the power arriving at the ULE cavity.

The remaining control seed light, ≈ 40 mW, is amplified using a TA, producing ≈ 1 W seed power, which is then coupled in to a ring cavity with a SHG crystal inside it, as seen in figure 4.24. At the exit of the ring cavity we can typically produce ≈ 450 mW at 479 nm. This is sent through a double passed AOM (Gooch & Housego 3080-125), driven at ≈ 80 MHz, used for fast shuttering and on-the-fly detuning of the light. After, the light is coupled in to a fiber and sent to the science chamber. Due to the various losses we are able to obtain a maximum power of ≈ 100 mW of 479-nm light at the atomic cloud.

Ideally, we would perform a similar cavity filtering technique as we have with




Figure 4.24: Schematic of control laser setup. Wavelength for red colored beams is 958 nm, while blue colored beams are at 479 nm.

the probe for the control light, to eliminate any phase noise outside the lock bandwidth. However, it was determined that we would be unable to pass sufficient power through the ULE cavity to seed the TA, which in turn would have a severe impact on the output 479 nm power. It should be noted that the SHG ring cavity has a specified linewidth on the order of ≈ 1 MHz which potentially performs some filtering, but outside of this no further measures were able to be taken to remove the servo bumps from the control field. Presently, work is ongoing to take the 958 nm light exiting the cavity to injection lock another SHG system, which is likely to be beneficial in reducing phase noise on the light sent to the chamber.

Another minor issue with the present control field setup is that of arbitrary

Rydberg state frequency addressability. This problem has two components, firstly, the output of the seed ECDL steers slightly as its coarse grating is adjusted. Consequently, as we adjusted the seed frequency to address different Rydberg states the light must be re-coupled into the TA. However, this is not easily possible past a certain point as you exhaust the thread of one of the fixed mirrors within the laser system. Without making more significant changes to the SHG system, this means there is a Rydberg state below which the available 479-nm power is significantly lowered. As an example we have observed the power exiting the SHG cavity to address the $5P_{3/2}$ to $54S_{1/2}$ transition drops to ≈ 340 mW. This sets a threshold on the minimum Rydberg state we can easily work with. In addition, as shown in figure 4.23, the seed ECDL frequency is detuned from the various cavity resonances using the aforementioned ≈ 1 GHz AOM. Due to power constraints this AOM is run in single pass configuration, which affords a detuning bandwidth of $\approx \pm 200$ MHz. This bandwidth is insufficient to span the ≈ 1.5 GHz FSR of the cavity meaning that we are unable to arbitrarily set the frequency of the seed, and thus unable to arbitrarily address any Rydberg state.

As a final point of note for the control laser, we have historically had some issues associated with getting the ring cavity to lock to the seed light, which is required to get high efficiency frequency conversion to 479 nm. The problem appears to be very sensitive to the frequency of the seed light. From discussion with Toptica the belief is that this issue is related to absorption of the 958-nm light by water vapor. The ring cavity is designed to be used with a desiccant, however, the presence of the desiccant did not appear to be sufficient to alleviate the problem. We have taken to

releasing dry nitrogen into the case of the Toptica TA-SHG Pro and have seen no further resurgence of this issue.

4.2.5 General Laser System Information

We have already discussed some of the specifics of the various laser systems, but there are some more general properties that they all share. As we have noted, for the purpose of fast shuttering of light to the chamber we use AOMs. In all cases we additionally use slower mechanical laser shutters, based on the design in the JQI wiki⁷. The use of both allows for the AOM to be driven to be almost constantly, which is necessary to prevent drifts due to thermal effects as the AOM heats and cools.

All light sent to the chamber is delivered along polarization maintaining fiber, in order to mitigate unwanted polarization drifts. A further precaution against polarization drifts is taken for all light, other than that for the MOT beams, in a set of waveplates and PBS after the fiber but prior to the chamber. This effectively converts any unwanted polarization drifts into power drifts which we have found to be less problematic.

4.3 Experimental Sequence

In the previous section a brief outline of the experimental sequence was given. Here I discuss in more detail the steps involved in preparation of the cloud for

⁷[https://jqj-wiki.physics.umd.edu/d/documentation/electronics/hdd_shutter?s\[\]=shutter](https://jqj-wiki.physics.umd.edu/d/documentation/electronics/hdd_shutter?s[]=shutter)




Figure 4.25: Waveforms for various parameters during the experimental sequence. Note that the different segments of the temporal axis (x-axis) are not necessarily to scale, however, each segment is shown to scale.

probing. Figure 4.25 shows waveforms for various experimental parameters for the steps up to, and including the optical molasses.

4.3.1 Magneto-Optical Trap

The first stage in the experimental cycle is the MOT, loaded directly from the background rubidium vapor. For this, current is run through the quadrupole coils to generate a quadrupole field, with a gradient of ≈ 15 G/cm in the z -direction. The cooling power in each of the three MOT beams is ≈ 40 mW, which is $\approx 2\pi \times 20$ MHz red detuned from resonance, with the x - y plane beams additionally having ≈ 1 mW of on-resonance D2-repump power.

To keep the total experimental cycle time short we typically load the MOT for 250 ms. The atom number in the MOT saturates after several seconds of loading, with the load time used being in the regime where the atom number linearly increases with load time. However, we have observed that the atom number eventually loaded into the optical dipole trap saturates on a much faster timescale. Therefore, limited gain is obtained by increasing the loading time much beyond 250 ms, with a significant hit to the cycle rate of the experiment.

After the initial MOT loading stage we transition to a compressed magneto-optical trap (CMOT) by ramping up the quadrupole field from ≈ 15 to ≈ 32 G/cm over 10 ms. During the field ramp the dipole trapping beams are also ramped on. The magnetic field is held at this value for a further 12 ms. The intention here is to shrink the volume of the cloud to better match that of the optical dipole

trap. This must be done on a long enough timescale that compression occurs, but quickly enough that the losses associated with a large atomic density do not lead to a significant depletion of the atom number.

Additional compression of the cloud is performed using a so-called “temporal dark spot”. Here, the D2-repump and cooling power is decreased to $\approx 20 \mu\text{W}$ and $\approx 8 \text{ mW}$ respectively in each of the relevant beams. Additionally, the cooling light is further red detuned to $2\pi \times 100 \text{ MHz}$. Due to the small amount of D2-repump power the atoms are quickly pumped into the $F=1$ manifold, and are thus dark to the cooling light. This dramatically reduces the outward radiation pressure, causing the cloud to briefly collapse in on itself. However, as most of the atoms are not in a state addressed by the cooling light, the cloud is not particularly well trapped, therefore, we perform this step only for 8 ms.

4.3.2 Optical Molasses

At the end of the MOT stages the cloud is still relatively hot, 100s of μK , which is a similar order of magnitude to the depth of dipole trap. Therefore, in order to effectively load into the dipole trap we further cool the cloud using a gray optical molasses [81]. Here, the quadrupole field is extinguished and the magnetic bias fields are all tuned to zero out any stray magnetic field. The D2-repump and cooling powers are both increased slightly to $\approx 50 \mu\text{W}$ and $\approx 9 \text{ mW}$ respectively in each of the relevant beams. Additionally, the cooling detuning is increased such that it is $\approx 2\pi \times 215 \text{ MHz}$ red detuned, and the D2-repump light is $\approx 2\pi \times 51 \text{ MHz}$ blue

detuned from their transitions. In this configuration the D2-repump and cooling light are Raman resonant giving rise to dark states, which slow moving atoms are able to adiabatically follow. Meanwhile, fast moving atoms are not able to adiabatically follow the dark states and scatter photons, which results in them undergoing standard optical molasses cooling. This allows us to achieve phase space densities several times that which are achievable with a standard optical molasses. Empirically we find 15 ms of gray molasses cooling to be optimal for dipole trap loading, resulting in in-trap atom densities on the order of 10^{11} cm^{-3} and temperatures $\approx 10 \text{ } \mu\text{K}$.

4.3.3 Optical Pumping

Post molasses the atoms are distributed across the ground state manifold. As mentioned before, we wish to pump them into the $|F = 2, m_f = 2\rangle$ stretched state. To this end, we apply a bias field of $\approx 1.5 \text{ G}$ along the y-axis and turn on both D1-pump and D2-repump light in the optical pumping beam, with a power of $\approx 1 \text{ mW}$ and $\approx 250 \text{ } \mu\text{W}$. The optical pumping is performed in-situ in the optical dipole trap, with the D1-pump light $\approx 2\pi \times 80 \text{ MHz}$ blue detuned from resonance, and the D2-repump $2\pi \times 14 \text{ MHz}$ red detuned from resonance. We find that the pumping fidelity reaches near unity for pumping times greater than $\approx 10 \text{ ms}$. However, this step is performed for a duration of 20 ms to ensure experimental drifts in the pumping power do not decrease the state preparation fidelity.

At the end of the pumping duration the D1-pump and D2-repump light is

turned off. The bias field angle is then rotated linearly through $\approx 90^\circ$ over 25 ms till it is parallel with the probe beam. This has been empirically found to be sufficiently slow that the atoms adiabatically follow the rotating B-field.

4.4 Computer Control

As may have become apparent when we discussed the experimental cycle, we have a requirement for accurate and precise temporal control of experimental parameters. As we will see later, this timing is often required to be good to the order of 100s of nanoseconds. To that end, we take a similar approach to several others labs at JQI where a ‘primary’ device (Spincore PulseblasterUSB) is used as a clock for other ‘minion’ devices. We run a forked version of the JQI’s homegrown ‘Setlist’⁸ which provides a LabVIEW-based GUI which allows users to easily write a chronological table of instructions for each device. Extensive documentation for Setlist already exists, both in the original repository⁹ and on the JQI wiki¹⁰, so here I shall just discuss the specific details of the computer control pertaining to our experiment.

As mentioned, we use a Pulseblaster as our primary device, which has twenty-four¹¹ 3.3 V LVTTTL outputs that either serve as clock for minion devices, or are directly connected to pieces of equipment, e.g. AOM drivers, shutter drivers etc. A set of National Instrument cards (two NI PCI-6733 and two NI USB-6363) pro-

⁸<https://github.com/RbRy/SetList>

⁹<https://github.com/JQIamo/SetList>

¹⁰<https://jqj-wiki.physics.umd.edu/d/documentation/software/computercontrol/setlist/home>

¹¹unfortunately, three of these have met an untimely end in the line of duty

vide twenty¹² analog output channels, with an output range ± 10 V, and a further forty-eight¹³ 5 V TTL outputs. For the production of low frequency tones, up to 170 MHz, we use a DDS (Novatech 409B), which has two ‘static’ and a further two channels capable of cycling through a clock-able frequency list. A signal generator (Agilent E4426B) provides a method of generating tones up to 4 GHz, however, the functionality to alter this frequency mid-cycle has not yet been implemented in Setlist.

Although the control stack thus far described has enabled the relatively painless implementation of experimental procedures, it is not without its quirks. For one, the maximum current output of the Pulseblaster, and digital and analog outputs of the National Instrument devices mean that they are unable to drive LVTTTL, TTL, and the full range of voltage outputs respectively into 50Ω loads. This then leaves us with the unhelpful situation that either we try to have well impedance-matched loads, but with a voltage droop, or non-impedance matched loads, leading to ringing of signals which could potentially cause timing issues. To combat this we have taken an ad-hoc approach using unity gain buffers (Texas Instruments BUF634) on channels where issues have been discovered.

Prior to each experimental cycle some communication occurs between Setlist and the various devices it controls. The time it takes to do this is non-negligible, for example for a procedure including only the Pulseblaster this can be ≈ 100 ms, whereas for a procedure with all the devices listed above this time can be upwards

¹²minus two dead

¹³at time of writing these are all still alive and well

of a second. This can be problematic for us, where a typical experimental cycle time, ignoring this ‘dead’ time, is on the order of a few hundred milliseconds, as the real cycle time can end up being several times longer. For long data taking runs we have mitigated this dead time problem by programming certain devices outside of Setlist, with the Novatech and Agilent being the main offenders, and manually adding the clocking triggers to those devices into the Setlist procedure. One does, however, need to be careful as those devices can easily become de-synchronized from the Setlist procedure. As another solution, in some instances we have made use of the dead time to perform the initial MOT loading, although this has the downside that there is some fluctuation in the dead time, and thus a fluctuation in the MOT load time and therefore final atom number loaded into the dipole trap.

The Pulseblaster we use has a design feature/bug that, at the end of the experimental cycle, it sets all of its digital outputs low. This would normally mean that the lines would stay low during the dead time, which is not ideal if they are hooked up to equipment like AOM drivers where, for reasons already discussed, the driver wants to be on near-constantly. Helpfully, the Pulseblaster has the ability to program it to wait for an external trigger before continuing. Therefore, a kludgy fix we implement is, as the final instruction in the procedure, to tell the Pulseblaster to wait for a trigger that never comes, thereby causing it to retain the final digital output values till the next experimental cycle. Although this works for a large amount of the dead time, there is a small period, a few milliseconds, where Setlist causes the Pulseblaster to reset and its outputs are driven low. For something like the dead MOT loading this would mean that the MOT light would disappear briefly,

which is not ideal. For devices where uninterrupted service is required we, therefore, use the National Instruments card digital lines.

A final quirk of note is related to loops within the experiment. Often for Rydberg ensemble experiments you want to repeat some small section of the experimental cycle many times¹⁴. However, our version of the Pulseblaster has a memory capable of storing only four-thousand chronological instructions, which is clearly insufficient. Thankfully the Pulseblaster, and Setlist, have the ability to perform loops within the procedure, which enables us to perform these repeating sections of the cycle without exhausting the Pulseblaster memory. The downside is that Setlist renders non-primary channels, i.e. any that are not from the Pulseblaster, unchangeable during these loops. To get around this issue, we take the approach of removing from the Setlist procedure any non Pulseblaster devices that need changing during the loop, then externally programming those devices and setting their clocking triggers within Setlist.

4.5 Diagnostic Tools

For determining whether things are working as they should be, either when implementing/re-tooling one of our experimental schemes, or just for day-to-day operation, there are a number of tools/techniques we have at our disposal. Here, I want to touch on how these diagnostic tools operate and where they have proved useful.

¹⁴later we shall discuss procedures where we repeat a 5 μ s instruction set $\sim 10^5$ times per experimental cycle




Figure 4.26: Example of EIT transmission spectra. Fitted parameters here indicate an optical depth $OD \approx 37$, control Rabi frequency $\Omega_c \approx 2\pi \times 16$ MHz and ground-Rydberg dephasing rate $\gamma \approx 2\pi \times 500$ kHz.

4.5.1 Spectroscopy

This is by far our most commonly used diagnostic. Here, we take spectra, monitoring the transmission of the probe light through the cloud as a function of its frequency. Typically this is performed by taking the light exiting the chamber and coupling it into a multimode fiber which is connected to a single photon avalanche detector (SPAD). To obtain a value for the cloud transmission, three separate measurements are taken, one where the atoms and probing light are present, one where the atoms are not, and a final one with neither present, all taken within the same experimental cycle. Exact details of the full data collection process are unnecessary

for the discussion here, but are documented in section 4.6.

Generally we take the obtained spectra and fit it, as we have in figure 4.26, with the EIT transmission expression derived in equation 3.14. This allows us to simultaneously extract information about the optical depth, telling us whether our loading procedure is operating correctly, and the control Rabi frequency and ground-Rydberg dephasing rate, which together give information on problems coupling to or problems with the Rydberg states. Slightly more qualitatively, poorness of the fit, often in the form of some asymmetry in the spectra at the large detunings, can be indicative of some mis-alignment.

Due to the non-linearities associated with the Rydberg state it is necessary to perform these spectra at low probe photon flux. Otherwise, the EIT window can appear suppressed which leads to artificially inflated fitted values of γ . In a similar vein, even in the absence of the Rydberg control field, where we can ignore the Rydberg non-linearities, one must still be careful with the probe photon flux so as not to induce heating within the cloud.

As mentioned, the spectra are taken by coupling the probe light exiting the chamber into a multi mode fiber. Historically we had attempted to instead couple the light into a single mode fiber due to the increased spatial filtering of unwanted light sources, such as that of the dipole trapping light or control field. However, we noticed that spectra measured with the single mode fiber tended to have a higher degree of asymmetry, in addition to a larger fitted optical depth, to that measured with a multi mode fiber. It is believed that this is due to a lensing effect [87] present due to the high atomic density and the curvature of the cloud. This lensing causes

the probe light to not only be absorbed but deflected which, for the single mode fiber, manifests as loss as the fiber coupling is reduced. Therefore, we have moved away from and advise against extracting parameters from spectra obtained using single mode fibers in systems similar to ours.

4.5.2 Probe Transmission Imaging

While it is true that the transmission spectroscopy has wide utility in diagnosing problems within the cloud, it lacks information about the transverse spatial degree of freedom of the transmitted light which can occasionally be helpful. We have found images of the transverse mode of the probe transmitted through the cloud to be helpful in identifying potential misalignments of the probe with the atomic cloud and/or control field, as they tend to manifest in asymmetries in the mode structure.

To obtain this information, rather than sending the probe light post-chamber to a fiber, we instead direct the light through a set of lenses¹⁵ such that the in-chamber probe focal plane is imaged on to an EMCCD (Andor iXon 3). The high quantum efficiency of these devices mean that even with the low probe photon flux required to perform spectra, information about the transverse mode of the probe light transmitted through the cloud can be obtained often in a single experimental cycle.

Although imaging of the transverse mode of the probe has been presented

¹⁵apologies for the vagueness here, this is not commonly performed, and the imaging system changes almost every time it is

here as a diagnostic tool, there has been some discussion of using it to probe useful physics. For instance if one were able to in some way gate the camera, we could imagine performing an experiment where we looked at the correlations of the transverse spatial mode of the photons leaving the Rydberg ensemble medium, which to the best of my knowledge remains an experimentally unexplored problem.

4.5.3 Absorption Imaging

Unlike many other atomic physics experiments, images of the atomic cloud do not form a primary source of data in the physics we are investigating. However, we have found absorption imaging to be a useful diagnostic tool, mainly for the purpose of alignment, some cloud optimizations, and thermometry.

Primarily we perform imaging along the z-axis using light which follows the same beam path as the z-axis MOT light, but with the opposite polarization. After the chamber a PBS allows us to separate out the cooling and imaging light. Post-chamber we use a $\approx 1 : 1$ imaging system to form an image of the x-y plane on a CCD camera (Point Grey Flea3 FL3-GE-28S4M-C). When necessary, imaging along the y-axis can also be achieved by sending imaging light along the same path as the optical pumping beam. A pair of lenses post-chamber form an image of the x-z plane on another CCD camera (Point Grey Flea3 FL3-GE-28S4M-C). We follow standard absorption imaging practices to determine the transmission through the cloud, taking three images, one with the atoms and imaging light, one with just the imaging light, and one with neither.




Figure 4.27: Schematic for generalized photonic data collection. Only three outputs are shown for the optical processing here, however, in general this could be any number.

4.6 Photonic Data Collection

For all the projects discussed within this thesis, the raw data collected is that of a series of absolute timestamps, or time tags, of photon detection events on a set of SPADs. Here, I describe the experimental apparatus and software used for collecting this data.

Generally speaking data is collected using a scheme like that shown in figure 4.27. The post-chamber probe light is taken, and passed through some optical processing scheme, such as a Hanbury Brown and Twiss [88] or Hong-Ou-Mandel [89] interferometer. We couple the outputs of the optical processing into a set of fibers which are each connected to a SPAD (Excelitas SPCM-NIR- $\#$ -FC¹⁶). The SPADs operate in Geiger mode where, with some high probability, photons incident on the SPAD face cause an electron avalanche event. The avalanche events cause the

¹⁶Here, $\#$ refers to the dark count rate of the detector. We use both 13 and 14 models

Channel	Purpose
1	Defines measurement ‘windows’
2	Electronic clock/temporal reference
3-8	SPADs

Table 4.2: Time tagger connections.

SPAD to produce a TTL pulse which we send to a time tagging device (Roithner Laser TTM8000). In addition to the SPADs, we attach several lines from the Pulseblaster to the time tagger, as shown in table 4.2, in order to define various temporal references within the experimental cycle.

For actually collecting data the time tagger is physically connected, via gigabit Ethernet cable, to a computer upon which we run a small executable¹⁷ written with the time tagger’s C/C++ API. A LabVIEW front-end¹⁸ provides a more user-friendly interface for the executable for setting data collection parameters. We always operate the time tagger in ‘Continuous-I’ mode¹⁹, recording all rising edge events on the SPAD and clock channels specified in the LabVIEW GUI, and rising and falling edge events on the ‘window’ channel, defined in table 4.2. The host-computer-based software continuously reads from the tagger, writing to file the tags that occur during periods where the ‘window’ channel is high. In this way, we can filter the tags that occur during some experimentally relevant time period.

Strictly speaking the tags for each event are 64-bit words, with a large number of the bits reserved for recording the timestamp of the tag, with a temporal resolution

¹⁷Github repository for this software: <https://github.com/acraddoc91/timeTaggerRbRy>

¹⁸Github repository for front-end: <https://github.com/acraddoc91/DataAcquisitionRbRy>

¹⁹see user manual: http://www.roithner-laser.com/datasheets/accessories/ttm8000_manual.pdf

Bit	Purpose
0	High/Low Word Indicator
1-31	Temporal Payload

Table 4.3: Anatomy of time tagger high words.

Bit	Purpose
0	High/Low Word Indicator
1-27	Temporal Payload
28	Slope
29-31	Channel Number

Table 4.4: Anatomy of time tagger low words.

of ≈ 82.3 ps. To gain nearly a factor of two compression the tags are actually recorded as 32-bit ‘high’ and ‘low’ words with structures as defined in tables 4.3 and 4.4 respectively, where the low word is recorded for every event, and the high word only when it changes.

In ‘Continuous-I’ mode the time tags are referenced such that $t = 0$ corresponds to the time that the time tagger is initialized. Since this time is pretty arbitrary, it is, in general, not a very helpful reference point. Therefore, in post processing we usually re-reference tags to the start of their respective measurement ‘window’ by subtracting off the time tag corresponding to the rising edge on channel 1 of the tagger.

This data collection scheme is not free from its own oddities. Most notable of these is related to the time tagger resolution. Through use, we have noticed a peculiar property of the tags produced by the tagger, which is illustrated by figure 4.28. Here, we have taken a sample data set with a couple hundred thousand




Figure 4.28: Fraction of tags from time tagger with a temporal payload bit value equal to one, for several time tagger channels. For the data shown the tags on each channel are due to SPAD detection events for continuous wave light.

tags, and are looking at the values for each bit of the temporal payload as a function of the bit and channel number. For almost all the bits the fraction of the tags for which the bit value is one is about fifty percent, as would be expected for continuous wave light where the photons are arriving at the detectors randomly. However, for the zeroth bit there is a clear channel-dependent bias. From discussion with the manufacturer this is apparently a known problem with the time-to-digital converter chip. Therefore, we have, and we recommend, ignoring the least-significant bit of the tagger. While this changes the resolution from ≈ 82.3 ps to ≈ 164.6 ps, this has not been an issue as the SPADs themselves have a temporal resolution larger than this, on the order of a few hundred picoseconds.

A further issue relates to gating. The SPADs themselves are rather delicate instruments, saturating at a few tens of million detection events per second, with damage possible at higher rates. Therefore, it is often desirable to temporarily gate the detectors, for example to protect them from bright light sources. The SPADs themselves do have a hardware gate. However, when this has been used we have seen that a detection event is observed with higher than expected probability when the SPAD is gated on, with the problem getting worse as a function of the light intensity falling on the SPAD during the time it was gated off. As a result hardware gating, to protect the SPADs from high intensity light, has been implemented in the form of AOMs used as shutters. For the situation where gating wants to be performed not to protect the SPADs, but for temporal filtering, we typically use a software based approach, where certain regions of time are ignored during post-processing.

4.7 Experimental Techniques

In the process of building (and re-building) the apparatus we have developed a number of techniques which have greatly accelerated certain tasks. Here, I document some of these in the hope that it may be helpful for students working on this, and other projects.

4.7.1 Spatial Alignment

For the apparatus to function properly, we require several tightly focused beams to all intersect within the chamber. This is non-trivial and a a non-negligible

fraction of my time as a graduate student has been spent performing intricate alignments. However, I have picked up a few tricks that speed up the process that I wish to pass on.

4.7.1.1 Probe

For our setup we are relatively free to pick an absolute reference within the chamber, to which everything shall be aligned. The in-vacuum lenses do, however, provide some constraint as a beam sent off-axis can potentially cause the focused light to be aberrated. This can be problematic for a beam, such as the probe, where we desire the transverse mode to be contained within a Rydberg blockade radius. Therefore, we send our probe beam through the center of the in-vacuum lenses, with the light entering the chamber collimated. The focus of the probe in-vacuum then serves as our absolute reference for all the other beams within the experiment.

When sending the probe through the chamber several reflections off the various surfaces can be seen, albeit faintly as they are all AR coated at the probe wavelength. Off the two windows we see a set of reflections that are approximately collimated and, due to the angle of the lenses to the windows, should subtend some angle to the injected light when correctly aligned. The reflections off the two in-vacuum lenses are all diverging outside the chamber, with the exception of the reflection off the curved face of the second encountered lens, which focuses ≈ 2 cm from the input window. Once the four reflections from the in-vacuum lenses are identified, they can be used for alignment by walking the probe beam to overlap the input and reflected

light.

4.7.1.2 Dipole Trap

We align the dipole trapping light such that the trapped atomic cloud is centered on the focus of the probe beam. This is most straightforwardly accomplished by red detuning the probe from resonance, usually $\approx 2\pi \times 500$ GHz. Here, with a few hundred microwatts of power the probe beam forms a deep potential capable of optically dipole trapping atoms, as seen in figure 4.29. Imaging the trapped cloud allows us to precisely locate the probe focus both transversely, due to its narrow waist, and longitudinally, due to its short Rayleigh range. By similarly imaging the trapped atoms in the three dipole beams we are straightforwardly able to identify and adjust their transverse and longitudinal alignments to overlap with the focus of the probe.

This technique works well with our two, orthogonal, imaging systems for aligning the dipole arms in the three dimensions. However, there is an issue with the dimple beam in that it propagates along the same direction as the imaging light when imaging along the y-axis, preventing us from using this technique to set the z-axis alignment of this beam. Therefore, in order to correctly align the dimple beam in the z-axis we typically first align the arms, then tune the z-axis alignment of the dimple beam till it intersects with that of the arms, which can readily be seen using the z-axis imaging system. Once this coarse alignment is achieved, finer alignment is usually performed by monitoring the near-resonance probe transmission as




Figure 4.29: Absorption image of the trapped atoms formed with the 1004 nm optical dipole trapping light (top) and the far detuned probe light (bottom). Absorption image taken with z-axis imaging system. Note the color scales for the two images are different.

a function of the height of the dimple beam, with the ‘correct’ alignment taken to be where the absorption is maximized.

When adjusting the alignment of the dipole beams we not only want to ensure their overlap with the probe focus, but also want to ensure all the dipole beams propagate in the x-y plane, so as not to distort the cloud shape. For the arms we can verify straightforwardly through imaging if this is the case, while for the dimple the condition is practically satisfied owing to the geometry of the system. Additionally, we want the crossing angle to be symmetric about the probe for the

two arms, which is guaranteed by requiring the two dipole beams enter the chamber parallel to and equidistant from the probe. Alignment is thus performed by walking the beams to meet these additional constraints, along with that of intersecting the probe.

4.7.1.3 Control Field

For optimal operation the control and probe beams must be overlapped and counter-propagating. To perform this alignment we first align light in a tracer beam, of wavelength intermediate between the probe and control fields, usually generated by a fiber pen with a output wavelength ≈ 650 nm, and emanating from a single mode fiber. The tracer beam co-propagates with the probe, and is aligned to it by optimizing the coupling of both beams post-chamber into the same single mode fiber. We then align the control field light by coupling it into the fiber launch of the tracer beam. After this procedure the probe and control fields are usually suitably well overlapped that a Rydberg EIT signal can be observed. Finer alignment of the control field is then performed by maximizing the doubly resonant ($\delta = \Delta = 0$) transmission through the cloud.

4.7.2 Polarization Alignment

In addition to spatial alignments, correct polarization alignment at the atoms for many of the beams is also important. Given that the optical pumping fidelity is heavily influenced by the polarization purity of the pump beam, its polarization

needs to be carefully aligned. We do this by increasing the pumping time far past the saturation time, usually to hundreds of milliseconds. As scattering only occurs if the light is not perfectly σ -polarized, one can perform the polarization alignment by performing the pumping and monitoring the atom loss in the cloud. Both the magnetic field direction during pumping and waveplates that set the pump field's polarization in-chamber are adjusted to minimize the atom loss.

A similar method is also performed for fine probe polarization alignment. The probe field is tuned to the $F = 2$ to $F' = 2$ transition of the D2 line and atom loss monitored as a function of the magnetic field direction during probing and the probe waveplates, when probing for a long time. Although atom loss can still occur with perfect polarization purity, due to off-resonant scatter on the $F = 2$ to $F' = 3$ transition, we have found the technique still provides enough resolution to be able to tune waveplates and magnetic fields to within a few degrees.

The control field polarization is somewhat less critical than the other two, in that poor alignment will only lead to a reduction in the control Rabi frequency, as the π and incorrect σ have nothing to couple to. However, we do want to make the most of the control power we have. Alignment here is usually accomplished, as with the fine control beam alignment, by monitoring the doubly resonant ($\delta = \Delta = 0$) EIT transmission and trying to maximize it. Note here we only adjust the waveplates associated with control light as the magnetic field direction is fixed when the probe polarization alignment is performed.




Figure 4.30: Coincidences measured between two SPADs as a function of delay time off a generalized Hanbury Brown and Twiss setup, for probe light obtained by driving the double pass AOM at two tones. The frequency difference between the tones here is 10 MHz. The blue points represent actual data, while the orange line is a fit of the form $a \cos(\omega(\tau + t_{\text{off}})) + b$, with the fitted values being $t_{\text{off}} \approx 4.2$ ns, $a \approx 3000$ and $b \approx 6000$.

4.7.3 Temporal Synchronization

As we shall see later, for a lot of our projects we shall be interested in extracting information about the system by looking at the coincidence photon detection events of two or more detectors. Further, we are in general going to be interested in these coincidence events as a function of detection time between the various SPADs. Therefore, we want to try and mitigate any inherent unwanted delays that may exist, be they from optical or electronic path length differences.

By applying two different tones to the probe double pass AOM, the resulting

light arriving at the chamber has an intensity which oscillates at twice the frequency difference of the tones. Provided the light passes through the chamber unhindered, i.e. with no atoms present, and the optical processing resembles that of a generalized Hanbury Brown and Twiss [88] setup, this beating can be seen in the coincidence detection events between pairs of SPADs, as seen in figure 4.30. In the absence of any delay between the detectors a maxima in the coincidence should occur at zero delay time between the detectors. Any inherent delay will cause an offset in coincidence space that can be compensated for, which we usually do in software/post-processing.

Care does have to be taken as the temporal offset can be hard to determine when it becomes of similar, or greater, order to the beat period due to aliasing. To ensure the temporal offset is correctly identified, we usually take several measurements with progressively increasing beat frequencies.

Chapter 5: Measurement and Calculation of Photonic Correlation Functions

Within quantum optics, light is commonly characterized using a set of normalized correlation functions $\{g^{(n)}\}$, where $g^{(n)}$ is referred to as the n th degree of temporal coherence [90].

5.1 First Order Coherence

The degree of first order coherence quantifies the electric field fluctuations of the light being measured, and is defined by [90]

$$g^{(1)}(t, \tau) = \frac{\langle E^-(t)E^+(t+\tau) \rangle_t}{\sqrt{\langle E^-(t)E^+(t) \rangle \langle E^-(t+\tau)E^+(t+\tau) \rangle_t}}, \quad (5.1)$$

where E^+ and E^- correspond to the rotating and counter-rotating terms of the electric field, and $\langle \rangle_t$ denotes a time averaged quantity. Practically speaking, measurement of $g^{(1)}$ differs from that of the higher order coherences. Furthermore, we shall not be concerned with $g^{(1)}$ for the light discussed in the rest of this thesis and, as such, shall not explore it further.

5.2 Second Order Coherence

5.2.1 Formal Definition

The degree of second order coherence is formally defined as [90]

$$g^{(2)}(t, \tau) = \frac{\langle E^-(t)E^-(t+\tau)E^+(t+\tau)E^+(t) \rangle_t}{\langle E^-(t)E^+(t) \rangle_t \langle E^-(t+\tau)E^+(t+\tau) \rangle_t}. \quad (5.2)$$

It quantifies intensity fluctuations of the light, and is an important quantity in the field of quantum optics, as we shall see later.

5.2.2 Classical Limit

In the classical limit, electric fields commute and we are able to write equation 5.2 in terms of the intensity, $I(t) \propto E^-(t)E^+(t)$, of the light

$$g_{\text{classical}}^{(2)}(\tau) = \frac{\langle I(t)I(t+\tau) \rangle_t}{\langle I(t) \rangle_t \langle I(t+\tau) \rangle_t}. \quad (5.3)$$

It should be noted that in equation 5.3 we have dropped the explicit dependence of $g^{(2)}$ on t , as we shall do for the remainder of this section. This is strictly only correct to do for stationary light. However, it is not uncommon for the t dependence to be ignored for non-stationary light, for example pulsed light as in references [41, 42, 91].

The classical commutativity of the electric field has some important conse-

quences for $g^{(2)}$. Given the positivity of the variance

$$\text{Var}(A) = \langle A^2 \rangle - \langle A \rangle^2 \geq 0 \quad (5.4)$$

then it can be seen that

$$g_{\text{classical}}^{(2)}(0) = \frac{\langle I^2(0) \rangle_t}{\langle I(0) \rangle_t^2} \geq 1 \quad (5.5)$$

Additionally, from the Cauchy-Schwarz inequality, $\langle AB \rangle^2 \leq \langle A^2 \rangle \langle B^2 \rangle$, we can show

$$g_{\text{classical}}^{(2)}(\tau) \leq \frac{\sqrt{\langle I(t) \rangle_t} \sqrt{\langle I(t+\tau) \rangle_t}}{\langle I(t) \rangle_t \langle I(t+\tau) \rangle_t} = g_{\text{classical}}^{(2)}(0). \quad (5.6)$$

These two inequalities together preclude the classical production of important quantum optical phenomenon, such as single photons.

5.2.3 Quantum Mechanically

Quantum mechanically we can write the electric field in terms of the raising, \hat{a}^\dagger , and lowering, \hat{a} , operators of the electromagnetic field

$$\hat{E}^+ \propto \hat{a}, \quad (5.7)$$

$$\hat{E}^- \propto \hat{a}^\dagger, \quad (5.8)$$

allowing us to write equation 5.2 in terms of these quantum operators

$$g^{(2)}(\tau) = \frac{\langle \hat{a}^\dagger(t) \hat{a}^\dagger(t+\tau) \hat{a}(t+\tau) \hat{a}(t) \rangle_t}{\langle \hat{n}(t) \rangle_t \langle \hat{n}(t+\tau) \rangle_t}, \quad (5.9)$$

where $\hat{n} = \hat{a}^\dagger \hat{a}$ is the number operator.

5.2.4 Probabilistic definition

In addition to its formal definition in terms of fields or intensities, I have often found it intuitive to think about $g^{(2)}$ in a probabilistic way. Working from the quantum mechanical definition, let us first consider $g^{(2)}(\tau = 0)$. We can make use of the commutation relation, $[\hat{a}(t), \hat{a}^\dagger(t)] = 1$, to write $g^{(2)}(\tau = 0)$ purely in terms of the number operator

$$g^{(2)}(0) = \frac{\langle \hat{n}(t)(\hat{n}(t) - 1) \rangle_t}{\langle \hat{n}(t) \rangle_t^2}. \quad (5.10)$$

Considering some arbitrary state of light

$$|\psi\rangle = \sum_{n=0}^{\infty} c_n |n\rangle, \quad (5.11)$$

where $|n\rangle = (\hat{a}^\dagger)^n |0\rangle$ is an n -photon Fock state, we see that

$$g^{(2)}(0) = \frac{\langle \sum_{n=2}^{\infty} n(n-1) |c_n|^2 \rangle_t}{\langle \sum_{n=1}^{\infty} n |c_n|^2 \rangle_t^2}. \quad (5.12)$$

In the limit that $|c_n|^2 \gg |c_{n+1}|^2$, we find

$$g^{(2)}(0) \approx \frac{2 \langle P_2 \rangle_t}{\langle P_1 \rangle_t^2}, \quad (5.13)$$

where $P_i = |c_i|^2$ is the probability of observing i photons. We can see that $g^{(2)}(0)$ essentially quantifies how likely we are to observe a two-photon event, relative to a

single photon one. Therefore, $g^{(2)}(0)$ is a useful property for characterizing a single photon source, for which we want the likelihood of producing multiple photons to be small.

For sufficiently large τ we can take $[\hat{a}(t), \hat{a}^\dagger(t + \tau)] = 0$, and thus we can write

$$g^{(2)}(\tau) = \frac{\langle \hat{n}(t)\hat{n}(t + \tau) \rangle_t}{\langle \hat{n}(t) \rangle_t \langle \hat{n}(t + \tau) \rangle_t}. \quad (5.14)$$

Making the assumption that $P_n(t) \gg P_{n+1}(t)$, we can write this in terms of probabilities

$$g^{(2)}(\tau) \approx \frac{\langle P_1(t)P_1(t + \tau) \rangle_t}{\langle P_1(t) \rangle_t \langle P_1(t + \tau) \rangle_t} = \frac{\langle P(t + \tau | t) \rangle_t}{\langle P_1(t + \tau) \rangle_t}, \quad (5.15)$$

where $P(t + \tau | t)$ is the probability of a photon being present at $t + \tau$ given that one was present at t . Reiterating this in words, $g^{(2)}(\tau)$ tells us how likely we are to observe a photon at a time $t + \tau$, given that we have seen one at t , relative to random chance.

5.2.5 Measurement

For a classical photon source, with a large photon flux $g^{(2)}$ can be readily measured using a single standard photodiode up to the bandwidth of the detector, by measuring the intensity as a function of time and straightforwardly using equation 5.3. At low enough fluxes, like those we shall be concentrating on in this thesis, one needs to work with detectors which have resolution at the single photon level. Ideally a photon-number resolving detector would allow us to use the same mea-




Figure 5.1: Hanbury Brown and Twiss setup.

surement process as for the bright light source. However, photon-number resolving detectors typically have a low bandwidth and are not widely commercially available, making them unsuitable and impractical for all but a narrow subset of applications. In contrast, single photon detector (SPD) technology is mature enough that high bandwidth, high quantum efficiency SPDs are commercially available. Generally speaking, however, SPDs suffer from an issue known as ‘dead time’. That is the time, once a detection event has occurred, in which a further detection event is impossible. This makes measurement of $g^{(2)}$ at short delay times impossible with a single SPD.

To get around this a Hanbury Brown and Twiss setup, shown in figure 5.1, is commonly employed, utilizing a beamsplitter, which is assumed not to alter the underlying statistics of the light, and a pair of SPDs. By counting the coincidences between the two detectors, $N_{coincidence}$, as well as the the single detection events on the two detectors, $N_{1,2}$, one can then approximate $g^{(2)}$ [92]

$$g^{(2)}(\tau) \approx \frac{T}{\Delta\tau} \frac{N_{coincidence}(\tau)}{N_1 N_2}, \quad (5.16)$$




Figure 5.2: Generalized Hanbury Brown and Twiss setup.

where T is the total measurement time, and $\Delta\tau$ is the temporal bin width used for determining coincidences. This approximation is valid in the regime that $P_1(t) \gg P_2(t) \gg P_{n>2}(t)$, which shall be the case for the work in this thesis.

5.3 Higher Order Coherences

For higher order coherences of light, we can generalize most of the expressions for the second order coherences. The formal definition of the n -th order coherence, where $n > 1$, is given by

$$g^{(n)}(t, \tau_1, \tau_2 \dots) = \frac{\langle E^-(t) \prod_{i=1}^{n-1} E^-(t + \tau_i) \prod_{i=n-1}^1 E^+(t + \tau_i) E^+(t) \rangle_t}{\langle E^-(t) E(t) \rangle_t \prod_{i=1}^{n-1} \langle E^-(t + \tau_i) E^+(t + \tau_i) \rangle_t}, \quad (5.17)$$

which can be written in its classical, quantum, and probabilistic form

$$g_{classical}^{(n)}(\tau_1, \tau_2 \dots) = \frac{\langle I(t) \prod_{i=1}^{n-1} I(t + \tau_i) \rangle_t}{\langle I(t) \rangle_t \prod_{i=1}^{n-1} \langle I(t + \tau_i) \rangle_t}, \quad (5.18)$$

$$g_{quantum}^{(n)}(\tau_1, \tau_2 \dots) = \frac{\langle \hat{a}^\dagger(t) \prod_{i=1}^{n-1} \hat{a}^\dagger(t + \tau_i) \prod_{i=n-1}^1 \hat{a}(t + \tau_i) \hat{a}(t) \rangle_t}{\langle \hat{n}(t) \rangle_t \prod_{i=1}^{n-1} \langle \hat{n}(t + \tau_i) \rangle_t}. \quad (5.19)$$

$$g_{probabilistic}^{(n)}(0, 0 \dots) \approx \frac{n! \langle P_n \rangle_t}{\langle P_1 \rangle_t^n}, \quad (5.20)$$

$$g_{probabilistic}^{(n)}(\tau_1, \tau_2 \dots) \approx \frac{\langle P(t + \tau_{n-1} | t, t + \tau_1, t + \tau_2, \dots) \rangle_t}{\langle P(t + \tau_{n-1}) \rangle_t}, \quad (5.21)$$

where we have taken the probabilistic definitions to lowest order, and made the assumption $P_n(t) \gg P_{n+1}(t)$. Measurement of the higher order coherences can be performed with a generalized Hanbury Brown and Twiss setup, shown in figure 5.2, where n SPDs are required to measure $g^{(n)}$.

5.4 Note on Losses

A contributing reason for the ubiquity of the correlation functions, $\{g^{(n)}\}$, for the measurement of classical and quantum optical photon statistics is their robustness to loss. This robustness is readily seen classically. Loss simply applies a scale factor to the light intensity, $I' = \mathcal{T}I$, where \mathcal{T} is the transmission probability through some lossy medium. The classical correlation function is modified

$$\begin{aligned} g_{classical}^{(n)}(\tau_1, \tau_2 \dots) &= \frac{\langle \mathcal{T}I(t) \prod_{i=1}^{n-1} \mathcal{T}I(t + \tau_i) \rangle_t}{\langle \mathcal{T}I(t) \rangle_t \prod_{i=1}^{n-1} \langle \mathcal{T}I(t + \tau_i) \rangle_t} \\ g_{classical}^{(n)}(\tau_1, \tau_2 \dots) &= \frac{\mathcal{T}^n \langle I(t) \prod_{i=1}^{n-1} I(t + \tau_i) \rangle_t}{\mathcal{T}^n \langle I(t) \rangle_t \prod_{i=1}^{n-1} \langle I(t + \tau_i) \rangle_t} \\ g_{classical}^{(n)}(\tau_1, \tau_2 \dots) &= \frac{\langle I(t) \prod_{i=1}^{n-1} I(t + \tau_i) \rangle_t}{\langle I(t) \rangle_t \prod_{i=1}^{n-1} \langle I(t + \tau_i) \rangle_t}, \end{aligned} \quad (5.22)$$

which is identical to that without loss.

We can also see this robustness against loss in the quantum world by consid-

ering an arbitrary state

$$|\psi\rangle = \sum_{n=0}^{\infty} c_n |n\rangle = \sum_{n=0}^{\infty} c_n \frac{(\hat{a}^\dagger)^n}{\sqrt{n!}} |0\rangle. \quad (5.23)$$

Quantum mechanically we can model loss by imagining passing the light through a beamsplitter with transmission coefficient $t = \sqrt{\mathcal{T}}$. We shall assume, without loss of generality, that $t \in \mathbb{R}$, therefore \mathcal{T} is identical to the classical transmission probability we previously defined. For such a beamsplitter the raising operator transforms as

$$\hat{a}^\dagger = t\hat{b}^\dagger + \sqrt{1-t^2}\hat{c}^\dagger, \quad (5.24)$$

where \hat{b}^\dagger is the raising operator associated with the transmitted mode, and \hat{c}^\dagger that for the lost photons. Our state thus transforms

$$|\psi\rangle \rightarrow |\phi\rangle = \sum_{n=0}^{\infty} c_n \frac{(t\hat{b}^\dagger + \sqrt{1-t^2}\hat{c}^\dagger)^n}{\sqrt{n!}} |0\rangle. \quad (5.25)$$

Making use of the binomial expansion this can be written

$$|\psi\rangle \rightarrow |\phi\rangle = \sum_{n=0}^{\infty} c_n \sum_{k=0}^n \sqrt{\frac{n!}{k!(n-k)!}} t^k (\sqrt{1-t^2})^{n-k} |k, n-k\rangle, \quad (5.26)$$

which can also be written in density matrix form

$$\begin{aligned} \hat{\rho} = |\phi\rangle\langle\phi| &= \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} c_n c_m^* \sum_{k=0}^n \sum_{l=0}^m \sqrt{\frac{n!}{k!(n-k)!}} \sqrt{\frac{m!}{l!(m-l)!}} \\ &\quad \times t^{k+l} (1-t^2)^{\frac{n+m-k-l}{2}} |k, n-k\rangle\langle l, m-l|. \end{aligned} \quad (5.27)$$

Given that the we can only care about the transmitted photons we trace out the lost photon mode

$$\begin{aligned} \hat{\rho}_t &= \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} c_n c_m^* \sum_{k=0}^n \sqrt{\frac{n!}{k!(n-k)!}} \sqrt{\frac{m!}{(m-n+k)!(n-k)!}} \\ &\quad \times t^{2k+m-n} (1-t^2)^{n-k} |k\rangle\langle m-n+k|, \end{aligned} \quad (5.28)$$

to obtain the density matrix for the transmitted mode. Let us now evaluate

$$\begin{aligned} \left\langle \left(\hat{b}^\dagger \right)^d \left(\hat{b} \right)^d \right\rangle &= \text{Tr} \left[\left(\hat{b}^\dagger \right)^d \left(\hat{b} \right)^d \hat{\rho}_t \right] \\ &= \sum_{n=0}^{\infty} |c_n|^2 \sum_{k=d}^n \frac{n!}{k!(n-k)!} t^{2k} (1-t^2)^{n-k} \frac{k!}{(k-d)!} \\ &= t^{2d} \sum_{n=d}^{\infty} |c_n|^2 \frac{n!}{(n-d)!}. \end{aligned} \quad (5.29)$$

From which we can calculate the d-th order coherence

$$\begin{aligned}
g^{(d)} &= \frac{\langle (\hat{b}^\dagger)^d (\hat{b})^d \rangle}{\langle \hat{b}^\dagger \hat{b} \rangle^d} \\
&= \frac{t^{2d} \sum_{n=d}^{\infty} |c_n|^2 \frac{n!}{(n-d)!}}{t^{2d} \left(\sum_{n=1}^{\infty} |c_n|^2 n \right)^d} \\
&= \frac{\sum_{n=d}^{\infty} |c_n|^2 \frac{n!}{(n-d)!}}{\left(\sum_{n=1}^{\infty} |c_n|^2 n \right)^d},
\end{aligned} \tag{5.30}$$

which is independent of the loss term.

5.5 Note on Normalization

As seen in the previous sections, the photonic correlation functions contain a normalization factor. These factors are crucial as they allow for the direct comparison of sources of light without needing to worry about the relative intensities or brightness of the sources themselves. In the above we have made assumptions about the stationarity of the light being measured. However, experimentally speaking this is often not the case either deliberately, as is the case for a pulsed light source, or as a result of imperfections, as for experimental drifts. Therefore, we need to take care in how we normalize our correlation functions, as it can have undesired consequences.

As a demonstration-by-example of this phenomenon let us consider an experiment of the form described in reference [42]. Here, the experimentalists were interested in measuring the $g^{(2)}$ correlation function of probe light that has passed




Figure 5.3: $g^{(2)}$ calculated for coherent light with a 50% duty cycle square pulse envelope using different normalization methods. Equation 5.16 used for blue curve and equation 5.31 used for orange curve.

through a Rydberg-EIT medium that caused attractive interactions between the polaritons. In order to avoid unwanted differential light shifts from the trapping light, the probe and trapping light were strobed out of phase with each other. If they were to perform the measurement scheme outlined in section 5.2.5 and utilize equation 5.16 to determine $g^{(2)}$, they would find it would be the product of the polaritonic correlations of interest, along with an envelope, similar to that shown in figure 5.3, due to the pulsed nature of the probe light. This is clearly undesirable, as the uninteresting pulse-shape-induced correlations could obscure the interesting ones that originate from the physics being probed.

As an alternative, one can instead calculate

$$g^{(2)}(\tau) = \frac{N_{\text{coincidence}}(\tau)}{N_{\text{coincidence}}(\tau + T)}, \quad (5.31)$$

where T is some fixed time that is chosen to appropriately account for any undesired structure. As an example, for a pulsed source, T could be chosen to be an integer multiple of the pulse repetition rate, as has been done in the orange curve in figure 5.3. Or for an experiment where there is a long time scale experimental drift, T could be chosen such that it exceeds the time scale of the physics-of-interest but is smaller than that of the drift. In the limit that the light being measured is stationary then if T is sufficiently large, the arrival of photons should be uncorrelated with one another, thus

$$N_{\text{coincidence}}(\tau + T) = \frac{\Delta\tau}{T} N_1 N_2, \quad (5.32)$$

and we can see that equation 5.16 and 5.31 are equivalent.

This normalization technique, however, is not without its drawbacks. For instance, one has to be careful in picking T . If it becomes comparable to the timescale of the correlations of interest, it can potentially wash them out when $g^{(2)}$ is calculated. Further, the coincidence rate is generally going to be significantly smaller than the singles count rates on the SPDs. As a result, the uncertainty in the measured value of $g^{(2)}$ for a given measurement time is increased compared the traditional normalization method. However, this can be somewhat mitigated by

sampling several values of T and averaging.

While the normalization technique presented here is specific to calculating $g^{(2)}$, it is generalized to higher order correlation functions

$$g^{(n)}(\tau_1, \tau_2 \dots) = \frac{N_{\text{coincidence}}(\tau_1, \tau_2 \dots)}{N_{\text{coincidence}}(\tau_1 + T_1, \tau_2 + T_2, \dots)} \quad (5.33)$$

where the T_i 's should be chosen such that $T_i \neq T_j \forall i, j$.

5.6 Calculation

As has been seen in the previous section, the calculation of photonic autocorrelation functions, and indeed any correlation function calculated from discretely sampled data, is essentially a problem of counting coincidences. At face value this seems like a simple problem. However, it is one that can be computationally expensive. Throughout the course of my graduate studies I have devoted a not insignificant amount of time into developing various algorithms and related code for the counting of coincidences. Here, I present a semi-chronological report of this work so that future students may learn from my mistakes.

5.6.1 Problem to be Solved

For what follows, we shall assume that we have performed some measurement using a pair of SPDs for which we have obtained a time-tagged pair of vectors, \mathbf{u}^1 and \mathbf{u}^2 , containing the photon arrival times at each detector. Further, we shall assume, due to the time-tagged nature of the data, that both vectors are chrono-

logically ordered such that $u_{i+1} > u_i$. From this data we wish to calculate two-fold coincidences¹ between the two detectors, for some set of delay times, $\{\tau_i\}$, and some time binning, Δt_{bin} , of the time-tags.

5.6.2 Methods

In the following I present the methods trialed for calculating coincidences. A brief outline of each method precedes a formal description of the algorithm., which is then followed by a short discussion on the relative merits and deficiencies of the approach.

5.6.2.1 Per-Photon Histogramming

In this approach, for each photon from detector 1 we calculate the relative times of the photons from detector 2 and produce a histogram. Algorithmically:

1. Take tag u_1^1 and perform a histogram of $\mathbf{u}^2 - u_1^1$, with desired Δt_{bin} , to obtain a histogram vector \mathbf{H}_1
2. Repeat for all u_i^1
3. Calculate $\mathbf{coinc} = \sum_i \mathbf{H}_i$

This algorithm is relatively simple to code and conceptually is most similar to how one thinks about coincidences. A single pass through of the algorithm gives us information for all desired τ 's. This was the first method we used for calculating

¹while the following specifically focuses on two-fold coincidences the techniques are readily scalable to higher-order

coincidences, but was found to be prohibitively slow. The algorithm requires the calculation of $N(\mathbf{u}^1)$ histograms², which have a computational complexity equal to the sum of the length of the vector and the number of histogram bins, and is likely the main limitation on calculation speed. As such the total complexity of the algorithm scales as $O(N(\mathbf{u}^1) * N(\mathbf{u}^2))$. This scaling becomes worse for the higher order correlation functions as the complexity scales multiplicatively in the length of the vectors involved. For a calculation of n-fold coincidences the complexity goes as $O(\prod_i^n N(\mathbf{u}^i))$.

5.6.2.2 Double Histogram

In this approach a pair of histograms for the arrival times of photons from both detectors are created. Correlations are then calculated from the histograms. Algorithmically:

1. Histogram, with desired Δt_{bin} , \mathbf{u}^1 and \mathbf{u}^2 , from $t = 0$ to $t = \max(\mathbf{u}^1, \mathbf{u}^2)$, to obtain a new set of vectors \mathbf{h}^1 and \mathbf{h}^2
2. For $\tau = n \times \Delta t_{bin}$, calculate $\text{coinc}(\tau) = \sum_i h_i^1 h_{i+n}^2$
3. Repeat step 2 for each desired τ

This algorithm is also relatively straightforward to implement and was the second one attempted for calculating coincidences. However, it was also found to be restrictively slow and was soon discarded in favor of a faster method. Here, the limiting

²Henceforth $N(\mathbf{u})$ shall be used as shorthand for the number of elements in a vector \mathbf{u} .

factors in computation speed are the calculation of the vector dot products and the calculation of the two histograms. As such, the algorithm has portions with computational complexities scaling as $O(N(\mathbf{u}^1))$, $O(N(\mathbf{u}^2))$, and $O(N(\{\tau_i\})N(\mathbf{h}^1)N(\mathbf{h}^2))$. It follows that for n th-order coincidences the algorithm will have parts with complexities scaling with $O(N(\mathbf{u}^i))$ and $O((N(\{\tau_i\}))^{n-1} \prod_i^n N(\mathbf{h}^i))$.

A further problem with this method is that it is not uncommon for the flux of photons at each detector to be significantly lower than $\frac{1}{\Delta t_{bin}}$. As a result, \mathbf{h}^i can be incredibly sparse and thus inefficient objects to store in memory.

Finally, it is worth noting that this method is less exact than the “per-photon histogram” method as one picks up rounding errors on both channels when performing the histograms. However, provided Δt_{bin} is chosen to be smaller than any structure seen in the correlation function the difference between the two methods is small.

5.6.2.3 Reduced Double Histogram

This approach makes use of the time-tagged nature of the data by performing a sorted search to determine the number of instances where the time difference between photons on the two detectors is equal to τ . Algorithmically:

1. For the desired Δt_{bin} calculate a new set of vectors $\mathbf{v}^1 = \mathbf{u}^1 \frac{\Delta t_{tagger}}{\Delta t_{bin}}$ and $\mathbf{v}^2 = \mathbf{u}^2 \frac{\Delta t_{tagger}}{\Delta t_{bin}}$, where Δt_{tagger} is the width of the least significant bit of the time tagger. Here, the elements of \mathbf{v}^i should be integers, so the scaling should be performed then rounded

2. Determine $n = \frac{\tau}{\Delta t_{bin}}$
3. Create a pair of pointers, $i = 0$ and $j = 0$, and a running total of the coincidences, $\text{coinc}(\tau)$
4. Perform the following operations until $i > N(\mathbf{v}^1)$ or $j > N(\mathbf{v}^2)$:
 - (a) If $v_i^1 > v_j^2 + n$ then increment j by 1
 - (b) If $v_i^1 < v_j^2 + n$ then increment i by 1
 - (c) If $v_i^1 = v_j^2 + n$ then increment $\text{coinc}(\tau)$, i and j by 1

Note: For simplification the above assumes that there is never more than one tag within a given Δt_{bin} . If this is not the case then an extra step is added to the algorithm where, when $v_i^1 = v_j^2 + n$, one finds $s_i = \sum_k \delta(v_i^1 - v_k^1)$ and $s_j = \sum_k \delta(v_j^2 - v_k^2)$. The relative quantities are then incremented: $\text{coinc}(\tau)$ by $s_i \times s_j$, i by s_i and j by s_j

5. Repeat steps 2-4 for each τ for which $\text{coinc}(\tau)$ is to be evaluated

This method is essentially the same as the “double histogram” method, and is one of the those currently in implementation for calculating of coincidences in the lab. It makes use of the fact that the calculated histograms are likely to be sparse. Rather than performing the dot product between the two histogram vectors, instead one looks for the common elements between the two, scaled, time tag vectors, making use of the fact that the vectors are chronologically ordered.

The limiting factor in computation here is generally going to be the intersection-like operation which scales additively in the length of the two vectors. This results

in an overall complexity which scales as $O((N(\{\tau_i\}))^{n-1} \sum_i^n N(\mathbf{u}^i))$ for n-fold coincidences. Given the additive, rather than multiplicative scaling, of the complexity with the number of tags, this method should perform well compared to the “per-photon histogram” method. Additionally, under the assumption the photon flux is low, this method should also experience a speedup relative to the “double histogram” method since $N(\mathbf{h}^i) \gg N(\mathbf{u}^i)$.

As with the “double histogram” method this procedure suffers from being less exact than the “per-photon histogram” approach due to the rebinning of both sets of tags during step 1 of the algorithm. However, as was noted above, a suitable choice Δt_{bin} can mitigate this issue.

5.6.2.4 Reduced Per-Photon Histogramming

Similar in nature to the “per-photon histogram” method, here we make use of an additional assumption that the set of delay times, $\{\tau_i\}$, are contiguous. As with the “Reduced double histogram” method we also make use of the chronological nature of the vectors by performing ordered searches. However, this time a search is performed for each element in \mathbf{u}^1 to find the range of elements in \mathbf{u}^2 that fall within the range of τ for which we are measuring coincidences. Algorithmically:

1. For the desired Δt_{bin} calculate a new set of vectors $\mathbf{v}^1 = \mathbf{u}^1 \frac{\Delta t_{tagger}}{\Delta t_{bin}}$ and $\mathbf{v}^2 = \mathbf{u}^2 \frac{\Delta t_{tagger}}{\Delta t_{bin}}$, where Δt_{tagger} is the width of the least significant bit of the time tagger. Here, the elements of \mathbf{v}^i should be integers, so the scaling should be performed then rounded

2. Create a pair of pointers, $i = 0$ and $j = 0$, and a vector to hold the running number of coincidences, **coinc**
3. Calculate $n_{min} = \min(\{\tau_i\})/\Delta t_{bin}$ and $n_{max} = \max(\{\tau_i\})/\Delta t_{bin}$
4. Increment i until $v_1^1 \leq v_i^2 + n_{min}$
5. Increment j until $v_1^1 \leq v_j^2 + n_{max}$
6. For each k in the range $[i, j - 1]$ increment $\text{coinc}(v_1^1 - v_k^2 + n_{min})$ by 1
7. Repeat steps 4-6 for each successive element in \mathbf{v}^1

This method was born out of discussion with collaborators at NIST and is one of the two presently being utilized in the lab. It makes use of the fact that we usually want to calculate coincidences for a contiguous set of delay times when dealing with photonic autocorrelation functions. In this situation, we can perform two ordered searches for each element in the first vector, to find the range of elements in the second vector which have elements that will produce relevant coincidences. Furthermore, we can make use of the chronological nature of the vectors to provide the starting point for the successive ordered searches.

The computational complexity mainly arises here in both the pick, steps 4-5, and the place operations, step 6. The pick steps are essentially just a pair of intersection-like operations between the two vectors, and have a complexity that scales as such. The place operation scales linearly in the number of elements of \mathbf{v}^1 but has a slightly more complicated relationship to that of \mathbf{v}^2 . In general, $\tau_{max} - \tau_{min}$ is likely to be small relative to the experimental duration, T_{exp} .

Therefore, for a given element of \mathbf{v}^1 we expect, for i, j found in steps 4-5, that $j - 1 - i \ll N(\mathbf{v}^2)$. Hence, we expect the n-fold coincidence calculation is going to have computational complexities which scale something like $O\left(\left(\sum_i^n N(\mathbf{u}^i)\right)\right)$ and $O\left(\left(\frac{\tau_{max}-\tau_{min}}{T_{exp}}\right)^{n-1} N(\mathbf{u}^1) \prod_{i=2}^n N(\mathbf{v}^i)\right)$ for the pick and place parts of the algorithm respectively.

As with the other methods where the re-binning is performed on both sets of tags in step 1 of the algorithm, this method is slightly less exact. However, we have already discussed how such an issue can be mitigated.

5.6.3 Computational implementation

From the previous section it can be seen that both the “per-photon histogram” and “double histogram” methods have poor scaling properties, with both scaling multiplicatively in the size of vectors which are not necessarily small. Initial investigations with both of these methods demonstrated them to be prohibitively slow, potentially taking order(s) of magnitude longer to process the data than to collect it. In this section we shall, therefore, concentrate on the computational implementation of the two “reduced” methods. For both I shall describe and benchmark some optimizations made in the computational implementation of their algorithms.

It should be noted that comparisons of different implementations can be a little tricky, given that the different methods scale differently as a function of things like the size of $\{\tau_i\}$, or the various vectors. However, to give some sense of how the different methods, and computational implementations compare, I shall present

Component	Name
CPU	AMD RYZEN Threadripper 1950x
GPU	2x Gigabyte GeForce GTX 1080 Windforce OC
Motherboard	Gigabyte X399 AORUS Gaming 7
RAM	3x G.Skills Ripjaws V Series 8GB
HDD	WD Blue 500GB SSD

Table 5.1: Specifications for computer used for correlations code benchmarking.

benchmarks on a set of 201 data files, whose form is described in section 4.6 of the previous chapter. The set of files is representative of typical data obtained in our lab, with each file consisting of a set of tags taken for light sent through a generalized Hanbury Brown and Twiss configuration, with the outputs monitored by three SPADs. For each file, data was accumulated over a period of 100 ms, with the experimental cycle time (including MOT load, optical molasses etc.) being ≈ 500 ms. The data was taken where the light incident on the SPADs had statistics $g^{(2)}(\tau) \approx 1$ for all τ , with there being $\approx 6 \times 10^4$ recorded tags per detector. All benchmarks have been performed using a bin size, $\Delta t_{bin} \approx 1$ ns, on a computer with specifications as described in table 5.1. The code for the various computational implementations, as well as that for benchmarking, and the benchmark results is all contained in a Github repository³ For all the implementations the code consists of a “C/C++” backend, which does the heavy lifting of the computations, which is called as a DLL/shared library by a Python frontend, which exposes the backend in a more user-friendly way.

³https://github.com/acraddoc91/RbRy_Correlations

5.6.3.1 Reduced Double Histogram

As-written, the algorithm in section 5.6.2.3 is sequential. However, there is no reason one may not simultaneously calculate the number of coincidences separately for each τ_i and/or file. This, therefore, falls into the category of “embarrassingly parallel” problems [93]. As such, we would expect to see near a one-to-one speedup in computation as a function of the number of threads used in a multi-threaded implementation. This behavior is readily in figure 5.4, which shows the time required for calculating coincidences for 2001 τ_i ’s as a function of the number of threads used. Here, OpenMP parallel for loops [94] have been used to allow for parallel calculation of the coincidences for different τ_i ’s, for a given number of processing threads. For low thread number we observe a one-to-one speedup per thread. Past ≈ 8 threads the speedup per thread decreases slightly, and at > 32 threads no additional speedup is seen when increasing the thread count. The former of these effects is likely due to the overhead associated with creating threads, along with that to read the files from disk and put them in a format suitable for processing. The latter is related to the CPU itself, which has 16 cores, each capable of running two simultaneous threads.

Although we can achieve an $\approx 20x$ speedup just by fully utilizing the threads on the CPU, we have the issue that the compute time scales linearly with the size of the set $\{\tau_i\}$. Extrapolating from the data in figure 5.4, we would expect something like calculating the three-fold coincidences, with a similar range of τ_i ’s, to take on the order of tens of hours for the benchmark data, which is considerably longer than the time taken to collect the data itself, ≈ 100 s.




Figure 5.4: Benchmarks for the compute time as a function of threads for a multi-threaded implementation of the reduced double histogram coincidence calculation method. Speedup calculated relative to single threaded performance.

Fortunately, the computational algorithm lends itself well to general purpose GPU computing. This is potentially helpful given the propensity for GPUs to have significantly more cores than a CPUs, albeit with those cores being slightly limited compared to their CPU counterparts. The CUDA API [95] was used to implement the “reduced double histogram: method on the nVidia GPUs stated in table 5.1. With GPU computation, optimization often requires specific tweaks that maximize the utilization of the underlying hardware. I do not believe discussion of those to be helpful here, given how hardware specific they are. However, several more general purpose optimizations were found for the CUDA based implementation which warrant a brief discussion.

One of these optimizations was “compute stream concurrency”. The idea here is that within CUDA, commands are usually queued and executed by the GPU from a FIFO “stream”. The default behavior is that each GPU has a single stream. However, occasionally this resulted in sub-optimal performance if a command in

the stream did not fully utilize the resources of the GPU. I found that creating multiple streams (16 at time of writing) was able to eliminate this issue with minimal downside.

Further gains were found by utilizing “pinned memory transfer”. On modern OSs memory allocated to programs is usually “paged”, a kind of hybrid virtual memory that uses both the main memory, usually the RAM, and a second, usually slower, form of storage e.g. a hard disk. This causes issues under CUDA as data transfer to the GPU from paged memory requires a blocking call on the GPU, which temporarily suspends computation. With “pinned” memory, which has a well defined location in the RAM, data transfer is able to occur asynchronously, and may be performed in one stream concurrently with computation on a different stream. The use of pinned memory for host-GPU data transfer, combined with the use of multiple streams, was found to give a performance boost.

A final general tweak that was found to significantly boost performance was that of “mini-batching”. Normally on the GPU, data, the photon time tags in this case, resides in the GPU RAM. In the algorithm, each loop requires the comparison of two (or more for higher order correlations) time tags. On the GPU, this means there are many threads all trying to simultaneously read from the GPU RAM, with the memory access being sparse, resulting in the memory becoming a bottleneck in computation. To get around this we make use of the memory hierarchy of the GPU, where, for each thread, we load a small buffer of time tags (four for each larger time tag vector at time of writing) into “shared” memory. The access times for the shared memory is significantly faster than the RAM due to its on-chip location,




Figure 5.5: Benchmarks for the CPU and GPU based implementations of the reduced double histogram method. Speedup here is calculated relative to single threaded CPU performance.

although it is smaller in size and more limited in terms of its accessibility. Step four of the algorithm in section 5.6.2.3 is performed on the time tags within these smaller buffers, with the buffers being refreshed when they are exhausted. It might not seem obvious that this approach would result in a performance gain as the same number of tags must be read from the GPU RAM. However, this technique also makes use of what is called “memory coalescing”, which means that the reading of sequential values in the RAM can occur with similar speeds to single values, i.e. reading the four tags to the shared memory buffer is likely to be nearly as fast as only reading one of those time tags.

Figure 5.5 shows the result of benchmarks performed using the GPU computational approach, as well as that with the multithreaded CPU implementation, for a range of sizes of $\{\tau_i\}$. In all cases, for a small size of $\{\tau_i\}$ the computation time approaches some constant related to the overhead in reading and preparing files from disk, as well as initializing threads/the GPUs. While for a large $\{\tau_i\}$

the compute time scales linearly, as expected. For typical sizes of $\{\tau_i\}$ in our lab, $\approx 10^3 - 10^4$ for calculating $g^{(2)}$, the GPU implementation is at least an order of magnitude faster than for the multithreaded CPU approach. We see that the speedup compared to the multi-threaded CPU approach saturates at around $30\times$ for a single GPU, which should carry through when computing higher order correlations where the multi-dimensionality of the problem means a large number of calculations need to be performed. We see that the speedup for a pair of GPUs saturates at around twice that of a single GPU, as expected given the massively parallel nature of the underlying problem. Therefore, it is anticipated that further speedups could be gained with the addition of more GPU processing power.

5.6.3.2 Reduced Per-Photon Histogram

As with the “reduced double histogram” method, the algorithm as-written in section 5.6.2.4 is sequential. However, unlike that method there is not as much inherently in the algorithm that can be obviously be performed in parallel. Early on in the code writing process I had expected there to be potential gains to be made from splitting up the vector \mathbf{v}^1 and performing steps 2 through 7 in parallel on these smaller vectors⁴. However, this turned out to not be overly fruitful. Therefore, for this method the only real computational optimization I have performed is parallel processing of data files which, as before, has been implemented using OpenMP parallel for loops [94]. Figure 5.6 shows the benchmarks for calculating coincidences

⁴to get the initial pointer values a binary search algorithm was also implemented which prevented issues associated with step 4 and 5 becoming slow for the chronologically later smaller vector segments.




Figure 5.6: Benchmarks for the compute time as a function of threads for a multi-threaded implementation of the reduced per-photon histogram coincidence calculation method. Speedup calculated relative to single threaded performance.

for 2001 τ_i 's as a function of the number of threads used. We can see that the improvement, when increasing the number of threads, deviates very rapidly from the one-to-one behavior we have previously seen. My speculation is that this is related to the fact that the HDF5 libraries⁵ used did not allow parallel file reading, which would make reading the files from disk a potential bottleneck, at least in this regime.

A major issue that arises in software development, and in our lab, is the “it works on my machine” problem, where software fails to operate correctly on a machine other than the one used for development/testing. This can occur for any number of reasons, differences in available libraries, operating systems, CPU architectures etc. Docker [96] provides a remedy to this issue by allowing developers to write “images”, which can be run by users in a virtualized “container”. Therefore, a Dockerized version of the software was created, which has proven to be less

⁵Note that implementation of this method was performed in Linux, whereas the “reduced double histogram” was written in Windows. I believe there may be a difference in the read time of the HDF5 files in the two cases.




Figure 5.7: Benchmarks for the Docker and native based implementations of the reduced per-photon histogram method.

problematic in terms of portability and debugging. One might expect to see some performance hit due to virtualization. However, as can be seen in figure 5.7, the Dockerized and natively compiled version of the software perform similarly in the regimes benchmarked.

5.6.3.3 Closing Remarks

Given the differing scaling of the computational complexity for the two methods, we do not necessarily expect there to be a one-size-fits-all fastest solution. In the benchmarks shown in figure 5.8 we see that this is in fact the case⁶, with the “reduced double histogram: method being faster when $\{\tau_i\}$ is small, and the “re-

⁶Although it is possible some of this behavior may be related to the conjectured differing HDF5 file read speeds in the two cases




Figure 5.8: Comparison of the benchmarks for the optimized implementations of the two methods.

duced per-photon histogram” method quicker when it is large. In practice, because of its superior compatibility and lesser hardware requirements, we have tended to favor using the Dockerized “reduced per-photon histogram” software. However, both software approaches remain in active development.

We have largely been concerned with reducing the processing time required to calculate coincidences for some benchmark data. However, it is perhaps important to bear in mind what sent me down this path in the first place. That is that the more basic methods of calculation took far longer to perform than taking the data itself. This is problematic in our experiment as we typically need long data runs, multiple hours, to build up good statistics. Long processing relative to accumulation times therefore mean a long wait for results, in addition to an inability to know at some intermediate stage whether the data being taken is worthwhile. In an ideal

situation we want to perform data processing as fast, if not faster, than the rate at which it is accumulated. When we consider that the benchmark data took ≈ 100 s to obtain then, from figure 5.8, we see that the typical $\{\tau_i\}$ size we use of $\approx 10^3 - 10^4$ falls within the regime where real-time processing is possible. We have found this to be the case in normal lab use, with there rarely being issues with computation speed when calculating either⁷ $g^{(2)}$ or $g^{(3)}$.

⁷I can not comment on higher order correlation functions as we have never needed to calculate them.

Chapter 6: On-demand, Pure and Indistinguishable Single Photon Source Based on a Rydberg Ensemble

Sources that produce on-demand single photons with a high rate, efficiency, purity, and indistinguishability have a wide range of applications in the fields of linear optical quantum computing [97], boson sampling [98], quantum networking [14] and quantum metrology [99]. The forefront of this field is dominated by sources based on solid state platforms such as quantum dots [100–105], spontaneous-parametric downconversion [106, 107], and multiplexed sources [108, 109]. However, none of these sources fulfill all the desired criteria. Additionally, these sources in general, produce photons which are spectrally incompatible with atomic systems. While this is by no means essential, there are applications where such compatibility is likely to be desirable, e.g. hybrid quantum networking.

In addition to solid state sources, there also exists a body of work on developing sources using atomic systems, for example single atoms [110] or ions [111] in free space, and cavity coupled single atoms [112–114] and ensembles [115]. While these have the benefit that they produce photons that are inherently atomically compatible, they tend to lag behind the solid state sources in terms of their other properties.




Figure 6.1: Sequence for on-demand single photon generation.

In this chapter, I report on an efficient single photon source based on a Rydberg ensemble. This chapter is based on work [39] that has recently been accepted for publication. Dalia Ornelas-Huerta, the first author of the paper, and I are simultaneously writing our theses, which both contain chapters based on this publication. In an effort to minimize overlap, I shall be reporting the main results from the paper, while adding a description of additional work and analysis performed that did not make it into the paper due to length constraints.

6.1 General Method

With the atoms loaded into the dipole trap, the process of which has been described in detail in chapter 4, we use a write and retrieve protocol to produce photons from the ensemble, similar to that described in reference [64]. The pulse




Figure 6.2: Level scheme for production of on-demand photons.

sequence for this, shown in figure 6.1, is repeated many times per experimental cycle.

For the write stage, the probe and control fields, shown in figure 6.2, are detuned such that Δ is much greater than Γ , the $5P_{3/2}$ linewidth, but with $\delta \approx 0$. By choosing a Rydberg state with a sufficiently high principal quantum number, and by tuning the shape of the dipole trap, we can work in a regime where the blockade radius becomes larger than the volume of the system, suppressing multiple excitations. In this way the write process results in only a single collective excitation stored in the ensemble.

After the spin wave is written, the control field is brought close to resonance and turned back on mapping the Rydberg spin wave excitation to a Rydberg polariton, which exits the cloud as photon [73]. As discussed in chapter 3, the spin wave nature of the excitation means that the photon is preferentially retrieved into a single mode. Here, the same beam configuration for the control field is used during both writing and retrieval, so the photon is preferentially retrieved in the forward




Figure 6.3: Experimental setup for production and characterization of single photons.

direction. Given that the transverse mode of the written spin wave should be set by the probe, the produced photon should exit the cloud in a similar mode to that of the probe itself.

After the chamber, the photon is passed through a set of interference filters (Alluxa 780-1 OD6 and Semrock Brightline 780/12) and a PBS, shown in figure 6.3, to perform spectral and polarization filtering respectively. An AOM after the filtering stage allows us to gate the light to prevent saturation of the SPADs during the write stage. We couple the first-order diffracted light from the AOM into a short, few meter long, polarization-maintaining fiber (PMF) (Thorlabs PM780-HP) and a long, ≈ 1 km, single mode fiber (SMF). The power between the two fibers can be adjusted using the prior HWP and PBS shown in figure 6.3. In order to ensure the transverse spatial mode exiting the short and long arms is the same the km SMF is butt-coupled to a near-identical fiber to the short PMF. The outputs of both fibers

are passed through a further polarization filter, not pictured, and directed towards a 50:50 beamsplitter, whose outputs are coupled into a pair of SMF each connected a SPAD. We use a HWP to adjust the relative polarization of the two input arms of the 50:50 beamsplitter. By adjusting the power in the long SMF and short PMF we are able to switch between performing a Hanbury Brown and Twiss [88] and self-Hong-Ou-Mandel [89] measurement of the light produced by the system.

6.2 Source Optimizations

The general method, outlined above, has several processes with various parameters which can be adjusted, changing the properties of the photon source. This can be a double edged sword, on the one hand it grants us the freedom to alter the source characteristics, but on the other it makes finding optimal values of the parameters difficult, given the multi-dimensionality of the optimization problem. As such, an exhaustive search of the parameter space is near impossible. Nevertheless, I discuss here some of the reasoning behind the value of the parameters used in the paper upon which this chapter is based [39].

6.2.1 Single Photon Purity

Tautologically, a single photon source should be a source of single photons. That is, it should have an output with a $g^{(2)}(0) \approx 0$. For our source the suppression of multi-photon states in the output relies on only a single collective excitation being written to the ensemble. In experiment, if the Rydberg state principal quantum




Figure 6.4: Coincidences, as a function of delay time, at the outputs of a Hanbury Brown and Twiss setup for light produced using $n = 120$ with two different axial cloud radii (radius along the probe propagation direction). Coincidences have been normalized by dividing by the maximum coincidence value. Blockade radius calculated to be $r_B = (C_6/\Omega_{2\text{-ph}})^{1/6} \approx 45 \mu\text{m}$. Both figures show the same data, with the right-hand figure zoomed in around $\tau = 0$ to better show the peak there.

number being addressed is too low, then the blockade radius can be small enough, relative to the cloud length, that multiple excitations are simultaneously stored in the cloud during the write stage. This was seen in preliminary work performed at $n = 120$, where $g^{(2)}(0)$ was observed to have a clear dependence on the axial radius¹ of the cloud, as seen in figure 6.4. Therefore, in order to suppress multiple excitations, we use a principal quantum number $n = 139$, for which the blockade radius is typically² on the order of $r_B = (C_6/\Omega_{2\text{-ph}})^{1/6} \approx 60 \mu\text{m}$, where $\Omega_{2\text{-ph}}$ is the effective two photon Rabi frequency during the write stage. Additionally we use a shortened cloud, for which the axial radius, $\sigma_a \approx 27 \mu\text{m}$, is obtained by reducing the power in the dipole trap arms, with the power in the first arm $\approx 150 \text{ mW}$. Although, as will be seen later, this did not result in a perfectly pure source, further reduction of the axial radius of the cloud and/or working with a higher Rydberg state would

¹This is the radius along the probe propagation direction

²Note there is some variation due to the dependence of the two-photon write Rabi frequency

likely have had negative impacts on the other source properties and not been worth the compromise.

In a previous work [64] on a similar photon production process there was a suggestion that a dephasing process, during the time the excitations are stored, could contribute to the suppression of multi-photon output events. We would therefore expect $g^{(2)}(0)$ to decrease as the storage time of the excitations are increased. However, this was, within experimental uncertainties, not observed to be the case in our system. As a result, we believe the purity of the source to be solely a function of blockade in the write process, with $g^{(2)}(0)$ being essentially “locked-in” once the write has been performed.

6.2.2 Source Efficiency

In addition to having a high purity output, it is desirable that a single-photon source produce photons with high efficiency, i.e. upon request the source produces a photon, in the desired mode, with some high probability. For our source the efficiency of photon production is a combination of the efficiency of writing the initial spin wave, η_w , the efficiency of the storage of that spin wave, η_s , and the retrieval efficiency of the spin wave into a single photonic mode, η_r . In practice, we were only able to measure the total photon production efficiency, and not each of the individual efficiencies. However, the processes are sufficiently distinct from each other, both in practice and in theory, that it makes sense to talk about them separately.

6.2.2.1 Writing

Given our choice of Rydberg state and cloud dimensions, practically the entire excitation volume is contained within a single blockade radius. Therefore, we can effectively think of the ensemble as a so-called “superatom”, with a set of Dicke-like energy levels [71]. For the most part, we can ignore the intermediate $|5P_{3/2}\rangle$ state, which can be adiabatically eliminated as the write procedure is performed far from intermediate resonance, $\Delta \gg \Gamma$. As such, we can consider two Dicke states, $|G\rangle$ a state where all the atoms are in the ground state, and $|S\rangle$ a symmetric superposition state with exactly one Rydberg excitation, coupled with a two-photon Rabi frequency, $\Omega_{2\text{-ph}} = \frac{\Omega_p \Omega_c}{2|\Delta|}$. Due to the collective nature of the excitation, the probe Rabi frequency is enhanced by a factor³ [58] $\sqrt{N} \approx 20$, related to the atom number, $\Omega'_p = \sqrt{N}\Omega_p$.

In the absence of any dissipation, the spin wave write procedure could be performed with perfect efficiency. However, the finite spin wave dephasing rate, $\gamma \approx 2\pi \times 90$ kHz, and, to a lesser extent, the linewidth of the intermediate state⁴, $\Gamma \approx 2\pi \times 6.8$ MHz, set an upper bound on the writing efficiency which can be obtained. To maximize the fidelity, the write procedure should be performed as quickly as possible, for which we want to use as large a two-photon Rabi frequency as is obtainable. However, the control Rabi frequency is limited to $\Omega_c \approx 2\pi \times 7$ MHz by the control power available, ≈ 100 mW due to the necessity to use a double pass

³Note that this factor is omitted when calculating the blockade radius.

⁴This is the measured value, and is broadened relative to the atomic value, primarily due to the dipole trap.

AOM to allow for the dynamic adjustment of the control frequency. For the probe we do not have such a constraint, but off-resonant scatter from the intermediate state starts to become an issue at higher probe powers causing the cloud to be rapidly heated. While this can be mitigated by further detuning from the intermediate state, the bandwidth of the control double-pass AOM set a limit on how far this could be pushed.

Optimization of the write process was performed by measuring, as a function of the write parameters, the overall photon production efficiency averaged over many, of order several tens of thousand, photon production cycles. The maximum control Rabi frequency of $\Omega_c \approx 2\pi \times 7$ MHz, along with a probe Rabi frequency of $\Omega_p \approx 2\pi \times 1$ MHz, with an intermediate state detuning of $\Delta \approx 2\pi \times 50$ MHz, pulsed for ≈ 370 ns were empirically found to be optimal. The ideal two-photon detuning was found to be $\delta \approx -2\pi \times 2$ MHz, consistent with its theoretically expected value due to the imbalance of the two excitation Rabi frequencies, Ω'_p and Ω_c .

6.2.2.2 Storage

Given the experimental realities of our finite spin wave dephasing, we want to keep the storage time as short as possible, with the ideal situation being zero storage time. At present, a technical lower bound on the storage time is set by the switching time of the double-pass AOM used for the control light. During writing the control frequency should be far detuned from resonance, but for retrieval it needs to be close to resonance, a frequency difference of $\approx 2\pi \times 50$ MHz. This shift in frequency is

achieved by using an RF switch (Mini-Circuits ZASWA-2-50DR) to toggle between two different RF tones being sent to the AOM. Limitations of the AOM system mean that the minimum time required to make this switch is on the order of ≈ 300 ns. A storage time of 350 ns is used, with a slight overhead, compared to the switching time, added to ensure there are not issues resulting from experimental jitter.

6.2.2.3 Retrieval

For an infinitely long-lived spin wave the retrieval efficiency should only depend on the optical depth of cloud [73]. However, as with the other two stages of the photon production process, the retrieval efficiency also suffers as a result of spin wave decay, causing there to be an additional dependence of the efficiency on the control Rabi frequency and detuning during retrieval.

Given the finite dephasing, it is always preferable to minimize the amount of time spent in the Rydberg state. Therefore, we perform retrieval using the maximum available control Rabi frequency. We observe a curious asymmetry in the retrieval efficiency as a function of the control field detuning during retrieval, as seen in figure 6.5. This behavior is not expected theoretically [73], and its origin remains a mystery. To determine the control field detuning for photon production, a similar curve to that in figure 6.5 was taken, but not normalizing the control field power to account for the control AOM diffraction efficiency, which peaks at a detuning of $\approx 2\pi \times 20$ MHz. From that, we find a control field detuning of $\approx 2\pi \times 7$ MHz during retrieval to be optimal for photon production.




Figure 6.5: Normalized photon production efficiency as a function of the control field detuning during retrieval. All other experimental parameters kept constant. Control power here is 30% of its max value for all detunings.

For the dipole trap configured to suppress multiple excitations, and using the standard MOT loading time of 250 ms we obtain an optical depth $OD \approx 13$. Given the control Rabi frequency available, and the dephasing rate, this OD should theoretically be close to optimal. Experimentally, only minor gains in the production efficiency, $\approx 10\%$, were observed when increasing the MOT loading time by an order of magnitude to increase the optical depth, as seen in figure 6.6. For the rest of the work shown here we have used a 250 ms MOT loading time, and therefore $OD \approx 13$, which served as a good compromise between retrieval efficiency and experimental duty cycle.




Figure 6.6: Normalized photon production efficiency as a function of MOT load time/optical depth. Note the optical depth is *not* linear in the MOT load time.

6.3 Source Characterization

A single photon source is fully characterized by relatively few properties, its efficiency and production rate, and the purity, indistinguishability, single modedness, wavelength and bandwidth of the produced photons. Here, I report on the measurement of these properties for our source.

6.3.1 Efficiency and Rate

After performing the optimizations described in the previous section we observed peak in-PMF efficiencies up to 0.18(2), after accounting for the SPAD quantum efficiency. Accounting for the losses, listed in table 6.1, along the optical path

Element	Efficiency
Optics transmission	0.75(2)
AOM diffraction	0.79(2)
PMF coupling	0.75(2)
HOM interferometer	0.38(1)
SPAD quantum efficiency	0.67(1)

Table 6.1: List of the efficiencies along the post-cloud probe path.




Figure 6.7: Left: Average in-fiber efficiency as function of the duty cycle. Right: Effective fibered rate as a function of the duty cycle. Note, these were taken prior to full optimization of the photon generation process, hence the lower than reported fibered efficiency. Production period here is $5 \mu\text{s}$.

and in coupling into the PMF, where we have assumed the mode of the single photons to be identical to that of the probe, we determine the photon generation probability from the cloud to be 0.40(4). Theoretically modeling the photon production process using a Lindblad master equation approach to model the write and storage processes and the Maxwell-Bloch formalism of reference [73] for the retrieval process, we expect the source efficiency to be 0.42(3), which is consistent with our measurement. It should be noted that these experimental values are obtained for a relatively slow production cycle time, $200 \mu\text{s}$, and a small photon producing duty




Figure 6.8: Left: In-fiber efficiency as function of the production period. Right: Effective fibered rate as a function of the effective production attempt rate, Duty Cycle/ t_{period} . Duty cycle here is 0.2.

cycle⁵, 20%. Therefore, the effective rate, that is the rate accounting for the duty cycle of the experiment, of photons into the fiber is low, $\approx 200 \text{ s}^{-1}$.

An increase in the effective photon flux may be obtained by increasing the production duty cycle, by increasing the amount of time spent producing photons per dipole trap load. This comes with a reduction in the average per-attempt efficiency, as seen in figure 6.7. The trade-off exists as a result of the finite lifetime of the dipole trap, with the efficiency of both the write and retrieve stages being dependent on the atom number in the cloud.

In a similar vein, the production period may be reduced, to a minimum of a few microseconds, which can also increase the fibered photon rate. This again has the downside that it reduces the per-attempt efficiency, as seen in figure 6.8. The reason for this behavior is two-fold. Firstly, an increase in the repetition rate increases the scattering rate of the write photons from the cloud, which causes a

⁵here I am using duty cycle to mean the time spent in the photon producing part of the experimental cycle, where the other portion of the experimental cycle includes the MOT load, molasses, optical pumping etc.




Figure 6.9: On-demand single photon source purity measurement scheme.

decrease in the lifetime of the atomic cloud. The other reason is related to Rydberg contaminant states, which are themselves an interesting phenomenon that shall be covered in more detail in a later section.

For much of the remainder of the work, related to characterizing the source, we used a cycle time $\approx 5 \mu\text{s}$ and a duty cycle of 60%. This struck a good compromise between per-photon fibered efficiency, ≈ 0.1 , and effective in-fiber photon flux, $\approx 1.18 \times 10^4 \text{ s}^{-1}$.

6.3.2 Purity

The purity quantifies how likely it is that the output of the source is a single photon. We define the purity, in terms of the measurable quantity $g^{(2)}(0)$, as $1 - g^{(2)}(0)$. For a perfect single photon source $g^{(2)}(0)$ is zero and thus the purity is one. To measure $g^{(2)}$ for our source, we send all the light output from the chamber




Figure 6.10: Gating window, shown by black dotted lines, for on-demand single photon measurements.

to a single fiber, in this case the short PMF, with the other blocked, as shown in figure 6.9.

Although the post-chamber gating AOM prevents a large amount of the write pulse from reaching the SPADs, there is a small amount of leakage, that can be seen in figure 6.10. Therefore, for the calculation of coincidences, we apply a $1.4 \mu\text{s}$ gating window, also shown in the figure, to the time tags from both SPADs.

Figure 6.11 shows the coincidences, as a function of delay time, between the two SPADs for the source. As can be seen, the coincidences at zero time delay are greatly suppressed. Integrating the area for the raw coincidences around $\tau = 0$ and comparing it to the areas at a delay time $\tau = n \times t_{\text{period}}$, where here $t_{\text{period}} = 5 \mu\text{s}$, and n is sufficiently large enough for the coincidence events to have no undesired




Figure 6.11: $g^{(2)}$ measurement for single photon source. Both figures show the same data, with the right hand being zoomed in around $\tau \approx 0$. Orange curve shows the calculated coincidences due to background events.

correlations, we find $g_{raw}^{(2)}(0) = 0.0145(2)$. The major reason that $g_{raw}^{(2)}(0)$ is non-zero comes from coincidences due to background events. The dominant source of these accidental coincidence events in our system are due to coincident detection events of a produced photon with one unrelated to the photon production process, either a SPAD dark count, or from room light leakage into the SPAD. From the independently measured photon profile, background rate and photon rate, we determine the magnitude and profile of the coincidences due to background events, shown by the orange curve in figure 6.11. Subtracting the calculated background and performing the same analysis of integrating areas in coincidence space, we determine the background-subtracted $g^{(2)}(0) = 5.0(1.6) \times 10^{-4}$, indicating that our source has a near-perfect purity.




Figure 6.12: On-demand single photon source indistinguishability measurement scheme.

6.3.3 Indistinguishability

We use the indistinguishability of a source to quantify how similar the photons it produces are. This is an important property for many applications where the interference of photons is required. Experimentally we measure the ability of the source to produce indistinguishable photons using a self-Hong Ou Mandel interferometer. This is implemented by adjusting the power of the post-chamber HWP to roughly balance the photon flux at the output of the PMF and km SMF. The km SMF fiber acts as a $4.92 \mu\text{s}$ delay line. By tuning the production repetition period to match this delay time, as shown in figure 6.12, we are able to interfere photons from adjacently produced pulses on the 50:50 beamsplitter. The HWP immediately prior to the 50:50 beamsplitter allows us to adjust the polarization of the photons arriving at the two input ports to be either parallel or perpendicular.




Figure 6.13: Self Hong-Ou-Mandel indistinguishability measurements for single photon source. Both figures show the same data with the right being zoomed in around $\tau \approx 0$. Orange curve shows the calculated coincidences due to background events.

As with the purity measurements, a software gating window is applied to the time tags when calculating coincidences. We observe a suppression in the number of coincidences when the photons arriving at the input ports are parallel as opposed to perpendicular, as seen in figure 6.13. Integrating the area around $\tau = 0$ in both cases we find the raw visibility to be $V_{raw} = 1 - \text{Coincidences}_{\parallel} / \text{Coincidences}_{\perp} = 0.892(6)$. Accounting again for accidental coincidences, in addition to slight deviations from perfection in the transmission and reflection coefficients of the beamsplitter, we determine the corrected interference visibility to be $V = 0.980(7)$, suggesting our source produces near-indistinguishable photons.

In contrast to figure 6.11, the peaks at $\tau \neq 0$ in figure 6.13 have differing heights, and display a slight temporal asymmetry. This is expected behavior. The differing heights arise from the fact that the beamsplitter configuration in figure 6.12 forms several Hanbury Brown and Twiss setups, with differing delay times. While the temporal asymmetry occurs due to imperfections in the 50:50 beamsplitter.

These effects are not particularly relevant to the discussion here, but are explored further in appendix B.

6.4 Contaminants

As briefly touched upon in section 6.2, an interesting issue for our source is that of Rydberg contaminants, namely the creation of Rydberg states other than the one targeted by the excitation light. It should be noted that this is not a new phenomenon, and there is an existing body of experimental [116–118] and theoretical [91, 119–122] literature on the subject. However, it poses a challenge for Rydberg-ensemble based single-photon sources which, outside of the paper [39] this chapter is based on, remains unexplored.

We first identified contaminant states as a potential issue when we observed the dependence of the average photon production efficiency on the production cycle period, which we saw in figure 6.7. Further evidence of a problem came from the long timescale behavior of $g^{(2)}$ as well as production efficiency measurements made in pulse trains, both shown in figure 6.14. In this figure the $g^{(2)}$ data shown is the same as shown in figure 6.11, albeit looking at much larger values of τ , and where we have integrated the coincidences around the $\tau = n \times t_{period}$ peaks. For the displayed pulse train measurements we have used a production period of $2.5 \mu s$ for 75 production cycles before allowing a “cool down” of $200 \mu s$ between trains, where no production occurs. For the pulse train measurements, we see a clear decay in the production efficiency as a function of the pulse number within the train, indicative




Figure 6.14: Left: Area integrated $g^{(2)}$ showing long-time behavior. Right: Normalized production efficiency for a train of 75 pulses with a production period of $2.5 \mu\text{s}$, where the trains are separated by a $200 \mu\text{s}$ “cool down” time. In both cases the orange line is a fit, of form discussed in the text, to the data.

of some sort of build up which inhibits the production of photons. The fact that this decay is also present in $g^{(2)}$ suggests that the successful detection of a photon projects the cloud into a similar “clean” state that is present after a “cool down” has occurred in the pulse train measurement. From the data we infer that these contaminants are produced as a by-product of the photon production process, and that the presence of a contaminant within the cloud strongly suppresses the further production of photons.

We can gain further insight into the nature of the contaminant states through a simple model. We assume that contaminants are created in a clean medium with a probability P_c and that they have a lifetime in the cloud of τ_c . In addition, we shall assume that only a single contaminant may be present in the medium at any given time, i.e. contaminants blockade the production of further contaminants, and that contaminants fully blockade the photon production process. Then the probability

of having a contaminant for the n -th pulse within a pulse train is given by

$$P_n = P_{n-1}e^{-t_{\text{period}}/\tau_c} + (1 - P_{n-1})P_c. \quad (6.1)$$

Solving the recursion relation and using the initial conditions $P_0 = 0$, as the cloud is clean, and $P_1 = P_c$ then we find

$$P_n = \frac{P_c}{1 - e^{-t_{\text{period}}/\tau_c} + P_c} \left[1 - (e^{-t_{\text{period}}/\tau_c} - P_c)^n \right]. \quad (6.2)$$

For a given attempt in the train, the probability of producing a photon is given by

$$P_g(n) = P_{\text{max}}(1 - P_n) = P_{\text{max}} \left(1 - \frac{P_c}{1 - e^{-t_{\text{period}}/\tau_c} + P_c} \left[1 - (e^{-t_{\text{period}}/\tau_c} - P_c)^n \right] \right), \quad (6.3)$$

where P_{max} here is the maximum photon production probability. Fitting this to the pulse train data in figure 6.14, see orange curve, we are able to extract $P_c = 1.9(3) \times 10^{-2}$ and $\tau_c = 65(8) \mu\text{s}$.

In a similar way we can model the behavior of $g^{(2)}$ by recalling our probabilistic definition from chapter 5

$$g^{(2)}(m) \approx \frac{\langle P_g(n+m|n) \rangle_t}{\langle P_g(n+m) \rangle_t}, \quad (6.4)$$

where the notation should be taken to mean $g^{(2)}$ after summing over coincidences, and the probabilities as the generation probabilities. The denominator is the time averaged generation probability, which is the same as the generation probability as

State	Probability
138 $P_{3/2}$	0.217
139 $P_{3/2}$	0.212
138 $P_{1/2}$	0.111
139 $P_{1/2}$	0.104
137 $P_{3/2}$	0.036
140 $P_{3/2}$	0.030

Table 6.2: Decay probability, calculated with reference [61], of the 139 $S_{1/2}$ state due to blackbody induced transitions at $T = 300$ K. Only states with a probability higher than 2% are shown.

$n \rightarrow \infty$

$$\langle P_g(n+m) \rangle_t = P_{max} \left(1 - \frac{P_c}{1 - e^{-t_{period}/\tau_c} + P_c} \right). \quad (6.5)$$

While the numerator can be evaluated by recognizing that a photon produced in pulse n informs us that the cloud was clean at that point, and therefore simply

$$\langle P_g(n+m | n) \rangle_t = P_g(|m|) \quad \forall m \neq 0, \quad (6.6)$$

where the negative time behavior comes from time-symmetry arguments. Therefore, we can write

$$g^{(2)}(m \neq 0) \approx 1 + P_c \frac{(e^{-t_{period}/\tau_c} - P_c)^{|m|}}{1 - e^{-t_{period}/\tau_c}}. \quad (6.7)$$

Fitting this to data, see the orange curve in figure 6.14, we find values, $P_c = 7.2(2) \times 10^{-3}$ and $\tau_c = 75(2) \mu s$. Although these differ slightly from those found from the pulse train fits, most notably in the contaminant creation probability, the two data sets were taken almost a month apart, with various optimizations being performed in the interim. As such, the discrepancy is not necessarily surprising.

The initial suspicion was that the contaminants are the product of blackbody-




Figure 6.15: Normalized retrieval efficiency as a function of microwave time. Microwaves resonant with $139S_{1/2}$ to $139P_{3/2}$ transition. Orange curve is a model assuming a $2\pi \times 14$ MHz Rabi frequency and no dissipation.

induced transitions from the target Rydberg state, discussed in chapter 2 when we talked about the blackbody shortened lifetime of Rydberg states. As can be seen in table 6.2, the dominant blackbody decay path of the $139S_{1/2}$ state is to the $(139\pm 1)P$ states. To investigate this idea of blackbody contaminants further, we attempted to produce our own artificial contaminant states. This was accomplished by performing the write stage of the photon production, then using a resonant microwave field to transfer the excitation to a nearby P state. Confirmation of the ability to transfer the excitation to the P state was obtained by performing the retrieval part of the photon production procedure and observing clear Rabi flopping, seen in figure 6.15, in the retrieval probability as a function of the microwave duration.




Figure 6.16: Normalized photon production efficiency of second cycle as a function of the microwave time in the first cycle. Microwaves resonant with $139S_{1/2}$ to $139P_{3/2}$ transition. Orange curve to guide the eye is of form $1 - a \sin^2(\Omega t_{\text{microwave}}/2)$, where $\Omega = 2\pi \times 14$ MHz and $a \approx 0.4$.

The first in the series of experiments performed with the microwaves was to look at the behavior of subsequent photon production in the presence of the artificial P states. For this, we performed pairs of photon production cycles with a large “cool down”, $100 \mu\text{s}$, between pairs. In the first production cycle of the pair, we applied resonant microwaves during the storage time for a time $t_{\text{microwave}}$, whereas the second cycle was performed as normal. As seen in figure 6.16, the second cycle experiences a suppression in the production efficiency in phase with the flopping of the first cycle to the $139P_{3/2}$ state.

Again using the pairs of production cycles, we then probed the lifetime of the artificial contaminants. Here, we use $t_{\text{microwave}} = 100$ ns, to convert any stored target




Figure 6.17: Lifetime measurement of artificially created contaminant. Pairs of production cycles are performed where we use microwaves, resonant with $139S_{1/2}$ to $139P_{3/2}$ transition, to transfer the first Rydberg excitation to an artificial contaminant state. The second cycle is performed as normal. Plotted is the normalized production efficiency, as a function of the delay time between the two cycles. Orange curve is fit of form $a \left(1 - be^{-t_{\text{delay}}/\tau_{139P_{3/2}}}\right)$, where $a = 0.93(4)$, $b = 0.43(8)$ and $\tau_{139P_{3/2}} = 100(40) \mu\text{s}$.

excitation in the first production cycle to a $139P_{3/2}$ state, then the second production cycle is performed as normal after some delay, t_{delay} , with a “cool down” period, $500 \mu\text{s}$, between cycle pairs. As expected, we see an increase in the production efficiency of the second cycle as the delay time between the cycles is increased, which can be seen in figure 6.17. From this we extract a lifetime for the created $139P_{3/2}$ in the cloud of $\tau_{139P_{3/2}} = 100(40) \mu\text{s}$. This lifetime is consistent with that seen for the naturally occurring contaminants, however, it is considerably less than the calculated blackbody lifetime of the state, $\approx 850 \mu\text{s}$. Similar experiments performed




Figure 6.18: Contaminant production probability as a function of the peak atomic density of the cloud. $t_{storage}$ fixed at 350 ns.

with $138P_{3/2}$ corroborate the findings presented here.

Although there was a similarity in behavior of the naturally and artificially created contaminants states, a number of observations indicate that the natural contaminants do not arise from blackbody decay. Firstly, attempts were made to “clean out” the cloud after the photon production cycle by applying microwaves resonant with the prominent blackbody decay channels, along with the resonant control light. These were found to be unsuccessful in reducing the presence of contaminants within the cloud. Additionally, P_c , measured with the pulse train method, was seen to have a dependence on the cloud density, shown in figure 6.18. Finally, the dependence of P_c on the spin wave storage time, shown in figure 6.19, indicates a decay significantly faster, a few tens of microseconds, than would be expected for blackbody




Figure 6.19: Contaminant probability as a function of the spin wave storage time. Atomic density fixed at $\approx 4 \times 10^{11} \text{ cm}^{-3}$.

decay. Given this behavior we believe the natural contaminants to be produced by collisional processes between Rydberg and ground state atoms, similar to those described in reference [120]. In that work they discuss multiple collisional processes, some of which result in states that would not be cleaned in the manner attempted with microwaves. Due to the aforementioned issues with our state sensitive field ionization setup we are both unable to fully confirm this hypothesis, nor are we able to easily alleviate the issue.

6.5 Issues, Improvements and Outlook

We have already discussed the contaminant states at length, which cause there to be a trade-off between the production efficiency and rate. Eliminating this prob-

lem is therefore highly desirable. A field ionizing pulse after each production cycle is likely to be the most optimal solution, as it ensures a clean cloud for each production cycle. This is not feasible in our current setup, due to the limitations of our high voltage in-vacuum electronics, discussed in chapter 4. Failing this, the effect of the contaminants can be mitigated by reducing their production probability. This could be achieved either by reducing the time spent in the Rydberg state, or, if the contaminants are caused by collisional processes, by reducing the mean atom number encompassed by the Rydberg electron. For both of these, a reduction of the Rydberg principal quantum number would be helpful, in terms of increasing the control field matrix element and decreasing the Rydberg orbital radius. However, the trap geometry would need to be adjusted to ensure suppression of multiple excitations with the shrinking blockade radius.

Another issue relates to the electric field sensitivity of the Rydberg state. At $n = 139$ the polarizability is $\alpha \approx 62 \text{ GHz}/(\text{V}/\text{cm})^2$ [61]. This means a field of only a few mV/cm is necessary to shift the Rydberg state such that the probe and control are no longer at the optimal two-photon detuning. Annoyingly, the stray fields in the lab were found to drift on the order of tens of mV/cm over the course of a few hours. Therefore, for the longer data-taking runs required for the purity, indistinguishability and experiment in the following chapter, it was necessary to frequently adjust the applied field. This was done on-the-fly using a grad-student PID by monitoring the production efficiency and altering the applied fields to re-maximize it when it was seen to fall. Although this mode of operation worked fine for experiments shown in this thesis, it is unlikely to be practical for any longer




Figure 6.20: Photon production probability as a function of the spin wave storage time.

term operation, or for a lab where the stray fields drift more rapidly. An obvious solution would be to implement an automated PID to feedback and zero out the fields. However, the need for such a PID would rapidly become unnecessary if the principal quantum number were lowered, given the $(n^*)^7$ scaling of the polarizability.

A major limitation on the production efficiency is the finite lifetime of the Rydberg spin wave, which can be seen in figure 6.20, which affects every stage of the production cycle. The lifetime, on the order of microseconds, is significantly smaller than that of the Rydberg state itself, as a result of dephasing of the stored spin wave. There are a number of dephasing mechanisms at play in the system related to, motion of the atoms during storage, differential scalar light shifts, collisional shifts of the Rydberg state etc. From calculations, none of the known dephasing mechanisms seem to fully account for the observed lifetime. However, other exper-

iments performed with Rydberg spin waves have demonstrated significantly longer lifetimes [123, 124], so there is potential scope for improvement.

Rather than increasing the lifetime of the Rydberg spin wave, an alternative solution, and one that is being actively worked on, is the idea of shelving the Rydberg excitations. Here, after the write stage we would apply a second set of pulses with the control and a shelving beam, coupling the Rydberg and some other “shelving” ground state, different from the one the ensemble is initialized in. This would map the Rydberg spin wave into a spin wave that is a symmetric superposition state with exactly one atom in the shelving state. A benefit here is that ground state spin waves can be made to have lifetimes that are near infinite in relation to our experimental cycle times [70]. Additionally, the effective time spent in the Rydberg state would be restricted to the duration of the write and shelving processes, which is far smaller than at present. Further, the lack of need to change the control frequency mid-production cycle means we could increase the control Rabi frequency by working with a single pass, rather than double pass AOM.

While useful for the purpose of increasing the efficiency of the single photon source, the idea of spin wave shelving has a number of further uses. For one, it is a crucial step for modifying the source to generate photon-spin entanglement [31, 32], and therefore for the creation of a Rydberg-ensemble based quantum repeater. Another potential avenue of research would be in using the shelving for creating more exotic states of light. For example, in a similar vein to reference [125], we could imagine preparing arbitrary atomic spin wave Fock states. By applying a near-resonant shelving beam these could be read out as a photonic Fock state.

Along the same lines, we could even think about making arbitrary photonic states, in a similar manner to that in reference [126]. Slightly more ambitiously still, with the ability to shelve into multiple states of the ground manifold there is the potential to make more exotic and potentially useful photonic states, such as N00N states.

Chapter 7: Quantum Interference Between Photons from an Atomic Ensemble and a Remote Atomic Ion

Advances in the distribution of quantum information will likely require entanglement shared across a hybrid quantum network [36–38]. The complimentary strengths and functions of the different quantum systems gives heterogeneous networks an advantage over those consisting of identical nodes. Many protocols for generating remote entanglement require interference between photons produced by the different network nodes, which has largely prevented investigations into photonic-based hybrid entanglement, owing to the large differences in the spectral characteristics of single photons generated by different quantum systems [28,30,36]. Although this is not a physical limitation [127,128], vanishing entanglement generation rates, along with the necessity for detectors with bandwidths orders-of-magnitude greater than currently available has prohibited the linking of heterogeneous systems. Overcoming this spectral disparity will allow for the construction of hybrid networks with practical entanglement rates and expanded capabilities compared to a networks based only on homogeneous components [36].

Two of the leading systems in the field of quantum information are Rydberg atoms, which we have already discussed at length, and trapped ions. The strong

optical nonlinearity exhibited by neutral-atom Rydberg ensembles enables the construction of single-photon sources [64], gates [53], and transistors [129]. Strong light-matter interactions make them well suited as quantum memories [123], and for implementing quantum repeaters [31, 32]. Furthermore, arrays of Rydberg atoms are a powerful new platform for quantum simulation [46, 130]. Trapped ions are leading candidates for quantum computation, communication, and simulation with good matter-to-photon conversion [131–138]. Their continued success owes to long coherence and trapping lifetimes [135], high fidelity operations [136], and ease of generating ion-photon entanglement [137, 138].

In this chapter, I describe the observation of high-visibility Hong-Ou-Mandel (HOM) interference [89] between photons generated from a rubidium atomic ensemble and a trapped barium ion after closely matching their center frequencies via difference frequency generation (DFG) [139]. Recognizing that this is a key step towards the entanglement of these two systems, towards the end of the chapter we shall explore the feasibility of generating hybrid ion-atomic ensemble remote entanglement. This chapter has been adapted from a published work [40] which was a collaboration between our lab and Qudsia Quraishi’s trapped barium ion group.

7.1 Experimental Configuration

This experiment, shown in figure 7.1, spans two buildings. Building A contains a single trapped $^{138}\text{Ba}^+$ ion as well as two DFG setups. Building B contains the previously described ^{87}Rb atomic ensemble and a HOM interferometer to measure




Figure 7.1: Joint experimental layout. Here VBG and ULEC stand for volume Bragg grating and ultra low expansion cavity respectively.

two-photon interference. In this section I describe the various components that make up the total experimental setup.

7.1.1 Barium Trapped Ion Details

The configuration of the experiment in building A, and level scheme for the barium ion are shown in figure 7.2. Here, a single $^{138}\text{Ba}^+$ is confined and Doppler cooled using a radio-frequency Paul trap [140] and 493-nm light. An additional laser at 650-nm is used as a re-pumper. The ion emits single photons near 493 nm via spontaneous emission from the $6P_{1/2}$ excited state to the $6S_{1/2}$ ground state. Photon collection is performed with a custom coated, ex-vacuo 0.4 NA objective,




Figure 7.2: Experimental layout and energy level diagram for barium ion system.

corresponding to about 4% collection of light emitted by the ion. The total photon collection efficiency out of the SMF is $\approx 1\%$. The collected photons are spatially combined with the pump laser on a dichroic mirror before passing through a 20X objective and coupled into a wave guide in the periodically poled lithium niobate crystal, DFG-1 (SRICO Model: 2000-1005). Here, DFG converts the 493-nm photons to 780 nm, whilst preserving their quantum statistics [139, 141]. The converted and unconverted photons, as well as pump light, exit through a fiber butt-coupled and glued to the output of the wave guide. The light from DFG-1 is passed through a set of interference filters (two each of Semrock: LL01-780-25 and FF01-1326/SP-25) and a volume Bragg grating (OptiGrate BP-785) to filter out the pump, noise photons, and unconverted photons [142]. Finally, the remaining 780-nm single photons are passed through a polarization filter before being coupled into the PMF

connecting to Building B. The single-photon conversion efficiency, measured as the ratio between the number of output 780-nm photons after the Bragg grating and the number of 493-nm photons before combination with the pump, is $\approx 10\%$ on average, with fluctuations that are attributed to photo refractive effects caused by the high intensity pump light [143–145]. Due to these effects, the experimental run time is limited to about 10 hr before the pump must be turned off for an extended period of time (10 hr) to allow the crystal to recover.

To ensure the converted ion-produced photons are at a similar frequency as those produced by the atomic ensemble, a second frequency conversion setup, DFG-2, is used. Laser light at 493-nm, with a known detuning from the photons emitted by the ion (± 10 MHz), is combined with the same pump light used in DFG-1, producing continuous wave light at 780-nm. The 780-nm light from DFG-2 is combined with frequency-locked 780-nm light from building B onto a fast photodetector (Electro-Optics Technology ET-2030A), with which we measure an optical beat note. We use the beat note to produce an error signal, which is then used to feed back to the pump laser’s frequency control. We stabilize and set the frequency of the output 780-nm light from both DFG setups. Uncertainties in the center frequency of the converted 780-nm single photons were present in the experiment due to uncertainties in the ion spectroscopy, and drifts in the 493-nm and 650-nm laser wavemeter locks. This affects the two-photon interference and is investigated in a later section.




Figure 7.3: Experimental layout and energy level diagram for atomic ensemble system.

7.1.2 Atomic Ensemble Details

As the primary focus of this thesis, the atomic ensemble experiment has already been discussed in detail. However, for the sake of completeness, I include the brief description of the experiment from the published work.

The atomic-ensemble source uses Rydberg blockade [146] to produce single photons, utilizing a typical Rydberg polariton experimental layout [41, 129]. The ground, $|5S_{1/2}, F = 2, m_F = 2\rangle$, and Rydberg states, $|nS_{1/2}, J = 1/2, m_J = 1/2\rangle$ are coupled using a two-photon transition, via an intermediate state, $|5P_{3/2}, F = 3, m_F = 3\rangle$, shown in figure 7.3. The 780-nm probe light that has passed through the cloud is collected and coupled ($\approx 70\%$ efficiency) into a PMF. We operate with Rydberg levels with principal quantum numbers, $n \geq 120$, where the blockade radius is significantly larger than the probe beam waist, making the medium effectively one dimensional [41]. The atomic ensemble has a lifetime of ≈ 1 s, limited by the background vapor pressure. Thus, to maintain reasonable atom numbers over the course




Figure 7.4: Experimental layout for HOM interferometer.

of the measurements, we periodically reload the ensemble.

7.1.3 HOM Interferometer

Light from both sources is transmitted to the interferometer setup, shown in figure 7.4, by PMF. At the output of each fiber we use a PBS, not pictured, to clean the polarization of the light before it passes to the 50:50 beamsplitter. For the light from the atomic-ensemble source, a HWP after the PBS allows us to adjust the relative polarization of the two sources at the 50:50 beamsplitter. We couple the output ports of the 50:50 beamsplitter to a pair of SMFs with similar mode field diameters to the input PMFs, which are connected to a pair of SPADs (Excelitas SPCM-780-13). Immediately prior to both output SMFs we use a bandpass filter (Semrock Brightline 780/12) to remove stray light. We use a time-tagger (Roithner-Laser TTM8000) to record timestamps for SPAD detection events, from which we use software to calculate coincidences.




Figure 7.5: $g^{(2)}(\tau)$ for stochastically produced photons from the ion source.

7.2 Interference for Stochastically Produced Photons

First we consider the case of interference where each source continuously produces single photons with stochastic arrival times.

7.2.1 Trapped Ion Stochastic Photon Production

For the case of stochastic photon generation, the ion is constantly Doppler cooled with Rabi frequencies of $\approx 2\pi \times 25$ MHz and $\approx 2\pi \times 11$ MHz, and detunings $\approx 2\pi \times (-99)$ MHz and $\approx 2\pi \times 29$ MHz, for the 493-nm and 650-nm beams respectively. The average count rate of converted photons throughout the experiment, R_{ion} , measured as the sum of counts on SPAD A and B in Building B, is ≈ 400 s^{-1} . Figure 7.5 shows $g_{\text{ion}}^{(2)}(\tau)$ for the resulting 780-nm photon stream. We measure




Figure 7.6: $g^{(2)}(\tau)$ for stochastically produced photons from the atomic ensemble source.

$g_{\text{ion}}^{(2)}(0) = 0.05(8)$ after background subtraction. Here, the $g_{\text{ion}}^{(2)}$ dip width is set by the effective Rabi frequency ($\approx 2\pi \times 100$ MHz) of the driving 493-nm light, which additionally dictates the emitted photon's bandwidth.

7.2.2 Atomic Ensemble Stochastic Photon Production

To produce a stochastic photon stream from the atomic ensemble source, we tune the probe and 479-nm control fields to their respective atomic resonances, $\Delta = \delta = 0$ (see figure 7.3). Rydberg EIT ensures that only single photons propagate through the medium without large losses [41]. In steady-state operation at a high Rydberg level, $n = 120$, and large optical depth, $\text{OD} \approx 30$, we observe a background subtracted $g_{\text{atom}}^{(2)}(0) = 0.119(7)$, shown in figure 7.6. This is significantly larger




Figure 7.7: Normalized coincidences for the cases where the relative polarization of the two sources at the interferometer are parallel, $n_{\parallel}(\tau)$, and perpendicular, $n_{\perp}(\tau)$.

than the value of $g^{(2)}(0)$ we saw in the previous chapter for the on-demand source. However, this is not necessarily surprising given the differing mechanisms which set $g^{(2)}(0)$ in both cases. Here we attribute the non-zero value to finite probe beam size and input photon flux effects [41, 91]. The width of the $g_{\text{atom}}^{(2)}$ dip is set by the Rydberg-EIT bandwidth [41]. However, in contrast to the barium ion, the majority of the photons exiting the medium have similar spectral bandwidths to the input probe field [147]. We measure an average photon count rate throughout the experiment, R_{atom} , of $\approx 10^4 \text{ s}^{-1}$, with an experimental duty cycle of 0.56, where the off time is used for reloading.

7.2.3 Stochastic Hong-0u-Mandel Results

The background-subtracted normalized coincidences for the HOM interference are shown in figure 7.7 for the cases where the relative polarization at the interferometer of the photons from the two sources are parallel, $n_{\parallel}(\tau)$, and perpendicular,




Figure 7.8: Beamsplitter input and output operators. Input photons are represented by the raising operators \hat{a}^\dagger and \hat{b}^\dagger for the atomic-ensemble and ion sources respectively, with \hat{x}^\dagger and \hat{y}^\dagger representing photons exiting each of the two output ports.

$n_\perp(\tau)$. We define the visibility of the interference

$$V = \frac{n_\perp(0) - n_\parallel(0)}{n_\perp(0)} \quad (7.1)$$

and observe $V = 0.43(9)$ using 1-ns bins.

For a perfect 50:50 beamsplitter two factors can contribute to a non-unity visibility: multiphoton events from either of the sources, quantified by $g^{(2)}(0)$, and distinguishability. Given that $g^{(2)}(0)$ is not perfectly zero for either of the sources, let us determine the suppression of V due to multiphoton events from either of the two sources. The relation between the input and output bosonic operators shown

in figure 7.8 is given by

$$\begin{pmatrix} \hat{a}^\dagger \\ \hat{b}^\dagger \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix} \begin{pmatrix} \hat{x}^\dagger \\ \hat{y}^\dagger \end{pmatrix}. \quad (7.2)$$

Initially, we consider the case where a single photon from the atomic ensemble, represented by $\hat{a}_{\text{atom}}^\dagger$, and a single photon from the ion, represented by $\hat{b}_{\text{ion}}^\dagger$, are present at separate inputs of beamsplitter. For such an input state we have:

$$\begin{aligned} |1_{\text{atom}}, 1_{\text{ion}}\rangle_{\text{in}} &= \hat{a}_{\text{atom}}^\dagger \hat{b}_{\text{ion}}^\dagger |0, 0\rangle_{\text{out}} \\ &\rightarrow \frac{1}{2} \left(\hat{x}_{\text{atom}}^\dagger + i\hat{y}_{\text{atom}}^\dagger \right) \left(i\hat{x}_{\text{ion}}^\dagger + \hat{y}_{\text{ion}}^\dagger \right) |0, 0\rangle_{\text{out}}. \end{aligned} \quad (7.3)$$

To take into account the distinguishability between the two input photons, we define

$$\begin{aligned} \hat{x}_{\text{atom}}^\dagger &= \hat{x}^\dagger \\ \hat{y}_{\text{atom}}^\dagger &= \hat{y}^\dagger \\ \hat{x}_{\text{ion}}^\dagger &= \sqrt{c} \hat{x}^\dagger + \sqrt{1-c} \hat{x}_n^\dagger \\ \hat{y}_{\text{ion}}^\dagger &= \sqrt{c} \hat{y}^\dagger + \sqrt{1-c} \hat{y}_n^\dagger, \end{aligned} \quad (7.4)$$

where c is a real number, $0 \leq c \leq 1$, that parameterizes the mode overlap of the two photons. Here \hat{x}_n^\dagger and \hat{y}_n^\dagger consist of all modes orthogonal to \hat{x}^\dagger and \hat{y}^\dagger , respectively,

i.e. $0 = \langle 0 | \hat{x}_n \hat{x}_n^\dagger | 0 \rangle = \langle 0 | \hat{y}_n \hat{y}_n^\dagger | 0 \rangle$. From equation 7.3 and 7.4 we find

$$\begin{aligned}
|1_{\text{atom}}, 1_{\text{ion}}\rangle_{\text{in}} &\rightarrow |\psi_{\text{out}}\rangle \\
|\psi_{\text{out}}\rangle &= \frac{i}{2} \left(\sqrt{2c} |2, 0\rangle_{\text{out}} + \sqrt{1-c} |1\&1_n, 0\rangle_{\text{out}} + \sqrt{2c} |0, 2\rangle_{\text{out}} + \sqrt{1-c} |0, 1\&1_n\rangle_{\text{out}} \right) \\
&\quad + \frac{\sqrt{1-c}}{2} (|1, 1_n\rangle_{\text{out}} - |1_n, 1\rangle_{\text{out}}),
\end{aligned} \tag{7.5}$$

where we use the notation $|1\&1_n, 0\rangle_{\text{out}}$ and $|0, 1\&1_n\rangle_{\text{out}}$ to denote instances where the photons exit the same port of the beamsplitter, but are otherwise in orthogonal modes. The probability, $P_{|1,1\rangle_{\text{in}} \rightarrow 1,1}$, of finding a photon at both output ports is then given by

$$\begin{aligned}
P_{|1,1\rangle_{\text{in}} \rightarrow 1,1} &= |\langle 1, 1 | \psi_{\text{out}} \rangle|^2 + |\langle 1, 1_n | \psi_{\text{out}} \rangle|^2 + |\langle 1_n, 1 | \psi_{\text{out}} \rangle|^2 \\
&= \frac{(1-c)}{2}.
\end{aligned} \tag{7.6}$$

From this, we can see that a value of $c = 1$ corresponds to perfect interference, with zero probability of finding photons at both output ports simultaneously. Similarly, a value of $c = 0$ corresponds to no interference between the photons, with equal probability for the photons to exit the same port, or separate ports.

We now repeat the same procedure for the situation where two photons from the atomic-ensemble source are present at the beamsplitter with none from the ion

source, i.e. $|2_{\text{atom}}, 0\rangle_{\text{in}}$

$$\begin{aligned}
|2_{\text{atom}}, 0\rangle_{\text{in}} &= \frac{1}{\sqrt{2}} \left(\hat{a}_{\text{atom}}^\dagger \right)^2 |0, 0\rangle_{\text{out}} \\
&\rightarrow \frac{1}{2\sqrt{2}} (\hat{x}^\dagger + i\hat{y}^\dagger)^2 |0, 0\rangle_{\text{out}} \\
&\rightarrow \frac{1}{2} \left(|2, 0\rangle_{\text{out}} - |0, 2\rangle_{\text{out}} + i\sqrt{2} |1, 1\rangle_{\text{out}} \right).
\end{aligned} \tag{7.7}$$

Using a similar procedure as for the $|1, 1\rangle$ input state, the probability of finding a photon at both output ports is $P_{|2,0\rangle_{\text{in}} \rightarrow 1,1} = 1/2$. From symmetry this is the same for the case of the $|0, 2_{\text{ion}}\rangle_{\text{in}}$ input state, i.e. $P_{|0,2\rangle_{\text{in}} \rightarrow 1,1} = 1/2$.

For near-single photon sources, such as those used in this work, input states with total photon number > 2 occur with negligible probability. We now calculate the coincidence rate, ignoring such terms

$$\mathcal{R}(\tau = 0) = \frac{P_{\text{atom}} P_{\text{ion}}}{\Delta\tau} P_{|1,1\rangle_{\text{in}} \rightarrow 1,1} + \frac{P_{2 \times \text{atom}}}{\Delta\tau} P_{|2,0\rangle_{\text{in}} \rightarrow 1,1} + \frac{P_{2 \times \text{ion}}}{\Delta\tau} P_{|0,2\rangle_{\text{in}} \rightarrow 1,1}, \tag{7.8}$$

where $P_{\text{atom(ion)}}$ and $P_{2 \times \text{atom(ion)}}$ are the probabilities of having a single and two photons from the specified source in a time interval, $\Delta\tau$, respectively.

For the case of continuously produced photons, we can rewrite the single photon probability in terms of the singles rates, $R_{\text{atom(ion)}}$, $P_{\text{atom(ion)}} = R_{\text{atom(ion)}} \Delta\tau$. Additionally we can make the approximation $g_{\text{atom(ion)}}^{(2)}(0) \approx 2P_{2 \times \text{atom(ion)}}/P_{\text{atom(ion)}}^2$ which is nearly exact in the limit of small photon flux [92], as is true for our exper-

iment. Equation 7.8 can then be written

$$\begin{aligned}\mathcal{R}(0) &= \Delta\tau R_{\text{atom}} R_{\text{ion}} \left[\frac{(1-c)}{2} + \frac{1}{4} \left(\frac{R_{\text{atom}}}{R_{\text{ion}}} g_{\text{atom}}^{(2)}(0) + \frac{R_{\text{ion}}}{R_{\text{atom}}} g_{\text{ion}}^{(2)}(0) \right) \right] \\ &= \Delta\tau R_{\text{atom}} R_{\text{ion}} \left[\frac{(1-c)}{2} + \frac{1}{4} \left(r g_{\text{atom}}^{(2)}(0) + r^{-1} g_{\text{ion}}^{(2)}(0) \right) \right],\end{aligned}\quad (7.9)$$

where $r = R_{\text{atom}}/R_{\text{ion}}$. From equation 7.9 we can calculate the expected normalized coincidences as a function of c , $g^{(2)}(0)$ for the two sources and r

$$\begin{aligned}n(0) &= \frac{\mathcal{R}(0)}{\mathcal{R}(0)|_{c=0, g_{\text{atom}}^{(2)}(0)=1, g_{\text{ion}}^{(2)}(0)=1}} \\ &= \frac{2(1-c) + r g_{\text{atom}}^{(2)}(0) + r^{-1} g_{\text{ion}}^{(2)}(0)}{2 + r + r^{-1}}.\end{aligned}\quad (7.10)$$

Assuming no interference ($c = 0$) in the perpendicular case, and making no assumptions about the overlap in the parallel case, we can calculate an expected visibility

$$\begin{aligned}V_{\text{exp}} &= \frac{n_{\perp}(0) - n_{\parallel}(0)}{n_{\perp}(0)} \\ &= \frac{n(0)|_{c=0} - n(0)}{n(0)|_{c=0}} \\ &= 1 - \frac{\left(2(1-c) + r_{\parallel} g_{\text{atom}}^{(2)}(0) + r_{\parallel}^{-1} g_{\text{ion}}^{(2)}(0) \right) (2 + r_{\perp} + r_{\perp}^{-1})}{\left(2 + r_{\perp} g_{\text{atom}}^{(2)}(0) + r_{\perp}^{-1} g_{\text{ion}}^{(2)}(0) \right) (2 + r_{\parallel} + r_{\parallel}^{-1})}.\end{aligned}\quad (7.11)$$

If we assume the rates are similar in the parallel and perpendicular measurements, $r = r_{\perp} = r_{\parallel}$, equation 7.11 reduces to

$$V_{\text{exp}} = c \left[1 + \frac{r g_{\text{atom}}^{(2)}(0) + r^{-1} g_{\text{ion}}^{(2)}(0)}{2} \right]^{-1}.\quad (7.12)$$

In the limit that $g_{\text{atom}}^{(2)}(0) = g_{\text{ion}}^{(2)}(0) = 0$ we see that the visibility is equal to the mode overlap of the photons from the two sources, c . However, given a finite $g^{(2)}(0)$ for either source the visibility is reduced, due to multi-photon events, by a factor

$$f_{mp} = \left[1 + \frac{r g_{\text{atom}}^{(2)}(0) + r^{-1} g_{\text{ion}}^{(2)}(0)}{2} \right]^{-1}. \quad (7.13)$$

Given the independently measured $g^{(2)}(0)$ for the sources and ratio, r , from equation 7.13 we determine $f_{mp} = 0.41(1)$, where we have assumed the flux from the two sources is constant throughout the experiment, which is a valid approximation for this data. The observed $V = 0.43(9)$ can thus be attributed entirely to multi-photon events, and therefore is consistent with perfect bunching of photons from the two sources. Additionally, we note that $n_{\parallel}(0)$ and $n_{\perp}(0)$ are in agreement with the values expected from the measured $g^{(2)}(0)$'s, shown by the bands in figure 7.7, which are calculated using equation 7.10. Due to the disparity in the spectral widths of the photons produced by the sources, the width of the HOM dip, seen in figure 7.7, is almost entirely determined by the temporally narrower ion-produced photon.

7.3 Interference for On-Demand Produced Photons

To be useful for quantum networking, the photons should arrive on demand in well-defined temporal modes [148]. To this end, we investigate two-photon interference in the case where a single photon from each source arrives at a known time relative to an experimental trigger shared between the two buildings.




Figure 7.9: Photon profile and area-integrated $g^{(2)}(\tau)$ for on-demand photons produced by the trapped ion source.

7.3.1 Trapped Ion On-Demand Photon Production

To produce on-demand single photons from the ion, we use a process similar to that described in [142]. First, the ion is pumped into the $D_{5/2}$ manifold using only the 493-nm laser for 750 ns. Due to the polarization of the 493-nm light, the ion is pumped, with equal likelihood, into any one of the Zeeman sub-levels. This light is then turned off, and we wait for 60 ns to ensure any laser scatter is not detected during the photon extraction phase. A 200 ns pulse of 650-nm light, separate from that used to Doppler cool, is then used to excite the ion to the $P_{1/2}$ manifold, from which a 493-nm photon may be emitted, $\approx 75\%$ branching ratio. This 650-nm light is linearly polarized, and propagates along the quantization axis, with a Rabi Frequency of $\approx 2\pi \times 30$ MHz and detuning of $\approx 2\pi \times 29$ MHz. The 650-nm light is then turned off, and after a short period, 60 ns, with no light, Doppler cooling of the ion is resumed, for a minimum of 500 ns.

We detect a photon at the output of the HOM interferometer with a proba-




Figure 7.10: Photon profile and area-integrated $g^{(2)}(\tau)$ for on-demand photons produced by the atomic-ensemble source.

bility of $\approx 2 \times 10^{-5}$ per attempt. From a Hanbury Brown and Twiss measurement performed on the 493 nm light prior to conversion we find $g_{\text{ion}}^{(2)}(0) = 0(1) \times 10^{-2}$ after background subtraction, as shown in figure 7.9. Photons are emitted from the ion with a nearly exponential decaying temporal profile, with a decay constant, ≈ 50 ns, set by the effective Rabi frequency of the 650-nm retrieval pulse.

Due to the magnetic bias field, ≈ 5 G, splitting the Zeeman states in the $6S_{1/2}$ and $5D_{3/2}$ levels, combined with the near-equal population distribution in the $5D_{3/2}$ manifold following pumping, the average photon spectrum consists of several distinct peaks. This affects the shape of the interference dip, and is discussed in more detail below.

7.3.2 Atomic-Ensemble On-Demand Photon Production

For the atomic-ensemble source, we generate on-demand photons using the write and retrieve protocol, documented in chapter 6. Although this has been thoroughly discussed in that chapter, I am again including the brief description of the

protocol from the published work for completeness.

A Rydberg spin wave is written to the cloud using coherent control and probe pulses, detuned far from intermediate resonance, $\Delta \gg \Gamma$, the linewidth of the intermediate state, and close to two-photon resonance, $\delta \approx 0$. Rydberg blockade during the write process ensures that a single Rydberg spin wave excitation is stored in the medium. The control field is tuned close to resonance and then turned on, retrieving the spin wave as a single photon with a spatial mode similar to the input probe light. After accounting for background coincidences, we measure¹ $g_{\text{atom}}^{(2)}(0) = 0(1) \times 10^{-4}$, with a per-attempt detection probability $\approx 3 \times 10^{-2}$ at the outputs of the HOM interferometer. The temporal profile of the retrieved photon is determined by the control Rabi frequency, $\approx 2\pi \times 7$ MHz, intermediate state detuning during retrieval, $\approx 2\pi \times 7$ MHz, and optical depth, ≈ 10 , of the cloud [149]. Figure 7.10 shows the temporal profile of the atomic-ensemble produced photon, also well approximated by a decaying exponential, with a decay constant ≈ 120 ns.

7.3.3 Synchronization and Measurement Scheme

Given the on-demand nature of the two sources, synchronization is necessary to control the arrival times of the photons at the interferometer. To synchronize the two experiments we operate in primary-minion configuration with the atomic-ensemble lab, in building B, as the primary and the ion lab, in building A, as the minion. In the ensemble lab we generate 1064-nm optical pulses using an AOM

¹Note that there is a small discrepancy between the value of $g^{(2)}(0)$ reported here compared to the previous chapter. This is likely due to the different gating methods used.

with laser light. These are sent over fiber to the ion lab where the optical pulse is converted to TTL, using a high bandwidth (Thorlabs PDA05CF2) photodiode, which triggers photon production. Due to drifts in the power of the 1064-nm optical pulse, we observe small drifts (≤ 20 ns over several hours) in the ion-produced photon arrival time relative to that generated by the atomic-ensemble.

To measure the visibility in a single experimental run, instead of using polarization to make the photons distinguishable, we use a procedure where the ion-produced photons alternately arrive simultaneously on the beamsplitter with the atomic-ensemble produced photons, with identical polarization, interleaved with pulses when their arrival times are not overlapped, depicted in figure 7.11. Our experimental sequence consists of requesting photons from the atomic ensemble at a rate of 200 kHz, while the ion produces photons at 400 kHz. To ensure the photon profiles overlap, even with the temporal drifts, we offset the average arrival time of the ion produced photon $\approx +40$ ns relative to atomic-ensemble produced photon. Calculations indicate that with this offset, such temporal drifts have negligible effect on the two-photon interference. We observed no measurable drift between the temporally overlapped and non-overlapped photons produced by the ion, which have a temporal separation of $2.5 \mu\text{s}$. Due to the finite lifetime of the atomic-ensemble trap, we operate at an experimental duty cycle of 0.6.

Along with events on SPAD A and B we additionally record timestamps for an electronic reference, which defines an absolute time reference within the $5\text{-}\mu\text{s}$ pulse cycle. This reference was provided by the same electronics that controlled the arrival time of the photons produced by the two systems. Throughout the experiment we




Figure 7.11: Photon profiles for the two sources for one measurement period.

observed no drift between the arrival time of the atomic-ensemble produced photon and the electronic reference.

7.3.4 On-Demand Hong-Ou-Mandel Results

Due to the relatively low flux of ion-produced photons at the output of the interferometer, we perform software gating to improve the signal-to-noise ratio of the coincidences for the on-demand interference. We calculate coincidences, $\mathcal{C}(t, \tau)$, between the two SPADs as a function of the relative time τ between events on SPAD A and B, and the time t between the event on SPAD A and the electronic clock, shown in figure 7.12. We take two windows, denoted by the regions between the dashed-black lines in figure 7.12, and calculate $\mathcal{C}(\tau) = \sum_{t \in t_w} \mathcal{C}(t, \tau)$, where t_w are the set of times in the windows. For our data, we chose a 120-ns window size, encompassing $\approx 80\%$ of the area of the ion-produced photon, that provided a good compromise between data accumulation rate and signal-to-noise. This method is




Figure 7.12: Photon profiles and three-fold coincidences for on-demand photons. Gating window shown by black dotted line. Coincidences are calculated between the two SPADs as a function of relative time between events on the two SPADs, τ , and the absolute time within the pulse sequence, t .

equivalent to a physical gate on the SPADs over the time period of integration.

Due to the large disparity, $\approx \times 10^3$, in the ion and atomic-ensemble-produced photon detection probabilities, coincidences between any two independent ion-produced photons are negligible. Ignoring background coincidences, which will be discussed later, features in $\mathcal{C}(\tau)$ around $\tau = 0$ arise from instances where the atomic-ensemble-produced and ion-produced photons are overlapped, while features around $\tau = \pm(2.5 + 5k)$ μs , for $k \in \mathbb{Z}$, arise from instances where the two photons are not overlapped. To directly compare the cases where the photons are overlapped to the case where they are non-overlapped, we temporally shift the $\mathcal{C}(\tau)$ curve. Additionally, we average together several of these temporally shifted curves in order to reduce our uncertainties for the non-overlapped case. Procedurally, we take a set of $\mathcal{C}(\tau')$ curves and shift each curve by an amount $\tau_k = (2.5 + 5k)$ μs to obtain a set $\{\mathcal{C}(\tau' + \tau_k)\}$. We then average to obtain the non-overlapped curve, $\langle \mathcal{C}(\tau = \tau' + \tau_k) \rangle_k$, shown in figure 7.13. For the data presented in figure 7.14, the non-overlapped curve was




Figure 7.13: Method for determining temporally overlapped and non-overlapped coincidences. Upper plots show coincidences after gating. These are then time shifted and averaged to give lower plot.

constructed using $-5 \leq k \leq 4$.

A non-negligible number of background coincidence events were recorded during the data taking. These are predominantly due to coincidences between the atomic-ensemble produced photon and background events on the SPADs either due to dark counts, or ambient light leakage. We calculate the expected background coincidence curve, shown in figure 7.13, from the independently measured experimental singles and background rates from the two SPADs. For this calculation




Figure 7.14: Normalized coincidences for on-demand interference measurements.

the same gating described above is used. To obtain the final coincidence curves, shown in figure 7.14, the expected background is subtracted from the overlapped and non-overlapped coincidences and the resulting curves scaled by the same factor. Using equation 7.1 for the visibility, where n_{\parallel} and n_{\perp} correspond to the temporally overlapped and non-overlapped coincidences respectively at zero time delay, we calculate a visibility of $V = 1.1(2)$, indicating perfect two-photon interference.

However, the visibility on its own does not tell the whole story. The width of the dip in coincidence space contains additional information about the interference between the photons, in addition to setting a limit on the usability of the interference, as we shall see later. Therefore, it is important to understand what sets the width of the interference dip. To theoretically model the coincidence profile we shall consider the case where we have two single photons incident on two different input ports of a 50:50 beamsplitter, as in figure 7.8. We assume that the photons have the same transverse-spatial and polarization mode but may have different temporal modes, ζ_i . The probability of detecting a photon at time t_0 in one detector followed

by a detection in the other detector at time $t_0 + \tau$ is given by [127]

$$P(t_0, \tau) = \frac{1}{4} |\zeta_{\text{atom}}(t_0)\zeta_{\text{ion}}(t_0 + \tau) - \zeta_{\text{atom}}(t_0 + \tau)\zeta_{\text{ion}}(t_0)|^2. \quad (7.14)$$

Under the assumption that the photon is transform limited we can write

$$\zeta_i(t) = a_i(t)e^{-i\omega_i t}, \quad (7.15)$$

where $a_i(t)$ is given by the temporal envelope of the photon and ω_i the center frequency. Without loss of generality, we assume $a_i(t) \in \mathbb{R}, \forall t$. We thus have

$$P(t_0, \tau) = \frac{1}{4} [a_{\text{atom}}^2(t_0)a_{\text{ion}}^2(t_0 + \tau) + a_{\text{atom}}^2(t_0 + \tau)a_{\text{ion}}^2(t_0) - 2 \cos(\Delta\omega \tau)a_{\text{atom}}(t_0)a_{\text{atom}}(t_0 + \tau)a_{\text{ion}}(t_0)a_{\text{ion}}(t_0 + \tau)], \quad (7.16)$$

where $\Delta\omega = \omega_{\text{ion}} - \omega_{\text{atom}}$ is the center frequency difference between the two photons. In general, we are interested in the coincidence profile as a function of the relative detection time on two detectors. This is determined from equation 7.16 by integrating over t_0

$$\mathcal{C}_{\text{theory}}(\tau) \propto \int dt_0 P(t_0, \tau), \quad (7.17)$$

where the integral is taken with limits such that any temporal gating is accounted for. For the remainder of this work we take the integral to be over a region that encompasses $\approx 80\%$ of the area of the ion-produced photon, to reproduce the effect of the gating of the SPAD in the experiment. Note that it is this gating which is responsible for the asymmetric shape seen in the experimental and the theoretical




Figure 7.15: Comparison to data of time resolved HOM theory, including accounting for non-transform limited barium photon due to the branching ratio of the $6P_{1/2}$ state.

coincidence curves shown.

Using equation 7.16 and 7.17, along with the profiles for the two photons shown in figures 7.9 and 7.10, we compare the expected and observed shapes of the HOM dip, shown in figure 7.15, assuming the photons are transform limited and have identical center frequencies. To obtain the curve where the two photons do not interfere, we note that in the limit $\Delta\omega \rightarrow \infty$ the final term in equation 7.16 oscillates rapidly as a function of τ and will average to zero with a finite detection bandwidth. A qualitative discrepancy is seen between the experimental observations and what is theoretically expected in the case where $\Delta\omega = 0$. However, the experimental data, outside of the dip around $\tau \approx 0$, matches what we would expect for non-interfering photons.

There are several explanations why the experimentally observed dip is narrower

than that expected by the theory discussed above. Thus far we have assumed $\Delta\omega = 0$ and that the single-photon pulses are transform limited. These assumptions are broken in the experiment in the following ways:

- for the ion, decay from the $6P_{1/2}$ manifold to the $5D_{3/2}$ manifold and subsequent re-excitation during the extraction phase destroys the transform limited character of the produced photons
- experimental uncertainties in the detunings of the 493-nm and 650-nm laser frequencies from their corresponding resonance frequencies causes uncertainties in $\Delta\omega$
- drift on the ion laser locks produces a corresponding drift in $\Delta\omega$
- the scheme used to pump the ion into the $5D_{3/2}$ manifold, as well as Zeeman splitting, causes the ion to emit photons at multiple frequencies

Let us explore the effect of these on the interference profile.

Following the initial excitation from the $5D_{3/2}$ manifold to the $6P_{1/2}$ manifold, the ion can decay to either to the ground state with $\approx 75\%$ probability, emitting a 493-nm photon, or back to the $5D_{3/2}$ manifold with $\approx 25\%$ probability.

If the latter occurs, re-excitation to the $6P_{1/2}$ manifold will result in emission of a 493-nm photon delayed relative to the photon that would have been emitted if the decay to the $5D_{3/2}$ manifold did not occur. This creates a temporally lengthened observed photon shape but leaves the spectrum unchanged [150], resulting in a non-transform limited average temporal profile. To account for this we calculate

the transform limited pulse shape for the ion-produced photons by numerically solving optical Bloch equations. To determine the expected coincidence curve when no scatter back to the $5D_{3/2}$ state occurs, we use equation 7.16 along with the ion's transform-limited pulse shape. For the case where a single scattering event back to the $5D_{3/2}$ state occurs, equation 7.16 is again used along with the ion-produced photon's transform-limited profile, but here we add a probabilistic temporal displacement to account for the photons delayed emission. The total non-transform limited theory curve, shown in figure 7.15, is the sum of these two coincidence curves weighted by the branching ratio. We do not account for higher-order processes, where the ion scatters back to the $5D_{3/2}$ state more than once. As we see in figure 7.15, taking into account the non-transform limited character of the ion-produced photon only marginally alters the expected coincidence curve from that where we assume the two photons to be transform limited.

While an effort was made to ensure that the photons produced by the two sources had identical center frequencies, there were uncertainties in the actual value of $\Delta\omega$, predominantly due to the limited resolution of the ion spectroscopy and drifts in the ion laser locks. From equation 7.16 it can be seen that a non-zero value of $\Delta\omega$ gives rise to an oscillating envelope to the final term. This can cause a reduction in the width of the HOM dip, as seen in figure 7.16, where we have used a value of $\Delta\omega = 2\pi \times 20$ MHz.

Throughout the experiment the ion was subjected to a magnetic bias field (≈ 5 Gauss), lifting the degeneracy of the Zeeman sub-levels for the three states, as shown in figure 7.17. Additionally, the ion is pumped such that the Zeeman




Figure 7.16: Comparison to data of time resolved HOM theory, when accounting for constant center frequency difference, $\Delta\omega = 2\pi \times 20$ MHz, between the photons from the two sources.

states of the $5D_{3/2}$ are populated equally. Thus the 493-nm photons were emitted with a frequency shift given by the differential shift between the initial and final Zeeman sub-levels. While the spectrum of a single ion-produced photon is practically monochromatic with a narrow spectral bandwidth, the average spectrum of the ion source consists of several spectral peaks separated by these differential shifts and weighted by their likelihood. To account for this we modify equation 7.16

$$\begin{aligned}
 P(t_0, \tau) = \frac{1}{4} & \left[a_{\text{atom}}^2(t_0)a_{\text{ion}}^2(t_0 + \tau) + a_{\text{atom}}^2(t_0 + \tau)a_{\text{ion}}^2(t_0) \right. \\
 & \left. - 2a_{\text{atom}}(t_0)a_{\text{atom}}(t_0 + \tau)a_{\text{ion}}(t_0)a_{\text{ion}}(t_0 + \tau) \sum_i c_i \cos((\Delta\omega_i + \Delta\omega)\tau) \right],
 \end{aligned}
 \tag{7.18}$$

where c_i and $\Delta\omega_i$ are the weighting and differential shifts due to Zeeman splitting, and $\sum_i c_i = 1$. Given the population spread across the $5D_{3/2}$ sub-levels, as well




Figure 7.17: Barium Photon Frequencies Including Zeeman Splittings. In our excitation scheme, all possible Zeeman sublevels of $5D_{3/2}$ are equally populated, and only σ excitation transitions are used. **a** Possible excitation and decay paths involving the $|6P_{1/2}, m = -1/2\rangle$ state of $^{138}\text{Ba}^+$. **b**, possible excitation and decay paths involving the $|6P_{1/2}, m = 1/2\rangle$ state of $^{138}\text{Ba}^+$. **c**, shows the polarization and detuning of the resulting 493-nm photons relative to the center of all possible emission frequencies.

as the polarization and propagation direction of the 650-nm excitation light, we expect photons to be emitted from the ion at several frequencies around a mean value with near-equal probability, shown in figure 7.17. As seen in figure 7.18, this type of probabilistic spectrum gives rise to a narrowed HOM dip with subsequent oscillations in coincidence space appearing less pronounced.

Due to the way the DFG light was frequency stabilized, it is likely that the ion-produced photon frequency drifted relative to that of the atomic-ensemble produced




Figure 7.18: Comparison to data of time resolved HOM theory, when accounting for the probabilistic spectrum of the ion-produced photon due to the Zeeman splitting of the $5D_{3/2}$ and $6S_{1/2}$ manifolds, and equal initial state population across the $5D_{3/2}$ manifold.

photon. To include this drift in the theory we modify equation 7.16

$$\begin{aligned}
 P(t_0, \tau) &= \int_{-\infty}^{\infty} d\Delta\omega \frac{e^{-\frac{(\Delta\omega - \Delta\omega_0)^2}{2\sigma_{\Delta\omega}^2}}}{\sigma_{\Delta\omega} \sqrt{2\pi}} P(t_0, \tau) \\
 &= \frac{1}{4} \left[a_{\text{atom}}^2(t_0) a_{\text{ion}}^2(t_0 + \tau) + a_{\text{atom}}^2(t_0 + \tau) a_{\text{ion}}^2(t_0) \right. \\
 &\quad \left. - 2 \cos(\Delta\omega_0 \tau) e^{-\frac{1}{2}\sigma_{\Delta\omega}^2 \tau^2} a_{\text{atom}}(t_0) a_{\text{atom}}(t_0 + \tau) a_{\text{ion}}(t_0) a_{\text{ion}}(t_0 + \tau) \right],
 \end{aligned} \tag{7.19}$$

where we have assumed a Gaussian profile to the drift with an average detuning, $\Delta\omega_0$, and variance, $\sigma_{\Delta\omega}$. Using values, $\Delta\omega_0 = 0$ and $\sigma_{\Delta\omega} = 2\pi \times 10$ MHz in figure 7.19, we see that the theoretical HOM dip narrows while not exhibiting the large oscillations characteristic of a static frequency offset between the two photons.




Figure 7.19: Comparison to data of time resolved HOM theory, when accounting for relative drift of the center frequencies of the two photons over the course of the experiment due to laser drift. For the plot, $\sigma_{\Delta\omega} = 2\pi \times 10$ MHz and $\Delta\omega_0 = 0$.




Figure 7.20: Comparison to data of time resolved HOM theory, when accounting for the combination of experimental imperfections. Theory curve calculated from equation 7.20 showing the combination of the separate experimental imperfections described in figures 7.15, 7.16, 7.18 and 7.19.

To account for all the effects discussed above, equation 7.16 becomes

$$P(t_0, \tau) = \frac{1}{4} \left[a_{\text{atom}}^2(t_0) a_{\text{ion}}^2(t_0 + \tau) + a_{\text{atom}}^2(t_0 + \tau) a_{\text{ion}}^2(t_0) - 2\mathcal{A}(t_0, \tau) \sum_i c_i \cos(\Delta\omega_i + \Delta\omega_0) \tau e^{-\frac{1}{2}\sigma_{\Delta\omega}^2 \tau^2} \right], \quad (7.20)$$

where we have modified the original $a_{\text{atom}}(t_0)a_{\text{atom}}(t_0+\tau)a_{\text{ion}}(t_0)a_{\text{ion}}(t_0+\tau)$ term from equation 7.16 to include the non-transform limitedness of the ion produced photon, denoted by $\mathcal{A}(t_0, \tau)$. We use equation 7.20, with the values used for the individual plots in figures 7.15, 7.16, 7.18 and 7.19, which represent reasonable experimental estimates, to produce the theory curve in figure 7.20, which more closely matches the data. The factors discussed which pose the largest problems, namely effects which cause the center frequencies of the two sources to be different either through experimental drift, constant offsets or differential Zeeman shifts, are correctable. Therefore, there is scope to increase the HOM dip width which, as we shall see, would be beneficial for future projects relating to entanglement generation.

7.4 Outlook

Having observed interference between photons generated from two fundamentally different quantum sources, a natural next step is to think about entangling these two systems. There are multiple ways in which this may be achieved, but let us focus on two that are likely to be experimentally viable, and explore their feasibility.

Let us first discuss the entanglement scheme considered in the published work,

which is similar to that in references [30, 137]. Here, we would extend the HOM interferometer into a Bell-state analyzer, shown in figure 7.21a, requiring the addition of a pair of PBSs and two extra SPADs. An additional requirement for this scheme is that the ion and atomic-ensemble produce photons which are polarization qubits, entangled with their respective sources. In the case of the ion source this is relatively straightforward requiring the minimal change during the production process of pumping into one of the stretched states of the $5D_{3/2}$ manifold [151]. Although this is slightly more complicated for the atomic-ensemble source, as we have discussed in previous chapters, proposals do exist [31, 32] for producing polarization-ensemble entangled photons. For the purposes of analysis, we will assume for both sources that we can produce polarization qubit photons perfectly entangled with an internal degree of freedom of their source, with their otherwise being no change to the source properties. Furthermore, we shall assume that the detectors have negligible background counts, which is potentially realizable by changing from using SPADs to single nanowire single photon detectors (SNSPDs).

As derived in the published work [40], the fidelity for this entanglement generation scheme is given by

$$F = \frac{1 + V}{2}, \quad (7.21)$$

where V is the HOM visibility. The entanglement generation rate may be estimated from the number of coincidences observed in the $\tau = 0$ bin for the non-overlapped case of the on-demand measurement, divided by the experimental runtime, ≈ 21 hours. For the entanglement scheme in figure 7.21, only half the Bell




Figure 7.21: Setup to herald entanglement between distant matter qubits. **a**, incoming photons have their polarizations entangled with their corresponding matter qubit’s internal states. A 50:50 BS is used to interfere the two photons, allowing for the heralding of entanglement between the matter qubits after detection using PBSs and SPDs. **b**, different combinations of detector clicks correspond to the detection of certain photonic bell states. Depending on which set of detectors click (labeled in **a**), different entangled states between the matter qubits can be heralded. In the case of any individual detector clicking, the Φ_+ and Φ_- photonic bell states cannot be distinguished from one another, resulting in a failed attempt to entangle the matter qubits. All other combinations of detector clicks not shown should not be possible in the case of perfect two-photon interference, and are ignored in the case of imperfect interference.

states may be heralded, hence we would expect the entanglement generation rate to be half the coincidence rate observed in experiment. However, in the experiment only half the ion produced photons sent to the DFG setup are converted due to polarization filtering. For the entanglement generation scheme both polarizations would be required to be converted and the factor of two, lost from the heralding process, would be recovered. In the previous section the data shown used a 5-ns bin-width, for which we saw a visibility of $V = 1.1(2)$ with ≈ 40 coincidences in the $\tau = 0$ bin in the non-overlapped case. This allows us to extrapolate an entanglement




Figure 7.22: Inferred entanglement fidelity and generation rate as a function of the coincidence bin width.

generation rate of $\approx 5 \times 10^{-4} \text{ s}^{-1}$ with perfect fidelity. By increasing the bin width used for the heralding process, the entanglement generation rate can be increased, at the expense of reducing its fidelity, as shown in figure 7.22. Although nowhere near high enough for any practical purposes, the values we predict are comparable to those achieved in the first experiments using the same scheme with pairs of homogeneous trapped ion qubits [152, 153].

For this avenue of entangling the two systems, there exists several potential improvements. Most straightforwardly, there are improvements that can be made on the ion experiment that could increase the photon rate including increasing the collection efficiency with a higher NA lens, increasing the initial fiber coupling efficiency, improving the DFG efficiency, and reducing optical losses present from fiber butt couplings and in the the optics in the DFG filtering stage. For the interferometer an increase in the quantum efficiency of the detectors, by transitioning from using SPADs to SNSPDs, would also be a simple beneficial experimental change. Based on conservative estimates, we anticipate all these together having the poten-




Figure 7.23: Theoretical entanglement fidelity and rate speedup as a function of bin width assuming imperfection free interference. Speedup is relative to 5-ns bin width.

tial to increase entanglement generation rates by almost two orders of magnitude, $\approx \times 50$. In addition, resolving the issues described in the previous section that cause the narrowing of the interference dip is also likely to be beneficial, as it would significantly reduce the fidelity hit seen when using large bin widths. Figure 7.23 shows the fidelity and rate speed-up, relative to a 5-ns bin width, for the theory where there are no dip-narrowing experimental imperfections. We can see that if we are able to eliminate the issues, we would be able to use large coincidence bin widths, gaining us more than an order of magnitude in rate with reasonable entanglement fidelity. With these improvements, our rates could start to approach 1 s^{-1} , moving us into a regime where we could start thinking about interesting applications of such an entangled hybrid system.

An alternative to the method discussed above/in the published work, and one that we have been actively pursuing, is the so called “shoot and catch” method, that makes use of the atomic ensembles ability to act as a quantum memory. Here, the

ion source would generate ion-entangled photons that would be sent to the atomic-ensemble. With the ensemble we would store the photon, most likely using an EIT based protocol that maps the photonic qubit states into different spin wave excitations within the cloud, similar to what is described in reference [154]. This scheme confers multiple benefits over the previously described one. Firstly, the generation of photon-ensemble entanglement is not required, which is potentially non-trivial to implement. Additionally, without the need for the optical non-linearities derived from the Rydberg states, the storage could be performed using a λ -like EIT scheme where the spin wave coherence would be between two states from the ground state manifold, for which high efficiency and long-lived storage has been demonstrated [69, 70]. Furthermore, no projective Bell-state measurement is required, which inherently limits the entanglement generation probability of the previous scheme by a factor of two or more, due to optical losses, detector quantum efficiencies, and the imperfect heralding efficiency. While we expect this entanglement generation scheme to be superior, our implementation still remains in its infancy and, therefore, it is difficult to make any concrete predictions as to entanglement rates or fidelities.

Chapter 8: Conclusion and Outlook

In this thesis I have described an experimental apparatus that has been constructed for the purpose of trapping a high optical density ensemble of ultracold atoms, which are addressed by a pair of beams for the production of Rydberg excitations. I have described a project where we have used this apparatus to develop an on-demand high efficiency, high purity and high indistinguishability single photon source, and a further project where we demonstrated Hong-Ou-Mandel interference between photons from our source and photons from a collaborators trapped barium ion. Both of these experiments serve as a proof-of-concept for more ambitious future projects.

In the case of the single-photon source, we have been actively working on trying to get some shelving system working, whereby we can transfer the written Rydberg spin-waves to a ground state spin-wave with a much longer lifetime. As relates to the source itself, this is likely to provide an increase in our photon production efficiency. Additionally, it will open up a number of other avenues of research. For one, it will enable us to transform our source into a quantum repeater node that can produce photon-ensemble entanglement. Thinking further down the road, this would allow us to start to explore the possibility of creating a Rydberg-ensemble

based quantum network. Shelving should allow us to create arbitrary photonic states, and potentially even interesting entangled photonic states, which have a number of applications in fields such as quantum optics and quantum metrology.

For the barium ion collaboration, the obvious next step is to try and entangle the ion and atomic ensemble. There is a couple of ways of achieving this, one of which relies on us being able to develop the single photon source into the repeater node capable of generating photon-ensemble entanglement. The other method, which seems most straightforward from preliminary discussions, involves having the ion “shoot” an ion-entangled photon, and then having the atomic ensemble “catch” it, storing the photon while mapping its qubit states to some ensemble degrees of freedom. With that accomplished, we can start to think about demonstrating quantum networking protocols between the two disparate network nodes, like quantum teleportation or quantum key distribution, and potentially even some practical uses of a hybrid ion-Rydberg ensemble system.

Appendix A: Detailed C_6 Coefficient Calculations

Here, I provide a more detailed calculation of the C_6 coefficients discussed in chapter 2. There, we were considering the dipole-dipole operator

$$\hat{V}_{dd} = \frac{1}{4\pi\epsilon_0 R^3} \left[\hat{d}_1 \cdot \hat{d}_2 - 3(\hat{d}_1 \cdot \vec{n})(\hat{d}_2 \cdot \vec{n}) \right], \quad (\text{A.1})$$

with the $\hat{d}_{1,2}$ being the dipole operator for atom 1 and 2 respectively, and \vec{n} the unit vector along the inter-atomic axis. We can define the spherical operators

$$\hat{d}_0 := \hat{d}_z \quad (\text{A.2})$$

$$\hat{d}_+ := -\frac{\hat{d}_x + i\hat{d}_y}{\sqrt{2}} \quad (\text{A.3})$$

$$\hat{d}_- := \frac{\hat{d}_x - i\hat{d}_y}{\sqrt{2}}. \quad (\text{A.4})$$

Using these definitions, along with assuming, without loss of generality, that \vec{n} subtends an angle θ to the quantisation axis¹ z , we can write the dipole-dipole

¹there's an implicit assumption here that the azimuthal angle $\phi = 0$. However, a non-zero value results in the matrix elements have a phase factor which is irrelevant as we're interested in the absolute value of the matrix elements.

potential as

$$\begin{aligned} \hat{V}_{dd} = \frac{1}{4\pi\epsilon_0 R^3} & \left[\frac{1 - 3\cos^2\theta}{2} (\hat{d}_{1,+}\hat{d}_{2,-} + \hat{d}_{1,-}\hat{d}_{2,+} + 2\hat{d}_{1,0}\hat{d}_{2,0}) \right. \\ & + \frac{3\sin\theta\cos\theta}{\sqrt{2}} (\hat{d}_{1,+}\hat{d}_{2,0} - \hat{d}_{1,-}\hat{d}_{2,0} + \hat{d}_{1,0}\hat{d}_{2,+} - \hat{d}_{1,0}\hat{d}_{2,-}) \\ & \left. - \frac{3\sin^2\theta}{2} (\hat{d}_{1,+}\hat{d}_{2,+} + \hat{d}_{1,-}\hat{d}_{2,-}) \right], \end{aligned} \quad (\text{A.5})$$

where the first number in the subscript for \hat{d} indicates the atom the operator corresponds to, and the second symbol indicates the type of the spherical operator. In this thesis we've been concerned with the Rydberg blockade arising from van der Waals interactions between a pair of identical Rydberg states. Therefore, let us concentrate on matrix elements of the form

$$\langle r, r | \hat{V}_{dd} | r', r'' \rangle, \quad (\text{A.6})$$

where $|r\rangle^{(','')} = |n^{(','')}, l^{(','')}, j^{(','')}, m_j^{(','')}\rangle$ are some Rydberg states. Making use of Clebsch-Gordan coefficients

$$C_{m_j, q, m'_j}^{j, 1, j'} = \langle jm_j 1 q | j' m'_j \rangle, \quad (\text{A.7})$$

we can write the dipole-dipole matrix elements in terms of reduced dipole matrix elements [60]

$$\begin{aligned}
\langle r, r | V_{dd} | r', r'' \rangle &= \frac{\langle r | \hat{d} | r' \rangle \langle r | \hat{d} | r'' \rangle}{4\pi\epsilon_0 R^3} \times \\
&\left[\frac{1 - 3 \cos^2 \theta}{2} (C_{m_j, 1, m'_j}^{j, 1, j'} C_{m_j, -1, m''_j}^{j, 1, j''} + C_{m_j, -1, m'_j}^{j, 1, j'} C_{m_j, 1, m''_j}^{j, 1, j''} + 2C_{m_j, 0, m'_j}^{j, 1, j'} C_{m_j, 0, m''_j}^{j, 1, j''}) \right. \\
&+ \frac{3 \sin \theta \cos \theta}{\sqrt{2}} ((C_{m_j, 1, m'_j}^{j, 1, j'} C_{m_j, 0, m''_j}^{j, 1, j''} + C_{m_j, 0, m'_j}^{j, 1, j'} C_{m_j, 1, m''_j}^{j, 1, j''}) - (C_{m_j, -1, m'_j}^{j, 1, j'} C_{m_j, 0, m''_j}^{j, 1, j''} + C_{m_j, 0, m'_j}^{j, 1, j'} C_{m_j, -1, m''_j}^{j, 1, j''})) \\
&\left. - \frac{3 \sin^2 \theta}{2} (C_{m_j, 1, m'_j}^{j, 1, j'} C_{m_j, 1, m''_j}^{j, 1, j''} + C_{m_j, -1, m'_j}^{j, 1, j'} C_{m_j, -1, m''_j}^{j, 1, j''}) \right].
\end{aligned} \tag{A.8}$$

As established in chapter 2, the C_6 coefficient for a pair state is defined in terms of the perturbative second-order energy shift

$$\Delta E_{r,r}^{(2)} = \sum_{|r', r''\rangle \neq |r,r\rangle} \frac{|\langle r, r | V_{dd} | r', r'' \rangle|^2}{2E_r^{(0)} - E_{r'}^{(0)} - E_{r''}^{(0)}} = \frac{C_6}{R^6}. \tag{A.9}$$

Let us now calculate the C_6 coefficient for pair states which are experimentally relevant to this thesis, namely $|r\rangle = |n, S, 1/2, m_j\rangle$. For the $|r, r\rangle$ pair state, we can see from equation A.9 that the only pair states that will contribute to the C_6 coefficient are those where both $|r'\rangle$ and $|r''\rangle$ are P states. Here, we can make use of the fact that we can further reduce the dipole matrix elements

$$\langle n, l, j | \hat{d} | n', l', j' \rangle = (-1)^{j'+l+1+s} \sqrt{(2j'+1)(2l+1)} \begin{Bmatrix} l & l' & 1 \\ j' & j & s \end{Bmatrix} \langle n, l | \hat{d} | n', l' \rangle, \tag{A.10}$$

allowing us to write the energy shift of the $|r, r\rangle$ pair state in terms of a sum over

the principal quantum numbers of the P states

$$\Delta E_{r,r}^{(2)} = \sum_{n',n''} \frac{|\langle n, S | \hat{d} | n', P \rangle|^2 |\langle n, S | \hat{d} | n'', P \rangle|^2}{81 (4\pi\epsilon_0 R^3)^2} \times \left[\frac{(25 - 3 \cos 2\theta)}{2E_{nS_{1/2}} - E_{n'P_{3/2}} - E_{n''P_{3/2}}} + \frac{(7 - 3 \cos 2\theta)}{2E_{nS_{1/2}} - E_{n'P_{1/2}} - E_{n''P_{1/2}}} + \left(\frac{11 + 3 \cos 2\theta}{2E_{nS_{1/2}} - E_{n'P_{1/2}} - E_{n''P_{3/2}}} + \frac{11 + 3 \cos 2\theta}{2E_{nS_{1/2}} - E_{n'P_{3/2}} - E_{n''P_{1/2}}} \right) \right]. \quad (\text{A.11})$$

For rubidium, the dominant contribution to the sum comes from the terms where $n' = n$, $n'' = n - 1$ and vice versa. Therefore, the C_6 coefficient for a pair of $|r\rangle = |n, S, 1/2, m_j\rangle$ states is given by

$$C_6 \approx \frac{2}{81 (4\pi\epsilon_0)^2} |\langle n, S | \hat{d} | n, P \rangle|^2 |\langle n, S | \hat{d} | n - 1, P \rangle|^2 \times \left[\frac{(25 - 3 \cos 2\theta)}{2E_{nS_{1/2}} - E_{nP_{3/2}} - E_{(n-1)P_{3/2}}} + \frac{(7 - 3 \cos 2\theta)}{2E_{nS_{1/2}} - E_{nP_{1/2}} - E_{(n-1)P_{1/2}}} + \left(\frac{11 + 3 \cos 2\theta}{2E_{nS_{1/2}} - E_{nP_{1/2}} - E_{(n-1)P_{3/2}}} + \frac{11 + 3 \cos 2\theta}{2E_{nS_{1/2}} - E_{nP_{3/2}} - E_{(n-1)P_{1/2}}} \right) \right]. \quad (\text{A.12})$$

Although not utilized in the experiments presented in this thesis, the $|r\rangle = |n, D, j, m_j\rangle$ states are also accessible with our setup. Calculation of C_6 coefficients for these $|r, r\rangle$ states is slightly more complicated. Firstly, the C_6 coefficient is different for each j, m_j pair. Additionally, there are a larger number of couplings that need to be accounted for in performing the calculation. As a point of comparison

to the S states, and because they are likely to be the states that would be addressed experimentally, let us focus our efforts on the stretched $|r\rangle = |n, D, 5/2, 5/2\rangle$ states. Again making use of equation A.10, we can write the second order energy shift of the $|r, r\rangle$ pair states

$$\begin{aligned} \Delta E_{r,r}^{(2)} = & \frac{1}{(4\pi\epsilon_0 R^3)^2} \sum_{n',n''} \left[\frac{|\langle n, D | \hat{d} | n', P \rangle|^2 |\langle n, D | \hat{d} | n'', P \rangle|^2}{4} \frac{9 \sin^4 \theta}{2E_{nD_{5/2}} - E_{n'P_{3/2}} - E_{n''P_{3/2}}} \right. \\ & + \frac{|\langle n, D | \hat{d} | n', F \rangle|^2 |\langle n, D | \hat{d} | n'', P \rangle|^2}{784} \left(\frac{(72 + 48 \cos 2\theta) \sin^2 \theta}{2E_{nD_{5/2}} - E_{n'F_{5/2}} - E_{n''P_{3/2}}} + \frac{305 + 180 \cos 2\theta + 75 \cos 4\theta}{2E_{nD_{5/2}} - E_{n'F_{7/2}} - E_{n''P_{3/2}}} \right) \\ & \left. + \frac{|\langle n, D | \hat{d} | n', F \rangle|^2 |\langle n, D | \hat{d} | n'', F \rangle|^2}{57624} \left(\frac{1680 + 780 \cos 2\theta - 300 \cos 4\theta}{2E_{nD_{5/2}} - E_{n'F_{7/2}} - E_{n''F_{5/2}}} + \frac{39025 - 28500 \cos 2\theta + 1875 \cos 4\theta}{2E_{nD_{5/2}} - E_{n'F_{7/2}} - E_{n''F_{7/2}}} \right) \right]. \end{aligned} \quad (\text{A.13})$$

This can be simplified by noting that the hyperfine splitting of the nF states is small for large n ($E_{nF_{5/2}} - E_{nF_{7/2}} < 1$ MHz for $n > 55$), allowing us to set $E_{nF_{7/2}} \approx E_{nF_{5/2}} = E_{nF}$

$$\begin{aligned} \Delta E_{r,r}^{(2)} \approx & \frac{1}{(4\pi\epsilon_0 R^3)^2} \sum_{n',n''} \left[\frac{|\langle n, D | \hat{d} | n', P \rangle|^2 |\langle n, D | \hat{d} | n'', P \rangle|^2}{4} \frac{9 \sin^4 \theta}{2E_{nD_{5/2}} - E_{n'P_{3/2}} - E_{n''P_{3/2}}} \right. \\ & + \frac{|\langle n, D | \hat{d} | n', F \rangle|^2 |\langle n, D | \hat{d} | n'', P \rangle|^2}{112} \frac{47 + 24 \cos 2\theta + 9 \cos 4\theta}{2E_{nD_{5/2}} - E_{n'F} - E_{n''P_{3/2}}} \\ & \left. + \frac{|\langle n, D | \hat{d} | n', F \rangle|^2 |\langle n, D | \hat{d} | n'', F \rangle|^2}{8232} \frac{5815 + 3960 \cos 2\theta + 225 \cos 4\theta}{2E_{nD_{5/2}} - E_{n'F} - E_{n''F}} \right]. \end{aligned} \quad (\text{A.14})$$

For rubidium, the dominant contribution to the sum comes from the $nDnD \rightarrow (n+1)P(n-1)F$ and $nDnD \rightarrow (n+2)P(n-2)F$ channels. Therefore, we can




Figure A.1: Angular dependence of C_6 coefficient for different $|r, r\rangle$ pair states. Values calculated using reference [61].

approximate the C_6 coefficient for a pair of $|r\rangle = |n, D, 5/2, 5/2\rangle$ states

$$\begin{aligned}
 C_6 \approx & \frac{2}{(4\pi\epsilon_0)^2} \frac{47 + 24 \cos 2\theta + 9 \cos 4\theta}{112} \left(\frac{|\langle n, D | \hat{d} | n-1, F \rangle|^2 |\langle n, D | \hat{d} | n+1, P \rangle|^2}{2E_{nD_{5/2}} - E_{(n-1)F} - E_{(n+1)P_{3/2}}} \right. \\
 & \left. + \frac{|\langle n, D | \hat{d} | n-2, F \rangle|^2 |\langle n, D | \hat{d} | n+2, P \rangle|^2}{2E_{nD_{5/2}} - E_{(n-2)F} - E_{(n+2)P_{3/2}}} \right). \tag{A.15}
 \end{aligned}$$

From equations A.12 and A.15 we can see that there is an angular dependence to the C_6 coefficient for both the pairs of nS and pairs of nD states. However, as can be seen in figure A.1, the variation is small for the nS state pair. Here the asymmetry arises due to the hyperfine splitting of the nP states, with the C_6 coefficient becoming spherically symmetric in the limit that the splitting goes to

zero. For the pair of nD states, and non-zero orbital angular momentum states in general, the angular variation is *not* the result of some small splitting, and can be large, possibly even exhibiting zeroes.

Appendix B: Calculation of Expected Coincidences for Self-Hong-Ou-Mandel Measurement




Figure B.1: Schematic for self-HOM measurement.

Here, I derive expressions for the expected behavior of a self-HOM taken with a pulsed source. The experimental setup in mind here is the one shown in figure B.1, where we have a pulsed source incident on a beamsplitter, whose outputs are sent to a long and a short delay arms. Both delay lines are directed towards a further beamsplitter, whose outputs we're going to monitor for coincidences.

In order to perform the calculations there's some assumptions we need to make and some parameters we want to define. To start, we'll assume that our source produces photons with an efficiency P_{ph} , with a period of t_{pulse} . Additionally,

we'll assume that the source is reasonably good in terms of purity, so that we can ignore high order multi-photon events and make the approximation $g^{(2)}(0) \approx \frac{2P_2}{P_1^2}$, where P_n is the probability of the source producing an n -photon pulse. We'll use the probability P_{short} to denote the probability of a photon incident on the first beamsplitter propagating to the second beamsplitter along the short arm, and similarly for P_{long} along the long arm¹. This allows us to treat the general case for the splitting ratio at the first beamsplitter, along with accounting for loss incurred by traversing the long and short arms. We'll assume that the second beamsplitter is a perfect 50:50 splitter. Finally, when we discuss coincidences, we'll assume that we're performing a sum over peaks in coincidence space, so that we're agnostic to the shape of the photons.

Let us consider the situation where the delay and the pulse spacing are matched, as was the case for the self-HOM measurements taken in chapter 6. Here, we expect to get coincidences at delay times between the detectors of $\tau = kt_{pulse}$, where k is an integer. It is, perhaps, easiest to determine the magnitude of coincidences as a function of k by considering the permutation of outcomes for pairs of photons arriving at the first beamsplitter. I'll describe photons in terms of the chronological order in which they arrive at the first beamsplitter. First, let's consider what happens if we have two photons in the same pulse, which happens with a probability $g^{(2)}(0)P_1^2/2$:

¹Note that P_{short} and P_{long} do not necessarily sum to one due to loss along the two arms.

First Photon	Second Photon	Probability	Coincidence Time
Short arm	Short arm	$\frac{P_{short}^2}{2}$	0
Long arm	Long arm	$\frac{P_{long}^2}{2}$	0
Short arm	Long arm	$\frac{P_{long}P_{short}}{2}$	$\pm t_{pulse}$
Long arm	Short arm	$\frac{P_{long}P_{short}}{2}$	$\pm t_{pulse}$

Table B.1: Outcomes for two photons in the same pulse.

Note that here I've talked about the 'first' and 'second' photon but in reality they are incident on the first beamsplitter simultaneously. Next when the two photons come from nearest neighbor pulses:

First Photon	Second Photon	Probability	Coincidence Time
Short arm	Short arm	$\frac{P_{short}^2}{4}$	$\pm t_{pulse}$
Long arm	Long arm	$\frac{P_{long}^2}{4}$	$\pm t_{pulse}$
Short arm	Long arm	$\frac{P_{long}P_{short}}{4}$	$\pm 2t_{pulse}$
Long arm	Short arm	$(1 - c)\frac{P_{long}P_{short}}{2}$	0

Table B.2: Outcomes for two photons in the adjacent pulses.

Here, I have used c to parameterize the mode overlap of the two photons, as we did in chapter 7. Then for next-nearest neighbors:

First Photon	Second Photon	Probability	Coincidence Time
Short arm	Short arm	$\frac{P_{short}^2}{4}$	$\pm 2t_{pulse}$
Long arm	Long arm	$\frac{P_{long}^2}{4}$	$\pm 2t_{pulse}$
Short arm	Long arm	$\frac{P_{long}P_{short}}{4}$	$\pm 3t_{pulse}$
Long arm	Short arm	$\frac{P_{long}P_{short}}{4}$	$\pm t_{pulse}$

Table B.3: Outcomes for two photons in the adjacent but one pulses.

We could continue on considering next-next-nearest neighbors, next-next-next-nearest

neighbors and so on. However, the table for these will look largely the same as for the next-nearest neighbors, albeit with the coincidence time increased. Collating these together we can write down the expected number of coincidences

$$\mathcal{C}(0) \propto \frac{g^{(2)}(0)P_1^2}{2} \left(\frac{P_{short}^2}{2} + \frac{P_{long}^2}{2} \right) + (1-c)P_1^2 \frac{P_{long}P_{short}}{2} \quad (\text{B.1})$$

$$\mathcal{C}(\pm t_{pulse}) \propto \frac{g^{(2)}(0)P_1^2}{2} \frac{P_{long}P_{short}}{2} + P_1^2 \left(\frac{P_{short}^2}{4} + \frac{P_{long}^2}{4} + \frac{P_{long}P_{short}}{4} \right) \quad (\text{B.2})$$

$$\mathcal{C}(|\tau| \geq 2t_{pulse}) \propto P_1^2 \left(\frac{P_{short}^2}{4} + \frac{P_{long}^2}{4} + \frac{P_{long}P_{short}}{2} \right). \quad (\text{B.3})$$

It is useful to normalize by dividing by $\mathcal{C}(|\tau| \geq 2t_{pulse})$

$$\bar{\mathcal{C}}(0) = \frac{g^{(2)}(0) (P_{short}^2 + P_{long}^2) + 2(1-c)P_{long}P_{short}}{(P_{short} + P_{long})^2} \quad (\text{B.4})$$

$$\bar{\mathcal{C}}(\pm t_{pulse}) = \frac{(g^{(2)}(0) - 1)P_{long}P_{short}}{(P_{short} + P_{long})^2} + 1. \quad (\text{B.5})$$

In the self-HOM experiment performed in chapter 6 we tried to adjust the splitting at the first beamsplitter such that the power arriving at the second beamsplitter for the two arms was approximately equal, $P_{short} \approx P_{long}$, which explains why the height of the $\pm t_{pulse}$ peaks in figure 6.13 are approximately 75% of the $\pm t_{pulse}$ peaks.

The above math, however, does not explain the asymmetry we see between the heights of the $\pm t_{pulse}$ peaks. This arises from the fact that the second beamsplitter isn't perfectly 50:50, combined with the way the coincidence delay time is defined in terms of detection times on the two SPADs, $\tau = t_A - t_B$. To see why this is the case, let's consider the case where the transmission, T , of the second beamsplitter is

slightly higher than 50%, and thus the reflection slightly lower. Looking at figure B.1 this means that a photon coming from the short arm is more likely to end up going to SPAD-A, and a photon from the long arm to SPAD-B. So we're more likely to get coincidences, if the two photons took different arms, from the short arm photon going to SPAD-A and the long arm photon going to SPAD-B than the other way around. Now let's go ahead and consider a pair of photons that are temporally separated by n pulse periods, and to see the origin of the asymmetry we only need to consider what happens when the two photons take different arms. If the first photon takes the short arm, and the second the long arm then we can get a coincidence at $\tau = (n + 1)t_{pulse}$ with probability $\propto T^2$ or at $\tau = -(n + 1)t_{pulse}$ with probability $\propto (1 - T)^2$. But if the first photon takes the long arm, and the second the short arm we can get a coincidence at $\tau = -(n - 1)t_{pulse}$ with probability $\propto T^2$ or at $\tau = (n - 1)t_{pulse}$ with probability $\propto (1 - T)^2$. If we're looking at large $\tau = kt_{pulse}$ all this doesn't matter as we get contributions from pulses separated by $k - 1$ and $k + 1$, so we get something which looks like $\propto (T^2 + (1 - T)^2) = 1$. However, for $k = 1$ the $k - 1$ term would be where the two photons came from the same pulse, which is suppressed as $g^{(2)}(0) \approx 0$ for a good source, with a similar argument able to be made for $k = -1$ with the $k + 1$ term. This leads to the slight asymmetry between the number of coincidences seen in figure 6.13, where the direction and magnitude of asymmetry depends on the SPAD, fiber and beamsplitter configuration.

Bibliography

- [1] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.
- [2] David Deutsch and Roger Penrose. Quantum theory, the Church–Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985. Royal Society.
- [3] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992. Royal Society.
- [4] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [5] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
- [6] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493–R2496, 1995.
- [7] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 452(1954):2551–2577, 1996. Royal Society.
- [8] Charles Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:175–179, 01 1984.
- [9] Artur K. Ekert. Quantum cryptography based on Bell’s theorem. *Physical Review Letters*, 67(6):661–663, 1991.

- [10] Charles H. Bennett, Gilles Brassard, and N. David Mermin. Quantum cryptography without Bell's theorem. *Physical Review Letters*, 68(5):557–559, 1992.
- [11] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69(20):2881–2884, 1992.
- [12] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70(13):1895–1899, 1993.
- [13] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [14] Juan Yin, Yuan Cao, Yu-Huai Li, Sheng-Kai Liao, Liang Zhang, Ji-Gang Ren, Wen-Qi Cai, Wei-Yue Liu, Bo Li, Hui Dai, Guang-Bing Li, Qi-Ming Lu, Yun-Hong Gong, Yu Xu, Shuang-Lin Li, Feng-Zhi Li, Ya-Yun Yin, Zi-Qing Jiang, Ming Li, Jian-Jun Jia, Ge Ren, Dong He, Yi-Lin Zhou, Xiao-Xiang Zhang, Na Wang, Xiang Chang, Zhen-Cai Zhu, Nai-Le Liu, Yu-Ao Chen, Chao-Yang Lu, Rong Shu, Cheng-Zhi Peng, Jian-Yu Wang, and Jian-Wei Pan. Satellite-based entanglement distribution over 1200 kilometers. *Science*, 356(6343):1140–1144, 2017.
- [15] Juan Yin, Yu-Huai Li, Sheng-Kai Liao, Meng Yang, Yuan Cao, Liang Zhang, Ji-Gang Ren, Wen-Qi Cai, Wei-Yue Liu, Shuang-Lin Li, Rong Shu, Yong-Mei Huang, Lei Deng, Li Li, Qiang Zhang, Nai-Le Liu, Yu-Ao Chen, Chao-Yang Lu, Xiang-Bin Wang, Feihu Xu, Jian-Yu Wang, Cheng-Zhi Peng, Artur K. Ekert, and Jian-Wei Pan. Entanglement-based secure quantum cryptography over 1,120 kilometres. *Nature*, 582(7813):501–505, 2020.

- [16] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, New York, 2000.
- [17] T. E. Northup and R. Blatt. Quantum information transfer using photons. *Nature Photonics*, 8(5):356–363, 2014.
- [18] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014.
- [19] Liang Jiang, Jacob M. Taylor, Anders S. Sørensen, and Mikhail D. Lukin. Distributed quantum computation based on small quantum registers. *Physical Review A*, 76(6):062323, 2007.
- [20] J. J. Bollinger, Wayne M. Itano, D. J. Wineland, and D. J. Heinzen. Optimal frequency measurements with maximally correlated states. *Physical Review A*, 54(6):R4649–R4652, 1996.
- [21] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. *Physical Review A*, 46(11):R6797–R6800, 1992.
- [22] Zachary Eldredge, Michael Foss-Feig, Jonathan A. Gross, Steven L. Rolston, and Alexey V. Gorshkov. Optimal and Secure Measurement Protocols for Quantum Sensor Networks. *Physical Review A*, 97(4):042337, 2018.
- [23] Shamsolah Salemian and Shahram Mohammadnejad. Analysis of polarization mode dispersion effect on quantum state decoherence in fiber-based optical quantum communication. In *Proceedings of the 11th International Conference on Telecommunications*, pages 511–516, 2011.
- [24] Hyang-Tag Lim, Kang-Hee Hong, and Yoon-Ho Kim. Effects of polarization mode dispersion on polarization-entangled photons generated via broadband pumped spontaneous parametric down-conversion. *Scientific Reports*, 6(1):25846, 2016.
- [25] Sheng-Kai Liao, Wen-Qi Cai, Wei-Yue Liu, Liang Zhang, Yang Li, Ji-Gang Ren, Juan Yin, Qi Shen, Yuan Cao, Zheng-Ping Li, Feng-Zhi Li, Xia-Wei Chen, Li-Hua Sun, Jian-Jun Jia, Jin-Cai Wu, Xiao-Jun Jiang, Jian-Feng Wang, Yong-Mei Huang, Qiang Wang, Yi-Lin Zhou, Lei Deng, Tao Xi, Lu Ma, Tai Hu, Qiang Zhang, Yu-Ao Chen, Nai-Le Liu, Xiang-Bin Wang, Zhen-Cai Zhu, Chao-Yang Lu, Rong Shu, Cheng-Zhi Peng, Jian-Yu Wang, and Jian-Wei Pan. Satellite-to-ground quantum key distribution. *Nature*, 549(7670):43–47, 2017.
- [26] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.

- [27] L.-M Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, 2001.
- [28] Nicolas Sangouard, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33–80, 2011.
- [29] L.-M. Duan and H. J. Kimble. Efficient Engineering of Multiatom Entanglement through Single-Photon Detections. *Physical Review Letters*, 90(25):253601, 2003.
- [30] Christoph Simon and William T. M. Irvine. Robust Long-Distance Entanglement and a Loophole-Free Bell Test with Ions and Photons. *Physical Review Letters*, 91(11):110405, 2003.
- [31] Bo Zhao, Markus Müller, Klemens Hammerer, and Peter Zoller. Efficient quantum repeater based on deterministic Rydberg gates. *Physical Review A*, 81(5):052329, 2010.
- [32] Neal Solmeyer, Xiao Li, and Qudsia Quraishi. High teleportation rates using cold-atom-ensemble-based quantum repeaters with rydberg blockade. *Physical Review A*, 93(4):042301, 2016.
- [33] Nicolas Sangouard, Christoph Simon, Jiří Minář, Hugo Zbinden, Hugues de Riedmatten, and Nicolas Gisin. Long-distance entanglement distribution with single-photon sources. *Physical Review A*, 76(5):050301, 2007.
- [34] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance. High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network. *Physical Review Letters*, 124(11):110501, 2020.
- [35] R. Stockill, M. J. Stanley, L. Huthmacher, E. Clarke, M. Hugues, A. J. Miller, C. Matthiesen, C. Le Gall, and M. Atatüre. Phase-Tuned Entangled State Generation between Distant Spin Qubits. *Physical Review Letters*, 119(1):010503, 2017.
- [36] Nicolas Maring, Pau Farrera, Kutlu Kutluer, Margherita Mazzera, Georg Heinze, and Hugues de Riedmatten. Photonic quantum state transfer between a cold atomic gas and a crystal. *Nature*, 551(7681):485–488, 2017.
- [37] M. Lettner, M. Mücke, S. Riedl, C. Vo, C. Hahn, S. Baur, J. Bochmann, S. Ritter, S. Dürr, and G. Rempe. Remote Entanglement between a Single Atom and a Bose-Einstein Condensate. *Physical Review Letters*, 106(21):210503, 2011.

- [38] H. M. Meyer, R. Stockill, M. Steiner, C. Le Gall, C. Matthiesen, E. Clarke, A. Ludwig, J. Reichel, M. Atatüre, and M. Köhl. Direct Photonic Coupling of a Semiconductor Quantum Dot and a Trapped Ion. *Physical Review Letters*, 114(12):123001, 2015.
- [39] D. P. Ornelas-Huerta, A. N. Craddock, E. A. Goldschmidt, E. A. Goldschmidt, A. J. Hachtel, Y. Wang, P. Bienias, P. Bienias, A. V. Gorshkov, A. V. Gorshkov, S. L. Rolston, and J. V. Porto. On-demand indistinguishable single photons from an efficient and pure source based on a Rydberg ensemble. *Optica*, 7(7):813–819, 2020.
- [40] A. N. Craddock, J. Hannegan, D. P. Ornelas-Huerta, J. D. Siverns, A. J. Hachtel, E. A. Goldschmidt, J. V. Porto, Q. Quraishi, and S. L. Rolston. Quantum Interference between Photons from an Atomic Ensemble and a Remote Atomic Ion. *Physical Review Letters*, 123(21):213601, 2019.
- [41] Thibault Peyronel, Ofer Firstenberg, Qi-Yu Liang, Sebastian Hofferberth, Alexey V Gorshkov, Thomas Pohl, Mikhail D Lukin, and Vladan Vuletić. Quantum nonlinear optics with single photons enabled by strongly interacting atoms. *Nature*, 488(7409):57–60, 2012.
- [42] Ofer Firstenberg, Thibault Peyronel, Qi-yu Liang, Alexey V. Gorshkov, Mikhail D. Lukin, and Vladan Vuletić. Attractive photons in a quantum nonlinear medium. *Nature*, 502(7469):71–75, 2013.
- [43] Qi-Yu Liang, Aditya V. Venkatramani, Sergio H. Cantu, Travis L. Nicholson, Michael J. Gullans, Alexey V. Gorshkov, Jeff D. Thompson, Cheng Chin, Mikhail D. Lukin, and Vladan Vuletić. Observation of three-photon bound states in a quantum nonlinear medium. *Science*, 359(6377):783–786, 2018.
- [44] Benoît Vermersch, Tomás Ramos, Philipp Hauke, and Peter Zoller. Implementation of chiral quantum optics with Rydberg and trapped-ion setups. *Physical Review A*, 93(6):063830, 2016.
- [45] Przemyslaw Bienias, Michael J. Gullans, Marcin Kalinowski, Alexander N. Craddock, Dalia P. Ornelas-Huerta, Steven L. Rolston, J. V. Porto, and Alexey V. Gorshkov. Exotic photonic molecules via Lennard-Jones-like potentials. arXiv: 2003.07864.
- [46] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletic, and Mikhail D. Lukin. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 2017.
- [47] Henning Labuhn, Daniel Barredo, Sylvain Ravets, Sylvain de Léséleuc, Tommaso Macrì, Thierry Lahaye, and Antoine Browaeys. Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models. *Nature*, 534(7609):667–670, 2016.

- [48] Antoine Browaeys, Daniel Barredo, and Thierry Lahaye. Experimental investigations of dipole–dipole interactions between a few Rydberg atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(15):152001, 2016.
- [49] Hyosub Kim, YeJe Park, Kyungtae Kim, H.-S. Sim, and Jaewook Ahn. Detailed Balance of Thermalization Dynamics in Rydberg-Atom Quantum Simulators. *Physical Review Letters*, 120(18):180502, 2018.
- [50] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman. Demonstration of a Neutral Atom Controlled-NOT Quantum Gate. *Physical Review Letters*, 104(1):010503, 2010.
- [51] Daniel Tiarks, Steffen Schmidt-Eberle, Thomas Stolz, Gerhard Rempe, and Stephan Dürr. A photon–photon quantum gate based on Rydberg interactions. *Nature Physics*, 15(2):124–126, 2019.
- [52] David Petrosyan, Mark Saffman, and Klaus Mølmer. Grover search algorithm with Rydberg-blockaded atoms: quantum Monte Carlo simulations. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(9):094004, 2016.
- [53] Jeff D Thompson, Travis L Nicholson, Qi-Yu Liang, Sergio H Cantu, Aditya V Venkatramani, Soonwon Choi, Ilya A Fedorov, Daniel Viscor, Thomas Pohl, Mikhail D Lukin, and Vladan Vuletić. Symmetry-protected collisions between strongly interacting photons. *Nature*, 542:206–209, 2017.
- [54] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast Quantum Gates for Neutral Atoms. *Physical Review Letters*, 85(10):2208–2211, 2000.
- [55] Christopher L. Holloway, Matthew T. Simons, Abdulaziz H. Haddab, Carl J. Williams, and Maxwell W. Holloway. A “real-time” guitar recording using Rydberg atoms and electromagnetically induced transparency: Quantum physics meets music. *AIP Advances*, 9(6):065110, 2019.
- [56] David H. Meyer, Kevin C. Cox, Fredrik K. Fatemi, and Paul D. Kunz. Digital communication with Rydberg atoms and amplitude-modulated microwave fields. *Applied Physics Letters*, 112(21), 2018.
- [57] O. Firstenberg, C. S. Adams, and S. Hofferberth. Nonlinear quantum optics mediated by Rydberg interactions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(15):152003, 2016.
- [58] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Reviews of Modern Physics*, 82(3):2313–2363, 2010.
- [59] Thomas F Gallagher. *Rydberg atoms*. Cambridge University Press, 1994.
- [60] Daniel A Steck. Quantum and Atom Optics. Available at <http://steck.us/teaching> (revision 0.12.6, 23 April 2019).

- [61] N. Šibalić, J. D. Pritchard, C. S. Adams, and K. J. Weatherill. ARC: An open-source library for calculating properties of alkali Rydberg atoms. *Computer Physics Communications*, 220:319–331, 2017.
- [62] Klemens Hammerer, Anders S. Sørensen, and Eugene S. Polzik. Quantum interface between light and atomic ensembles. *Reviews of Modern Physics*, 82(2):1041–1093, 2010.
- [63] Nina Stiesdal, Jan Kumlin, Kevin Kleinbeck, Philipp Lunt, Christoph Braun, Asaf Paris-Mandoki, Christoph Tresp, Hans Peter Büchler, and Sebastian Hofferberth. Observation of Three-Body Correlations for Photons Coupled to a Rydberg Superatom. *Physical Review Letters*, 121(10):103601, 2018.
- [64] Y. O. Dudin and A. Kuzmich. Strongly interacting Rydberg excitations of a cold atomic gas. *Science*, 336(6083):887–889, 2012.
- [65] Robert Löw, Hendrik Weimer, Johannes Nipper, Jonathan B. Balewski, Björn Butscher, Hans Peter Büchler, and Tilman Pfau. An experimental and theoretical guide to strongly interacting Rydberg gases. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(11):113001, 2012.
- [66] Daniel A Steck. Rubidium 87 D Line Data. Available online at <http://steck.us/alkalidata> (revision 2.2.1, 21 November 2019).
- [67] M. Fleischhauer and M. D. Lukin. Dark-State Polaritons in Electromagnetically Induced Transparency. *Physical Review Letters*, 84(22):5094–5097, 2000.
- [68] Przemyslaw Bienias. *Few-body quantum physics with strongly interacting Rydberg polaritons*. PhD thesis.
- [69] Ya-Fen Hsiao, Pin-Ju Tsai, Hung-Shiue Chen, Sheng-Xiang Lin, Chih-Chiao Hung, Chih-Hsi Lee, Yi-Hsin Chen, Yong-Fan Chen, Ite A. Yu, and Ying-Cheng Chen. Highly efficient coherent optical memory based on electromagnetically induced transparency. *Physical Review Letters*, 120(18):183602, 2018.
- [70] Y. O. Dudin, L. Li, and A. Kuzmich. Light storage on the time scale of a minute. *Physical Review A*, 87(3):031801, 2013.
- [71] R. H. Dicke. Coherence in Spontaneous Radiation Processes. *Physical Review*, 93(1):99–110, 1954.
- [72] M. Saffman and T. G. Walker. Creating single-atom and single-photon sources from entangled atomic ensembles. *Physical Review A*, 66(6):065403, 2002.
- [73] Alexey V. Gorshkov, Axel André, Mikhail D. Lukin, and Anders S. Sørensen. Photon storage in λ -type optically dense atomic media. II. Free-space model. *Physical Review A*, 76(3), 2007.

- [74] Sergio H. Cantu, Aditya V. Venkatramani, Wenchao Xu, Leo Zhou, Brana Jelenković, Mikhail D. Lukin, and Vladan Vuletić. Repulsive photons in a quantum nonlinear medium. *Nature Physics*, 2020.
- [75] C. Klempt, T. van Zoest, T. Henninger, O. Topic, E. Rasel, W. Ertmer, and J. Arlt. Ultraviolet light-induced atom desorption for large rubidium and potassium magneto-optical traps. *Physical Review A*, 73(1):013410, 2006.
- [76] A. Gozzini, F. Mango, J. H. Xu, G. Alzetta, F. Maccarrone, and R. A. Bernheim. Light-induced ejection of alkali atoms in polysiloxane coated cells. *Il Nuovo Cimento D*, 15(5):709–722, 1993.
- [77] Christoph Tresp. *Rydberg polaritons and Rydberg superatoms - Novel tools for quantum nonlinear optics*. PhD thesis, Universität Stuttgart, 2017.
- [78] Daniel L Campbell. *Engineered Potentials in Ultracold Bose-Einstein Condensates*. PhD thesis, University of Maryland, College Park, 2015.
- [79] Karina Jimenez-Garcia. *Artificial Gauge Fields for Ultracold Neutral Atoms*. PhD thesis, University of Maryland, College Park, 2012.
- [80] Jürgen Appel, Andrew MacRae, and A. I. Lvovsky. A versatile digital GHz phase lock for external cavity diode lasers. *Measurement Science and Technology*, 20(5):055302, 2009.
- [81] Sara Rosi, Alessia Burchianti, Stefano Conclave, Devang S. Naik, Giacomo Roati, Chiara Fort, and Francesco Minardi. λ -enhanced grey molasses on the D 2 transition of Rubidium-87 atoms. *Scientific Reports*, 8(1):1–9, 2018.
- [82] L. Li, Y. O. Dudin, and A. Kuzmich. Entanglement between light and an optical atomic excitation. *Nature*, 498(7455):466–469, 2013.
- [83] E D Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics*, 69(1):79, 2001.
- [84] Benjamin Reschovsky. *Studies of Ultracold Strontium Gases*. PhD thesis, University of Maryland, College Park, 2017.
- [85] Sylvain de Léséleuc, Daniel Barredo, Vincent Lienhard, Antoine Browaeys, and Thierry Lahaye. Analysis of imperfections in the coherent optical excitation of single atoms to Rydberg states. *Physical Review A*, 97(5):053803, 2018.
- [86] Harry Levine, Alexander Keesling, Ahmed Omran, Hannes Bernien, Sylvain Schwartz, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-Fidelity Control and Entanglement of Rydberg-Atom Qubits. *Physical Review Letters*, 121(12):123603, 2018.

- [87] Stetson Roof, Kasie Kemp, Mark Havey, I. M. Sokolov, and D. V. Kupriyanov. Microscopic lensing by a dense, cold atomic sample. *Optics Letters*, 40(7):1137–1140, 2015.
- [88] R. Hanbury Brown and R. Q. Twiss. Correlation between photons in two coherent beams of light. *Nature*, 177(4497):27–29, 1956.
- [89] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. *Physical Review Letters*, 59(18):2044–2046, 1987.
- [90] Rodney Loudon. *The Quantum Theory of Light, 3rd Edition*. Oxford University Press, 2000.
- [91] Przemyslaw Bienias, James Douglas, Asaf Paris-Mandoki, Paraj Titum, Ivan Mirgorodskiy, Christoph Tresp, Emil Zeuthen, Michael J. Gullans, Marco Manzoni, Sebastian Hofferberth, Darrick Chang, and Alexey V. Gorshkov. Photon propagation through dissipative rydberg media at large input rates. *Physical Review Research*, 2:033049, 2020.
- [92] Martin J Stevens. Photon Statistics, Measurements, and Measurements Tools. In *Photon Statistics, Measurements, and Measurements Tools*, page 35. Academic Press, 2013.
- [93] Maurice Herlihy and Nir Shavit. *The art of multiprocessor programming*. Morgan Kaufmann, 2008.
- [94] Leonardo Dagum and Ramesh Menon. OpenMP: An Industry-Standard API for Shared-Memory Programming. *IEEE Computation Science & Engineering*, 5(1):46–55, 1998.
- [95] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. Scalable parallel programming with cuda. *Queue*, 6(2):40–53, 2008.
- [96] Dirk Merkel. Docker: lightweight Linux containers for consistent development and deployment. *Linux Journal*, 2014(239):2:2, 2014.
- [97] Jacques Carolan, Christopher Harrold, Chris Sparrow, Enrique Martín-López, Nicholas J. Russell, Joshua W. Silverstone, Peter J. Shadbolt, Nobuyuki Matsuda, Manabu Oguma, Mikitaka Itoh, Graham D. Marshall, Mark G. Thompson, Jonathan C. F. Matthews, Toshikazu Hashimoto, Jeremy L. O’Brien, and Anthony Laing. Universal linear optics. *Science*, 349(6249):711–716, 2015.
- [98] Hui Wang, Jian Qin, Xing Ding, Ming-Cheng Chen, Si Chen, Xiang You, Yu-Ming He, Xiao Jiang, L. You, Z. Wang, C. Schneider, Jelmer J. Renema, Sven Höfling, Chao-Yang Lu, and Jian-Wei Pan. Boson Sampling with 20 Input Photons and a 60-Mode Interferometer in a 10^{14} -Dimensional Hilbert Space. *Physical Review Letters*, 123(25):250503, 2019.

- [99] Sergei Slussarenko, Morgan M. Weston, Helen M. Chrzanowski, Lynden K. Shalm, Varun B. Verma, Sae Woo Nam, and Geoff J. Pryde. Unconditional violation of the shot-noise limit in photonic quantum metrology. *Nature Photonics*, 11(11):700–703, 2017.
- [100] Hui Wang, Yu-Ming He, T.-H. Chung, Hai Hu, Ying Yu, Si Chen, Xing Ding, M.-C. Chen, Jian Qin, Xiaoxia Yang, Run-Ze Liu, Z.-C. Duan, J.-P. Li, S. Gerhardt, K. Winkler, J. Jurkat, Lin-Jun Wang, Niels Gregersen, Yong-Heng Huo, Qing Dai, Siyuan Yu, Sven Höfling, Chao-Yang Lu, and Jian-Wei Pan. Towards optimal single-photon sources from polarized microcavities. *Nature Photonics*, 13(11):770–775, 2019.
- [101] N. Somaschi, V. Giesz, L. De Santis, J. C. Loredó, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Antón, J. Demory, C. Gómez, I. Sagnes, N. D. Lanzillotti-Kimura, A. Lemaître, A. Auffeves, A. G. White, L. Lanco, and P. Senellart. Near-optimal single-photon sources in the solid state. *Nature Photonics*, 10(5):340–345, 2016.
- [102] Juan C. Loredó, Nor A. Zakaria, Niccolò Somaschi, Carlos Anton, Lorenzo de Santis, Valerian Giesz, Thomas Grange, Matthew A. Broome, Olivier Gazzano, Guillaume Coppola, Isabelle Sagnes, Aristide Lemaître, Alexia Auffeves, Pascale Senellart, Marcelo P. Almeida, and Andrew G. White. Scalable performance in solid-state single-photon sources. *Optica*, 3(4):433–440, 2016.
- [103] Hui Wang, Yu He, Yu-Huai Li, Zu-En Su, Bo Li, He-Liang Huang, Xing Ding, Ming-Cheng Chen, Chang Liu, Jian Qin, Jin-Peng Li, Yu-Ming He, Christian Schneider, Martin Kamp, Cheng-Zhi Peng, Sven Höfling, Chao-Yang Lu, and Jian-Wei Pan. High-efficiency multiphoton boson sampling. *Nature Photonics*, 11(6):361–365, 2017.
- [104] Gabija Kiršanskė, Henri Thyrrestrup, Raphaël S. Daveau, Chris L. Dreeßen, Tommaso Pregnolato, Leonardo Midolo, Petru Tighineanu, Alisa Javadi, Søren Stobbe, Rüdiger Schott, Arne Ludwig, Andreas D. Wieck, Suk In Park, Jin D. Song, Andreas V. Kuhlmann, Immo Söllner, Matthias C. Löbl, Richard J. Warburton, and Peter Lodahl. Indistinguishable and efficient single photons from a quantum dot in a planar nanobeam waveguide. *Physical Review B*, 96(16):165306, 2017.
- [105] H el ene Ollivier, Ilse Maillette de Buy Wenniger, Sarah Thomas, Stephen C. Wein, Abdelmounaim Harouri, Guillaume Coppola, Paul Hilaire, Cl ement Millet, Aristide Lemaître, Isabelle Sagnes, Olivier Krebs, Lo ic Lanco, Juan C. Lored o, Carlos Ant on, Niccol o Somaschi, and Pascale Senellart. Reproducibility of High-Performance Quantum Dot Single-Photon Sources. *ACS Photonics*, 7(4):1050–1059, 2020.
- [106] Xi-Lin Wang, Luo-Kan Chen, W. Li, H.-L. Huang, C. Liu, C. Chen, Y.-H. Luo, Z.-E. Su, D. Wu, Z.-D. Li, H. Lu, Y. Hu, X. Jiang, C.-Z. Peng, L. Li,

- N.-L. Liu, Yu-Ao Chen, Chao-Yang Lu, and Jian-Wei Pan. Experimental Ten-Photon Entanglement. *Physical Review Letters*, 117(21):210502, 2016.
- [107] Han-Sen Zhong, Yuan Li, Wei Li, Li-Chao Peng, Zu-En Su, Yi Hu, Yu-Ming He, Xing Ding, Weijun Zhang, Hao Li, Lu Zhang, Zhen Wang, Lixing You, Xi-Lin Wang, Xiao Jiang, Li Li, Yu-Ao Chen, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan. 12-Photon Entanglement and Scalable Scattershot Boson Sampling with Optimal Entangled-Photon Pairs from Parametric Down-Conversion. *Physical Review Letters*, 121(25):250505, 2018.
- [108] C. Xiong, X. Zhang, Z. Liu, M. J. Collins, A. Mahendra, L. G. Helt, M. J. Steel, D.-Y. Choi, C. J. Chae, P. H. W. Leong, and B. J. Eggleton. Active temporal multiplexing of indistinguishable heralded single photons. *Nature Communications*, 7(1):10853, 2016.
- [109] F. Kaneda and P. G. Kwiat. High-efficiency single-photon generation via large-scale active time multiplexing. *Science Advances*, 5(10):eaaw8586, 2019.
- [110] Wenjamin Rosenfeld, Daniel Burchardt, Robert Garthoff, Kai Redeker, Norbert Ortengel, Markus Rau, and Harald Weinfurter. Event-Ready Bell Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes. *Physical Review Letters*, 119(1):010402, 2017.
- [111] P. Maunz, D. L. Moehring, S. Olmschenk, K. C. Younge, D. N. Matsukevich, and C. Monroe. Quantum interference of photon pairs from two remote trapped atomic ions. *Nature Physics*, 3(8):538–541, 2007.
- [112] Martin Mücke, Joerg Bochmann, Carolin Hahn, Andreas Neuzner, Christian Nölleke, Andreas Reiserer, Gerhard Rempe, and Stephan Ritter. Generation of single photons from an atom-cavity system. *Physical Review A*, 87(6):063805, 2013.
- [113] T. Wilk, S. C. Webster, H. P. Specht, G. Rempe, and A. Kuhn. Polarization-Controlled Single Photons. *Physical Review Letters*, 98(6):063601, 2007.
- [114] Peter B. R. Nisbet-Jones, Jerome Dilley, Daniel Ljunggren, and Axel Kuhn. Highly efficient source for indistinguishable single photons of controlled shape. *New Journal of Physics*, 13(10):103036, 2011.
- [115] James K. Thompson, Jonathan Simon, Huanqian Loh, and Vladan Vuletić. A High-Brightness Source of Narrowband, Identical-Photon Pairs. *Science*, 313(5783):74–77, 2006.
- [116] B. J. DeSalvo, J. A. Aman, C. Gaul, T. Pohl, S. Yoshida, J. Burgdörfer, K. R. A. Hazzard, F. B. Dunning, and T. C. Killian. Rydberg-blockade effects in Autler-Townes spectra of ultracold strontium. *Physical Review A*, 93(2):022709, 2016.

- [117] E A Goldschmidt, T Boulier, R C Brown, S B Koller, J T Young, A V Gorshkov, S L Rolston, and J V Porto. Anomalous Broadening in Driven Dissipative Rydberg Systems. *Physical Review Letters*, 116(11), 2016.
- [118] D. P. Sadler, E. M. Bridge, D. Boddy, A. D. Bounds, N. C. Keegan, G. Lochhead, M. P. A. Jones, and B. Olmos. Radiation trapping in a dense cold Rydberg gas. *Physical Review A*, 95(1):013839, 2017.
- [119] J. A. Aman, B. J. DeSalvo, F. B. Dunning, T. C. Killian, S. Yoshida, and J. Burgdörfer. Trap losses induced by near-resonant Rydberg dressing of cold atomic gases. *Physical Review A*, 93(4):043425, 2016.
- [120] Michael Schlagmüller, Tara Cubel Liebisch, Felix Engel, Kathrin S. Kleinbach, Fabian Böttcher, Udo Hermann, Karl M. Westphal, Anita Gaj, Robert Löw, Sebastian Hofferberth, Tilman Pfau, Jesús Pérez-Ríos, and Chris H. Greene. Ultracold Chemical Reactions of a Single Rydberg Atom in a Dense Gas. *Physical Review X*, 6(3):031020, 2016.
- [121] T. Boulier, E. Magnan, C. Bracamontes, J. Maslek, E. A. Goldschmidt, J. T. Young, A. V. Gorshkov, S. L. Rolston, and J. V. Porto. Spontaneous avalanche dephasing in large Rydberg ensembles. *Physical Review A*, 96(5):053409, 2017.
- [122] Jeremy T. Young, Thomas Boulier, Eric Magnan, Elizabeth A. Goldschmidt, Ryan M. Wilson, Steven L. Rolston, James V. Porto, and Alexey V. Gorshkov. Dissipation-induced dipole blockade and antiblockade in driven Rydberg systems. *Physical Review A*, 97(2):023424, 2018.
- [123] J. Lampen, H. Nguyen, L. Li, P. R. Berman, and A. Kuzmich. Long-lived coherence between ground and Rydberg levels in a magic-wavelength lattice. *Physical Review A*, 98(3):1–10, 2018.
- [124] Steffen Schmidt-Eberle, Thomas Stolz, Gerhard Rempe, and Stephan Dürr. Dark-time decay of the retrieval efficiency of light stored as a Rydberg excitation in a noninteracting ultracold gas. *Physical Review A*, 101(1):013421, 2020.
- [125] Matthew Ebert, Alexander Gill, Michael Gibbons, Xianli Zhang, Mark Saffman, and Thad G. Walker. Atomic Fock State Preparation Using Rydberg Blockade. *Physical Review Letters*, 112(4):043602, 2014.
- [126] C. K. Law and J. H. Eberly. Arbitrary Control of a Quantum Electromagnetic Field. *Physical Review Letters*, 76(7):1055–1058, 1996.
- [127] T. Legero, T. Wilk, A. Kuhn, and G. Rempe. Time-resolved two-photon quantum interference. *Applied Physics B*, 77(8):797–802, 2003.
- [128] A. M. Dyckovsky and S. Olmschenk. Analysis of photon-mediated entanglement between distinguishable matter qubits. *Physical Review A*, 85(5):052322, 2012.

- [129] H Gorniaczyk, C Tresp, J Schmidt, H Fedder, and S Hofferberth. Single-Photon Transistor Mediated by Interstate Rydberg Interactions. *Physical Review Letters*, 113(5):053601, 2014.
- [130] Sylvain De Léséleuc, Sebastian Weber, Vincent Lienhard, Daniel Barredo, Hans Peter Büchler, Thierry Lahaye, and Antoine Browaeys. Accurate Mapping of Multilevel Rydberg Atoms on Interacting Spin- 1 /2 Particles for the Quantum Simulation of Ising Models. *Physical Review Letters*, 120(11), 2018.
- [131] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe. Complete 3-Qubit Grover search on a programmable quantum computer. *Nature Communications*, 8(1):1918, 2017.
- [132] H. Häffner, C. F. Roos, and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, 2008.
- [133] R. Blatt and C. F. Roos. Quantum simulations with trapped ions. *Nature Physics*, 8(4):277–284, 2012.
- [134] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature*, 551(7682):601–604, 2017.
- [135] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-Lived Qubit Memory Using Atomic Ions. *Physical Review Letters*, 95(6):060502, 2005.
- [136] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit. *Physical Review Letters*, 113(22):220501, 2014.
- [137] D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. Monroe. Modular entanglement of atomic qubits using photons and phonons. *Nature Physics*, 11(1):37–42, 2015.
- [138] Matthias Bock, Pascal Eich, Stephan Kucera, Matthias Kreis, Andreas Lenhard, Christoph Becher, and Jürgen Eschner. High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion. *Nature Communications*, 9(1):1998, 2018.
- [139] James D. Siverns, John Hannegan, and Qudsia Quraishi. Neutral-Atom Wavelength-Compatible 780 nm Single Photons from a Trapped Ion via Quantum Frequency Conversion. *Physical Review Applied*, 11(1):014044, 2019.

- [140] James D. Siverns and Qudsia Quraishi. Ion trap architectures and new directions. *Quantum Information Processing*, 16(12):314, 2017.
- [141] Prem Kumar. Quantum frequency conversion. *Optics Letters*, 15(24):1476–1478, 1990.
- [142] J. D. Siverns, J. Hannegan, and Q. Quraishi. Demonstration of slow light in rubidium vapor using single photons from a trapped ion. *Science Advances*, 5(10):eaav4651, 2019.
- [143] David Eger, M. Arbore, M. Fejer, and M. Bortz. High intensity illumination effects in LiNbO₃ and KTiOPO₄ waveguides. *Journal of Applied Physics*, 82:998–1005, 1997.
- [144] O. Kashin, M. Homann, V. Matusevich, F. Setzpfandt, T. Pertsch, and R. Kowarschik. Change of the refractive index in PPLN waveguides due to the photorefractive effect. *Applied Physics B*, 104(3):547, 2011.
- [145] Takumi Fujiwara, Xiaofan Cao, Ramakant Srivastava, and Ramu V. Ramaswamy. Photorefractive effect in annealed proton-exchanged LiNbO₃ waveguides. *Applied Physics Letters*, 61(7):743–745, 1992.
- [146] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman. Observation of Rydberg blockade between two atoms. *Nature Physics*, 5(2):110–114, 2009.
- [147] Emil Zeuthen, Michael J. Gullans, Mohammad F. Maghrebi, and Alexey V. Gorshkov. Correlated Photon Dynamics in Dissipative Rydberg Media. *Physical Review Letters*, 119(4):043602, 2017.
- [148] M. Aspelmeyer, T. Jennewein, M. Pfennigbauer, W.R. Leeb, and A. Zeilinger. Long-distance quantum communication with entangled photons using satellites. *IEEE Journal of Selected Topics in Quantum Electronics*, 9(6):1541–1551, 2003.
- [149] Alexey V Gorshkov, Axel André, Michael Fleischhauer, Anders S Sørensen, and Mikhail D Lukin. Universal approach to optimal photon storage in atomic media. *Physical Review Letters*, 98(12), 2007.
- [150] Philipp Müller, Tristan Tentrup, Marc Bienert, Giovanna Morigi, and Jürgen Eschner. Spectral properties of single photons from quantum emitters. *Physical Review A*, 96(2):023861, 2017.
- [151] C. Crocker, M. Lichtman, K. Sosnova, A. Carter, S. Scarano, and C. Monroe. High purity single photons entangled with an atomic qubit. *Optics Express*, 27(20):28143–28149, 2019.

- [152] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe. Entanglement of single-atom quantum bits at a distance. *Nature*, 449(7158):68–71, 2007.
- [153] D. N. Matsukevich, P. Maunz, D. L. Moehring, S. Olmschenk, and C. Monroe. Bell Inequality Violation with Two Remote Atomic Qubits. *Physical Review Letters*, 100(15):150404, 2008.
- [154] Yunfei Wang, Jianfeng Li, Shanchao Zhang, Keyu Su, Yiru Zhou, Kaiyu Liao, Shengwang Du, Hui Yan, and Shi-Liang Zhu. Efficient quantum memory for single-photon polarization qubits. *Nature Photonics*, 13(5):346–351, 2019.