Single-Photon Generation via Four-Wave Mixing in a Thermal Rubidium Vapour at a High Magnetic Field

RENJU SATISH MATHEW

Single-Photon Generation via Four-Wave Mixing in a Thermal Rubidium Vapour at a High Magnetic Field

Renju S. Mathew

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Department of Physics
Durham University

21st February 2021

Single-Photon Generation via Four-Wave Mixing in a Thermal Rubidium Vapour at a High Magnetic Field

Renju S. Mathew

Abstract

We present investigations of the nonlinear optical phenomenon of four-wave mixing (4WM) in a thermal vapour at a high magnetic field. The $5S_{1/2}$, $5P_{1/2}$, $5P_{3/2}$, and $5D_{3/2}$ terms in rubidium-87 are used. In a 0.62 T field, the optical transitions become separated by more than their Doppler width, effectively isolating a four-level system. Spontaneous 4WM in this system is used to generate herald-signal photon pairs and to realise a heralded single-photon source, as demonstrated by a Hanbury Brown–Twiss (HBT) experiment giving a $g^{(2)}(0)$ value of 0.35 ± 0.02. We show that increased correlation between herald and signal leads to a $g^{(2)}(0)$ value closer to zero, and we investigate the parameters that increase this correlation. We additionally characterise the parameters in the seeded 4WM system and show that working in the hyperfine Paschen–Back (HPB) regime leads to a good agreement between theory and experiment. The splitting of the seeded 4WM lineshape is used to extract an excited state dipole matrix element. We also study the background photons in our system, showing that they are partly a result of a collisional transfer process. We build a bespoke etalon lens filter to aid in filtering out these photons and show that the use of this filter improves the correlation between herald-signal photon pairs. In order to perform HBT experiments that take over twenty-four hours to run, we devise a novel laser locking scheme, named STROLL, that simultaneously stabilises the frequencies of two lasers to a two-photon transition over that time period. Difficulty with optimising the alignment of photons into a fibre leads us to also develop an innovative method of using machine learning for automatically aligning laser beams; we implement this method on a physical device, which we have named the Pi Auto-aligner.

Declaration

I confirm that no part of the material offered has previously been submitted by myself for a degree in this or any other University. Where material has been generated through joint work, the work of others has been indicated.

Renju S. Mathew
Durham, 21st February 2021

For my parents

Contents

Abstract i

Declaration ii

Dedication iii

Acknowledgements viii

List of Figures x

1 Introduction 1

1.1 A brief history of quantum theory, leading to quantum optics 3

1.2 Thesis summary 5

1.3 Publications 7

1.4 Contributions 7

2 Theory: Atom-light interactions 9

2.1 Time evolution of a two-level system 9

2.2 Three-level system 16

2.3 Four-level system 18

2.4 Doppler broadening due to atomic motion 20

2.5 The atomic structure of rubidium 22

2.6 Application of a high magnetic field: the hyperfine Paschen–Back (HPB) regime 24

2.7 Conclusion 29

3 Experimental techniques and apparatus **30**
3.1 Setting up a tapered amplifier . . . 32
3.2 A lens etalon spectral filter . . . 37
3.3 Laser frequency calibration . . . 40
3.4 Conclusion . . . 40

4 Machine learning in quantum optics experiments **42**
4.1 Introduction . . . 42
4.2 Machine learning (ML) . . . 44
4.3 Implementation . . . 46
4.4 Using ML to maximise fibre coupling efficiency . . . 55
4.5 Using ML to maximise photon count rate . . . 59
4.6 Discussion . . . 63
4.7 Conclusion . . . 64

5 Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime **65**
5.1 Introduction . . . 66
5.2 Concept . . . 69
5.3 Experimental demonstration . . . 70
5.4 Results . . . 74
5.5 Conclusion . . . 78

6 Spectral Redistribution **79**
6.1 Introduction . . . 79
6.2 Experimental set-up . . . 83
6.3 Results & Discussion . . . 85
6.4 Conclusion . . . 97

7 Seeded four-wave mixing **98**
7.1 Introduction . . . 99
7.2 Theory . . . 102
7.3 Experimental set-up . . . 107
7.3.1 Initial alignment . . . 110

7.3.2 Realignment . . . 111
7.3.3 Detection . . . 112
7.4 Results . . . 114
7.4.1 Dependence on atomic number density . . . 114
7.4.2 Dependence on seed power . . . 114
7.4.3 Dependence on pump power . . . 116
7.4.4 Dependence on coupling detuning . . . 117
7.4.5 Dependence on coupling power . . . 120
7.5 Extracting a dipole matrix element . . . 123
7.5.1 Calculating $|\langle 5P_{3/2},\ m_J = 3/2||er||5D_{3/2},\ m_{J'} = 1/2\rangle|$ . . . 125
7.5.2 Analysing experimental data . . . 127
7.5.3 Estimating the error . . . 128
7.5.4 Experimental determination of the excited state dipole matrix element . . . 129
7.6 Conclusion . . . 133

8 Spontaneous four-wave mixing leading to a heralded single-photon source 134
8.1 Introduction . . . 135
8.2 Theory . . . 137
8.3 Experimental set-up . . . 139
8.4 Normalised cross-correlation, $g^{(2)}_{\mathrm{h,s}}(\tau)$ . . . 140
8.5 Simulating photon arrivals . . . 142
8.6 The Hanbury Brown–Twiss (HBT) experiment . . . 146
8.6.1 Obtaining a $g^{(2)}(\tau)$ measurement from experimental data . . . 146
8.6.2 Relationship between $g^{(2)}_{\mathrm{h,s}}$ and $g^{(2)}$ . . . 151
8.7 Towards a more ideal heralded single-photon source . . . 152
8.7.1 Dependence on temperature . . . 153
8.7.2 Dependence on pump detuning . . . 155
8.7.3 Dependence on pump power . . . 157
8.7.4 Dependence on coupling power . . . 158
8.7.5 Deepening the $g^{(2)}$ dip using an etalon lens filter . . . 159
8.8 Discussion . . . 162

8.9 Conclusion 164

9 Summary and Outlook **165**

A Full Optical Table Layout **169**

B Python Code **171**

Bibliography **172**

Acknowledgements

I would like to begin by thanking Ifan Hughes; you are an inspiring lecturer, teacher, and supervisor. I stayed in Durham to do my PhD principally so that I could work under your supervision. You have taught, mentored, and supported me from when I was an undergraduate, and I am deeply grateful for all that you have done.

I would also like to thank the following people for their help and support throughout my time at Durham:

Dani, your detailed comments on my thesis were invaluable. You and Liz are two of the best human beings that I know. Liz, Australia's gain was very much England's loss. Clare, thank you for being my main link to the group during the COVID crisis when I was relatively isolated, and for all the helpful conversations about this thesis. Francisco, you arrived just in time for us to be able to start figuring out the STROLL lock together. Lina, thank you for continuing on with the dipole matrix element work and for encouraging me to move onto the spontaneous four-wave mixing work. James & Dan, you both set such an exceptionally high bar; I learned an enormous amount from you both. Mark, the journey for me becoming an experimentalist started with you when we worked on Faraday filtering in rubidium. Thanks for putting me on the right track. Jake, Nick, and Lucy—we made a good cohort!

Kev, I gained much even from our brief conversations. Nick, you encouraged me to go out on a limb with the machine-learning work. None of us really knew whether it would be useful but it shows even more promise now than when we started working on it. Roshan, your work on the machine-learning project was exceptional. Steve,

thank you for encouraging me to give first-year undergraduate physics tutorials. David & Robert, thank you for keeping an eye on my progression through the years. Steve and the rest of the team in the mechanical workshop, you were always helpful despite often being very busy. I would also like to thank my friend, Will Parr, the python wizard, for essential last-minute python help.

Charles Adams, I gained an immense amount from our conversations. I would like to thank all the members of the QLM group and the Durham Physics department for creating such a supportive community. I would also like to acknowledge that this work was made possible by funding from the Engineering and Physical Sciences Research Council (EP/L023024/1, EP/R002061/1).

Last but, by no means, least, I would like to thank my parents; you have provided unflagging support from before I can remember. It is impossible to put into words my gratitude to you both.

List of Figures

2.1 Two-level, three-level, and four-level systems. 10
2.2 Absorption & dispersion of the seed beam in a two-level system. . . 16
2.3 Absorption & dispersion of the seed beam in a three-level system. . 17
2.4 Absorption & dispersion of the seed beam in a four-level system. . . 19
2.5 Doppler-broadened absorption profile for a four-level system. 21
2.6 Schematic diagram of ^{87}Rb atomic structure. 23
2.7 Theoretical transmission plots for 795 nm seed light on the rubidium D_1 line in magnetic fields of different strengths & associated transitions 27
2.8 Theoretical transmission plots for 780 nm pump light on the rubidium D_2 line in magnetic fields of different strengths & associated transitions 28

3.1 Optical table layout for maximising the output of the BoostTA tapered amplifier. 31
3.2 TA output power increases with TA current. 34
3.3 TA output power increases with TA current (log scale). 35
3.4 TA power increases with seed power and current. 36
3.5 Normalised transmission for the lens etalon filter 38

4.1 Machine learning (ML) feedback loop 45
4.2 Build 1.0: Photo of small stepper motor with steel connector for connection to stable mirror mount 47
4.3 Build 2.0: Photo of large stepper motor attached to stable mirror mount. 47
4.4 Schematic of the lab implementation of the concept 49

4.5 Illustration of the basic concept behind auto-aligning a laser beam . 50
4.6 Correcting for motor hysteresis 53
4.7 Reproducibility: returning the mirror to its starting position after repeated changes to its orientation 54
4.8 Direct output from M-LOOP of parameter landscape exploration . 56
4.9 Direct output from M-LOOP of the full parameter landscape 58
4.10 Explaining M-LOOP's "thought process" (1) 61
4.11 Explaining M-LOOP's "thought process" (2) 62

5.1 Schematic of the STROLL (Simultaneous Two-Photon Resonant Optical Laser Lock) experimental configuration. 68
5.2 Explanation of the spectra used to create STROLL (Part I) 72
5.3 Explanation of the spectra used to create STROLL (Part II) 73
5.4 Overlapping Allan deviation of the frequency measurements of the two lasers being locked 75
5.5 Overlapping Allan deviation of the temperature measurement of the locking cell 77

6.1 780 nm photons generated when 795 nm beam is inputted and vice versa 80
6.2 Experimental set-up diagram 82
6.3 Theoretical transmission plots for 780 nm pump light on the rubidium D_2 line & associated transitions 84
6.4 Experimental transmission plots for 780 nm light & associated transitions 86
6.5 795 nm photon generation increases with increasing pump power . . 87
6.6 795 nm photon generation associated with each 780 nm transition (low temperature) 88
6.7 795 nm photon generation associated with each 780 nm transition (high temperature) 89
6.8 Dependence of 795 nm photon count rate on atomic number density 91
6.9 Theoretical transmission plots for 795 nm seed light on the rubidium D_1 line & associated transitions 92

6.10 Experimental transmission plots for 795 nm light & associated transitions . . . 94
6.11 780 nm photon generation associated with each 795 nm transition . 95
6.12 Dependence of 780 nm photon count rate on atomic number density 96

7.1 Rubidium energy levels in the diamond configuration, with and without a magnetic field . . . 101
7.2 Example experimental four-wave mixing (4WM) spectra . . . 101
7.3 Diagram of laser set-up . . . 105
7.4 Diagram of experimental set-up . . . 106
7.5 Spectra of the rubidium D_1 (795 nm) line, with and without a magnetic field. Spectra with seed, seed plus pump, and seed plus pump plus coupling. . . . 109
7.6 4WM signal as a function of atomic number density and temperature. 115
7.7 4WM signal as a function of seed beam power. . . . 116
7.8 4WM signal as a function of pump beam power. . . . 117
7.9 4WM signal as a function of pump beam power for powers less than 1 μW. . . . 118
7.10 Heatmap showing that the theoretical 4WM signal has individual contributions from many velocity classes. . . . 119
7.11 4WM signal as a function of both coupling and seed beam detunings. 121
7.12 4WM signals for different coupling powers . . . 122
7.13 Dependence upon pump power of the scaling factor, m', between peak splitting and coupling Rabi frequency . . . 124
7.14 Experimental four-wave mixing signal with theoretical signal overlaid. 128
7.15 Extracted values of the splitting of the peaks, S, as a function of coupling beam power. Data superset 1. . . . 130
7.16 Extracted values of the splitting of the peaks, S, as a function of the square root of coupling-beam power. Data superset 1. . . . 130
7.17 Extracted values of the splitting of the peaks, S, as a function of the square root of coupling-beam power. Data superset 2. . . . 132

8.1 Diagram of experimental set-up . . . 141

8.2 The normalised cross-correlation between simulated arrivals at channels 0 and 1. . . . 145
8.3 Normalised double coincidence and normalised triple coincidence events. . . . 147
8.4 Hanbury Brown–Twiss experiment shows a $g^{(2)}(\tau = 0)$ dip below 0.5 which demonstrates the presence of single-photons. . . . 150
8.5 The $g^{(2)}$ dip of the Hanbury Brown–Twiss (HBT) experiment is lower for larger $g^{(2)}_{\mathrm{h,s-MAX}}$ . . . 151
8.6 Herald-signal pair rate increases with temperature but the max. cross-correlation decreases. . . . 154
8.7 Herald-signal maximum cross-correlation increases with pump detuning 156
8.8 Herald-signal maximum cross-correlation decreases with pump power 157
8.9 Maximum cross-correlation between herald and signal decreases with coupling power . . . 158
8.10 Ratio of correlation peaks (with and without etalon) . . . 160
8.11 Herald-signal cross-correlation with and without the lens etalon filter 161

A.1 Full Optical Table Layout . . . 170

"I am among those who think that science has great beauty... A scientist in his laboratory is not only a technician: he is also a child placed before natural phenomena which impress him like a fairy tale. We should not allow it to be believed that all scientific progress can be reduced to mechanisms, machines, gearings, even though such machinery also has its own beauty."

MARIE CURIE
1867–1934
[1]

Chapter 1

Introduction

In this work, we explore the optical phenomenon of four-wave mixing (4WM), in which three optical fields are coherently combined to produce a fourth field. This process in nonlinear optics has found use in a variety of applications, including the production of correlated photon pairs for use in quantum information protocols [2]; for transferring trans-spectral orbital angular momentum [3]; for precision spectroscopic measurements [4]; for the creation of entangled imaging systems [5]; and for observing collective quantum beats [6], to name but a few.

Atoms of a single isotope can be considered to be identical oscillators[1]. EM fields affect the state of an atom, allowing the atom to be used as a sensor for these fields [9–11]. Conversely, light incident upon an atom is changed by the atom, allowing atomic media to be used to manipulate light [12]. Traditionally—since laser cooling was developed [13]—cold atoms have been used for these purposes. However, thermal vapours have distinct advantages due principally to the absence of a need for sophisticated apparatuses for cooling and trapping. This leads to thermal atomic vapour experiments being simpler to make, cheaper, and more compact. The compactness allows for devices built using thermal vapour cells to be miniaturised [14,15]. Furthermore, with vapour cells, one can change the atomic

[1]This is true provided the atoms are observed under the same conditions (e.g., at rest in the same magnetic field). This then allows them to be used as ultra-precise clocks with long-term stability [7,8].

number density in a range that spans many orders of magnitude, simply by changing their temperature. The workhorses in the field are alkali metals because they have simple, well-established atomic structures and well-understood interactions with external fields. Throughout this work, we work with a thermal vapour of pure (i.e., at the 98% level) rubidium-87.

Atomic energy levels are very sensitive to externally applied magnetic fields [16], allowing them to be used as sensors for these fields [17]. In magnetic field strengths of ~1 T, the ground state alkali-metal atoms enter what is known as the hyperfine Paschen–Back (HPB) regime where the Zeeman shifts exceed the atomic hyperfine splittings. It is possible to purchase permanent magnets with the high field strengths required to reach this regime. As vapour cells can be small, good field homogeneity across the cell is also readily achievable with such magnets. In the HPB regime, each of the transitions are individually resolved and therefore separately addressable. This enables simplified spectroscopic measurements in thermal vapours [6, 18–25] as ideal two-, three- [24] and four-level [25] systems are realised.

We characterise a four-level seeded 4WM system in a diamond configuration in the hyperfine Paschen–Back regime and show how the 4WM signal can be used to extract an excited state dipole matrix element. We enter the arena of quantum nonlinear optics by showing how the particular case of spontaneous 4WM can be used to create a heralded source of single-photons in rubidium vapour. This is demonstrated by a Hanbury Brown–Twiss (HBT) experiment giving a $g^{(2)}(0)$ value of 0.35 ± 0.02. We find that to bring this number closer to zero, and thus the system closer to an ideal single-photon source, it is necessary to increase the correlation between the herald-signal photon pairs produced. We explore the various parameters that can increase this correlation.

Along the way to these results, we design and build a bespoke etalon filter to aid in the removal of background photons, and show that it has the effect of improving correlations between herald and signal photons. We perform a set of investigations into these background photons, confirming that they result from the collisional transfer of states. We devise a new laser locking method, called STROLL, that is effective at locking two lasers simultaneously to a two-photon transition in the

hyperfine Paschen–Back regime. We also design and build a novel machine-learning device to automatically align laser beams, which will, amongst other things, allow the maximum number of signal photons to be captured.

1.1 A brief history of quantum theory, leading to quantum optics

A short trip through the history of quantum theory will situate our work in quantum optics within the larger historical context. Since this is a long and interesting set of narratives, we will limit the discussion to points that the author found interesting and relevant to the work in this thesis.

Perhaps the best place to begin is with the rivalry between Newton and Huygens: in the early 1800s, theirs were the two competing theories of light. Huygens held to a wave theory of light and Newton to a particle ("corpuscular") theory. Both Huygens' wave concept and Newtons' particle concept differ considerably from the modern understanding of particles or waves. Huygens' wave concept lacked both a frequency or a wavelength; Newton's particles were four-sided [26]. Nevertheless, Newton was able to use his concept to explain the behaviour of light being split into ordinary and extraordinary beams by a calcite crystal, and it turns out that his explanation was essentially an elucidation of the polarisation properties of light.[2]

It seemed initially that the wave theory was the best explanation for Young's double-slit experiment of 1801 [27] and diffraction, as explained by Fresnel in 1815 [28]. Familiarity with Young's double-slit experiment and the ease with which the results can be explained by wave interference may make us overlook its strangeness from the point of view of light-as-particles: *adding* light (through a second hole) leads to *lower* light intensity (at some points on a screen)[3]. And so by the end

[2]The author finds it encouraging that it is possible to be wrong in the specifics and yet be right in the general or, alternatively, to be interestingly wrong such that one opens the way to the right answer.

[3]This strangeness only deepens with later experiments [29,30] showing that double-slit experiments with a single photon at a time still give these interference effects.

of that century, a little after the Maxwell-Heaviside equations showed wavelike solutions that travelled at the speed of light, the corpuscular theory seemed to have been undermined.

However, the corpuscular theory was swiftly revived in 1900 when Planck [31] hypothesised that black-body radiation could be explained as discrete packets of energy called quanta. In particular, he accounted for the spectral distribution of radiation from a thermal source by postulating that the energy of a harmonic oscillator is quantised [32, 33]. This solved the ultraviolet catastrophe problem that had been one of the outstanding problems in physics to that date, despite that specific problem not being the motivation for Planck's hypothesis [34]. Initially Planck did not believe this quantisation to be anything more than a mathematical trick [35]. Half a century later, only after quantum theory was fully developed would Planck write, "I now knew for a fact that the elementary quantum of action [Planck's constant] played a far more significant part in physics than I had originally been inclined to suspect..." [36]

In 1905, Einstein [37] used the nascent quantum theory to explain the photoelectric effect. Planck's and Einstein's work inspired much of the early work on the development of quantum mechanics. Yet those early experiments did not constitute direct experimental evidence that it was *light* that was quantised. Indeed, it can be shown that the photoelectric effect can be explained by assuming only that the *atom* has quantised energy levels. Throughout its history, there have been few experiments which directly show the quantum nature of light. This recurring theme appears in the experiments in this thesis. Proving the single-photon nature of our source took at least as long as building the source in the first place.

Experimental quantum optics proper can perhaps be said to begin with Taylor's experiment of 1909 [29]. This was a successor experiment to Young's with the light beam intensity attenuated such that there was only one quantum of energy in the apparatus at any given time.

Theoretical quantum optics proper began in the 1920s. The word "photon" was coined by Lewis in a letter to Nature in 1926 [38][4], and in 1927 Dirac [41] established

[4]To be exactly accurate the word "photon" was first used by Troland [39] but it meant something

his quantum theory of radiation. In the same decade, Dirac suggested that light is both a particle and a wave at the same time [42]. Until the 1950s, the subject of quantum optics, in theory and practice, was devoted to studying the optical spectra of atoms.

The game-changer of the field was the invention of the laser in 1960 ([43] & [44]). The laser allowed experimental explorations of quantum theory in multiple directions. Of particular interest to us is Glauber's 1963 [45] descriptions of states of light that have different statistical properties to classical light. In 1977, these non-classical properties were confirmed by Kimble, Dagenais, and Mandel [46] with the first demonstration of photon antibunching.

We finish by mentioning the work of Hanbury Brown and Twiss [47]. In 1956, they recorded correlations between starlight intensities detected on two separated detectors. These experiments investigated light intensity fluctuations on short-time scales which paved the way for photon statistics experiments that can only be explained by quantised light. It is by using a modern version of Hanbury Brown–Twiss experiment that we are able, in this thesis, to show we have built a single-photon source.

1.2 Thesis summary

Chapter 2 – We explicate the theoretical underpinnings of the basic atom-light interactions work underlying this work. Starting from a two-level system and computationally solving the Lindblad equation, extra levels are added to show how the absorptive and dispersive properties of the medium change. The changes to the energy-level structure of rubidium in the hyperfine Paschen–Back regime are explained.

Chapter 3 – We detail the ancillary experimental apparatus and techniques required for our experiments. We explain how to optimise the power output of a tapered amplifier and how to use an optical cavity to calibrate the frequency of a laser.

different to him; it was Lewis that used the word to mean something similar to what we mean now. The word remains controversial in some quarters [40].

The design of a bespoke etalon lens filter that we built to narrowly filter light is discussed.

Chapter 4 – We describe the design and implementation of an innovative machine-learning device for automating the alignment of laser beams. Both the first-generation (Build 1.0) and second-generation (Build 2.0) designs are discussed and their differences explained.

Chapter 5 – We describe a novel laser locking scheme that simultaneously stabilises the frequencies of the 780 nm pump laser and the 776 nm coupling laser to a two-photon transition in the hyperfine Paschen–Back regime.

Chapter 6 – We undertake an investigation of the background photons generated by collisional transfer. It is shown that the generation of 780 nm fluorescence light from 795 nm laser light in a rubidium vapour, and vice versa, is dependent on the atomic number density.

Chapter 7 – We characterise the seeded four-wave mixing (4WM) system, and show that the experimental data agree with the computational model. In particular, we explain that this works despite the model being a simple one because of the effective four-level system created in the hyperfine Paschen–Back regime. The 4WM signal is used to extract a dipole matrix element.

Chapter 8 – We describe how spontaneous four-wave mixing is used to create a single-photon source. The single-photon nature of the source is demonstrated by showing that a Hanbury Brown–Twiss experiment gives a value for $g^{(2)}(\tau = 0)$ below 0.5. All the experimental variables that bring this value closer to zero are investigated.

Chapter 9 – Each chapter of the thesis is summarised, and the outlook for the research is commented upon.

1.3 Publications

The following papers have come out of the work described in this thesis:

Renju S. Mathew, Francisco Ponciano-Ojeda, James Keaveney, Daniel J. Whiting, and Ifan G. Hughes, *Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime*, Optics Letters **43** 17, pp. 4204-4207 (2018). doi:10.1364/OL.43.004204 [48].

Renju S. Mathew, Roshan O'Donnell, Danielle Pizzey, and Ifan G. Hughes. *The Raspberry Pi Auto-aligner: Machine Learning for Automated Alignment of Laser Beams*, Review of Scientific Instruments **92** 015117 (2018). doi:10.1063/5.0032588 [49].

Daniel J. Whiting, Renju S. Mathew, James Keaveney, Charles S. Adams, and Ifan G. Hughes, *Four-wave mixing in a non-degenerate four-level diamond configuration in the hyperfine Paschen–Back regime*, Journal of Modern Optics **65** 713 (2018). doi:10.1080/09500340.2017.1377308 [25].

Clare R. Higgins, Danielle Pizzey, Renju S. Mathew and Ifan G. Hughes. *Atomic line versus lens cavity filters: a comparison of their merits.* OSA Continuum **3** 961 (2020). doi:10.1364/OSAC.390604 [50].

The following papers are in preparation:

Direct measurement of an excited-state dipole matrix element using four-wave mixing in the hyperfine Paschen–Back regime.

Single-photon generation in rubidium vapour via four-wave mixing in the hyperfine Paschen–Back regime.

1.4 Contributions

Ifan Hughes was the principal investigator for this project. Charles Adams was the secondary supervisor. Contributions from other researchers are detailed below.

Nicholas Spong was involved at the origin of Build 1.0 of the machine-learning project. He provided guidance and support in the early stages of that project.

Clare Higgins was involved at the end of that project and helped to test the device. Investigations were performed on the spontaneous four-wave mixing project with Clare's assistance. The etalon filter was built by Clare & Danielle Pizzey. Build 2.0 of the machine-learning project was built by Roshan O'Donnell. The building of the STROLL lock was done jointly with Francisco Ponciano-Ojeda. He was also involved at the beginning of the investigations into background photons. Seeded four-wave mixing (4WM) work, excluding the measurement of the dipole matrix element, was done jointly with Dan Whiting, with Dan doing much of the modelling. Postdoctoral support was provided by James Keaveney in the first third of the project, by Liz Bridge in the second third and Danielle Pizzey & Lina Hoyos-Campo in the last third. Any highly specific contributions are indicated in the relevant chapters.

Chapter 2

Theory: Atom-light interactions

In this chapter, we shall introduce the basic physics underlying the atom-light interactions described in the rest of the thesis. We shall begin by considering the dynamics of the simplest system, a hypothetical two-level atom coupled by a driving field (Fig. 2.1) and determine the steady-state solution. We shall then add levels until we reach the four-level system that is considered in most of this work. We then cover the energy level structure of the atom which we will use in all of our experiments, rubidium-87, and finish by explaining how the energy levels shift in the hyperfine Paschen–Back regime. More advanced, as well as more specific, theory will be covered in the relevant chapters.

2.1 Time evolution of a two-level system

Our discussion begins with a hypothetical atom with just two levels that are coupled by a monochromatic light field. The two-level atom has a ground-state $|g\rangle$ of energy E_g and an excited-state $|e\rangle$ of energy E_e, where $E_e - E_g = \hbar\omega_0$, and ω_0 is the atomic resonance frequency. We can describe the state of the two-level atom by using a density matrix [51]. The density matrix method allows the computation of the dynamics of quantum systems without requiring full information about all of the

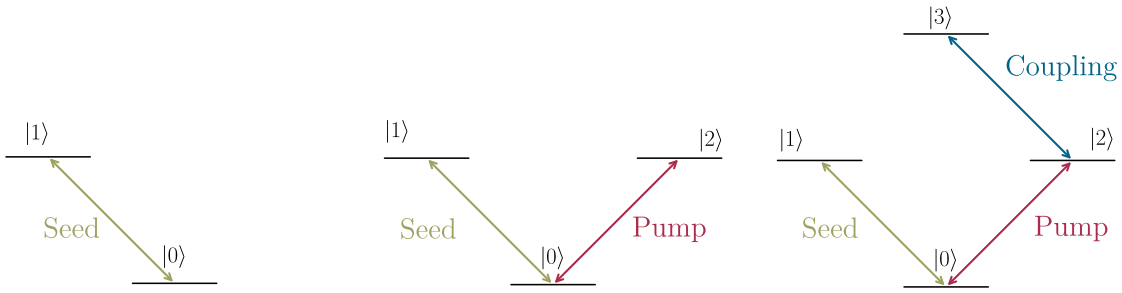

(a) Two-level system (b) Three-level system (c) Four-level system

Figure 2.1: (a) Two-level, (b) three-level, and (c) four-level systems coupled with seed, pump, and coupling lasers with Rabi frequencies Ω_s, Ω_p, and Ω_c respectively. In the main text, $|0\rangle$ & $|1\rangle$ are also referred to as $|g\rangle$ & $|e\rangle$, for ground and excited respectively.

system's underlying details [52]. The two-by-two density operator, $\hat{\rho}$, representing the two-level system is

$$\hat{\rho} = \begin{pmatrix} \rho_{gg} & \rho_{ge} \\ \rho_{eg} & \rho_{ee} \end{pmatrix}, \tag{2.1}$$

where the subscripts 'g' and 'e' refer to the ground and excited states, respectively. The diagonal elements represent the populations of the states, and the off-diagonal elements represent coherences between the states. This generalises so that an n-level system can be described with an n-by-n matrix, where the diagonal terms correspond to the populations of each level and the off-diagonal terms correspond to the coherences between levels.

We now consider an incident plane-wave optical field of angular frequency ω with an electric field

$$\vec{E}(\omega) = \hat{\mathbf{e}} E \cos(\omega t) = \frac{\hat{\mathbf{e}} E}{2} \left(e^{i\omega t} + e^{-i\omega t} \right), \tag{2.2}$$

where $\hat{\mathbf{e}}$ is a vector describing the polarisation of the field. In full generality, this field is written $\vec{E}(\omega, t) = \hat{\mathbf{e}} E \cos(\omega t - kz)$, where $k = 2\pi/\lambda$ is the wavenumber and z is the propagation direction, but we can make use of the dipole approximation [16] to allow us to neglect the spatial z dependence. This approximation can be made when the wavelength of the light is much larger than the spatial extent of

the atomic wavefunctions, as is the case for all the systems considered.

The field couples the two states via an electric dipole transition which has the dipole operator

$$\hat{d} = \begin{pmatrix} 0 & \vec{d}_{ge} \\ \vec{d}_{eg} & 0 \end{pmatrix},$$

where $\vec{d}_{ij} = \langle i|e\vec{r}|j\rangle$ are the dipole matrix elements, e is the electron charge, and $\vec{r}$ is the displacement vector.

The atom-field interaction can then be described by the following Hamiltonian: [16]

$$\hat{H}_{\text{int}} = -\hat{d} \cdot \vec{E}. \tag{2.3}$$

Liouville evolution

The time evolution of the system can now be calculated using the Liouville equation[1] [53]

$$\frac{\mathrm{d}\hat{\rho}}{\mathrm{d}t} = -\frac{\mathrm{i}}{\hbar}[\hat{H}, \hat{\rho}], \tag{2.4}$$

where the total Hamiltonian $\hat{H} = \hat{H}_0 + \hat{H}_{\text{int}}$ is the sum of the interaction Hamiltonian, $\hat{H}_{\text{int}}$ (Eq. 2.3), and the unperturbed atomic Hamiltonian, $\hat{H}_0$, which is described by

$$\hat{H}_0 = \begin{pmatrix} E_g & 0 \\ 0 & E_e \end{pmatrix}. \tag{2.5}$$

It is useful to perform a rotating frame transformation using the following unitary operator:

$$\hat{U} = \begin{pmatrix} \mathrm{e}^{-\mathrm{i}E_g/\hbar t} & 0 \\ 0 & \mathrm{e}^{-\mathrm{i}\omega t}\mathrm{e}^{-\mathrm{i}E_g/\hbar t} \end{pmatrix}. \tag{2.6}$$

This frame contains a term which oscillates with the driving field frequency, and it allows us to shift the energy scale such that the ground state is zero. This gives the transformed density matrix

[1] This is equivalent to using the time-dependent Schrödinger equation with a quantum wavefunction. Eq. 2.4 is often also referred to as the Liouville–von Neumann equation as it was von Neumann that first introduced the formalism of density operators and matrices [53].

$$\tilde{\rho} = \hat{U}^\dagger \hat{\rho} \hat{U} = \begin{pmatrix} \rho_{gg} & \rho_{ge} \mathrm{e}^{-\mathrm{i}\omega t} \\ \rho_{eg} \mathrm{e}^{\mathrm{i}\omega t} & \rho_{ee} \end{pmatrix},$$

and the transformed Hamiltonian

$$\begin{aligned} \tilde{H} &= \hat{U}^\dagger \hat{H} \hat{U} - \mathrm{i}\hbar \hat{U}^\dagger \frac{\mathrm{d}\hat{U}}{\mathrm{d}t} \\ &= \frac{1}{2} \begin{pmatrix} 0 & -\vec{d}_{ge} \cdot \vec{\epsilon} E[1 + \mathrm{e}^{-2\mathrm{i}\omega t}] \\ -\vec{d}_{eg} \cdot \vec{\epsilon} E[1 + \mathrm{e}^{2\mathrm{i}\omega t}] & 2[E_e - E_g - \hbar\omega] \end{pmatrix}. \end{aligned}$$

We now use the rotating-wave approximation (RWA) [16]: the rapidly oscillating terms can be discarded because their time dependence averages out over the relatively larger time scale required for the evolution of the atomic populations.

So, we have

$$\tilde{H}_{\mathrm{RWA}} = \frac{1}{2} \begin{pmatrix} 0 & -\vec{d}_{ge} \cdot \vec{\epsilon} E \\ -\vec{d}_{eg} \cdot \vec{\epsilon} E & 2[\hbar\omega_0 - \hbar\omega] \end{pmatrix}.$$

By making use of the definitions for the angular Rabi frequency, $\Omega = -\vec{d}_{ge} \cdot \vec{\epsilon} E/\hbar$, and the detuning, $\Delta = \omega - \omega_0$, we obtain the final Hamiltonian[2]

$$\tilde{H}_{\mathrm{RWA}} = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega \\ \Omega^* & -2\Delta \end{pmatrix}.$$

In preparation for adding more levels (See Fig. 2.1), let us write this as

$$\tilde{H}_{\mathrm{2level}} = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_{\mathrm{s}} \\ \Omega_{\mathrm{s}} & -2\Delta_{\mathrm{s}} \end{pmatrix}, \tag{2.7}$$

where the subscript 's' indicates the seed beam[3] that couples the two levels together.

[2] $\Omega = \Omega^*$ since $\vec{d}_{eg} = \vec{d}_{ge}$.

[3] This is conventionally referred to as the *probe* beam, but we refer to it as the seed beam as it will eventually be identified as the seed for the seeded four-wave mixing process.

Lindblad master equation

Even though we are considering the atom interacting with a single light field, there is some interaction of the atom with the vacuum, which leads to the emission of photons. This is the well-known phenomenon of spontaneous emission ([16] & Section 10.3.2 of [54]).

This decay is not accounted for by the Liouville–von Neumann equation (Eq. 2.4) because that only describes the coherent evolution of the density matrix. The decay *can* be included phenomenologically as a statistical process packaged into the decay matrix, $\hat{L}$, which is added to the right hand side of Eq. 2.4. This gives the Lindblad master equation ([55–57], Eq. 9 in [57])

$$\frac{\mathrm{d}\hat{\rho}}{\mathrm{d}t} = -\frac{\mathrm{i}}{\hbar}[\hat{H}, \hat{\rho}] + \hat{L}, \tag{2.8}$$

where $\hat{L} = \sum_n [2\hat{C}_n\hat{\rho}\hat{C}_n^\dagger - (\hat{\rho}\hat{C}_n^\dagger\hat{C}_n + \hat{C}_n^\dagger\hat{C}_n\hat{\rho})]/2$, $\hat{C}_n = \sqrt{\gamma_n}\hat{A}_n$ are the collapse operators, and $\hat{A}_n$ are the operators describing the coupling of the environment to the system with rates γ_n . For spontaneous decay from the excited state to the ground state (with a lifetime $\tau = 1/\Gamma_\mathrm{s}$), the collapse operator is $\hat{C} = \sqrt{\Gamma_\mathrm{s}}|g\rangle\langle e|$. We shall see that Γ_s is the linewidth of the transition.

Writing out each term in equation 2.8 leads to the optical Bloch equations [58]. For the two-level system, the steady-state solutions can be derived analytically. The coherence between the two states is,

$$\rho_{eg} = -\frac{\mathrm{i}\Omega/2}{\Gamma_\mathrm{s}/2 - \mathrm{i}\Delta}\left(\rho_{gg} - \rho_{ee}\right). \tag{2.9}$$

Taking the limit of ρ_{eg}/Ω as $\Omega \to 0$, i.e., in the limit of a weak incident electric field, we find that

$$\lim_{\Omega\to 0}\frac{\rho_{eg}}{\Omega} = \frac{1}{2\Delta + \mathrm{i}\Gamma_\mathrm{s}}. \tag{2.10}$$

Macroscopic properties

We now relate the ρ_{eg} term to the macroscopic properties of a medium containing many two-level atoms. The optical response of the system can be described in terms of its complex *electric susceptibility*, χ. The macroscopic polarisation, $\vec{P}$, is related to the individual (microscopic) electric dipole moment of each atom as follows. The polarisation is given by the number density of dipoles N multiplied by the average dipole moment,

$$\vec{P} = N\langle\hat{d}\rangle = \mathrm{Tr}(\hat{\rho}\hat{d})N = N(\vec{d}_{ge}\rho_{eg}e^{-\mathrm{i}\omega t} + \text{c.c.}), \tag{2.11}$$

where c.c. is the complex conjugate.

The polarisation of the macroscopic medium can also be written as [59]

$$\vec{P} = \epsilon_0\chi\vec{E} = \frac{1}{2}\epsilon_0\vec{E}(\chi e^{-\mathrm{i}\omega t} + \text{c.c.}), \tag{2.12}$$

where ϵ_0 is the permittivity of free space. The complex susceptibility is frequency-dependent, and the medium's dispersive and absorptive properties are described by the real and imaginary parts of χ, respectively. We can equate equations 2.11 and 2.12 in order to relate the microscopic and macroscopic quantities. Taking the dot product with $\vec{d}_{ge}$ of both sides gives

$$Nd_{ge}^2\rho_{eg} = \frac{1}{2}\epsilon_0\frac{\vec{E}\cdot\vec{d}_{ge}}{\hbar}\hbar\chi. \tag{2.13}$$

Rearranging, we can write the susceptibility in terms of the microscopic quantities as

$$\chi = -\frac{2d_{ge}^2}{\hbar\epsilon_0}\left(\frac{\rho_{eg}}{\Omega}\right)N, \tag{2.14}$$

where the definition of Rabi frequency $\Omega = -\vec{d}_{ge}\cdot\vec{\epsilon}E/\hbar$ has been used.

Now, combining Eq. 2.14 and Eq. 2.10 (i.e., the weak-probe regime [60] where the susceptibility is independent of the incident electric field strength) we have

$$\chi(\Delta) = \frac{\mathrm{i}Nd_{ge}^2}{\epsilon_0\hbar(\Gamma_\mathrm{s}/2 - \mathrm{i}\Delta)} \ . \tag{2.15}$$

The complex refractive index of the medium, $\underline{n}$ [4], is related to the susceptibility by $\underline{n} = \sqrt{1+\chi}$. In terms of the real and imaginary parts of the refractive index ($\underline{n} = n_\mathrm{R} + \mathrm{i}n_\mathrm{I}$), we can write the modification of the electric field due to the medium as $\hat{E}(z) = \hat{E}_0 e^{\mathrm{i}n_\mathrm{R}kz} e^{-kn_\mathrm{I}z}$. The real part of the complex refractive index, n_R, characterises the phase velocity, v, of the light in the medium: $n_\mathrm{R} = c/v$, where c is the speed of light in vacuum[5]. As the electric susceptibility is a function of detuning, n_R contains information on how the phase velocity varies with the frequency of the wave, i.e., its dispersion. The imaginary part of the refractive index, n_I, leads to an exponentially decaying amplitude which describes the *extinction* of the field. This is often referred to as absorption. The transmitted intensity through a medium of thickness l can be written as $T = \exp(-\alpha l)$ where the so-called absorption coefficient α depends linearly on the imaginary part of the refractive index $\alpha = 2kn_\mathrm{I}$. This is Beer's law [61], also known as the Beer-Lambert law.

Systems containing more than two levels cannot be easily solved analytically, and so computational methods are used. For this purpose, we have used, to good effect, the Quantum Toolbox in Python (`QuTiP`) module [57] which numerically solves the master equation (Eq. 2.8). Since we use `QuTiP` to model three- and four-level systems, we will, for consistency, also use it to model the two-level system just described.[6] The figures that follow are created using outputs from `QuTiP`. We are not here attempting to match theory to experiment—where relevant this will be done in the appropriate chapters of the thesis; rather, the figures illustrate how the absorptive and dispersive lineshapes vary as extra levels, and couplings between levels, are introduced. For the two-level system, Fig. 2.2 shows the absorptive and dispersive lineshapes. The absorptive line shows the characteristic Lorentzian lineshape of Eq. 2.10 with a full-width at half-maximum linewidth of Γ_s. The

[4] The underline indicates that it is a complex quantity.

[5] In everyday usage, often what is referred to as the "refractive index" is only the real part, i.e., $n_\mathrm{R} = n = c/v$.

[6] Although we used `QuTiP` to model the *weak* probe (seed) beams which we used in our experiment, `QuTiP` is not limited to low-intensity beams.

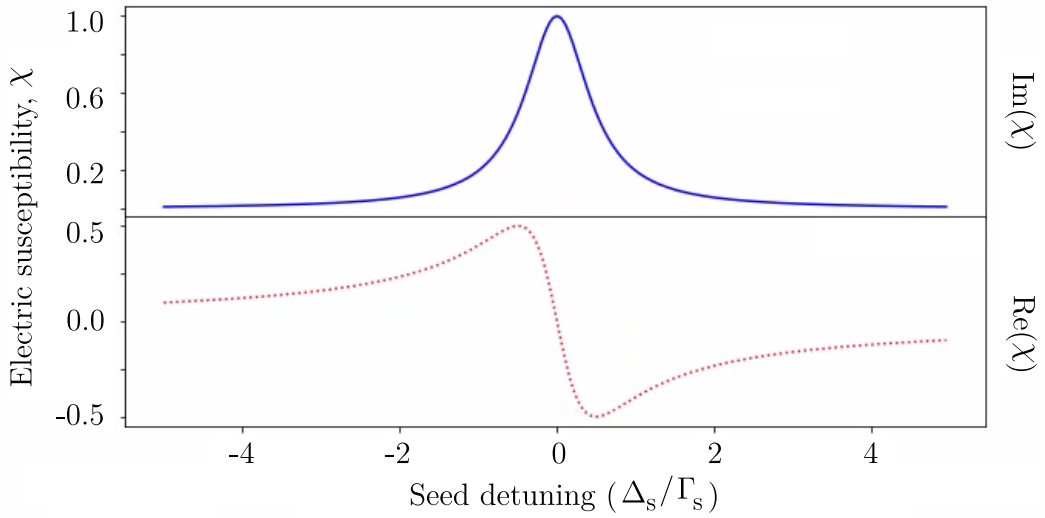

Figure 2.2: Two-level system. The imaginary (blue) and real (red dotted) parts of the normalised electric susceptibility, χ, are plotted, showing the characteristic absorptive and dispersive lineshapes respectively.

dispersive line has the characteristic dispersive lineshape, i.e., an odd function centred around resonance, positively valued below resonance and negatively valued above it.

2.2 Three-level system

Following the procedure detailed in Section 2.1, we can arrive at the Hamiltonian for a three-level system (Fig. 2.1), where the subscript 'p' indicates the pump laser that couples the ground state to the new excited state:

$$\tilde{H}_{3\text{level}} = \frac{\hbar}{2}\begin{pmatrix} 0 & \Omega_\text{s} & \Omega_\text{p} \\ \Omega_\text{s} & -2\Delta_\text{s} & 0 \\ \Omega_\text{p} & 0 & -2\Delta_\text{p} \end{pmatrix}. \tag{2.16}$$

By comparing elements, it can be seen that this naturally extends the $\tilde{H}_{2\text{level}}$ matrix (Eq. 2.7) and so could, in principle, be immediately written down without going through the entire procedure in Section 2.1.

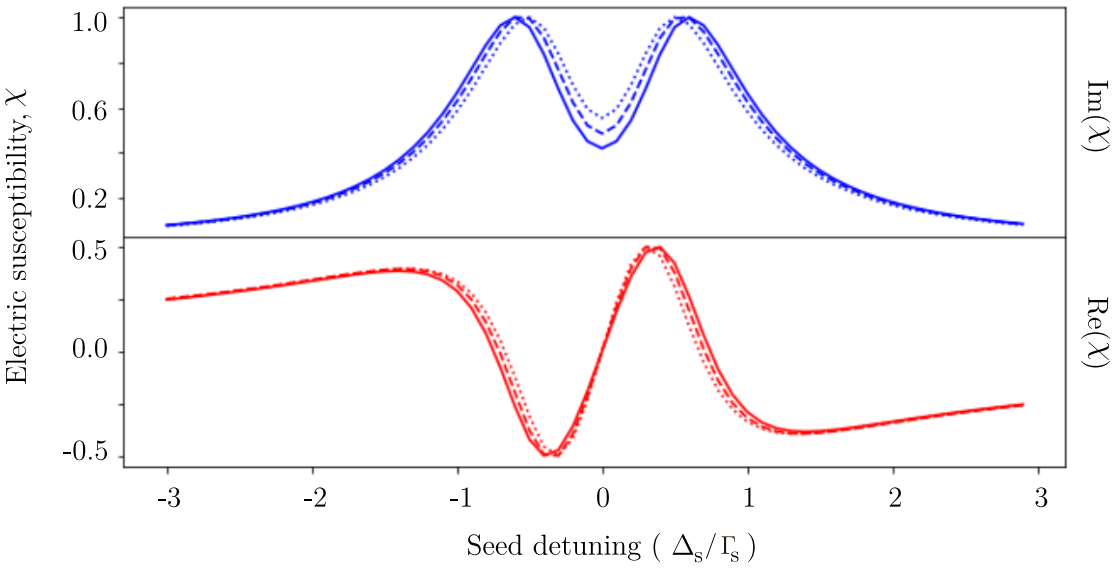

Figure 2.3: Three-level system. The real (red) and imaginary (blue) parts of the normalised electric susceptibility, χ, are plotted. The top panel shows that the EIT window increases in size as the pump beam Rabi frequency, $\Omega_p = \Omega_{02}$, increases from $\Omega_p = 1.6\ \Omega_s$ (dotted) to $\Omega_p = 1.8\ \Omega_s$ (dashed) to $\Omega_p = 2.0\ \Omega_s$ (solid).

We see in Fig. 2.1 that two resonant transitions ($|0\rangle \rightarrow |1\rangle$ & $|0\rangle \rightarrow |2\rangle$) are coherently coupled to a common state ($|0\rangle$). The presence of the strong pump laser creates an electromagnetically induced transparency (EIT) in the seed absorption, as seen in Fig. 2.3. This EIT window can be explained as follows: The strong pump laser shifts the energy levels of the bare atomic states (the AC Stark effect), "dressing" them. The new dressed states [58] *are* the eigenstates of the $\tilde{H}_{3\text{level}}$ Hamiltonian. Destructive interference between the resulting new paths leads to a reduction in the absorption of the weak seed laser[7], leading to an electromagnetically induced transparency (EIT) window within the absorption peak. This EIT window can be seen in Fig. 2.3, where our computational model shows how increasing the stronger pump laser leads to a larger EIT window. Also shown is that associated with the change in absorption, there is a concomitant modification of the refractive index, as seen in the change in the dispersive feature. In Chapter 5, this effect will

[7]In the case of one particular pump detuning, this is equivalent to saying that the relevant two probability amplitudes for excitation to the dressed states are equal but of opposite phases, and so cancel.

be exploited to create the STROLL laser locking system.

2.3 Four-level system

Following the procedure detailed in Section 2.1, we arrive at the Hamiltonian for a four-level system, where the subscript 'c' indicates the coupling laser that couples the excited state to the new doubly excited state ($|2\rangle \rightarrow |3\rangle$). Once again, it can be seen that this naturally extends the $\tilde{H}_{\text{3level}}$ matrix (Eq. 2.16):

$$\tilde{H}_{\text{4level}} = \frac{\hbar}{2}\begin{pmatrix} 0 & \Omega_{\text{s}} & \Omega_{\text{p}} & 0 \\ \Omega_{\text{s}} & -2\Delta_{\text{s}} & 0 & 0 \\ \Omega_{\text{p}} & 0 & -2\Delta_{\text{p}} & \Omega_{\text{c}} \\ 0 & 0 & \Omega_{\text{c}} & -2(\Delta_{\text{p}} + \Delta_{\text{c}}) \end{pmatrix}. \tag{2.17}$$

Addition of the strong coupling laser leads to the previously mentioned dressed states being symmetrically split by the AC Stark effect. Again, there is interference between paths, leading to the original EIT feature being split (as seen in Fig. 2.4).

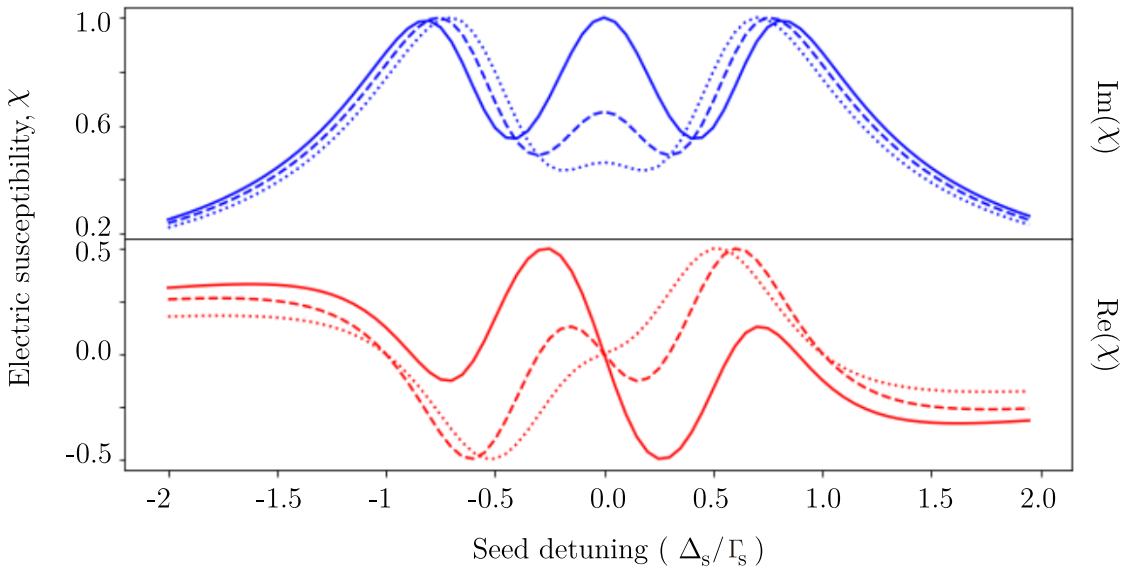

Figure 2.4: Four-level system. The real (red) and imaginary (blue) parts of the normalised electric susceptibility, χ, are plotted. $\Omega_p = 1.6\ \Omega_s$. The top panel shows that the splitting of the EIT window increases as the coupling beam Rabi frequency, $\Omega_c = \Omega_{23}$, increases from $\Omega_c = 1.9\ \Omega_s$ (dotted) to $\Omega_c = 2.6\ \Omega_s$ (dashed) to $\Omega_c = 3.2\ \Omega_s$ (solid).

2.4 Doppler broadening due to atomic motion

Since all the experiments covered in this thesis occur in thermal vapour, the thermal velocity of the atoms must be accounted for. The frequency of the light that the atom experiences is Doppler shifted by an amount

$$\Delta_D = -\vec{k} \cdot \vec{v}, \tag{2.18}$$

where $\vec{v}$ is the velocity of the atom and $\vec{k}$ is the wavevector of the incoming laser beam. We can choose $\vec{k}$ to lie along the z-axis and so $\Delta_D = -kv_z$. A Maxwellian distribution of velocities is expected so that the atomic density for a velocity class $N(v)$ is given by [62]

$$N(v) = N_0 \frac{e^{-v^2/u^2}}{u\sqrt{\pi}}, \tag{2.19}$$

where u is the one-dimensional root mean square (rms) speed of the atoms, $u = \sqrt{2k_{\mathrm{B}}T/m}$. In the case of ^{87}Rb atoms (of which more in the next section) at room temperature (20 °C), $u = 237$ m/s.

As an example, for the four-level system, we add Doppler broadening by replacing Δ_{i} with $(\Delta_{\mathrm{i}} - k_{\mathrm{i}}v)$ where $\mathrm{i} = \{\mathrm{p}, \mathrm{c}, \mathrm{s}\}$, with p, c and s referring to the pump, coupling and seed beams respectively. Assuming all the beams are parallel, this gives the modified Hamiltonian,

$$\tilde{H} = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_{\mathrm{s}} & \Omega_{\mathrm{p}} & 0 \\ \Omega_{\mathrm{s}} & -2(\Delta_{\mathrm{s}} - k_{\mathrm{s}}v) & 0 & 0 \\ \Omega_{\mathrm{p}} & 0 & -2(\Delta_{\mathrm{p}} - k_{\mathrm{p}}v) & \Omega_{\mathrm{c}} \\ 0 & 0 & \Omega_{\mathrm{c}} & -2(\Delta_{\mathrm{p}} - k_{\mathrm{p}}v + \Delta_{\mathrm{c}} - k_{\mathrm{c}}v) \end{pmatrix}. \tag{2.20}$$

This is the Hamiltonian used in the full model that we use later in this thesis and which we shall see agrees with our experimental data. Further details and

discussion of the full model will be found in Chapters 7 and 8, which discuss seeded and spontaneous four-wave mixing respectively. For the sake of completeness, a preview of the *theoretical* absorption profile generated by the full model for the rubidium four-level system is shown in this chapter in Fig. 2.5, for a temperature of 87 °C. The split EIT transmission window can be seen within the absorption profile. The key point to note is that the Doppler broadening is ~100 times that of the natural linewidth ($\Gamma_s/2\pi = 6$ MHz). In Chapter 7, we shall see the *experimental* transmission spectrum associated with this figure.

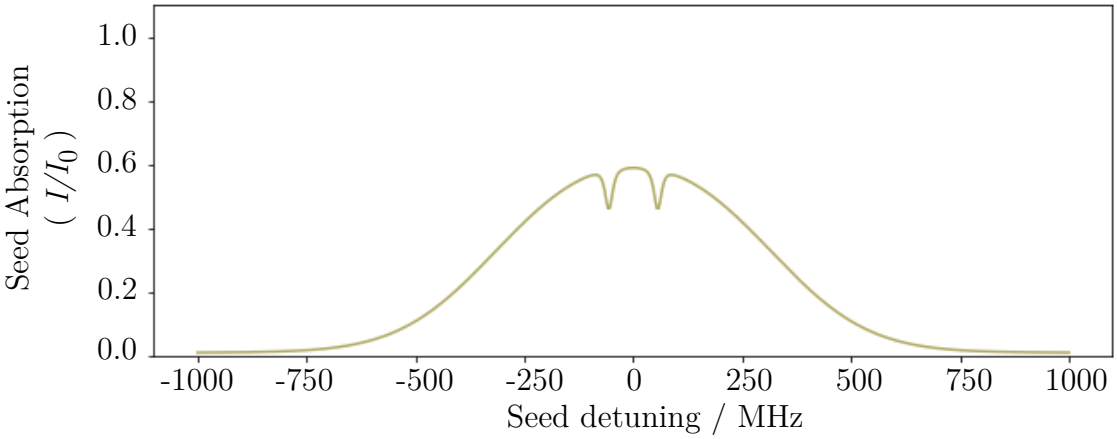

Figure 2.5: A Doppler-broadened absorption profile generated by the full computational model for rubidium vapour in a 2 mm-length cell at a temperature of 87 °C. The normalised intensity of the seed beam is plotted against the seed linear detuning from resonance. The Doppler broadening is ~100 times that of the natural linewidth (Γ_s).

2.5 The atomic structure of rubidium

We have hitherto focused on *hypothetical* atoms with two, three, and four levels. In all of the experiments covered in this thesis, the atom used is rubidium-87. Rubidium is a Group I alkali-metal atom with a single valence electron, occurring naturally as two isotopes: ^{85}Rb with a nuclear spin $I = 5/2$ and ^{87}Rb with a nuclear spin $I = 3/2$; rubidium-87 is used throughout this work. Figure 2.6 is a schematic diagram of its energy level structure. The gross structure is determined by the principal quantum number ($n = 5$). Fine-structure arises from the spin-orbit interaction, i.e., the interaction between the electron spin ($\vec{S}$) and its orbital angular momentum ($\vec{L}$). The resulting new energy eigenstates have total electron angular momenta $\vec{J} = \vec{L} + \vec{S}$. They are well identified by the quantum number $J = |\vec{J}|$, along with m_J, which is the projection of $\vec{J}$ onto a quantisation axis defined by a magnetic field. Hyperfine structure arises from the interaction between the total electron angular momentum ($\vec{J}$) and the nuclear spin ($\vec{I}$). The new energy eigenstates have total angular momenta $\vec{F} = \vec{I} + \vec{J}$. They are well identified by the quantum number $F = |\vec{F}|$, along with m_F, the projection of $\vec{F}$ onto a quantisation axis defined by a magnetic field. The energy level structure changes on application of a large magnetic field [63], to the point where, in the hyperfine Paschen–Back regime, we are able to address individual transitions and effectively treat the atom as a four-level system.

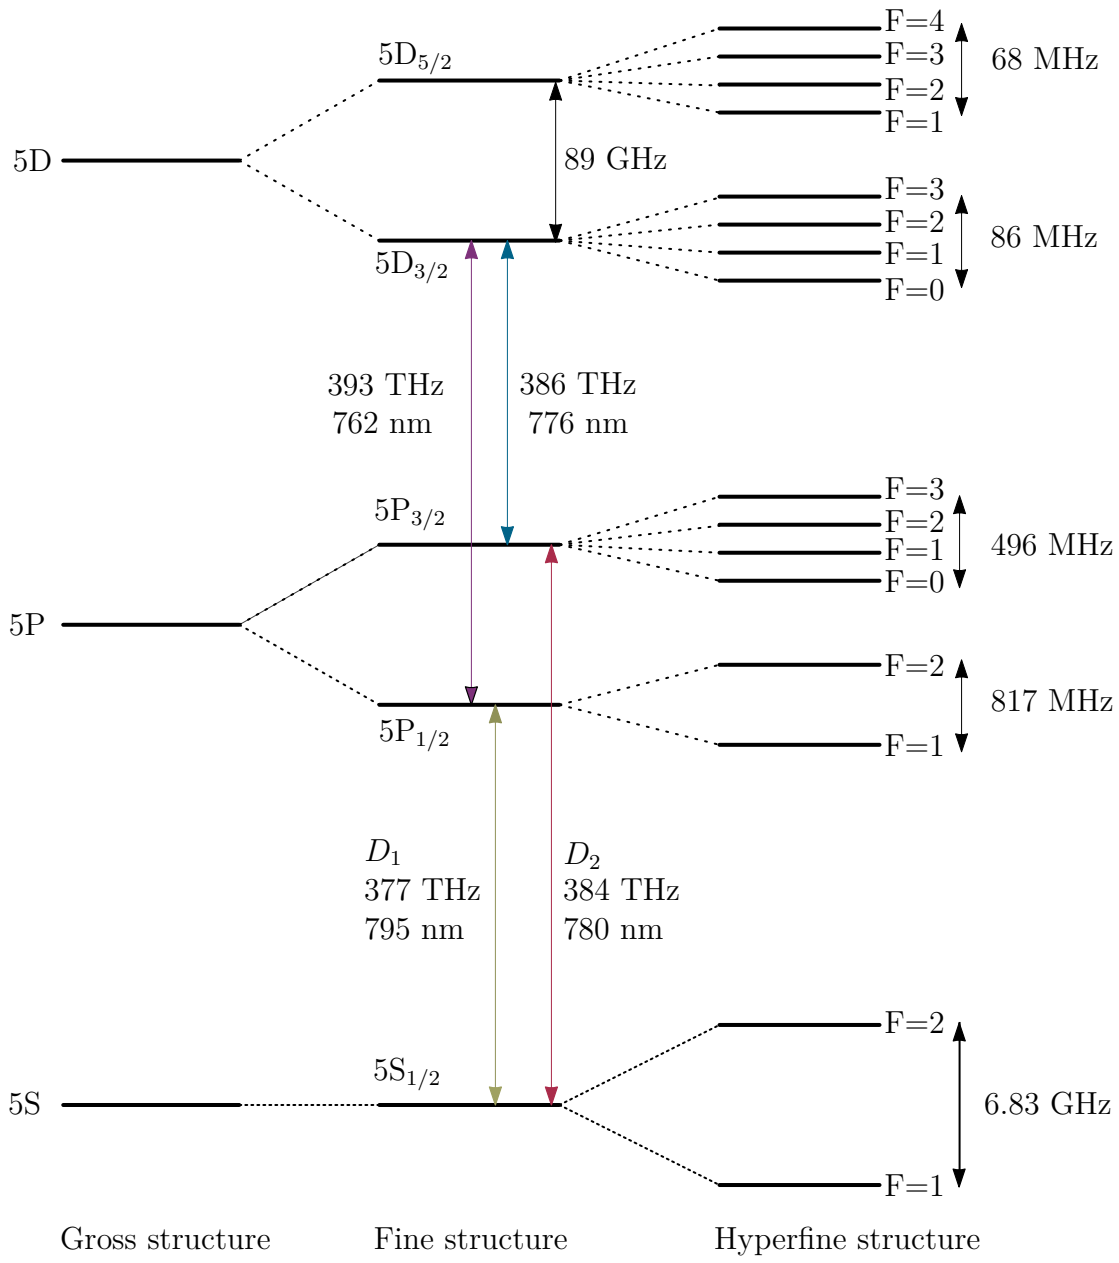

Figure 2.6: ^{87}Rb atomic structure. Data from [64] and [65].

2.6 Application of a high magnetic field: the hyperfine Paschen–Back (HPB) regime

The Zeeman effect [63] describes the splitting of sharp spectral lines into multiple lines when an external magnetic field is applied. It is caused by the Zeeman interaction in which the external magnetic field exerts a torque on magnetic dipoles; it can be described with the magnetic interaction Hamiltonian:

$$\hat{H}_{\text{Zeeman}} = -(\vec{\mu}_I + \vec{\mu}_L + \vec{\mu}_S) \cdot \vec{B}, \tag{2.21}$$

where $\vec{\mu}_I$, $\vec{\mu}_L$, and $\vec{\mu}_S$ are the magnetic moments associated with the nuclear spin, the orbital motion of the electron, and the electron spin respectively. The contribution due to nuclear spin can be ignored because the nuclear magneton is three orders of magnitude smaller than the Bohr magneton [51].

For weak magnetic fields, each hyperfine level is split into $2F+1$ levels symmetrically about the zero field level with the energy level shift equalling $\Delta E = g_F m_F \mu_B B$, where μ_B is the Bohr magneton and g_F is the Landé factor. This is known as the hyperfine linear Zeeman (HLZ) regime. In the HLZ regime, the $|F, m_F\rangle$ base describes the interaction best.

The so-called hyperfine Paschen–Back (HPB) regime is entered when the Zeeman interaction causes an energy shift equal to the hyperfine splitting. At fields on the order of ~1 Tesla, for alkali-metal atoms in the HPB regime, the nuclear, I, and total electron, J, angular momenta are decoupled. The energy splitting is then equal to $\Delta E = g_J m_J \mu_B B + g_I m_I \mu_B B$.

Fig. 2.7 (e) shows[8] the change in the energy level structure of ^{87}Rb with applied magnetic fields up to a strength of 0.62 T. In the absence of a magnetic field, the m_F states are degenerate. Application of a magnetic field lifts this degeneracy. By 0.33 T [65], the HPB regime is reached for the ground state. J equals 1/2, giving two branches, $m_J = +1/2$ and $m_J = -1/2$. Since $I = 3/2$, there are $2I + 1 = 4$ energy levels in each branch, labelled $m_I = -3/2, -1/2, +1/2, +3/2$. States

[8] The figure was created using the author's adaptation of code written by J. Keaveney; it makes use of the modules of the ElecSus computational model [66, 67]. Linear detunings are used.

with $m_J = +1/2, -1/2$ shift linearly with the magnetic field as $\Delta E = g_J m_J \mu_B B$ [9], where g_J is the Landé factor [68]. Since, in this case, the Landé factor $g_J = 2$ [10], the energy shift per magnetic field simply equals the value of the Bohr magneton in frequency units, ± 14 GHz/T (Table E. 2 of [68])[11]. We are stating shifts of frequency, $\Delta\nu$, as it is most useful for the experimentalist, but $E = h\nu$ can be used to convert between the two, where h is Planck's constant. Note that at a field of 0.62 T, J and I are not yet completely decoupled. Thus, for the ground states whose state decomposition has been explicitly stated at the bottom of the diagram, we see there is still some small amount of hyperfine mixing.

For the excited state $5P_{1/2}$, J equals 1/2. This again gives two branches, $m_J = \pm 1/2$, with each branch again containing four energy levels, $m_I = \pm 3/2,\ \pm 1/2$. In this case, the Landé factor $g_J = 2/3$. Thus the energy shift is equal to 1/3 of the Bohr magneton value[12]. In the case of the excited state $5P_{3/2}$, shown in Fig. 2.8 (e), J now equals 3/2 and so there are now four branches, $m_J = \pm 3/2,\ \pm 1/2$. Each branch again contains four energy levels, $m_I = \pm 3/2, \pm 1/2$. The Landé factor $g_J = 4/3$ and so the energy shift of the maximally stretched states ($m_J = \pm 3/2$) equals twice the Bohr magneton value, i.e., ± 28 GHz/T [13]. For the other two states ($m_J = \pm 1/2$), the slopes are one-third of this value.

The transmission spectra at different magnetic field strengths for the seed/795 nm beam illuminating the ^{87}Rb D_1 line are shown in Fig 2.7 (a–d), and the spectra for the pump/780 nm beam illuminating the ^{87}Rb D_2 line are shown in Fig. 2.8 (a–d). Both figures will be discussed together. The beam in both cases propagates along the direction defined by the magnetic field; it is linearly polarised and, as such, equally drives σ_+ ($\Delta m_J = +1$) and σ_- ($\Delta m_J = -1$) transitions. In the forward geometry, where the beams are propagating along the direction of the B field, there is no electric field in that direction (i.e., it is all in the x-y plane) and so π transitions are not driven. $\Delta m_I = 0$ in all cases as the light does not interact with

[9] The contribution due to the nuclear spin has been ignored due to it being three orders of magnitude smaller. It is, however, included in the code generating the Breit-Rabi diagrams.

[10] $g_J = \frac{3}{2} + \frac{S(S+1)-L(L+1)}{2J(J+1)}$, with $J = 1/2$, $L = 0$, and $S = 1/2$

[11] $\Delta E/B = g_J m_J \mu_B = 2 \times \frac{1}{2} \times \mu_B = \mu_B$

[12] $\Delta E/B = g_J m_J \mu_B = \frac{2}{3} \times \frac{1}{2} \times \mu_B = \frac{1}{3}\mu_B$

[13] $\Delta E/B = g_J m_J \mu_B = \frac{4}{3} \times \frac{3}{2} \times \mu_B = 2\mu_B$

the nuclear spin.

In (b) and (c), the excited state is in the HPB regime but the ground state is not, giving transitions that are only partially resolved. It is therefore difficult to identify which spectral lines are due to which transitions. When the HPB regime is reached in (d), we see that the lines are fully resolved and each spectral line can be associated with a particular transition. This will allow us to isolate a four-level system. If the weak transitions due to remaining hyperfine mixing are ignored, spectral lines in the negatively detuned regions correspond to light driving σ^- (purple) transitions, and those in the positively detuned regions correspond to light driving σ^+ (cyan) transitions.

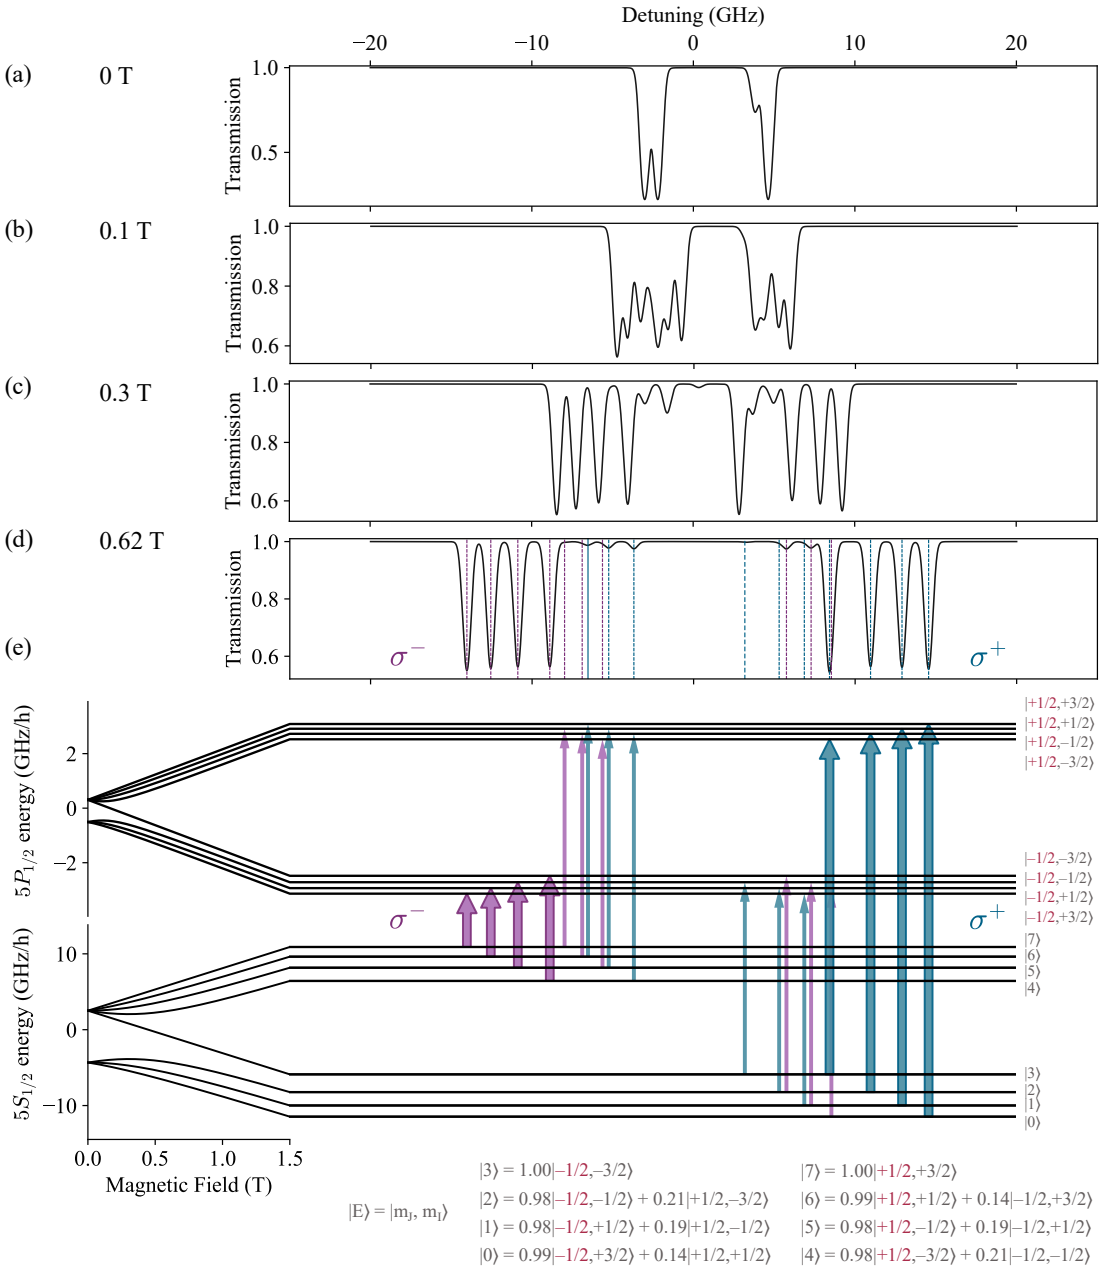

Figure 2.7: (a)–(d) Theoretical transmission plots for 795 nm light for a 2 mm length ^{87}Rb vapour cell on the D_1 line at 75 °C, for magnetic fields of different strengths, with the zero field case shown on the topmost panel. (e) The associated transitions for the 0.62 T case are shown on the Breit-Rabi diagram. For the geometry of this experiment, 795 nm light drives sigma minus, σ^- (purple), and sigma plus, σ^+ (cyan), transitions on the D_1 line. Weak transitions due to the residual hyperfine mixing are shown as thinner lines.

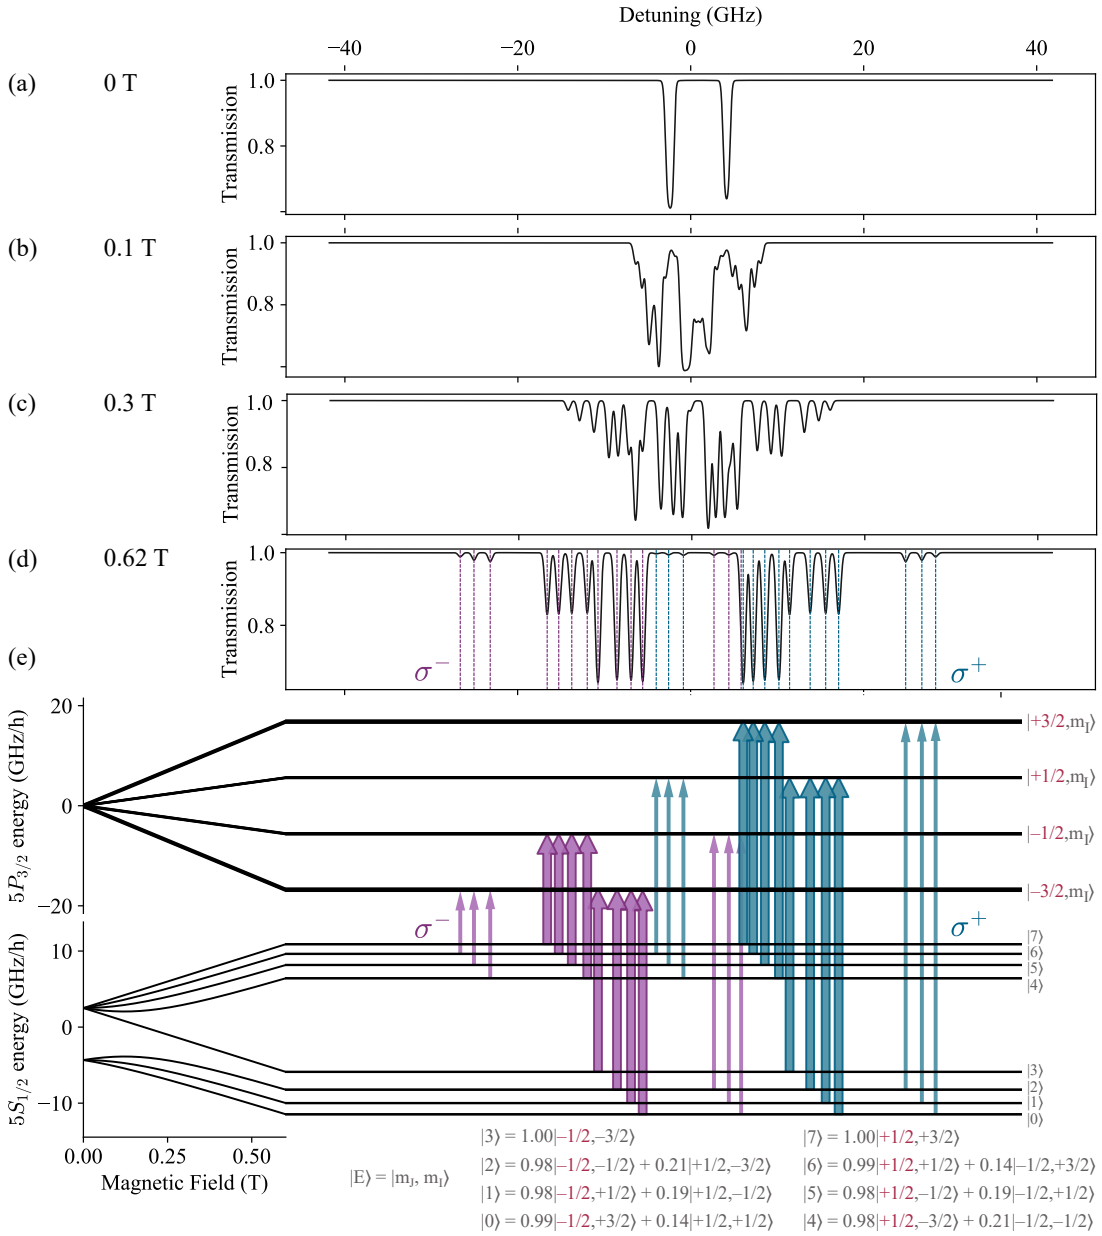

Figure 2.8: (a)–(d) Theoretical transmission plots for 780 nm light for a 2 mm length ^{87}Rb vapour cell on the D_2 line at 75 °C for magnetic fields of different strengths, with the zero field case shown on the topmost panel. (e) The associated transitions for the 0.62 T case are shown on the Breit-Rabi diagram. For the geometry of this experiment, 780 nm light drives sigma minus, σ^- (purple), and sigma plus, σ^+ (cyan), transitions on the D_2 line. Weak transitions due to the residual hyperfine mixing are shown as thinner lines. For the $5P_{3/2}$ state, the 4 $|m_I\rangle$ levels associated with each $|m_J\rangle$ are not visible at this scale.

2.7 Conclusion

In this chapter, we have covered the basic physics underlying atom-light interactions for two-, three-, and four-level systems. We have shown how the absorptive and dispersive features of the seed transmission vary with the addition of energy levels and couplings between levels. We finished by detailing the energy level structure of rubidium, our atom of choice, and explained how that structure changes on applying a magnetic field so that the HPB regime is reached. This will allow us to effectively treat the atom as a four-level system.

Chapter 3

Experimental techniques and apparatus

Section 3.2 of this chapter is based on the following publication:

C. R. Higgins, D. Pizzey, R. S. Mathew and I. G. Hughes. *Atomic line versus lens cavity filters: a comparison of their merits.* OSA Continuum **3** 961 (2020). [50]

In this chapter, we describe ancillary parts of the experimental set-up which were crucial to obtaining the best results from our experiments: (1) It was necessary to install a tapered amplifier (TA) because the power from our coupling/776 nm laser led to a photon pair rate in Chapter 8 which was too low to perform a Hanbury Brown–Twiss experiment. (2) The large number of background photons (See Chapter 6) in our experiments led to the implementation of a bespoke frequency filtration device. (3) A simple cavity is used to calibrate the frequency of the pump/780 nm and seed/795 nm lasers.

To maximise the number of photons collected into the fibres connected to our photon detectors, we designed and implemented a machine learning automation system (the Pi Auto-aligner); this will be covered in the next chapter.

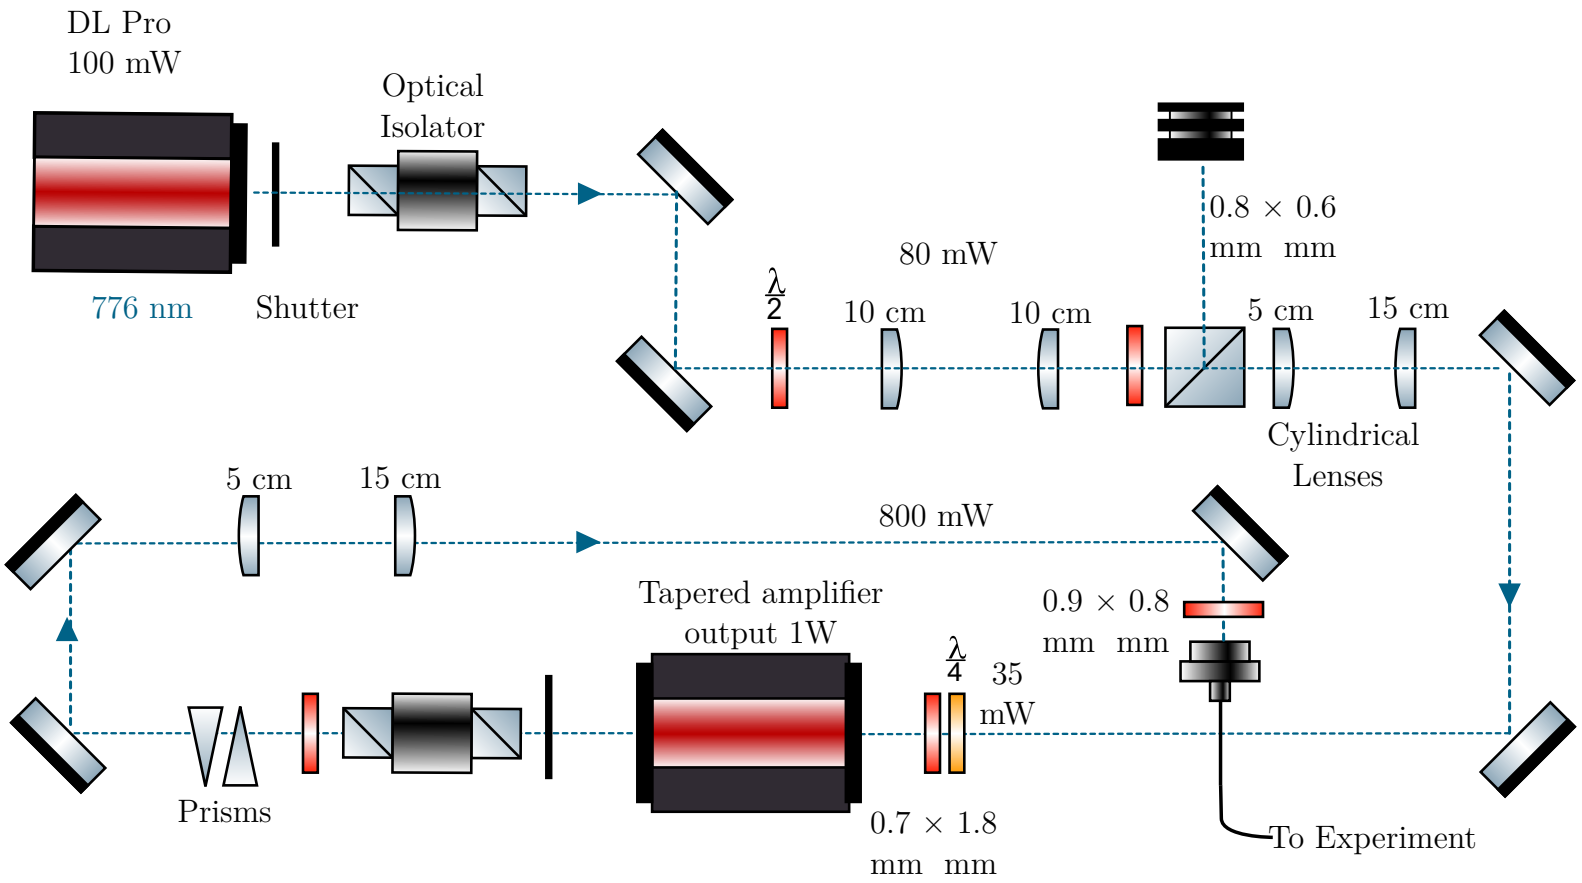

Figure 3.1: Optical table layout for maximising the output of the BoostTA tapered amplifier. The optical isolators after the DL Pro and the BoostTA have isolation levels of 40 dB (Isowave I-80-3L) and 60 dB (Isowave I-80U-4-L) respectively.

3.1 Setting up a tapered amplifier

We found that the Hanbury Brown–Twiss (HBT) experiment of Chapter 8 required a large photon pair rate and that a large coupling power was required to achieve this. We, therefore, installed the optical components needed to maximise the amplified power from a tapered amplifier (TA), the Toptica BoosTA.

A tapered amplifier is a semiconductor device in which a semiconductor chip with a tapered gain region amplifies seed light. The tapering prevents damage to chip from the high light intensities within. Unseeded, the TA behaves as a laser with amplified spontaneous emission (ASE) output from both the front and back of the chip. The beam emerging from the back of the chip can be used as a guide beam: to align the seed light into the TA, we ensure that the seed and guide beams are overlapped.

We achieve an absolute maximum of 1 W straight out of the tapered amplifier. However, to avoid damage from back-reflections, the TA should never be operated without an optical isolator, and so we will always quote the power output once the light has passed through an isolator with a 79% transmission; so, the effective maximum output is 820 mW. We shall see that losses throughout the path mean that by the time the light reaches the experimental cell, the maximum is only ~400 mW. Nevertheless, this is still greater than a tenfold increase compared to without the TA.

Detailed layout

A simplistic picture of TA operation would be that an input of x mW of seed light results in an output beam of $10x$ mW. There are several complications:[1]

1. The TA requires the input beam to have the same profile as the guide beam to ensure that the whole of the gain region is used. The beam must, therefore, be reshaped by a pair of cylindrical lenses before entering the TA.

2. The optical fibre taking the output beam to the experiment requires a particular

[1] The level of detail in this section (Section 3.1) is provided in the hope that it will be of help to future researchers setting up a TA.

profile. Therefore, a pair of anamorphic prisms and a pair of cylindrical lenses are placed before the input of the fibre.

Figure 3.1 shows the set-up required to obtain the maximum transmission from the TA. The light emitted from the DL Pro 776 nm laser passes through an optical isolator and is steered using two mirrors through a pair of 10 cm lenses to collimate the beam. The beam then passes through a $\lambda/2$ waveplate plus polarising beam splitter (PBS) combination to limit the beam power to <40 mW as a seed power above this risks damage to the TA. The beam next passes through two cylindrical lenses of focal length 5 cm and 15 cm respectively. This changes the beam waist from 0.8 mm $\times$ 0.6 mm before the lenses to 0.7 mm $\times$ 1.8 mm before the TA. The approximate 1:3 ratio matches the required input for the TA, as stated in its specification documentation. Before entering the TA, the beam is steered by two mirrors through a $\lambda/2$ and $\lambda/4$ pair as polarisation-matching into the TA is also required.

After being amplified, the beam exits the TA through an optical isolator with 79% transmission. A pair of anamorphic prisms and a pair of 5 cm and 15 cm cylindrical lenses reshape the beam from 1.9 mm $\times$ 0.6 mm to 0.9 mm $\times$ 0.8 mm, which is the size required for input into the optical fibre[2]. The fibre then carries the light to the experiment.

[2]Thorlabs P3-780PM-FC-2 optical fibre with the F220APC-780 fibre collimation package.

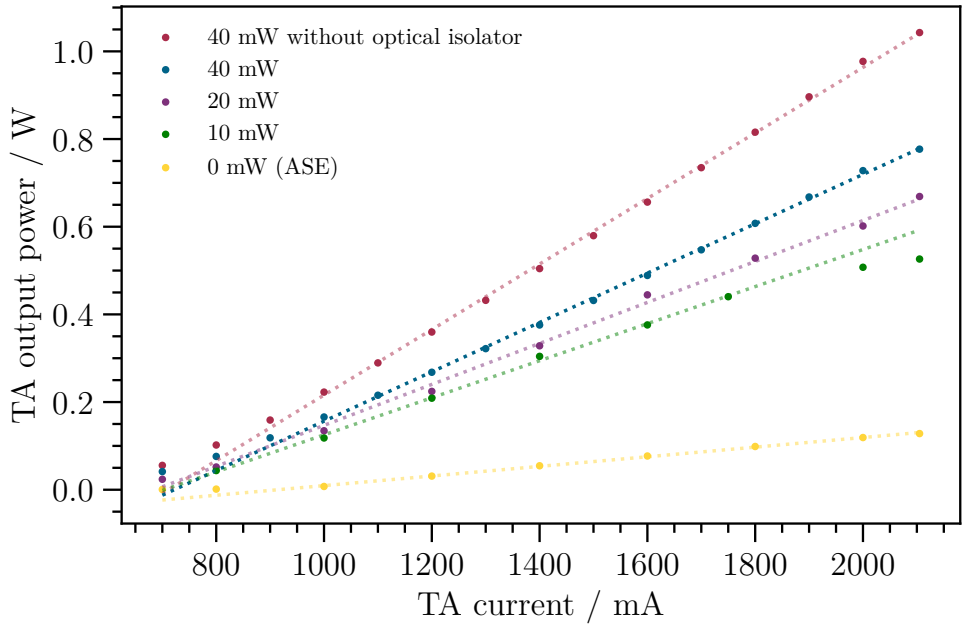

Figure 3.2: TA output power increases linearly with increasing current and increasing seed input power. The data are shown with straight-line fits for three different seed input powers: 40 mW (dark blue), 20 mW (purple), and 10 mW (green). We also show amplified spontaneous emission, ASE (yellow). The data without an optical isolator are also shown for reference (red), although the TA should generally not be operated without the optical isolator.

Characterisation of the TA

Figure 3.2 shows how the output power of the TA increases with current. Data are shown for different seed powers with fits to a straight line. In all cases, there is a linear increase from 800 mA upwards; in the case of 10 mW of seed power (green points), there is a deviation from linearity at the highest currents of 2000 mA and 2100 mA. The yellow points show that in the absence of seed light, the TA does indeed behave as a laser that outputs amplified spontaneous emission (ASE).

Figure 3.3 shows how the output power, P_{out}, of the TA increases with increasing seed power, P_{in}. Data are shown for different TA currents with fits to logarithmic lines, $P_{\text{out}} = m \times \ln(P_{\text{in}}) + c$. This is more evident on a graph with log scaling on the x-axis, as in Figure 3.4. The logarithmic scaling means that little is gained by

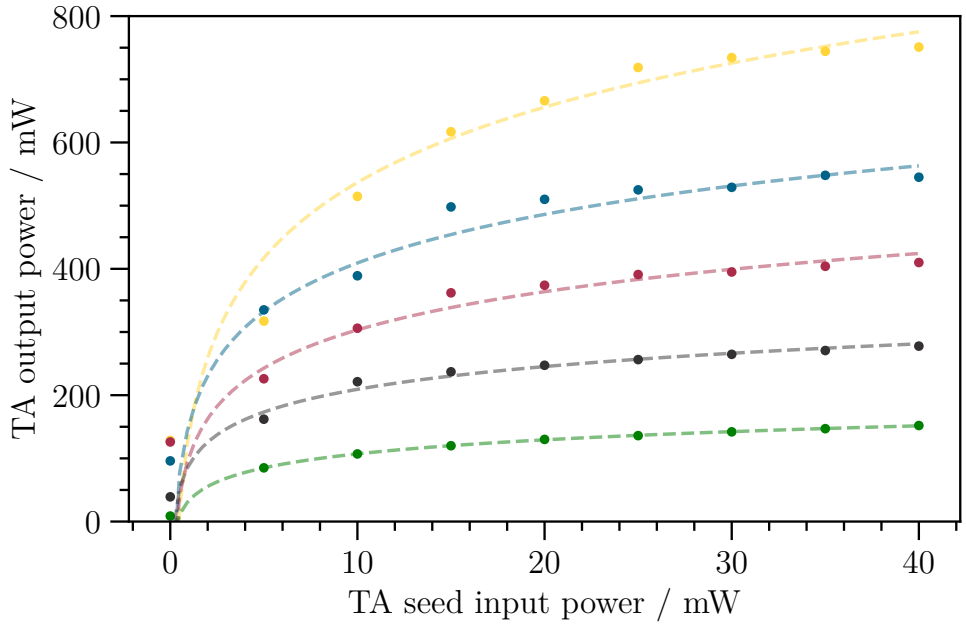

Figure 3.3: TA output power increases with increasing seed power and increasing current. The data are shown with fits to logarithmic lines, $P_{\text{out}} = m \times \ln(P_{\text{in}}) + c$. The TA currents are, from top to bottom, 2100 mA (yellow), 1750 mA (blue), 1500 mA (red), 1250 mA (grey), and 1000 mA (green).

increasing seed power, e.g., there is very little difference in output power between a seed input of 30 mW compared to 40 mW. In order to maximise output power, we, therefore, worked at maximum current (2100 mA) and limited the seed light input to 35 mW because, as previously mentioned, above 40 mW there is a risk of damage to the TA.

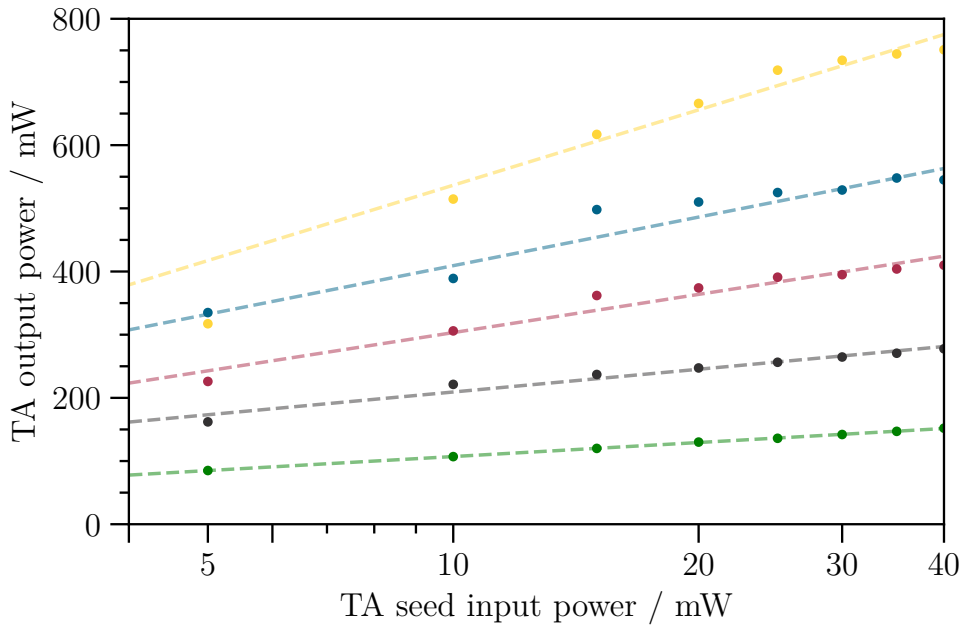

Figure 3.4: TA output power increases logarithmically with increasing seed power. These are the same data as in Figure 3.3, except that the logarithmic scaling is now immediately evident because of the log scale. The data are shown with fits to logarithmic lines, $P_{\text{out}} = m \times \ln(P_{\text{in}}) + c$. The TA currents are, from top to bottom, 2100 mA (yellow), 1750 mA (blue), 1500 mA (red), 1250 mA (grey), and 1000 mA (green).

3.2 A lens etalon spectral filter

Throughout this work, we will find it necessary to separate output light from input light. In most cases, we can use commercially available interference filters because the "noise light" differs from "signal light" by nanometres in wavelength. However, in Chapters 6 & 8, we shall see that our experimental situation creates a large number of background photons that are very close in frequency to the signal light. We therefore found it useful to implement a narrower, sub-nm, bandwidth filter.

We designed and implemented a cavity filter[3] consisting of two highly-reflective (HR) (reflectivity, R~99) dielectric-coated surfaces separated by a fixed length. Light entering the cavity only exits when the light is resonant with the cavity; the result is a periodic set of transmission peaks, as shown in Fig. 3.5. This periodicity can be seen in the spectrum for the 5.0 mm length lens (dotted cyan): both the leftmost and rightmost peaks on the figure are TEM_{00} modes, separated by the free spectral range (FSR) of 19.9 ± 0.1 GHz. The periodicity cannot be seen in the spectrum of 2.5 mm length lens (purple), which is the one we use in the main experiment: this is merely an artefact of not being able to scan the laser over the calculated FSR of 39.7 GHz. The free spectral range is given by

$$\mathrm{FSR} = \frac{c}{2nL}, \tag{3.1}$$

where c is the speed of light, n is the refractive index of the lens ($n = 1.51$), and L is its length (thickness).

Following the method of [69], we used a single high-reflection coated planoconvex lens to obtain the two highly-reflective (HR) surfaces. This design was chosen instead of the more familiar one that consists of separate HR coated mirrors attached to a spacer because our design obviates the need to lock mirror positions, and so is intrinsically stable. Compared to that other design, our planoconvex geometry also gives a higher cavity finesse, $\mathcal{F}$, where $\mathcal{F}$ is given by $(\pi\sqrt{R})/(1-R)$. The presence of the different TEM modes, as shown in Fig. 3.5, also demonstrates

[3]The terms lens "cavity filter" and lens "etalon filter" are used interchangeably.

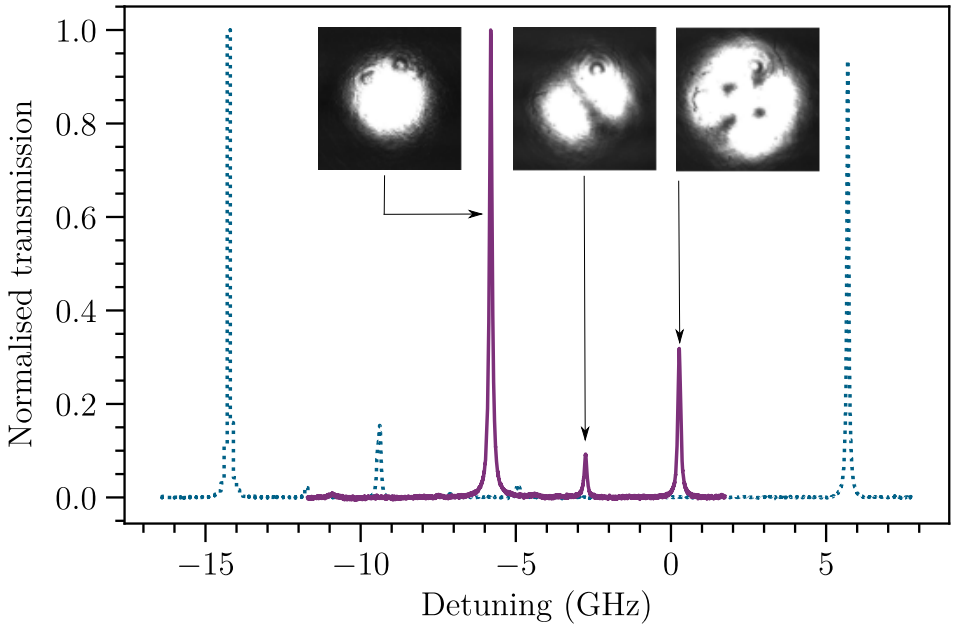

Figure 3.5: Normalised transmission spectrum (purple) for the lens etalon filter of length 2.5 mm used in the experiment. The spectrum (dotted cyan) for a 5.0 mm cavity is shown for comparison. Shown from left to right are CCD images of the TEM_{00}, TEM_{10}, and TEM_{20} modes. Alignment into the cavity is achieved by maximising the transmission of the TEM_{00} mode.[4]

the spatial mode filtering of our design. Any commercially available lenses can be used, provided they are given a high-reflectivity coating; we happened to use the services of Lambda Research Optics Inc.

We tested several lenses, as explained in [50]. For the work in this thesis, we used the 2.5 mm thickness lens, which gave transmission peaks of width 128 MHz, with the resonant frequency being tunable with temperature. The shift of the resonant frequency, ν, is given by

$$\frac{\mathrm{d}\nu}{\mathrm{d}T} \approx -\left(\alpha + \frac{1}{n}\frac{\mathrm{d}n}{\mathrm{d}T}\right)\nu, \tag{3.2}$$

where α and n are the thermal expansion coefficient and refractive index of the lens material respectively, and $\mathrm{d}n/\mathrm{d}T$ is calculated from the Sellmeier function [70][5].

[4]Data & photos taken by Clare Higgins.
[5]Eq. 10 of the technical report produced by SCHOTT (2016) [71] "TIE-19 Temperature Coefficient

We experimentally determined the resonant frequency change with temperature, $\mathrm{d}\nu/\mathrm{d}T$, to be -3.36 ± 0.06 GHz/K, which is in agreement with the value obtained with Eq. 3.2 of -3.37 GHz/K.

The corollary of tunability with temperature is that the appropriate temperature requires stabilisation. As given by Eq. 3.2, for a frequency stability of 10 MHz, we need a temperature stability of 3 mK, which we achieved as follows. The lens tube is screwed into a Teflon covered stainless steel block attached to a Peltier heater & thermistor, and the temperature is controlled with a Koheron TEC100L temperature controller.

Any stress in the cavity glass can lead to birefringence. As the polarisation of light changes, this birefringence induces a shift in the position and shape of the transmitted peaks [71]. Our mounting method minimises such stress, and so we measured negligible frequency shifts with changes in polarisation.

The mode of the laser beam must be matched to the fundamental mode of the cavity. For this purpose, a mode-matching lens is placed upstream of the cavity. It is selected so that the curvature of the wavefront matches the front (convex) surface of the cavity and the beam focuses at the back (planar) surface of the cavity. The beam waist of the fundamental cavity mode is $w_0 = \left(\frac{\lambda n L}{\pi}\left(\frac{r}{nL} - 1\right)^{1/2}\right)^{1/2}$ [72]. The focal length of the mode-matching lens needed is $f = \frac{w_1 w_0 \pi}{\lambda}$ [73], where w_1 is the initial beam waist[6].

It is necessary to carefully align the light into the cavity to maximise the transmission of the TEM_{00} mode. This is accomplished with the use of a pair of steering mirrors before the cavity. A camera may initially be needed to identify the TEM_{00} mode, but alignment thereafter can be performed by visualising the transmission spectrum on an oscilloscope[7].

of the Refractive Index" was used.

[6] Of course, lenses come in fixed focal lengths, and so perfect mode-matching is not possible.

[7] Because of laser beam drift over time, some periodic realignment of this part of the experiment is necessary to maintain maximum transmission through the etalon. Periodic and frequent realignment of laser beams is a recurring theme throughout this thesis. The next chapter, Chapter 4, describes a novel use of machine learning which we designed and implemented to help automate this humdrum activity.

3.3 Laser frequency calibration

There is another optical cavity in our experiment, which we use to calibrate the frequency of the pump/780 nm and seed/795 nm lasers. We follow the method of [74]; here, we describe the method in brief[8].

The laser beam passes through an optical cavity, and the signal is detected on a photodiode. When we "scan" the laser, i.e., change the frequency with time, the peaks in the cavity transmission signal must occur at fixed *frequency* intervals (as given by Eq. 3.1). However, we find that the cavity transmission peaks do not occur at fixed *time* intervals. This indicates that the laser frequency scan is nonlinear with time. However, our knowledge that the peaks occur at fixed frequencies allows us to fit a function which can thenceforth be used to convert from the time of the scan to the frequency of laser at that time. In this way, we linearise the frequency axis for our data. At this point, we still only have a relative frequency axis. To convert to an absolute frequency axis, we simply match our rubidium spectral dip features to the well-known atomic transition frequencies for rubidium vapour.

It should be noted that although we linearised the scan as a matter of course, in practice the original nonlinearity of the lasers is negligible for most purposes except when performing precise fitting with ElecSus. To give a measure of the difference when fitting, there is a difference of $<1\,^{\circ}\mathrm{C}$ in the temperature extracted from a spectrum which has had its nonlinearity corrected and one which does not. We emphasise that this will differ from laser to laser.

3.4 Conclusion

In this chapter, we have shown how to set up and optimise the power output of a tapered amplifier (TA). The TA will be of particular use in Chapter 7, when obtaining data for the measurement of a dipole matrix element, and in Chapter 8, when performing a Hanbury Brown–Twiss experiment. We have also shown how we designed the lens etalon spectral filter; this will be of benefit in creating 'purer' photons in Chapter 8. We ended with explaining the use of an optical cavity to

[8] A fuller description is given in Section 4.3 of [75].

calibrate the frequency of all of our lasers. In the next chapter, we discuss a novel laser-alignment device, which we designed and built, that uses machine learning to automatically optimise the alignment of laser beams.

Chapter 4

Machine learning in quantum optics experiments

The following paper is based on this chapter:

Renju S. Mathew, Roshan O'Donnell, Danielle Pizzey, and Ifan G. Hughes. *The Raspberry Pi Auto-aligner: Machine Learning for Automated Alignment of Laser Beams*, Review of Scientific Instruments **92** 015117 (2018). [49]

4.1 Introduction

One of the difficulties with quantum optics experiments involving single-photons is isolating the photons of interest from the countless number of background photons in the environment. In the spontaneous four-wave mixing experiment of Chapter 8, the photons of interest are the herald photons (762 nm) emitted on the ^{87}Rb transition $|5D_{3/2}, m_J = \frac{1}{2}\rangle \rightarrow |5P_{1/2}, m_J = -\frac{1}{2}\rangle$ and the signal photons (795 nm) emitted on the $|5P_{1/2}, m_J = -\frac{1}{2}\rangle \rightarrow |5S_{1/2}, m_J = \frac{1}{2}\rangle$ transition.

In order to obtain the normalised cross-correlation between herald and signal photons, $g^{(2)}_{\mathrm{h,s}}$, it is necessary to maximise the number of herald photons collected by

one optical fibre and to maximise the number of signal photons collected by another optical fibre. We perform this alignment manually for each fibre by adjusting the orientations of two mirrors[1] so that the photon count rate on our photon-counter is maximised. We first optimise the mirror orientations for the case of seeded four-wave mixing (Chapter 7), where the beams can be seen on an IR card or photodiode. For the purposes of this chapter, we shall call these beams "visible". Then we re-optimise the mirror orientations in the case of spontaneous four-wave mixing (Chapter 8), where the beams can no longer be seen. These beams we shall call "invisible".

Alignment is particularly problematic as we approach the single-photon level. In the case of a visible beam, the tacit knowledge of how to align a beam becomes second nature with practice. However, for an invisible beam, alignment becomes increasingly difficult as the beam becomes arbitrarily weak because of two issues that make the process non-intuitive even after much practice. Firstly, aligning photons onto a photon-counter requires the experimenter to wait, after each change to the mirror orientation, for the photon counts to be integrated. Secondly, the magnitude of the Poissonian noise on the signal can be approximately the same as the signal magnitude.

We demonstrate here that both cases, visible and invisible, can be done automatically via machine learning. For a machine, there is, in principle, no difference between maximising the coupling efficiency of a laser beam, as measured by a photodiode or power-meter, and maximising the coupling efficiency of a beam of photons, as measured by a photon-counter. Although there are several practical complications that must be addressed, our central idea is straightforward: Mirrors direct a beam of light into an optical fibre, and we attach motors to these mirrors; we give control of the motors to a computer, and a machine learning algorithm, M-LOOP [76], optimises[2] the orientation of the mirrors so that the power output of the fibre is maximised. We call the entire set-up a "beam auto-aligner" or "fibre auto-coupler".

[1] Each mirror has a pitch knob and a yaw knob, so throughout this chapter we are dealing with four mirror knobs. In other words, four independent parameters need to be optimised.

[2] For this to work, we must always begin by manually coupling some nonzero amount of light into the fibre.

Commercially available equipment to align just two mirrors costs ~£3000 [3] whereas our device can be assembled for a fraction of the cost[4]. Because commercial devices have hardware and software that are proprietary, the equipment must be used as a black-box, and so there is no room for customisation. Commercial software may need to be bought separately from the hardware and need periodic updating. Most relevantly, at the time of writing, commercial devices cannot align invisible (arbitrarily weak) beams.

4.2 Machine learning (ML)

Machine learning (ML) methods discover patterns in data without requiring any assumptions about the data's structure [77]. Performing research with machine learning began in earnest in the 1980s [78], and by 1992 ML methods were used to, for example, create non-intuitive laser pulse-sequences for excitation of rotational quantum states [79]. However, it is only in the last decade or so that ML methods have begun to be used more widely in the atomic, molecular, and optical (AMO) physics community. It has been used to create self-tuning, mode-locked lasers [80–82]; for automating the production of Bose-Einstein condensation [83]; and for maintaining doughnut-shaped beams in scattering media [84]. ML has recently even been used to create new quantum experiments: the system both learned to create a variety of entangled states and improved the efficiency of their realisation [85].

Surprisingly, to our knowledge, ML has not been used to achieve automatic beam alignment. Whilst it is entirely possible to imagine a science of the future still using waveplates, mirrors, and other such optical elements, it is difficult to imagine all of these being manipulated solely by hand. It would be disappointing if self-driving cars were in widespread use before automation appeared in the labs of laser physicists[5]. Automation is not new in other scientific disciplines, including other

[3] Polaris® Kinematic Mirror Mounts with Piezoelectric Adjusters (~£700 each) & Kinesis NanoTrak®Fiber Alignment Controllers (~£1000)

[4] The cost is dominated by the price of the Raspberry Pi (£65) and the motors (£35 each).

[5] Here we are adapting, and making more general, a comment made by Baumeister et al. [86] about their self-tuning laser.

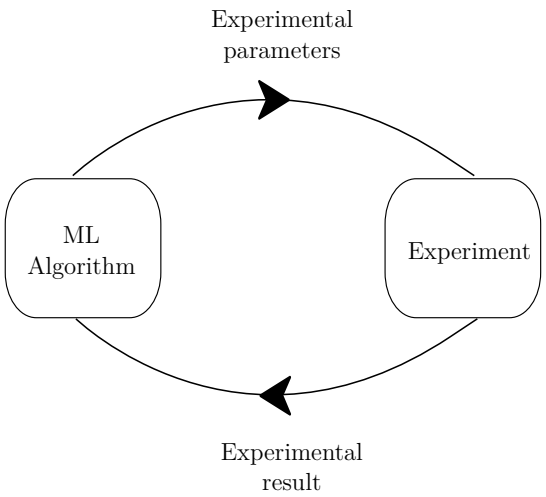

Figure 4.1: Machine learning (ML) feedback loop: The ML algorithm sends experimental parameters for one run of the experiment & the experimental result is returned to the algorithm. On each iteration of this loop, the ML algorithm builds a more accurate internal model so that the correct parameters for the "best" experimental result is found. The "best" case must be precisely defined by the experimenter.

branches of physics. For example, Adams [87] writes that in synchrotron radiation laboratories "the art of remote controlling a hundred-meter-long beamline has been carried to perfection." Recently in chemistry, a robotic lab assistant was developed to automate routine experiments [88]. It does seem a worthy goal to attempt to automate anything that does not require creativity [89]. More prosaically, but no less importantly, automation would be a helpful addition to laser safety protocols. Aligning beams with ML is a relatively simple case from which lessons can be learned about how to apply ML to more complicated cases. So, when thinking about the design of the apparatus, it is useful to think at a high level of abstraction, i.e., what are the most general terms in which to think about this problem? In the most general terms, to automate any part of an experiment, we require a feedback loop, Fig. 4.1, between the ML algorithm and the experiment. The output of the experiment is the input to the ML algorithm. From this input, the ML algorithm improves its internal model of the problem-situation and determines what new experiment would improve its model further. The output of the algorithm is the

set of parameters for the next experimental run. This loop cycles until the "best" set of parameters is found within some tolerance, where the best case is defined by the experimenter. All of this relies on efficient interfacing between the abstract algorithm and the physical experiment.

Many different types of ML algorithms, including genetic algorithms (GAs) [79, 90, 91] and deep neural networks (DNNs) [78, 86, 92] have been used in optics experiments. We used the Gaussian process (GP) algorithm in M-LOOP: M-LOOP was designed to be used in atomic physics experiments [76], and GPs can deal effectively with uncertainty in the model because GPs are a probabilistic technique based on Bayesian inference. This effectiveness with dealing with uncertainty should make them especially useful in beams dominated by Poissonian noise, as is the case of our experiments involving single-photons in Chapter 8.

The M-LOOP algorithm is minimising a cost function, similar to finding the minimum of an error surface. The internal model that M-LOOP builds becomes increasingly accurate with each iteration. Simultaneously, it tries to find the global minimum of the cost function landscape and ends when this is found, or some other halting conditions are met. Further details will be provided in Section 4.4. Unfortunately, GPs also have computational training times that scale cubically with data [93], although this is not a problem in our case as we only have four parameters to optimise.

4.3 Implementation

The implementation in the laboratory involves certain additional subtleties and complexities which we describe here. There are two different builds. The first, "Build 1.0", uses small motors (28BYJ-48 5V Stepper Motors with ULN2003 driver boards), Fig 4.2, and worked for coupling visible, weak beams into a fibre. It did not always work for invisible beams when those beams became arbitrarily weak, necessitating a second build, "Build 2.0", using large motors (17HD48002H-22B Stepper Motors with L298N driver boards), Fig 4.3. The testing of the second build is ongoing. The building of Build 2.0 was done as a master's level project by Roshan O'Donnell [94]; technical details for the build are found in the

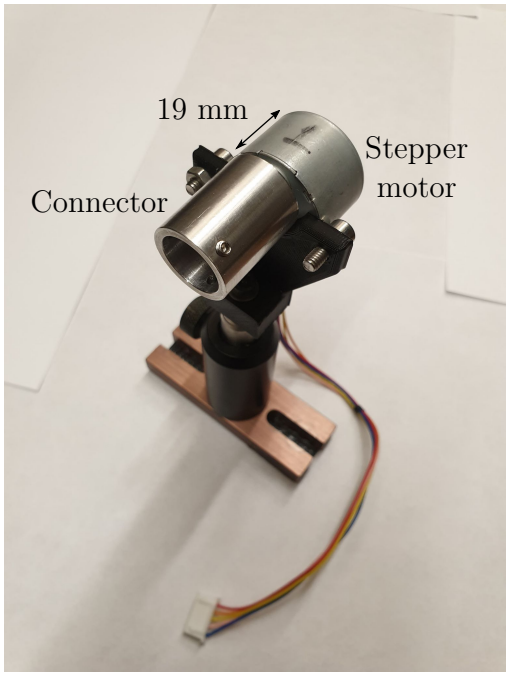

Figure 4.2: Build 1.0. Photo of the small 28BYJ-48 stepper motor (grey colour) with steel connector (silver colour) for connection to stable mirror mount (not shown). (Photo taken by Clare Higgins.)

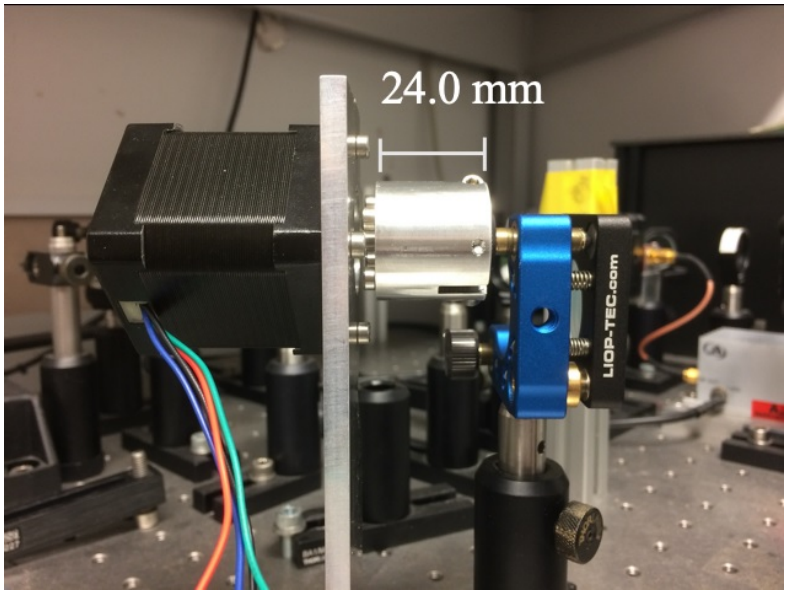

Figure 4.3: Build 2.0. Photo of large 17HD48002H-22B stepper motor (black, left) attached with steel connector (silver, centre) to a stable mirror mount (blue, right). (Photo taken by Roshan O'Donnell.)

supplementary material to the paper associated with this chapter [49]. At the time of writing, Build 2.0 had not yet been tested on invisible beams. To avoid later confusion, we will state here the principal difference between the two builds. Build 1.0 requires manual hysteresis correction; once this is performed, it works as required on visible beams. Build 2.0 requires no manual hysteresis correction because we tested different parts of the system individually, redesigning as required to minimise hysteresis. Build 1.0 was essentially a proof-of-principle construction which demonstrated that our central idea of using ML to automatically align lasers would work. The second version was an attempt to fine-tune the whole system. Nevertheless, Build 1.0 worked better than expected and, provided the hysteresis correction is done, works well for aligning visible beams.

Two motors are used to control the orientation of each mirror so as to change the alignment of the laser beam incident upon it. The computer that we use to both run the ML algorithm and to control the motors is a Raspberry Pi 3 Model B+. The Raspberry Pi has a row of GPIO (general-purpose input/output) pins that allows direct communication with the motors. Furthermore, it is a small device, with dimensions 8.5 cm $\times$ 5.6 cm $\times$ 1.7 cm, that can be placed directly on the optical table.

Unfortunately, because of software driver incompatibility, the Pi does not readily connect to either the power-meter (Thorlabs PM100D) or the photon-counter[6], so we had to use a desktop computer as an intermediary. Python sockets connect the desktop computer and the Pi: the computer (server) sends power-meter/photon-counter readings to the Pi (client) on request. The Pi and the computer can be connected wirelessly, although we did not implement this as we had a conveniently located ethernet port for a wired connection.

This slightly more complex implementation is diagrammatically represented in Fig. 4.4 and Fig. 4.5. We now discuss various specific details of the implementation.

[6]Excelitas SPCM-AQRH Single Photon Counting Module in combination with a SensL HRM-TDC High Resolution Time-to-Digital Converter Module; the timing module is required because the photon-detector cannot be directly connected to the Raspberry Pi.

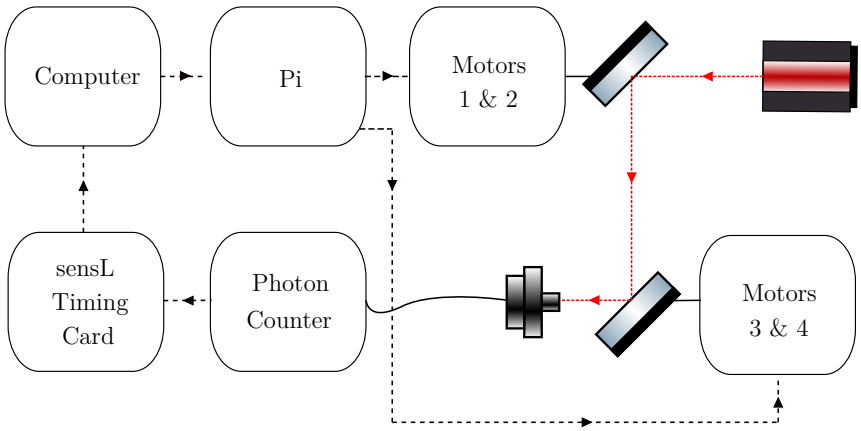

Figure 4.4: Schematic of laboratory implementation. A desktop computer acts as an intermediary between the photodiode/power-meter/photon-counter and the Raspberry Pi (which runs the ML algorithm). Stable mounts are required to secure the mirrors. Electrical paths are shown as black dashed lines; the laser beam path is shown in red.

Electronics

An analogue-to-digital converter (ADC), the 16-bit ADS1115 ADC, is used to convert the analogue output of a photodiode (PD) to the required digital signal input required by the Pi.

Shift registers (SRs) allow the relatively small number of GPIO (general-purpose input-output) pins on the Pi to control an arbitrary number of drivers/motors. The 74HC595 shift register is able to operate two stepper motors. Two shift registers are connected in series to operate four motors.

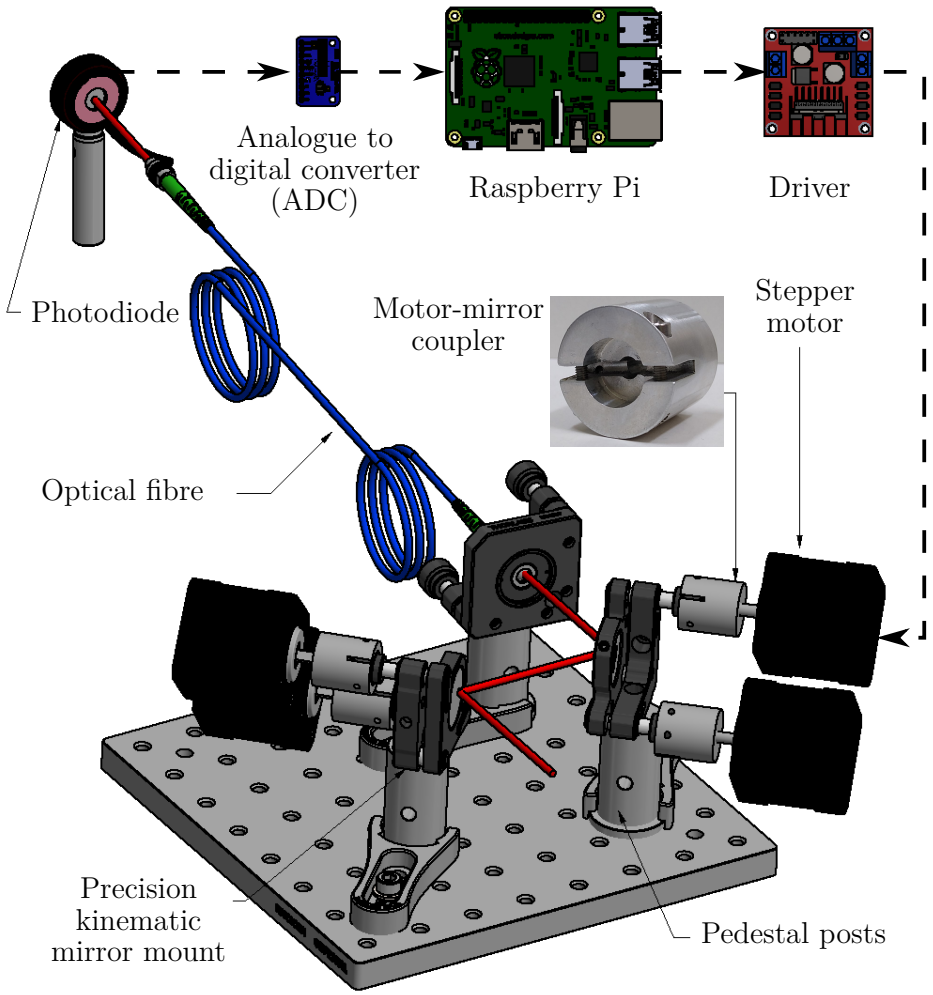

Figure 4.5: Illustration of concept. A Raspberry Pi controls stepper motors attached to mirrors, and a machine learning algorithm automatically aligns the laser beam so as to maximise the optical power coupled into a fibre. Also shown are the motor-mirror coupler connections to connect the mirror mount knobs securely to the motors. Inset is a photograph of the motor-mirror coupler. Support mount for the motor-mirror couplers are not shown. For scale, the breadboard that the beam auto-aligner is mounted on is 115 mm × 115 mm. Details of the different components are given in the main text.[7]

[7]Adapted from figure produced for paper by D. Pizzey.

Motors and mirrors

A motor is attached to the knob controlling the pitch of the mirror mount, and another is attached to the yaw knob. A set of pitch & yaw values sets the orientation of the mirror mount, hence changing the alignment of the laser beam incident upon it. A bespoke 3-D printed motor-to-mirror connector placed on the mirror-mount knob secures the motor in position. The 3-D printed components degrade within weeks of use, so our workshop machined the final versions in steel. We also created 3D-printed stands for the physical driver boards to electrically insulate them from the optical table.

We found that it was difficult to return the mirrors to their starting orientation if standard Thorlabs mirror mounts, posts, and bases were used. This problem was solved using stable mirror mounts (SR100-100-2-BU) from Photonics Technologies (LiOp-Tec) and stable pedestals (RDS-MNI-P-75) and holding forks (RDS-MNI-HF-M) from Radiant Dyes Laser.

Hysteresis

With Build 1.0, we found that we always needed to account for some amount of mechanical hysteresis: If the motor turns N steps[8] clockwise followed by N steps counterclockwise, the mirror mount knob does not return to its starting position. The coupling efficiency of the laser into the fibre changes correspondingly. This is because during the shift in direction (either clockwise to anticlockwise or vice-versa) there are effectively a number of steps where the motor turns, but the mirror knob does not[9].

For Build 2.0, we redesigned the interfaces between the motor-to-mirror so that the hysteresis is negligible. For Build 1.0, we need to calibrate each motor-knob pair to determine the number of extra steps required for a directional change. We then

[8]The relationship between the number of motor steps and the amount of mirror knob rotation depends on the particular mirror and motor used. Build 1.0 (with Thorlabs Mirrors): ~360 steps per 360° rotation; Build 2.0 (with LiOp-Tec mirrors): ~200 steps per 360°. In practice, it is always the number of steps that are the relevant quantity.

[9]The word "effectively" indicates a simplifying model: although the actual cause of the hysteresis may be more distributed across many parts of the apparatus, it can be treated and corrected by assuming the cause is at the motor alone.

correct for this when we send commands to the motor. This is a straightforward procedure: We begin by coupling some light into the fibre and measuring the power output using a power-meter. The Pi sends a command to the motor (via its driver) to turn a mirror knob by a full 360° clockwise rotation, in increments of 1 motor step. No changes are made to the other mirror knobs. Then, the Pi send a command to the motor to turn the original mirror knob in the reverse direction by the same number of steps. At each step, the power is recorded. The results are shown in Fig. 4.6, where the lack of overlap between the two cases indicates the hysteresis correction required. In the case of the data shown, nine motor steps needed to be added whenever the motor changes from clockwise to anticlockwise (and vice versa): this is the x distance between the peaks of the two curves. Once this is done, rerunning this experiment gives two curves which exactly overlap. This procedure is performed for each of the four mirror knobs separately.

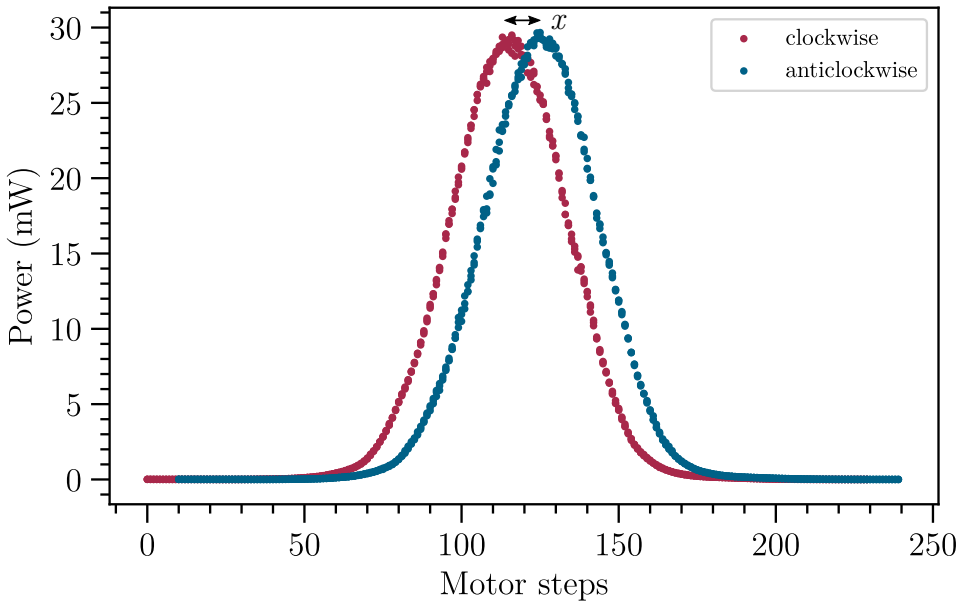

Figure 4.6: Before hysteresis correction, if the motor turns the mirror knob clockwise from 0 to 240 steps (in increments of 1 step) and then turns the mirror knob anticlockwise by the same number, the mirror does not return to the same position. By use of the curve above, hysteresis correction can be included in the code by adding extra steps equal to the distance, x, between the peaks each time the motor changes direction.

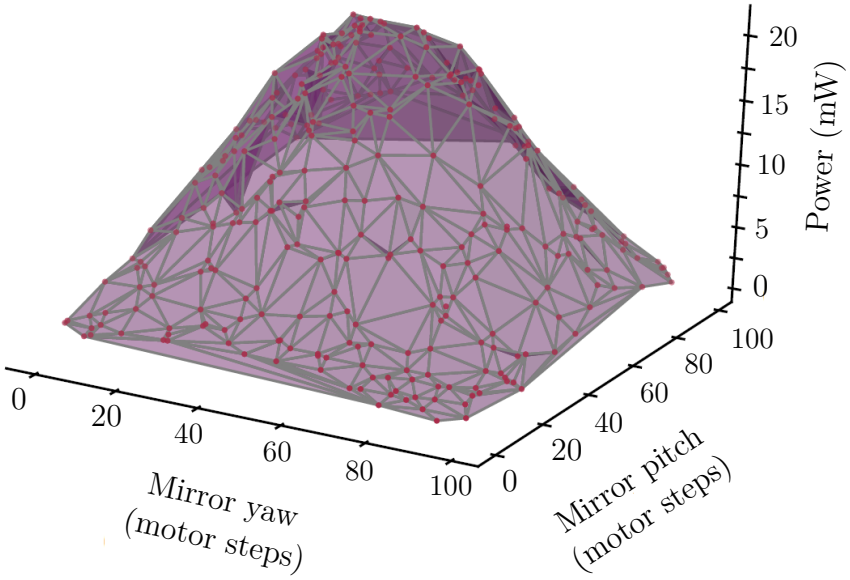

Figure 4.7: Starting by manually achieving maximum fibre coupling efficiency, the motor moves the mirror to a random orientation (a red point) and immediately moves it back to the starting position (the top of the landscape). This is repeated one thousand times. A landscape is mapped out, as shown by the gray lines and purple shading.

Reproducibility

It is necessary to be able to *reproducibly* put the mirror into an arbitrary orientation. To show this, we do a simple test whereby we cycle back and forth between the starting orientation and some random orientation for one thousand iterations. We test the reproducibility of this by using the mirror to couple light into a fibre. We first manually couple light into our fibre until maximum coupling efficiency has been achieved and use this as our starting mirror orientation. The results are shown in Fig. 4.7, where the top of the landscape is the starting point. The fact that this procedure has mapped out a Gaussian landscape of fibre output powers shows that the motor-mirror set-up is working as required. If the hysteresis correction steps are inaccurate, this landscape is not created because the mirror will, with each iteration, drift away from the starting orientation.

4.4 Using ML to maximise fibre coupling efficiency

To couple light into a fibre, we require a fibre and fibre collimator unit in a fixed position and two mirrors, each with a pitch knob and a yaw knob, to direct the light into the fibre. We begin by directing some amount of light into the fibre. Then, once we have connected four motors and determined and corrected the number of hysteresis steps for each motor-knob pair, we are in a position to allow M-LOOP to determine the best mirror orientations.

M-LOOP is used as follows. There is some "cost", which is simply a number that the algorithm attempts to *minimise*. We are interested in *maximising* the power out of the fibre, P_{out}, so we can simply set the cost to equal the negative of the power, $-P_{out}$. M-LOOP is initialised with the boundary values of the parameter space it should explore. This will initially require some trial and error. In the following, these values are normalised so that the parameters only ever take a value between 0 and 1. We have four parameters: a pitch (parameter 1, p_1) and yaw (parameter 2, p_2) for mirror 1 and a pitch (parameter 3, p_3) and yaw (parameter 4, p_4) for mirror 2. There is some unique set of values (p_1, p_2, p_3, p_4) where the coupling efficiency is at its maximum.

M-LOOP begins by outputting a file containing an initial set of parameters ($p^{(i)}{}_1$, $p^{(i)}{}_2$, $p^{(i)}{}_3$, $p^{(i)}{}_4$) for which it wants[10] the associated cost. The Pi sends commands to the motors to set each mirror knob to these values. The Pi then requests and receives from the server the power on the power-meter, P_{out}. The negative of this number ($-P_{out}$) is written to a cost file. M-LOOP waits until it detects this cost file and writes a file containing new parameters and the process repeats.

The initial set of parameter values is random. However, on each run of the above, M-LOOP begins to build an internal model of the parameter landscape. From this, it begins to test the parameter values most likely to minimise the cost.

[10] We will be using some anthropomorphic language throughout: there is no suggestion that M-LOOP is actually thinking; it is simply more intuitive to explain M-LOOP's behaviour *as though* it were an agent. This note is included as we are aware that such language has caused confusion in other branches of science. [95]

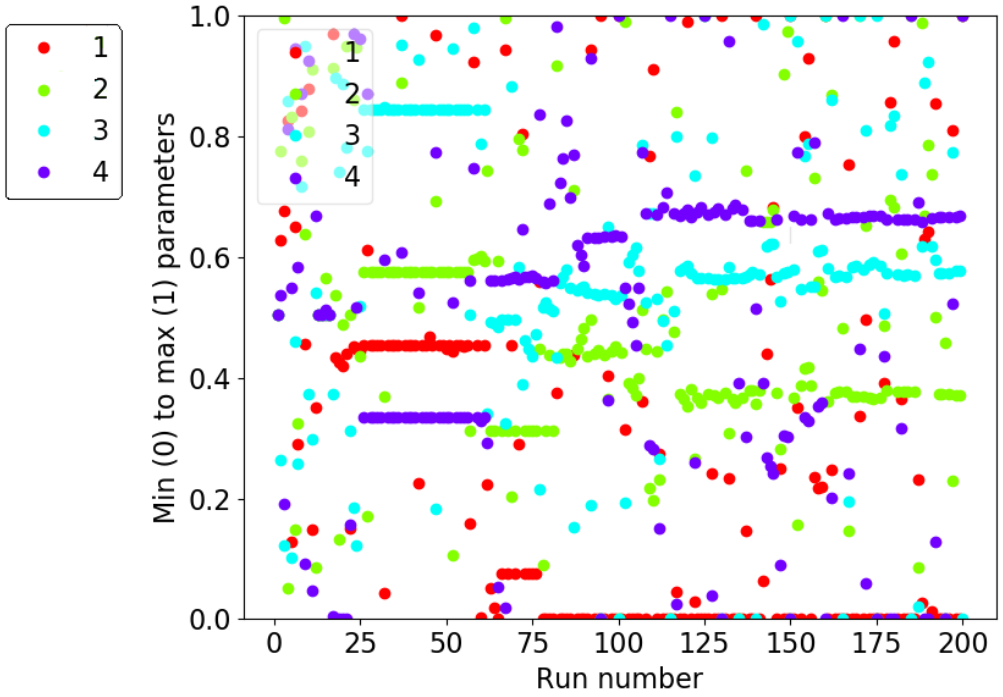

Figure 4.8: Direct output from the M-LOOP algorithm. The legend on the direct output is not easily visible; an extra legend has been added by the author (top left). The y-axis is the normalised number of turns on a mirror knob where 0.5 means no change from the original position. After a certain number of runs, M-LOOP homes in on the best parameters. For this dataset, the best value for parameter 1 (red) lay outside the range we allowed M-LOOP to explore.

M-LOOP finishes when one of three certain conditions are met; these are set by the experimenter before the first run and are as follows: (1) The maximum number of runs has been reached. (2) The minimum cost value has been reached. (3) A certain number of runs has elapsed without a lower cost being found.

Figure 4.8 is a direct output from the M-LOOP algorithm. Better formatted versions of these types of plots will appear later in this chapter, but it seemed to us that it might be helpful for those working with M-LOOP for the deciphering of a direct output to be explained as the direct output is difficult to read. A careful study of the output gives insight into what M-LOOP is doing at each run and gives indications on how to debug. For example, we can see that for the first 25

runs, a random set of parameters is tried. This is typical though the number of runs during this random search varies. By run 25, M-LOOP has homed in on an optimum set of parameter values, as shown by the plateaus for all four parameters. It continues testing the nearby parameter space but regularly tries completely different parameters to ensure that it is not stuck in a local cost minimum. Indeed, by run 70, it has discovered a different set of optimal parameters. However, by run 77, parameter 1 (red) is at the very edge of its constraint. The algorithm wants to explore values below 0, but our initial boundary conditions do not allow it.

This is also clear in Figure 4.9, which again is a direct output from the M-LOOP algorithm for the same experimental set, showing the predicted landscape of cost against parameter value. The minimum of the parameter 1 curve (red) seems to lie just outside of the range that we allowed M-LOOP to investigate, as can be seen if we imagine extrapolating the red curve into the region of the negative x-axis. Nevertheless, the smooth curves for each of the four parameters demonstrate that our program is working as expected; M-LOOP takes approximately 20 minutes from start to finish for 200 runs.

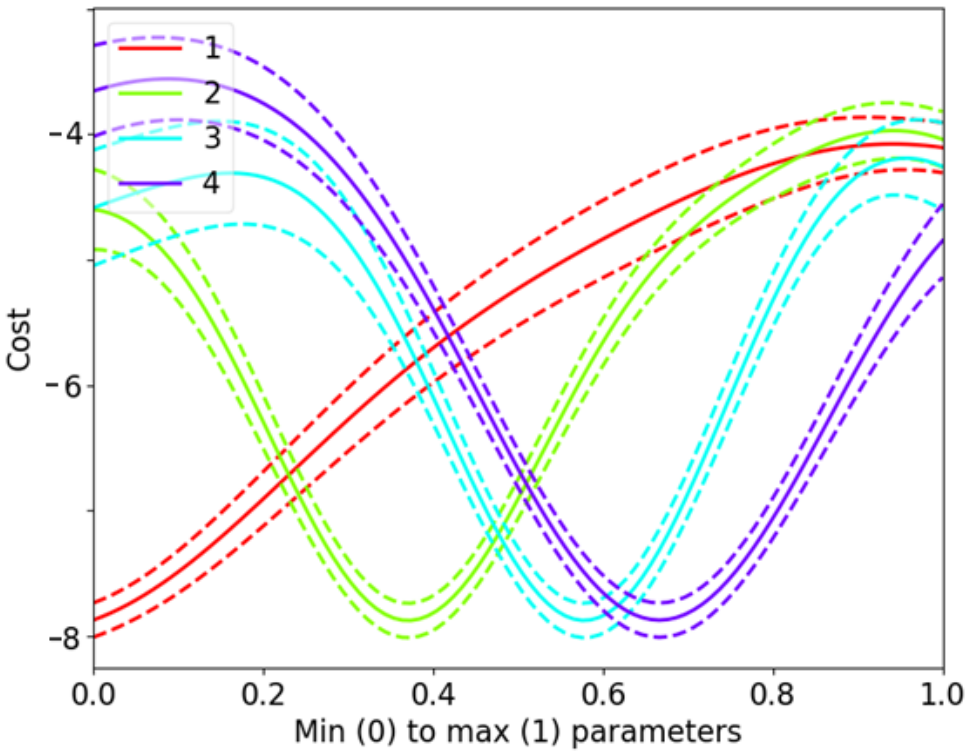

Figure 4.9: Direct output from the M-LOOP algorithm. With a sufficient number of runs, M-LOOP builds an internal model of the parameter landscape. The dotted lines indicate error boundaries as determined by M-LOOP. Smooth curves indicate that the device is functioning as expected. The cost is simply the negative of the power recorded on the power-meter. The minimum cost (maximum power) for all parameters except 1 have been found. The minimum cost for parameter 1 appears to lie outside the range we allowed M-LOOP to explore.

4.5 Using ML to maximise photon count rate

Once the system is optimised for visible beams incident on a power-meter, it is, in principle, straightforward to do the same for maximising counts on a photon-counter with an attenuated beam.

In practice, we found there was a greater proneness to failure in the case of using M-LOOP for photon counting. The figures in Figs. 4.10 and 4.11 show various scenarios.

Figures 4.10a and 4.10b show what happens when we start the machine learning after manually aligning the fibre. In Figure 4.10a, as expected, M-LOOP does a random exploration of the parameter space for 10 runs and quickly finds that the original location, {0,0,0,0}, is optimal. It nevertheless tests to see if there are any better values for these parameters; on run 55, it decides that indeed {0,0,0,0} is the best case. Although it continues to test further, it does not find better parameters. The fact that the points from runs 70 to 80 are not exactly overlapping is an indication of the mechanical creep between runs. Figure 4.10b is the case where the auto-coupler works exactly as it was designed to do. As before, we start at a position we *think* is the best. Between runs 15 to 50, M-LOOP agrees with us until it discovers an even better set of parameter values. In the remaining figures, we do *not* start at a position we think is best. Instead, we make sure there is enough light entering the fibre such that the photon-counters detect some light. Fig. 4.10c shows M-LOOP discovering three sets of good parameters, the first set between runs 10-30, the second between 80-140 and the third between 140-200. By checking the costs associated with each, we can confirm that the last set is the best, i.e., it gives maximum number of photons. Fig. 4.10d shows M-LOOP discovering a good set of parameters between runs 10 to 20, then looking for a better set, but not finding any before reaching the maximum number of runs, which was 100 in this case. M-LOOP takes approximately 10 minutes from start to finish for 100 runs.

Figs. 4.11 show failure cases where M-LOOP has to be reset and run again. In Fig 4.11a M-LOOP seems to find a good set between runs 15-20 but is not sure of the correct value for the parameter corresponding to blue. We find in these cases, as

with Fig 4.10d, that no changes need to be made to any M-LOOP configuration and simply rerunning the procedure will result in a good set being found.

In Fig 4.11b, however, it is clear that M-LOOP wants to look outside the boundary conditions we have given it for both red and blue. In this case, we must rerun M-LOOP having expanded the boundaries within which it is allowed to search.

It is not so clear in Fig 4.11c and Fig 4.11d why M-LOOP has been unable to find the optimal parameters. The solution is usually to restart the whole procedure but set the starting parameters to some sensible new values, e.g., the parameter set at 55 for 4.11c and run 20 for 4.11d.

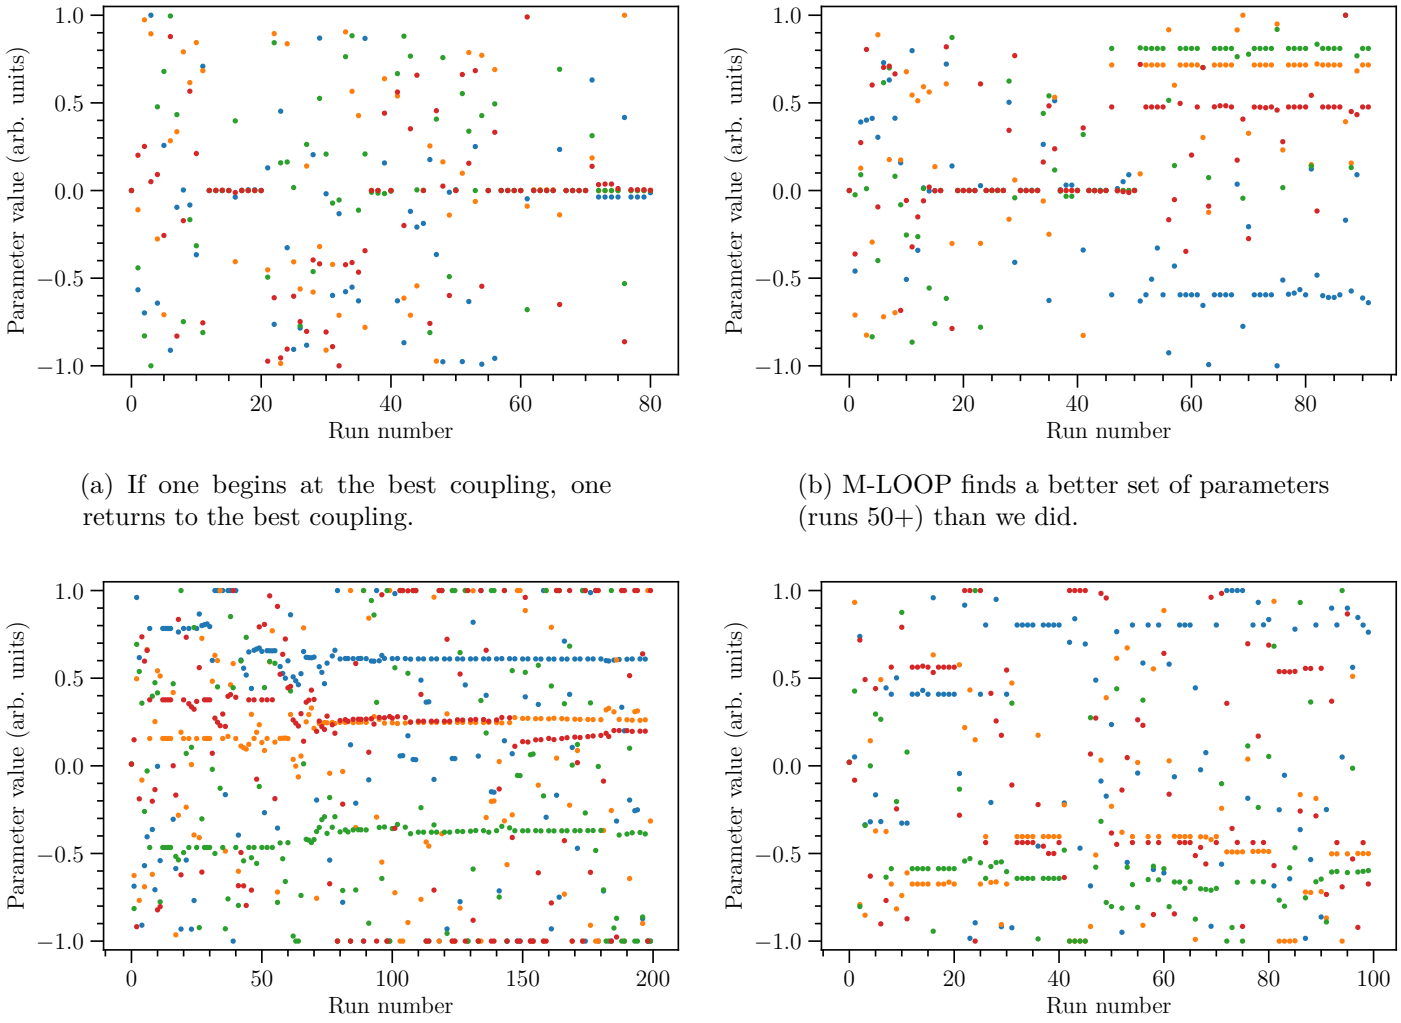

(a) If one begins at the best coupling, one returns to the best coupling.

(b) M-LOOP finds a better set of parameters (runs 50+) than we did.

(c) Three good parameter sets found.

(d) Parameters are found and then lost.

Figure 4.10: These figures illustrate M-LOOP in action. Parameters p_1, p_2, p_3, and p_4 are represented by blue, red, orange, and green, respectively.

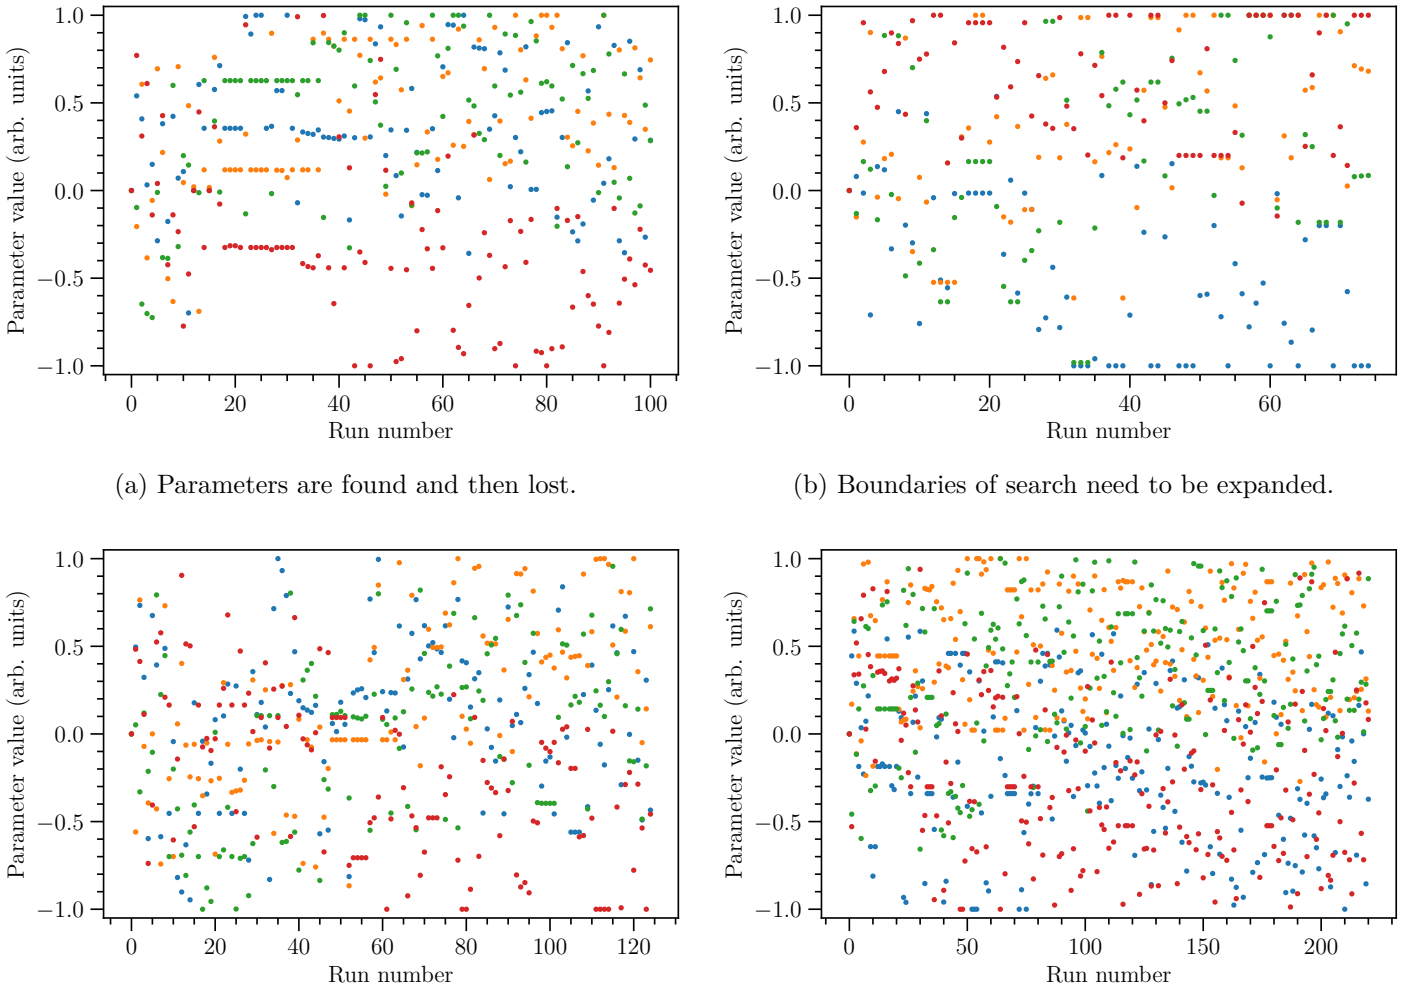

(a) Parameters are found and then lost.

(b) Boundaries of search need to be expanded.

(c) Not able to find optimal values for all parameters simultaneously.

(d) Optimum value for one parameter (red) cannot be found.

Figure 4.11: These figures help troubleshoot M-LOOP. Parameters p_1, p_2, p_3, and p_4 are represented by blue, red, orange, and green, respectively.

4.6 Discussion

Beam-alignment using four motors was consistently achieved within twenty minutes[11]. Both Build 1.0 and 2.0 are able to couple visible light. Build 1.0 occasionally fails for coupling invisible light and, indeed, this occasional failure was our motivation for creating the second build. The main difference between 1.0 and 2.0 is that the latter has, by design, a negligible amount of hysteresis. This was achieved by characterising each part of the new device individually and then redesigning parts as required. This involved rebuilding the motor-mirror connectors and mounting the motors more securely[12].

The general method we have used is that the output of the experimental signal becomes the input for the ML algorithm, and the output from the algorithm are the parameters for the next experiment. This procedure is iterated until the "best" signal is found. This method can be used for other experimental variables, provided they are electronically controllable. Electronic controllability has its weakest point at the physical interface between the electronics and the experimental equipment. We designed and built bespoke interfaces between motors and mirror mounts. If the described ML procedure were to be used to control voltage for something other than turning a motor, e.g., an AOM detuning or power, the interfacing should be less of a problem.

Build 2.0 suffers from a larger footprint than Build 1.0 because of the combined size of the motors and motor holders; nevertheless, it is comparable in size to the commercial alternative[13]. In comparison with the commercial product, our machine is easily customisable. All parts could be readily exchanged with alternatives. This includes changing the ML algorithm at the heart of the device. So, although we found the use of GPs useful because of the ease with which it deals with experimental uncertainty, GPs have the weakness that computational time increases considerably with increasing numbers of parameters. Genetic algorithms (GAs) do not have

[11] To avoid any confusion, we restate here that, to begin the procedure, some nonzero amount of light must first be coupled into the fibre manually.
[12] Technical details are available in the paper associated with this chapter [49].
[13] 121 mm×60 mm×47 mm for the Thorlabs Fiber Alignment device and 83 mm×100 mm×125 mm for Build 2.0 [94].

this particular problem and are a commonly-used alternative [79, 80, 90]. GAs work analogously to evolution by artificial or natural selection: they select the best parameter sets from a large population, recombining and creating a new generation, and then iterating the process.

For future versions of this machine, a piezoelectric device could be added to allow for even finer control of beam alignment. However, these devices require high voltage supplies which are bulky; they are also expensive, as is the device itself—these being the reasons we have not yet implemented them. The driver is more important for precision than the stepper motor, and so it is worth investing in better quality drivers. Our method is easily scalable and could also be used for continual auto-optimisation. It is much easier to return a signal to maximum once the signal has been found because only a small, local parameter space needs to be explored. So, for example, the device could be set to perform auto-alignment at a set time every morning before the experimenter enters the lab. Remote alignment could also be easily performed as we have shown that the Pi can be controlled via the intranet using the commonly used ssh protocol, which is inbuilt into the Pi's Linux operating system.

4.7 Conclusion

Aligning laser beams manually is both labour- and time-intensive. In this chapter, we have demonstrated how this process can be automated with our Beam Auto-aligner. Build 1.0 of our device works well enough to be used for aligning visible beams detectable by photodiodes. Build 2.0 is still being tested for use in aligning invisible, arbitrarily weak beams, detectable only by photon-counters. Depending on the number of optical components on the table, it is easy to underestimate how much researcher time is devoted to aligning and realigning lasers, so Build 1.0 alone is a useful addition to the lab.

Chapter 5

Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime

This chapter is based on the following publication:
Renju S. Mathew, Francisco Ponciano-Ojeda, James Keaveney, Daniel J. Whiting, and Ifan G. Hughes, *Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime*, Optics Letters **43** 17, pp. 4204-4207 (2018) DOI:10.1364/OL.43.004204

Many experiments in quantum optics require laser frequencies to be highly stable over an extended period of time; in particular, we find that Hanbury Brown–Twiss (HBT) experiments[1] require two lasers to have stable frequencies for at least 24 hours. However, because we need to stabilise the laser frequencies (commonly known as laser locking) by using an excited state transition in the hyperfine Paschen–Back (HPB) regime, none of the techniques in general use was immediately suitable for

[1] See Chapter 8.

our purpose. We therefore devised our own method of locking to a two-photon resonance.

In this chapter, we demonstrate a technique to simultaneously lock two laser frequencies to each step of a two-photon transition, in the presence of a magnetic field large enough to gain access to the hyperfine Paschen–Back regime. A ladder configuration with the $5S_{1/2}$, $5P_{3/2}$ and $5D_{5/2}$[2] terms in a thermal vapour of isotopically enriched (>98%) ^{87}Rb is used. Both lasers remain locked for more than 24 hours. For nearly all measured time scales, we measure a frequency instability for the sum of the laser frequencies of less than the rubidium D_2 natural linewidth of 6 MHz.

5.1 Introduction

Stabilisation of the optical output frequency of a laser, i.e., laser locking, is necessary in many research areas. It is especially important in atomic physics where the requisite absolute stability can often be sub-MHz because it is determined by the width of atomic resonance lines. A myriad of methods is available for on- or near-resonant locking. However, performing thermal vapour experiments in the hyperfine Paschen–Back (HPB) regime [6, 18–25], where the atomic resonances are Zeeman-shifted by tens of GHz[3], necessitates the development of new laser locking methods.

Currently available methods for on- or near-resonant locking include locking to wavelength meters [96], stable optical cavities [97], and beat-note locks [98, 99]. In atomic physics research, lasers can be stabilised to particular atomic resonance lines. For generating dispersive error signals with zero crossings, there are a variety of spectroscopic techniques that can be used. Those that require external modulation of the laser include frequency-modulation (FM) [100] and modulation transfer (MT) spectroscopy [101, 102]. Methods that do not require external modulation

[2]For four-wave mixing experiments, we use the $5D_{3/2}$ level. The only difference for the laser lock is that the error signal is a little smaller due to the weaker line strength, requiring tweaking of the laser PID (proportional-integral-derivative) controller setting. Alternatively, the coupling power can be increased.
[3]Chapter 2.6.

are simpler to implement. Such methods include dichroic atomic vapor laser locking (DAVLL) [103–106], polarisation spectroscopy [107], saturated absorption spectroscopy [108], and prismatic deflection [109]. The Faraday effect can also be used to form an off-resonance laser lock [110], as can Zeeman-shift based locking (ZSAR) [111]. Both of these have the advantage of being tunable over a wide range. At large detunings, frequency stabilisation can be achieved using saturation absorption spectroscopy [112] or with a low-quality cavity technique [113]. An interesting alternative to all of these schemes is to avoid external locking altogether by placing atomic media in the external cavity feedback of the laser, causing the laser to be intrinsically stabilised [114].

If there are several lasers, as in the case of atomic systems where multiple levels are coupled, the excited-state transitions can also be used as locking signals. These include locks based on fluorescence detection [115] and electromagnetically induced transparency (EIT) [116–118]. Excited-state polarisation spectroscopy can also be performed with [119] and without [120] a small magnetic field. As we perform our experiments in a large magnetic field, none of the techniques mentioned so far are immediately suitable for our purposes. Furthermore, quantum optics experiments [121–126], including those reported in Chapter 8, often require long integration times [2], meaning that lasers may need to remain locked for a day or more and many of the mentioned techniques are not suitable for such large time periods.

We shall present a technique to simultaneously lock two lasers to two transitions that form a ladder-type excitation scheme (see the top right of Fig. 5.1). We thus stabilise the sum of their frequencies over a timescale of hours. In the presence of the large magnetic field (0.62 T), both transitions are significantly Zeeman-shifted from their zero-field frequencies. Apart from being able to work in large fields, there are at least two other advantages of our scheme: (1) tunability on the first step of the excitation and (2) the lock compensating for drift in one laser by automatic adjustment of the other.

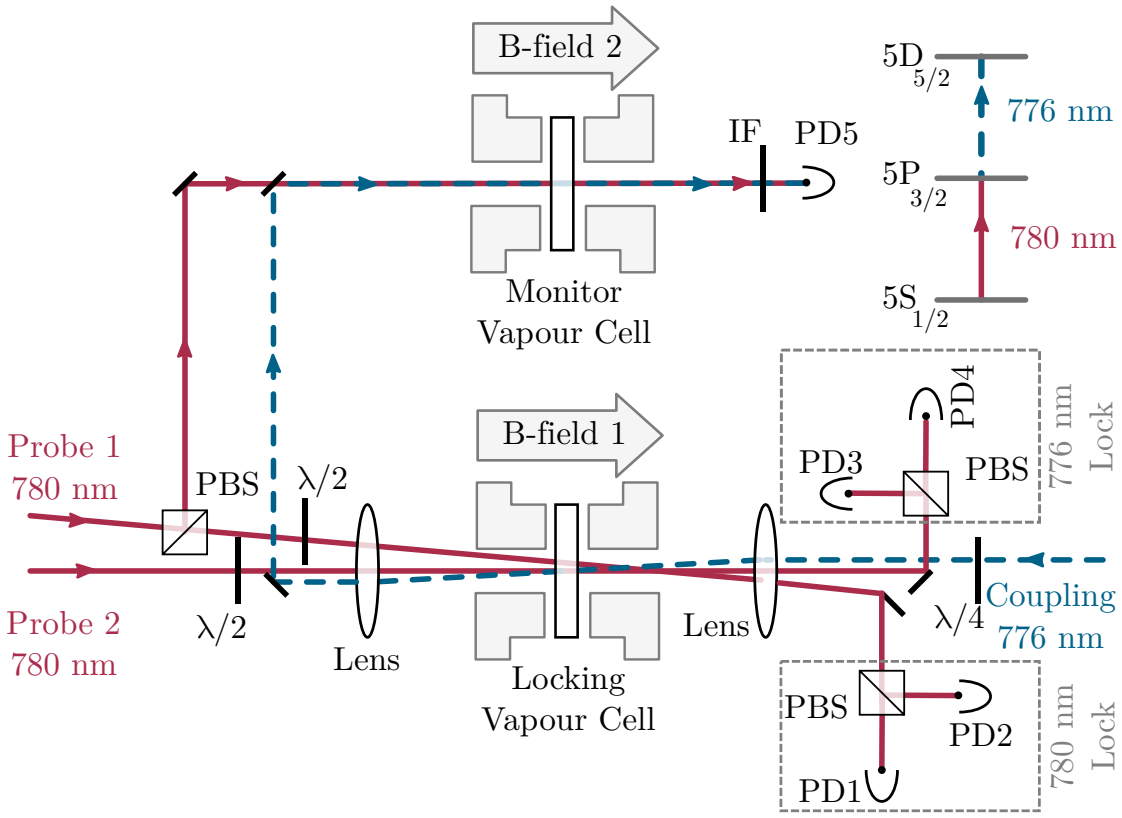

Figure 5.1: Schematic of the experimental configuration (read from left to right). Probe beams (solid red lines) and coupling beam (dashed blue line) are counter-propagated and focussed through the locking vapour cell. The 1 mm^3 cell contains isotopically enriched ^{87}Rb and is in a uniform magnetic field of up to 0.62 T along the optical axis. Only the coupling beam and probe beam 2 are overlapped within the cell. Angles are not to scale. Beam polarisations are set by half- and quarter-waveplates ($\lambda/2$ & $\lambda/4$). By subtracting the signals from the photodiodes (PD) at the output of a polarising beam splitter (PBS), the Stokes parameter S_1 is measured. PDs 1 & 2 are used for the 780 nm lock, and 3 & 4 are used for the 776 nm lock. The monitor vapour cell, interference filter (IF), and PD5 are used to monitor the atomic resonance of experimental interest. The relevant rubidium energy levels are shown on the top right.

5.2 Concept

We start by using the off-resonant Faraday-rotation method described in [110] to stabilise the 780 nm probe[4] laser. We then use a novel Faraday EIT method to stabilise the 776 nm coupling laser. The implementation of our simultaneous two-photon resonant optical laser lock (STROLL) is in the HPB regime in order to create a suitable crossing for the resonances of interest to us.

By measuring the optical rotation due to the Faraday effect near an atomic resonance, we obtain a dispersive error signal on the first probe beam (Probe 1). The probe laser can then be locked to this signal [110]. The origin of this signal is explained as follows. In an external axial magnetic field, an atomic medium has different refractive indices for right- and left-handed circularly polarised light. This is circular birefringence, and it leads to the rotation of the plane of polarisation of linearly polarised input light. The angle of rotation is proportional to the real part of the difference in refractive indices. That difference depends on the detuning of the light from the atomic resonances. The rotation is measured using the Stokes parameter, $S_1 = (I_x - I_y)/I_0$, where I_0 is the incident intensity, and I_x and I_y are the intensities of orthogonal linear polarisation components of the output light (see Fig. 5.1). Thus S_1 is the normalised difference in the intensities of the orthogonal linear polarisation components. The normalisation ensures that S_1 lies between -1 and 1, and so the lineshape of S_1 has zero crossings. An appropriate zero crossing in S_1 (see bottom panel of Fig. 5.2) is used as the error signal to the feedback loop of our probe laser PID (proportional-integral-derivative) controller.

Once the 780 nm probe laser is locked, the 776 nm coupling beam can then be locked. We use a second probe beam (Probe 2) that is overlapped with the coupling beam in the locking cell. It is important to ensure that the first probe beam is *not* overlapped with the coupling beam. This configuration of beams leads to the presence of an electromagnetically induced transparency (EIT) feature on Probe 2 but not on Probe 1. EIT is a well-known phenomenon that is often used in multi-

[4]For consistency with the paper associated with this chapter, we use the term *probe* instead of the term *pump* which we use in other chapters. There is no confusion of terms as both pump and probe are only ever used for the 780 nm beam.

level atomic systems [127, 128][5]. EIT is the reduction in the absorption of a weak probe beam when a strong coupling laser field is used to drive a resonant transition in a three-level atomic system, where the two resonant transitions are coupled coherently to a common state [58]. Associated with the change in absorption, there is a concomitant modification of the refractive index [129]. In the HPB regime, the EIT only changes the refractive index associated with one hand of polarisation. EIT thus causes additional birefringence and a change in the S_1 signal [110]. EIT appears on the S_1 signal in the form of a dispersive feature when the probe laser is scanning and the coupling beam is on and at a fixed frequency (see the highlighted region of Fig. 5.3 (b)). When the probe laser is locked and the coupling beam is scanning, we use this feature as the error signal (see Fig. 5.3 (c)) to the PID feedback loop of our 776 nm coupling laser controller.

5.3 Experimental demonstration

A schematic of the experimental set-up is shown in Fig. 5.1. Two weak, 50 μW, 780 nm probe beams are focussed to a beam ellipse with waists[6] of (83 $\pm$ 2) μm $\times$ (106 $\pm$ 2) μm through a 1 mm^3 vapour cell. This is the 'locking vapour cell' that contains the isotopically enriched rubidium. A strong, 16 mW,[7] 776 nm coupling beam is focussed to a beam ellipse with waists of (74 $\pm$ 2) μm $\times$ (80 $\pm$ 2) μm [130]. This beam is counterpropagated through the cell. Only the second probe beam and the coupling beam, which are resonant with the $|5S_{1/2}, m_J = \frac{1}{2}\rangle \rightarrow |5P_{3/2}, m_J = \frac{3}{2}\rangle$ and $|5P_{3/2}, m_J = \frac{3}{2}\rangle \rightarrow |5D_{5/2}, m_J = \frac{1}{2}\rangle$ transitions respectively, are overlapped within the cell. One advantage of our scheme is that, on exiting the cell, the strong coupling beam can be reused in further experiments. The raw intensity differences, i.e., $I_x - I_y$, are determined with the use of polarising beam splitters (PBS) and photodiodes: PD1 & PD2 are for the probe lock and PD3 & PD4 are for the coupling lock.

The probe and coupling light are also sent through a 2 mm length heated vapour

[5] Three-level EIT has been covered in Chapter 2, Section 2.4.
[6] The beam waist is measured using [130].
[7] A higher power of 40 mW is used if using the $5D_{3/2}$ level, to compensate for the weaker line strength.

cell—the 'monitor vapour cell'[8]—of pure (>98%) rubidium. By monitoring the absorption in this cell using a photodiode (PD5), we can choose where to lock the 780 nm laser. This is shown in Fig. 5.2: the zero crossing on the bottom panel is chosen depending upon the resonance of interest in the top panel. Both cells contain unknown buffer gas; this causes an additional broadening of 7 MHz on the D_2 line, as determined by a fit with ElecSus [66].

A magnetic field of up to 0.62 T is achieved across each vapour cell with the use of two cylindrical NdFeB magnets. Figure 5.1 shows a cross-sectional view of the top-hat profile of the magnets. There is a maximum variation of 1 mT across the 2 mm cell. The magnetic field is measured with a Hall probe and the value confirmed by fitting with ElecSus. The field profile is shown in Fig. 1 of reference [24]. Further details of the field uniformity and the magnet design are in [131] and [75]. By changing the separation of the respective magnets, the strength of each field can be varied. Changing the strength of the field across the locking cell gives us tunability for the lock-point of the 780 nm laser, although the STROLL will remain locked to the two-photon resonance.

[8]This also happens to be the experiment cell ('science cell') but, for the purposes of this chapter, it is simpler to consider it to be part of the laser locking set-up.

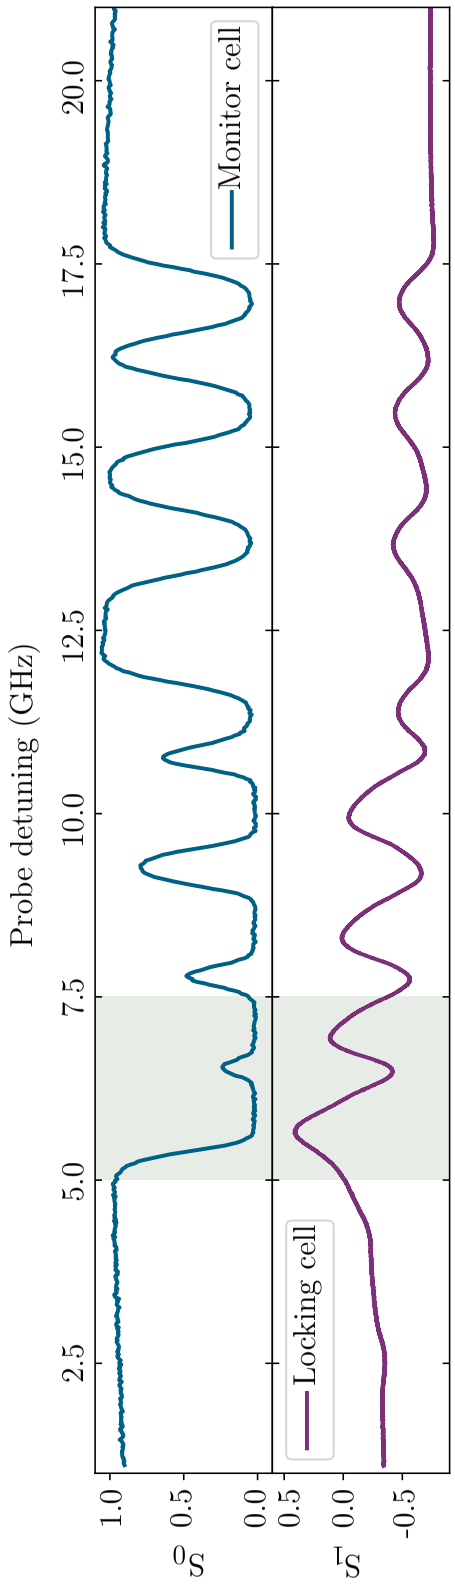

Figure 5.2: The monitor cell spectrum (top) is used to find the required zero crossing on the locking cell signal (bottom). Only probe beam 1 is on. The top panel displays the 780 nm spectrum in the HPB regime in the monitor cell at 106 °C; the bottom panel displays the S_1 signal in the locking cell at 100 °C. This temperature is chosen as a compromise between increased signal and optical depth & line broadening. Zero probe detuning is the weighted D_2 line centre of naturally abundant rubidium in zero magnetic field [132].

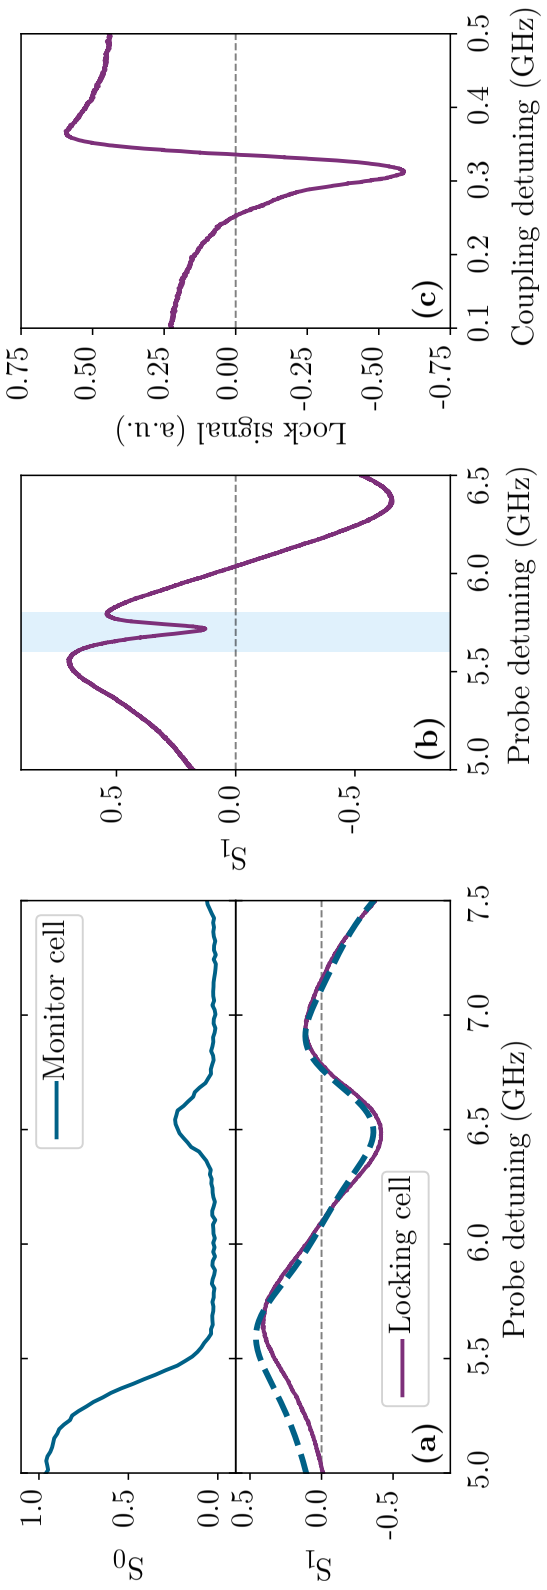

Figure 5.3: **(a)** The S_1 signal from probe beam 1 has a zero crossing that can be used in (c) to lock the probe laser. Shown is the shaded region from Fig. 5.2 where the dashed line is an ElecSus fit [66, 67, 133] . **(b)** The S_1 signal from probe beam 2 has an EIT feature (shaded region) when the coupling beam is turned on and resonant. **(c)** With the probe laser now locked using the method shown in part (a), scanning the *coupling* laser gives the error signal that is used to lock the coupling laser.

5.4 Results

To monitor the long-term stability of the locked lasers, we make use of the overlapping Allen Deviation, σ_{ODEV}, a measure typically used to quantify the stability of oscillators [134, 135]. The Overlapping Allan deviation as a function of averaging time, $m\tau_0$, can be defined via the following equation[9]:

$$\sigma^2_{ODEV}\left(m\tau_0\right) = \frac{1}{2\left(m\tau_0\right)^2\left(N-2m\right)} \sum_{n=1}^{N-2m} \left(x_{n+2m} - 2x_{n+1m} + x_n\right)^2, \tag{5.1}$$

where x_n is the time-series of frequency measurements spaced by the measurement interval τ_0, and N is the total number of measurements. We explain the meaning of equation 5.1 by first rewriting the terms in the brackets of the summation as $(x_{n+2m} - x_{n+1m}) - (x_{n+1m} - x_n)$. As an example, let $m = 1$ & $n = 1$, giving $(x_3 - x_2) - (x_2 - x_1)$. So, we have the change between the frequency difference of x_1 & x_2 and the frequency difference of x_2 & x_3. Doing this for all n up to $N-2$ means finding the change in frequency differences for every possible triplet of consecutive measurements. To complete the explanation, we turn to Fig. 5.4 which shows the overlapping Allan deviation of the concurrent frequency measurement of the 776 nm and 780 nm diode lasers. Averaging all the changes in frequency differences (and normalising) gives the frequency stability for a timescale equal to the measurement interval τ_0. The square root of this value gives a single point—the leftmost point—on Fig. 5.4. All other points are for timescales equal to integer multiples, m, of the measurement interval.[10] Thus we can tell how stable the frequency is across any timescale.

[9] We note, for clarity and completeness, that the square of the Allan deviation is the Allan variance (i.e., the left hand side of the equation) although the latter concept will not be used further in this work.

[10] The maximum value of the averaging time, $m\tau_0$, is half of the total measurement time, explaining why the figure shows only ~13 hours for a total measurement that took ~26 hours.

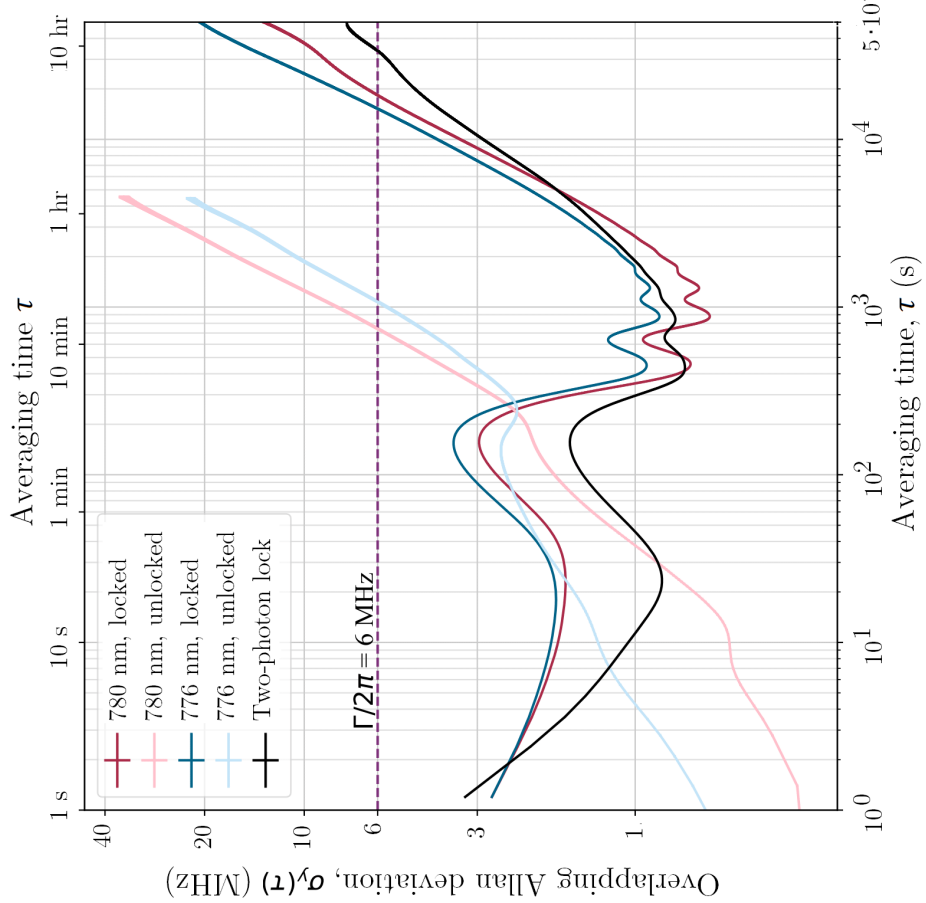

Figure 5.4: Overlapping Allan deviation of the frequency measurement of the 780 nm probe laser (red locked and pink unlocked), the 776 nm coupling laser (cyan locked, and light blue unlocked) and the summed frequency (black). $\Gamma/2\pi$ is the natural linewidth of the $5S_{1/2} \rightarrow 5P_{3/2}$ probe transition. The first maximal turning point is explained by temperature fluctuations (Fig. 5.5).

We used a High Finesse WS7 wavemeter with a switcher box to simultaneously monitor both laser frequencies over a period of 24 hours, where the measurement interval, $\tau_0 = 0.5$ s. The lasers stay locked for the whole of the 24-hour period.[11] The frequency instability of the sum of both lasers when locked is less than the natural linewidth of the probe transition of 6 MHz and of the EIT linewidth of 25 MHz. For most timescales, it is clear that the frequency instability of the sum is less than the frequency instability of either the 780 nm laser or the 776 nm laser alone. STROLL ensures that the lasers stay locked to the two-photon transition. Even if the frequency of one laser drifts, the frequency of the other laser changes accordingly to compensate. When the lasers are unlocked, they stay at an equivalent stability only for averaging times less than ~15 mins. This is of importance in quantum optics measurements where data must be accumulated over hours [6], e.g., for measurements of $g^{(2)}$ autocorrelations.

Our first attempt at making these measurements used two wavemeters, but the availability of the switcher box made the one-wavemeter method more convenient. In particular, when using two wavemeters it can be tricky to ensure that time is precisely synchronised between the two devices. (The solution is to write code so that a single program controls both wavemeters.)

The laser frequency stability that we achieve is adequate for our purposes. However, temperature sensitivity is a known issue with Faraday locking [110] with a temperature dependence of the zero crossing of < 1 MHz/°C. In future, if required, it is possible to make further improvements to the frequency stability by adding active temperature stabilisation. Figure 5.5 shows the overlapping Allan deviation for the temperature of the cell over the timescales of interest to us. The cell temperature was measured, using a Pico PT-104 Data Logger, to be stable to better than 1 °C over several hours. This measurement also revealed that the peaks at 150 s in Fig. 5.4 can be explained as being due to temperature variation since the frequency variation peaks coincide in averaging time with the temperature variation peaks.

[11] We find in Chapter 8 that, in practice, the lock can remain stable over a whole weekend period of 60 hours. The seeming loss of stability at long timescales is attributable to the fact that the reference wavemeter drifts at that timescale.

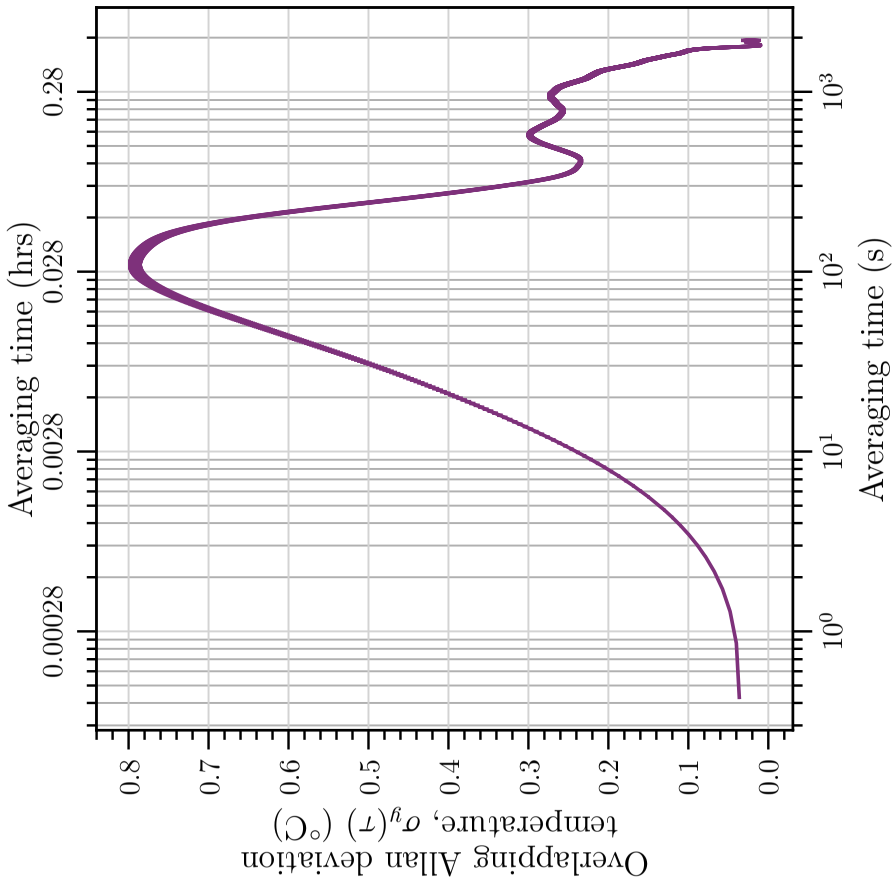

Figure 5.5: Overlapping Allan deviation of the temperature measurement of the locking cell: the cell temperature is stable to less than 1 °C over several hours. The maximum of the curve above coincides in averaging time (at 10^2 seconds) with the maximal turning point in Fig. 5.4, explaining the feature seen there.

The strength of the magnetic field sets the tunability of the probe laser lock, giving several GHz of freedom. There is further freedom from the fact that the Zeeman shift exceeding the Doppler width gives rise to many possible zero crossings in the S_1 signal, as seen in the bottom panel of Fig. 5.2.

5.5 Conclusion

In this chapter, we demonstrated a technique to simultaneously lock two laser frequencies to the two-photon transition $5S_{1/2} \rightarrow 5P_{3/2} \rightarrow 5D_{5/2}$ in ^{87}Rb, in the presence of an applied magnetic field that is sufficiently large to gain entry to the HPB regime. When the lasers are locked simultaneously, we showed a frequency instability for the sum frequency of less than 6 MHz for nearly all of the measured timescales. Whilst the specific application here happens to be in rubidium, the STROLL concept is easily transferrable to three-level ladder systems in other alkali metals. We shall find it necessary to use the STROLL lock to perform the quantum optics experiments in Chapter 8 which require two lasers to be locked in the HPB regime for timescales up to a day.

Chapter 6

Spectral Redistribution

6.1 Introduction

In the spontaneous four-wave mixing (4WM) experiments of Chapter 8, there appears to be a large number of 795 nm "background" photons present. An initial investigation into this phenomenon is described here as it is both experimentally and conceptually simpler than the experiments in the next two chapters. In the experiments of this chapter, only one laser is incident upon the rubidium atoms at any one time. We find that switching on only the 780 nm pump beam and passing it through heated rubidium vapour leads to the generation of 795 nm photons (Fig. 6.1 (a)). As the pump power is increased, so too are the number of 795 nm photons that are generated.[1] Since these photons have not been released via a 4WM process, they merely create a background of uncorrelated photons. We know that these photons are not produced by some decay mechanism from the higher-lying P state ($5P_{3/2}$) to the lower-lying P state ($5P_{1/2}$) because the opposite case is also observed: we *do* observe 780 nm photons when only the 795 beam is on[2]. A simple decay process is, therefore, not an adequate explanation. As we shall show in the rest of the chapter, the number of background photons (780 nm or 795 nm) depends in a quantifiable way on the temperature. As the temperature

[1] This shall be referred to as the 780→795 experiment(s).
[2] This shall be referred to as the 795→780 experiment(s).

increases, so does the atomic number density and the number of atomic collisions. This suggests the transitioning between the P states is due to the collision between atoms.[3]

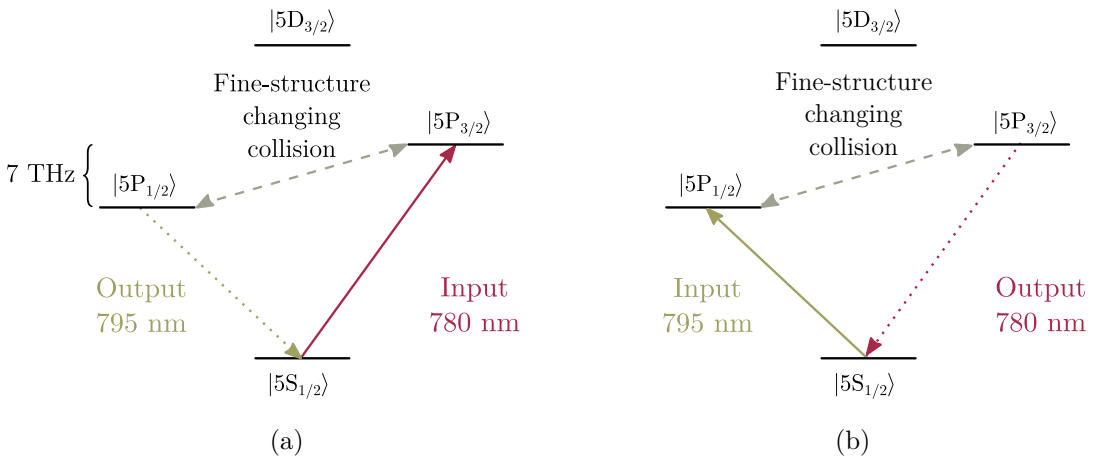

Figure 6.1: Rubidium energy levels for (a) 780→795 experiments: When a 780 nm beam is inputted, 795 nm photons are generated, and (b) 795→780 experiments: When a 795 nm beam is inputted, 780 nm photons are generated. This is due to collisions transferring population between the two P states shown. (The $5D_{3/2}$ state plays no part in this process but is included as a reminder that the above process always plays a part in all the 4WM experiments in later chapters where this state is involved.)

Investigation of this effect has a long history. Some of the earliest experiments in atomic physics in the 20$^{\text{th}}$ century were performed on this very phenomenon: In 1918, Wood & Mohler [136] illuminated the D_2 line in sodium and observed that light was produced on the D_1 transition. It was determined that this effect was due to collisions between sodium atoms or between sodium atoms and hydrogen molecules. At that time this phenomenon was referred to as "sensitised fluorescence", a term which has since fallen out of favour. There seems to be no single term used to describe this phenomenon now: throughout the years it has been variously referred to in the literature as "fine-structure changing collisions" [137], "collisional excitation" [138], "collisional transfer of excitation" [139, 140], "redistribution of

[3]We note that the energy gap between the two P states is on the order of $k_{\text{B}}T$. which is ~7 THz in frequency units.

resonance radiation" [141], and "spectral redistribution" [142], amongst others. We shall generally use the last term, "spectral redistribution" although the other terms will occasionally be employed as they emphasise different aspects of the phenomenon. Another reason for the lack of an accepted common term seems to be that this phenomenon appears in many different contexts. Towards the end of the last century [143], there was renewed interest in this topic as these fine-structure changing collisions were leading to atomic loss from cold-atom traps. Most recently, there has been interest in this phenomenon as a possible route by which an alkali laser [144–146] could be made. At high densities for the non-Rb species, the phenomenon can be due to both two-body and three-body collisions [137]. For many gases, and across a wide range of pressures, it is now well-known that as buffer gas densities increase, so does the amount of light generated [139, 147, 148]. As suggested by the variety of terms used to describe spectral redistribution, the literature is vast and the topic deep. We shall be limiting ourselves to performing some simple experiments in this chapter instead of attempting a full investigation, as our principal interest is in confirming that the background photons are indeed due to spectral redistribution.

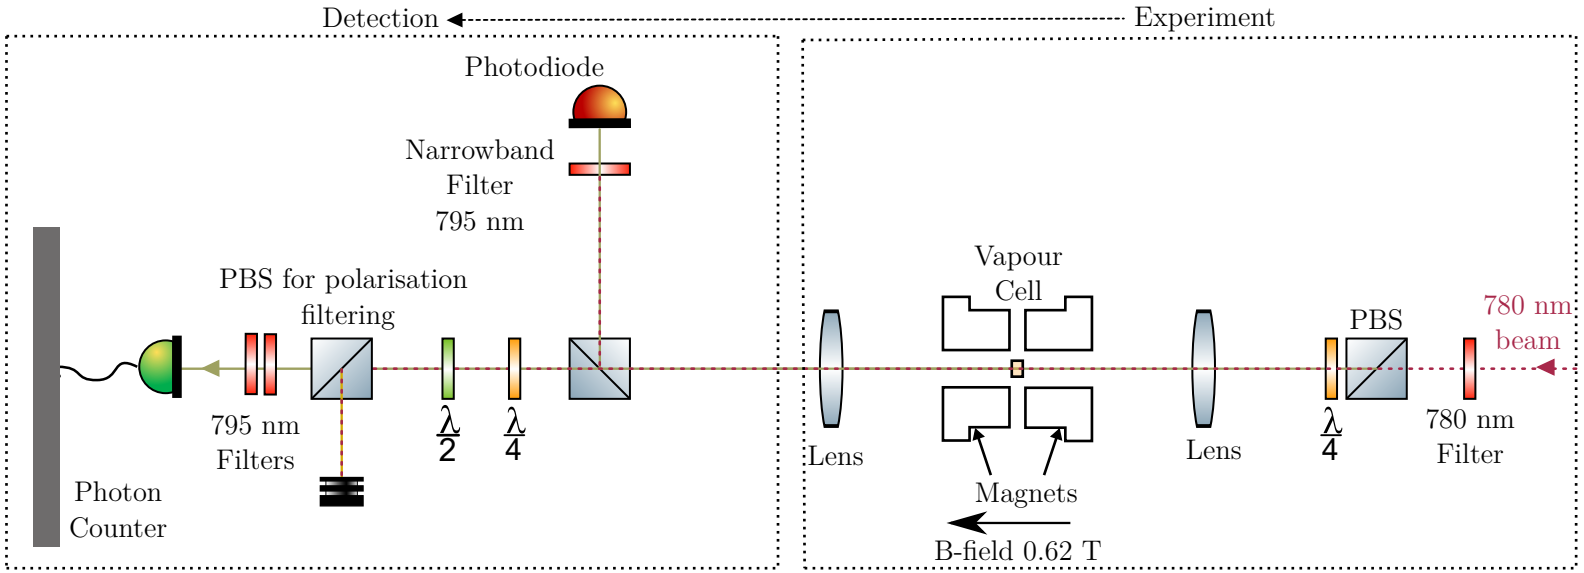

Figure 6.2: Diagram of the experiment, read from right to left. The pump/780 nm beam passes through a 780 nm interference filter to remove stray 795 nm photons. The quarter-waveplate ($\lambda/4$) sets the pump polarisation to be left-hand circularly polarised. The beam then passes through a heated 2 mm long vapour cell containing 98% rubidium-87. The cell is sandwiched between two magnets that produce a field of 0.62 T in the direction of the 780 nm beam propagation. After passing through the cell, the 780 nm transmission is measured by the photodiode. The 780 nm is then filtered out, allowing the generated 795 nm photons to be measured by the photon counter.

6.2 Experimental set-up

Figure 6.2 shows the experimental set-up. It is instructive to compare the set-up here to the one in the next chapter, Fig. 7.4: the principal difference is the absence of all lasers but one. As we are interested in this phenomenon in the context of the 4WM experiments performed in the next two chapters, it is reasonable to keep the set-up as similar as possible. For the 780→795 experiment, the pump/780 nm laser is passed through a heated 2 mm length vapour cell containing 98% rubidium-87. The cell is placed between two NdFeB magnets which create a magnetic field of 0.62 T [75] across the vapour cell. The circularly polarised 780 nm pump beam drives the σ^+ transitions between $5S_{1/2}$ and $5P_{3/2}$. A photodiode detects the transmitted pump beam and a photon detector detects the generated 795 nm photons along the axis of the pump beam. It is important to allow only 780 nm light into the cell and, conversely, to disallow 780 nm light onto the detector by using appropriate filters.

For the 795→780 conversion experiment, only a few changes are needed. The pump/780 nm laser is replaced with the seed/795 nm laser, and the 780 nm interference filters in the experiment are swapped with 795 nm interference filters and vice versa.

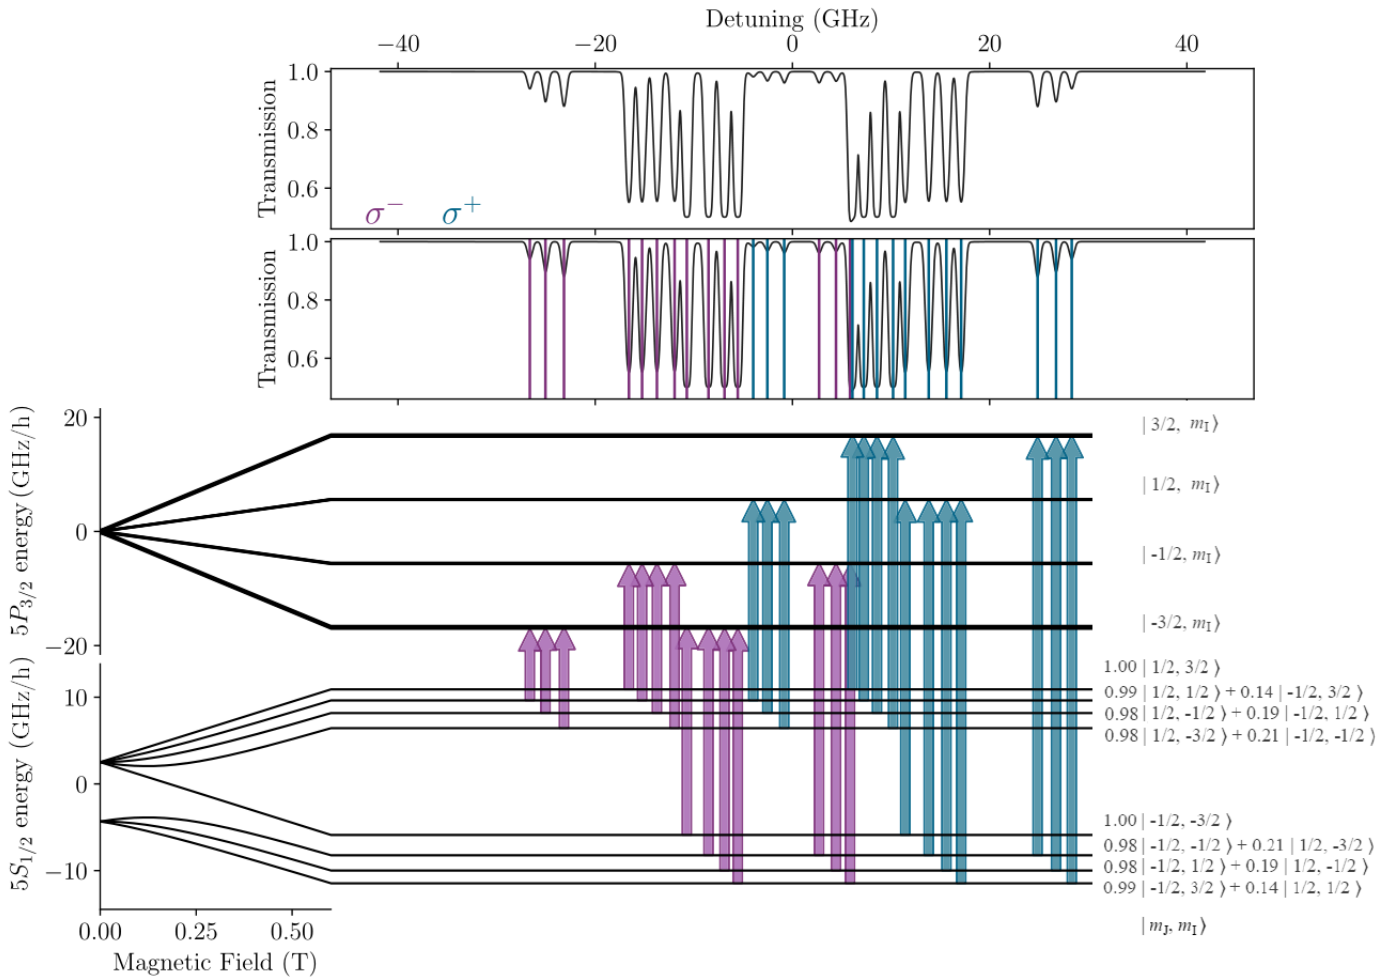

Figure 6.3: Theoretical transmission plots for 780 nm light (top) for an ^{87}Rb vapour cell at 100 °C in a 0.62 T magnetic field, and the associated transitions (bottom). For the geometry of this experiment, 780 nm light drives sigma minus, σ^- (purple), and sigma plus, σ^+ (cyan), transitions on the D_2 line. For the $5P_{3/2}$ state, there are 4 $|m_I\rangle$ levels associated with each $|m_J\rangle$ but they are not visible at this scale.

6.3 Results & Discussion

780 nm → 795 nm Spectral Redistribution

Figure 6.3 shows all the possible transitions for this geometry that are caused by the 780 nm/pump beam at 0.62 T; σ^- transitions are shown in purple and σ^+ transitions are shown in cyan. The top panels show the theoretical spectral features associated with each transition. The lowest level shown in the Breit-Rabi diagram is $0.99|-1/2, 3/2\rangle + 0.14|1/2, 1/2\rangle$ because at the magnetic field of 0.62 T the levels are not fully decomposed into pure $|m_J, m_I\rangle$ states. This leads to more transitions than would be expected if the bottom level were pure $|-1/2, 3/2\rangle$. The diagram is simplified in Fig 6.4, where only the major transitions for σ^+ are shown; the spectrum at the top of the diagram is experimental data for the transmission of 780 nm light. Setting the pump beam on resonance with the leftmost transition on the diagram[4] and increasing the power of the pump beam gives an increased number of generated 795 nm photons (Fig. 6.5). This is as expected: if fine-structure changing collisions involving P states are the underlying mechanism, we first have to promote atoms into one of the P states, and the higher the power, the more atoms are promoted. In order to maximise the signal-to-noise ratio, we chose to work at 16 mW of power, i.e., an intensity of 2.0×10^6 W/m^2 for a waist of 50 μm.

For 780 nm light driving σ^+ transitions, there are four transitions from $|5S_{1/2}, m_J = \frac{1}{2}\rangle$ to $|5S_{3/2}, m_J = \frac{3}{2}\rangle$, and a further four transitions from $|5S_{1/2}, m_J = -\frac{1}{2}\rangle$ to $|5S_{3/2}, m_J = \frac{1}{2}\rangle$. We find that associated with each transition, there is generation of 795 nm photons, as shown in Fig. 6.6.[5] For high temperatures, there is an apparent splitting of the peaks of the generated photons, perhaps due to reabsorption on-resonance.

Comparing the data for low (Fig. 6.6) and high temperatures (Fig. 6.7), we can already see that there are more 795 nm photons produced for higher temperatures. This is as expected since collisions are the underlying mechanism, and the rubidium

[4] $|5S_{1/2}, m_J = \frac{1}{2}, m_I = \frac{3}{2}\rangle \rightarrow |5P_{3/2}, m_J = \frac{3}{2}, m_I = \frac{3}{2}\rangle$

[5] Approximately double the number of photons are produced from the maximally stretched state (m_J=3/2, m_I=3/2), but this appears to be a high-power effect that disappears when the pump beam is in the weak probe regime.

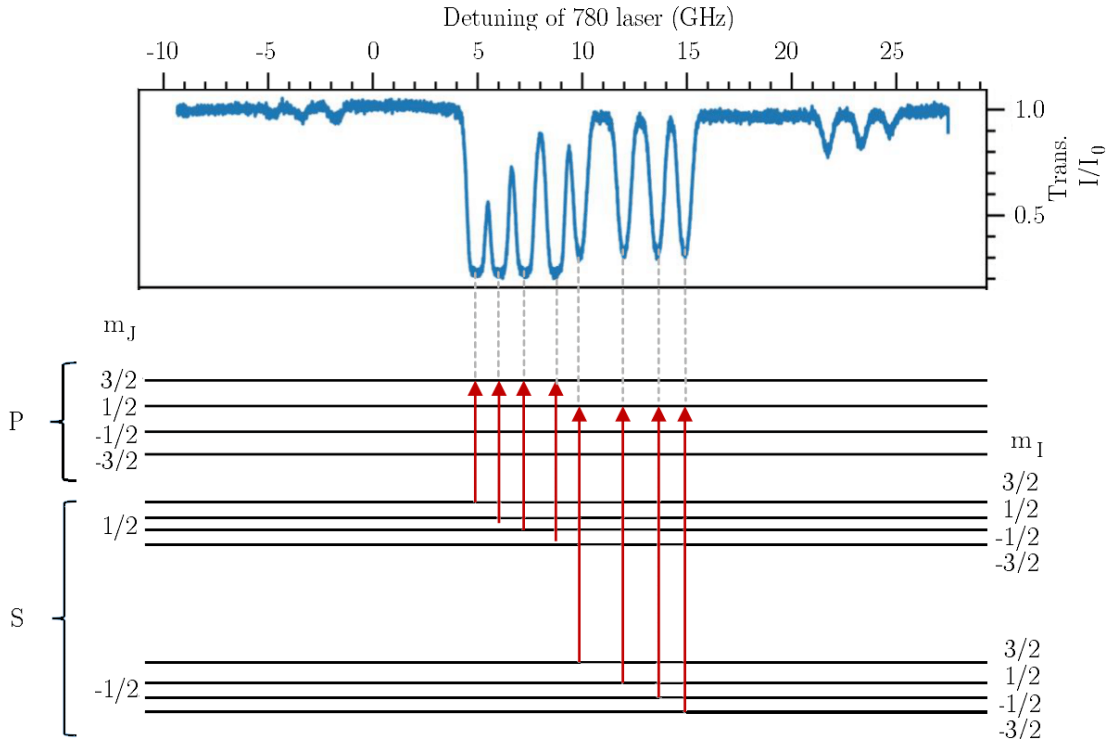

Figure 6.4: Experimental transmission plots for 780 nm light (top) for an ^{87}Rb vapour cell at 100 °C in a 0.62 T magnetic field, and the associated transitions (bottom). For 780 nm light driving σ^+ transitions, there are four transitions from $|5S_{1/2}, m_J = \frac{1}{2}\rangle$ to $|5P_{3/2}, m_J = \frac{3}{2}\rangle$ (the four transitions on the left), and a further four transitions from $|5S_{1/2}, m_J = -\frac{1}{2}\rangle$ to $|5P_{3/2}, m_J = \frac{1}{2}\rangle$ (the four transitions on the right).

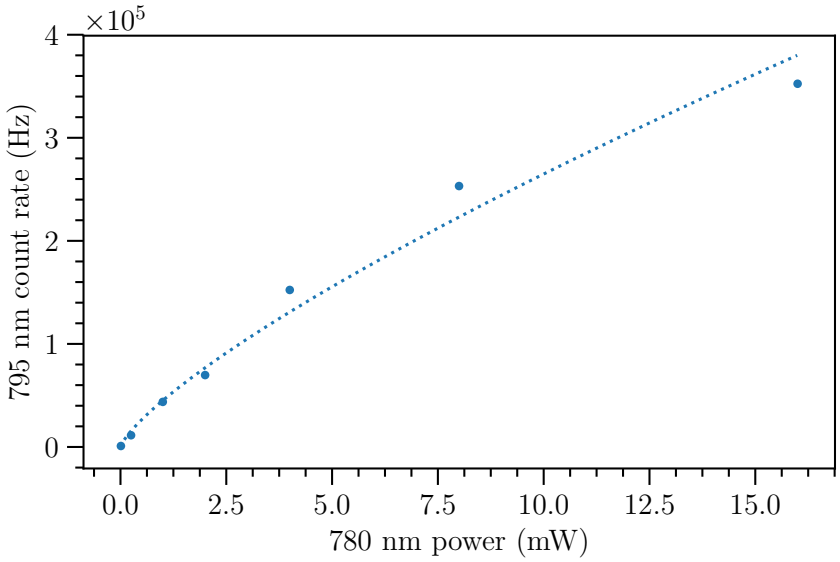

Figure 6.5: 795 nm photon count rate ($\dot{C}$) increases with increasing pump power (P). The dotted line is a fit to a function of the form $\dot{C} = aP^b$, where fitting gives $a = 4.5 \pm 0.6$ and $b = 0.77 \pm 0.6$. The Poissonian error bars are too small to be seen.

number density increases exponentially with temperature. So throughout the following, we need to convert temperature to number density. First, the temperature in degrees kelvin, T[K], is converted to vapour pressure in units of atmosphere, p[atm], and then the vapour pressure is converted to atomic number density, n [66]:

$$\log_{10}(p[\text{atm}]) = A + B(T[\text{K}])^{-1} + C\log_{10}(T[\text{K}]). \tag{6.1}$$

For solid Rb, A=4.857, B=−4215, and $C = 0$; for liquid Rb A=8.316, B=−4275, and C=−1.3102 [149]. The melting point of rubidium is 39.50 °C [150]. The atomic number density is found from the pressure by assuming an ideal gas:

$$n = \frac{p}{k_{\text{B}}T}. \tag{6.2}$$

We performed a series of experiments at different temperatures to determine the

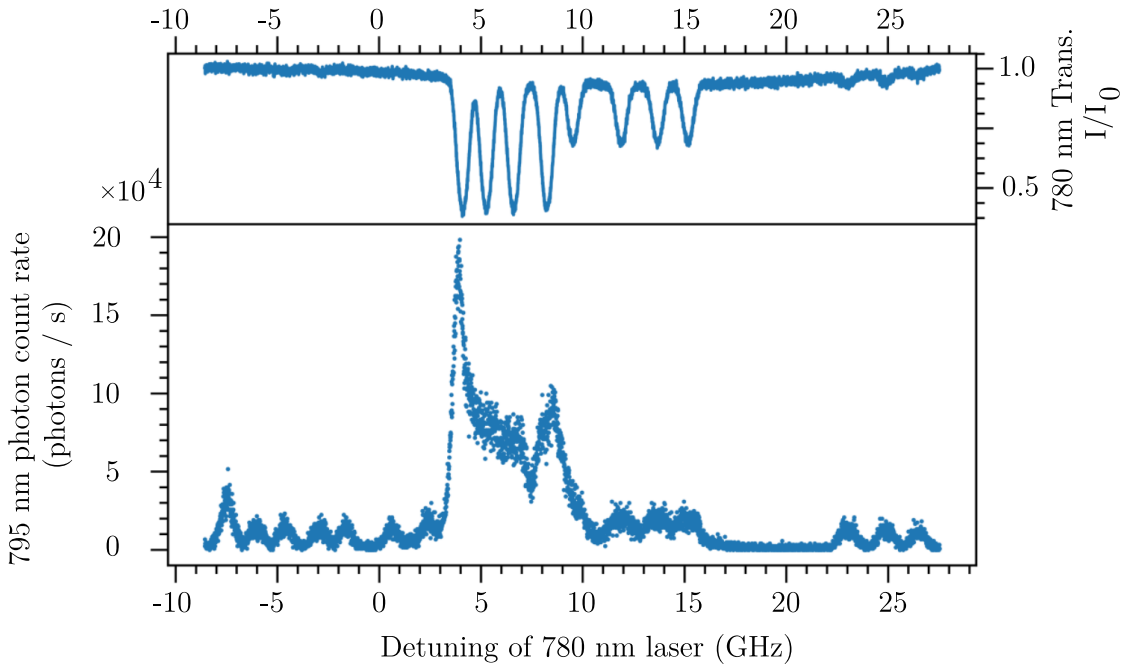

Figure 6.6: 780→795 experiment. Low temperature (73 °C). For each 780 nm transition (top panel), there is an associated peak of 795 nm photon generation (bottom panel). Approximately double the number of photons is produced from the maximally stretched state compared to the others, an effect that disappears at low input power.

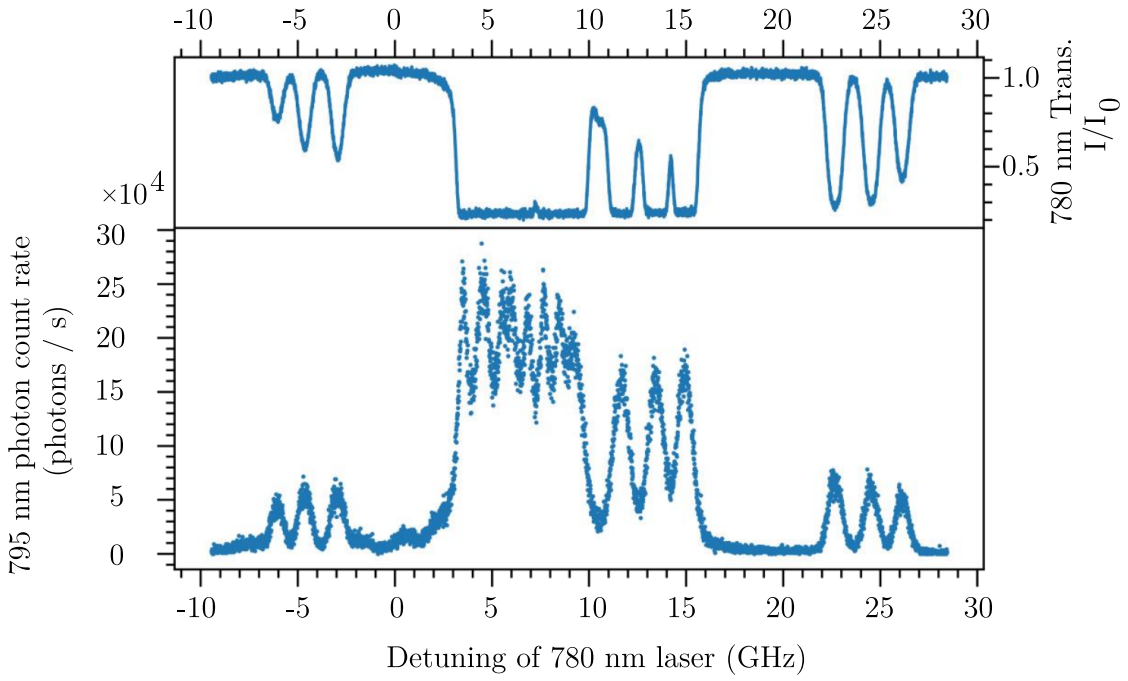

Figure 6.7: 780→795 experiment. High temperature (140 °C). For each 780 nm transition (top panel), there is an associated peak of 795 photon generation (bottom panel). The splitting of some of the peaks may be due to reabsorption. There is a greater photon generation rate at higher temperatures.

exact temperature dependence of the count rate of the generated 795 nm photons. As there was ample signal, for each data-point we needed to count for only 60 seconds. We focused our efforts on the peak used in the 4WM experiments, i.e., the leftmost peak in Fig. 6.6. Initially, the experiments were performed by setting the cell to a particular temperature, producing a figure like Fig. 6.6, and then repeating for a different temperature. For each temperature, we recorded the value of the count rate associated with the transition of interest. Whilst this procedure was adequate, it was also slow. We designed a much faster method[6] by making use of Newton's Law of Cooling [151]:

$$T(t) = T_{\text{env}} + (T(0) - T_{\text{env}})\, e^{-rt}, \tag{6.3}$$

where $T(t)$ is the temperature of the cell at time t, T_{env} is the temperature of the environment, $T(0)$ is the initial cell temperature, and r is the positive constant that characterises the system. The cell is heated to a high temperature (150 °C), and the heaters are switched off. Every minute, a weak probe spectrum is acquired, and the data are fitted with ElecSus [66] to determine the temperature, T, at that time, t. We are then able to fit these ($T(t)$, t) data points using Eq. 6.3 to determine r, and so obtain a calibrated cooling curve that gives temperature as a function of time.

We again heat the cell to the same high temperature, turn off the heaters and record the count rate recorded by the photon-counters as a function of time. We use the calibrated cooling curve to convert this to count rate as a function of temperature. Lastly, we convert to count rate versus number density using equations 6.1 and 6.2. The data obtained using this faster method matched that from the slower method, and we were able to obtain many more data points as the photon-counter is continuously recording.

We attempt to fit to the data, Fig. 6.8, a function of the form $\dot{C} \propto n^a$ where $\dot{C}$ is the count rate, n is the rubidium number density and the exponent a is allowed to float. For the data in the "low-temperature" range (31 °C — 86 °C), we obtain

[6] This reduced the time taken to obtain data from one day to one hour.

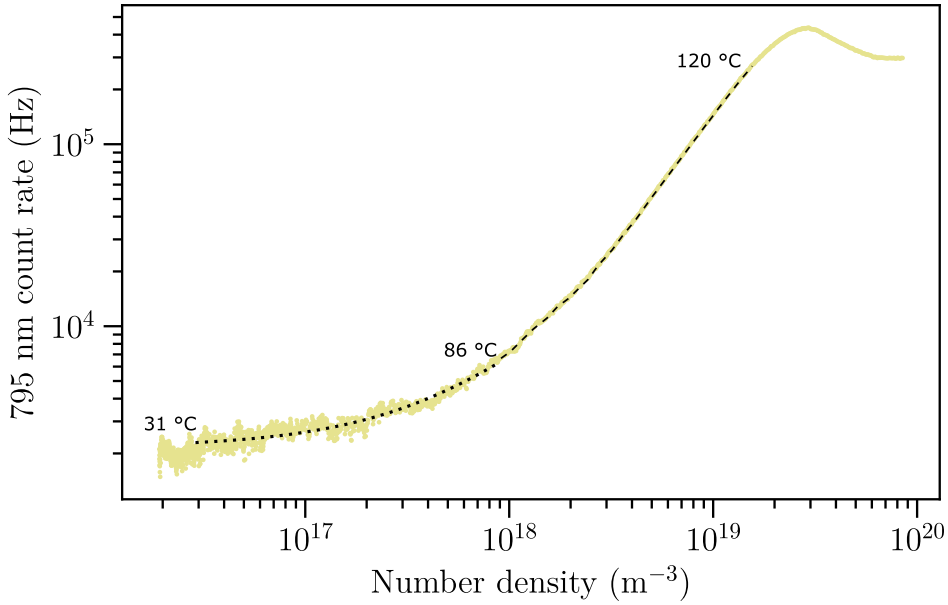

Figure 6.8: 780→795 experiments. In the low-temperature region (31 °C — 86 °C), where the fit to the data (yellow points) is shown as a dotted line, the photon generation rate is linearly proportional to the number density, n^1. In the high-temperature region (86 °C — 120 °C), where the fit is shown as a dashed line, the photon generation rate is proportional to $n^{1.6}$.

$a = 1.03 \pm 0.04$ with a reduced chi-squared, χ^2_ν, of 1.3, indicating a good fit [133]. For the data in the "high-temperature" range (86 °C — 120 °C), we obtain $a = 1.6 \pm 0.1$ with a reduced chi-squared, χ^2_ν, of 1.3, indicating a good fit. We performed the same experiment for all the different peaks. For the high-temperature region, a always lies between 1.6 and 1.8. Finally, at the highest temperatures (>120 °C) saturation occurs, which we surmise is due to 795 nm photon reabsorption in the dense atomic vapour. Recalling that these generated photons lead to uncorrelated background noise in our spontaneous 4WM experiments, their dependence on n gives us some control over their number. To keep the number of background photons to a minimum, it is best to work at lower temperatures. More exactly, we would like to work at the high end of the linear regime (~86 °C) because the genuine signal increases as n^2 with temperature as will be seen in the next chapter.

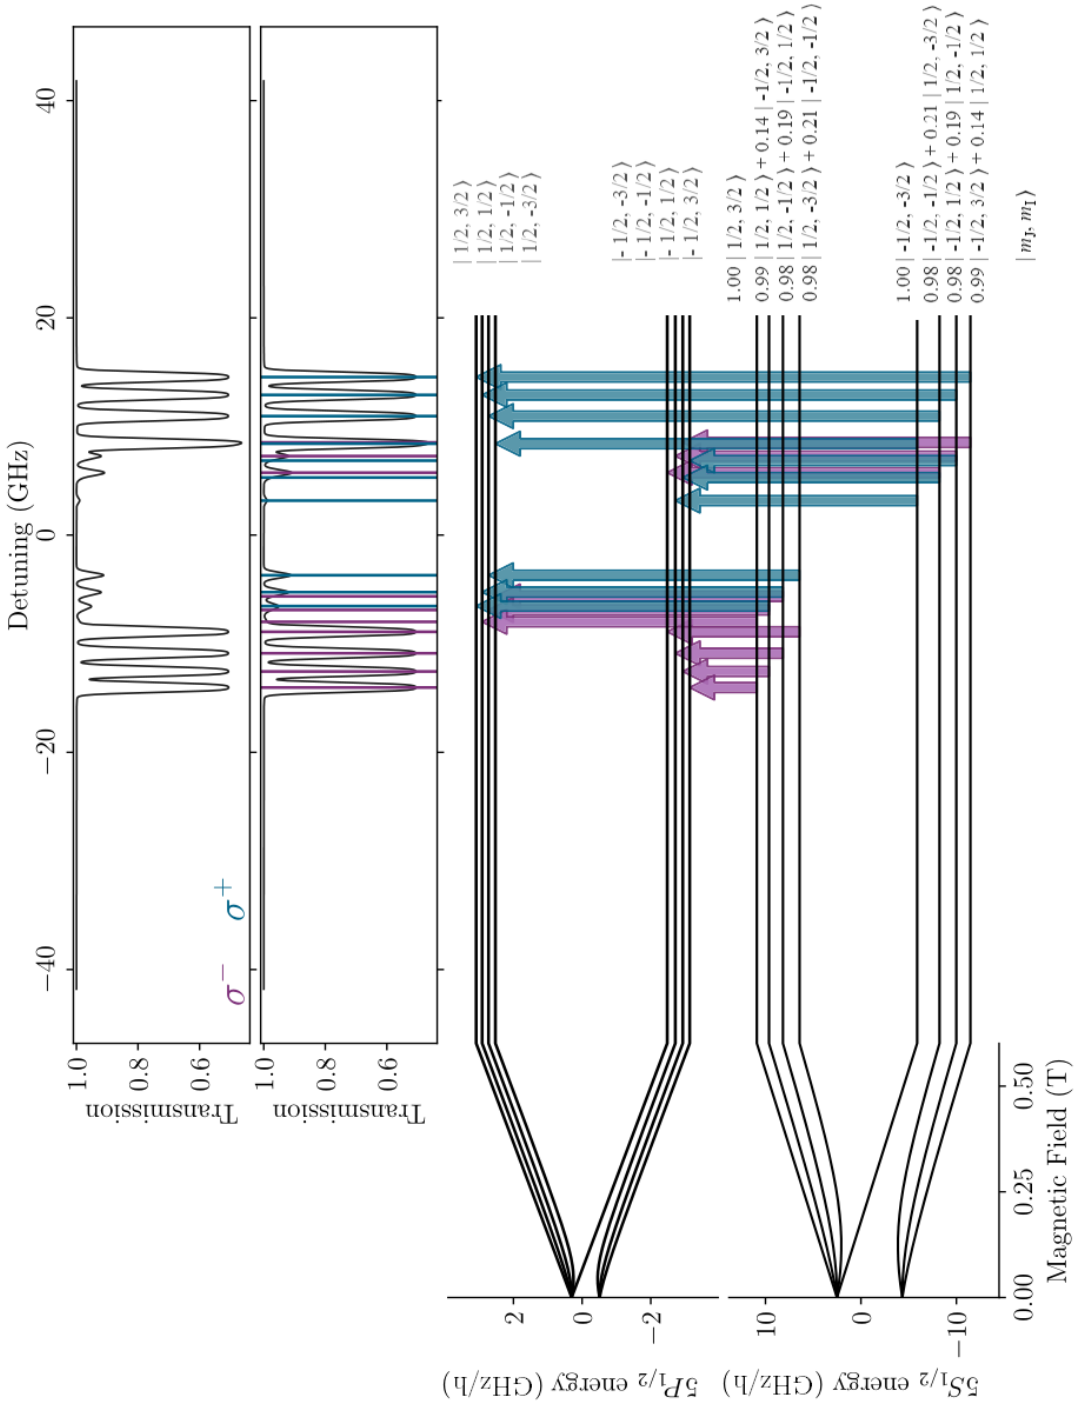

Figure 6.9: Theoretical transmission plots for 795 nm light (top) for an ^{87}Rb vapour cell at 100 °C in a 0.62 T magnetic field. For the geometry of this experiment, 795 nm light drives sigma minus, σ^- (purple), and sigma plus, σ^+ (cyan), transitions on the D_1 line.

795 nm → 780 nm Spectral Redistribution

With the seed/795 nm laser in place of the pump/780 nm laser, for each 795 nm transition (Fig. 6.10) 780 nm photons are generated (Fig. 6.11). Figure 6.9 shows, for this geometry, all the possible transitions that are caused by the seed beam at 0.62 T; σ^- transitions are shown in purple, and σ^+ transitions are shown in cyan. The top panels show the theoretical spectral features associated with each transition. A simplified version of the diagram is presented in Fig 6.10, where only the major transitions for σ^- are shown; the spectrum at the top of the diagram shows experimental data for the transmission of 795 nm light.

A short investigation gave similar results (Fig. 6.12) to the 780 nm → 795 nm experiment. We find that associated with each transition, there is the generation of 780 nm photons, as shown in Fig. 6.11.[7] There is a low-temperature region (31 °C — 85 °C) where the 780 nm photon generation rate has an n^1 dependence. There is a high-temperature region (85 °C — 120 °C) where the 780 nm photon generation rate has an $n^{1.8}$ dependence. Above a certain temperature, saturation occurs.

Spectral Redistribution: 780 nm → 795 nm & 795 nm → 780 nm

Both the 780→795 case and the 795→780 case suggest the following. At low temperatures where the n^1 dependence is evident, the rubidium atoms are colliding with some buffer gas, X, in the cell:

For 780→795,

$$\mathrm{Rb}(5\mathrm{P}_{3/2}) + \mathrm{X} \rightarrow \mathrm{Rb}(5\mathrm{P}_{1/2}) + \mathrm{X}. \tag{6.4}$$

For 795→780,

$$\mathrm{Rb}(5\mathrm{P}_{1/2}) + \mathrm{X} \rightarrow \mathrm{Rb}(5\mathrm{P}_{3/2}) + \mathrm{X}. \tag{6.5}$$

In the 795→780 case, this requires the conversion of the kinetic energy of the collisions ($k_\mathrm{B}T \sim 7$ THz) to the energy difference between the two P states (= 7 THz).

[7] We also looked for generated photons in the forward and backward direction to rule out any significant difference from looking from different directions.

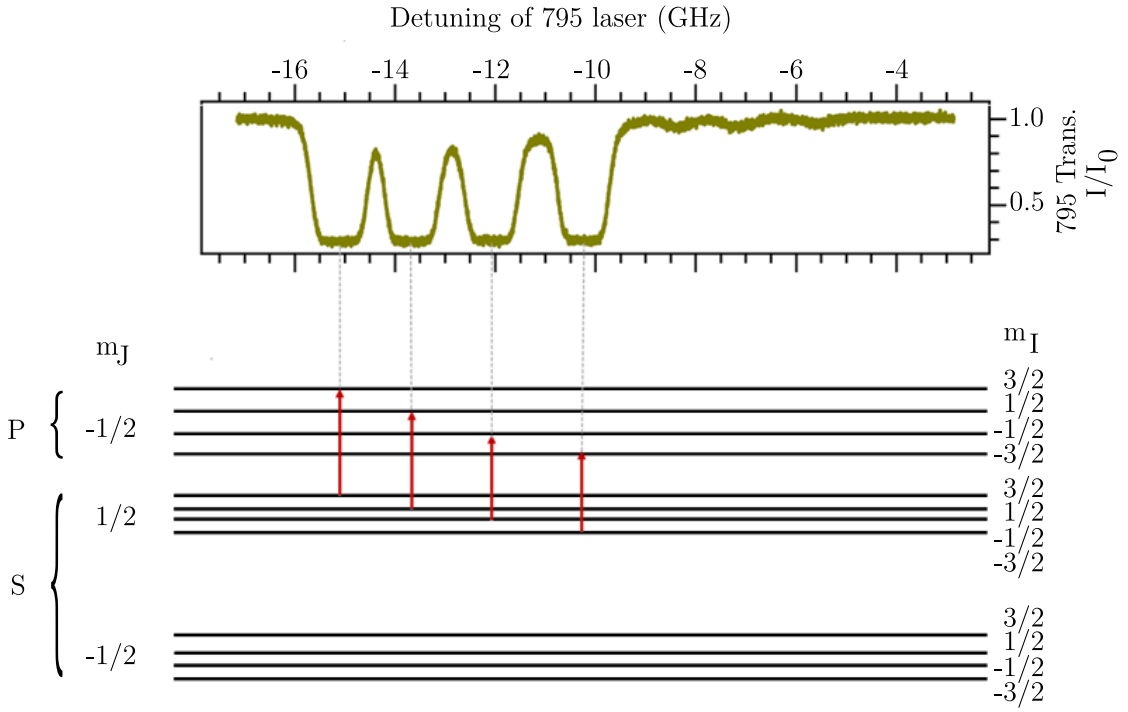

Figure 6.10: Experimental transmission plots for 795 nm light (top) for an ^{87}Rb vapour cell at 100 °C in a 0.62 T magnetic field, and the associated transitions (bottom). For 795 nm light driving σ^- transitions, there are four transitions from $|5S_{1/2}, m_J = \frac{1}{2}\rangle$ to $|5P_{1/2}, m_J = -\frac{1}{2}\rangle$.

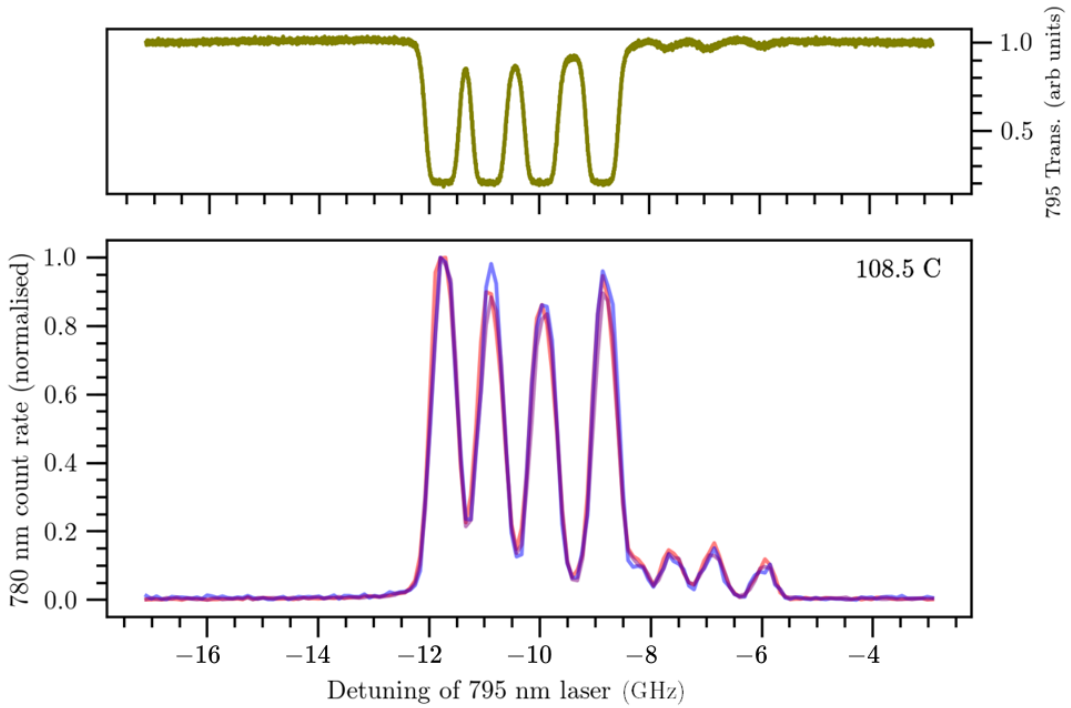

Figure 6.11: 780 generation associated with each 795 transition. Photons were detected in the backward (blue & purple) and forward (red) directions, as shown in the bottom panel.

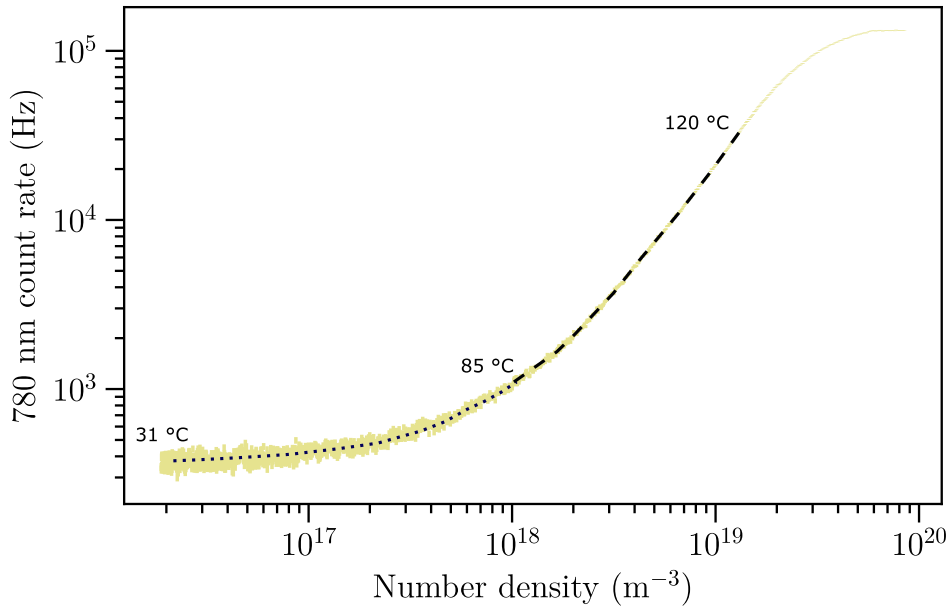

Figure 6.12: 795→780 experiments. In the low-temperature region (31 °C — 85 °C), where the fit to the data (yellow points) is shown as a dotted line, the photon generation rate is linearly proportional to the number density, n^1. In the high-temperature region (85 °C — 120 °C), where the fit to the data is shown as a dashed line, the photon generation rate is proportional to $n^{1.8}$.

Although the cell was not made with the intention of containing buffer gas, we know from previous measurements [24] that it does contain some residual gas, of unknown identity, from the manufacturing process.

At high temperatures, the exponent of n changes from 1 to approximately 2 indicating an extra mechanism at work. A possible explanation is that it is not only inter-species collisions (between Rb and buffer gas) but also intra-species collisions (between two Rb atoms) that occur in the high-temperature regime:

For 780→795,

$$\mathrm{Rb}(5\mathrm{P}_{3/2}) + \mathrm{Rb} \rightarrow \mathrm{Rb}(5\mathrm{P}_{1/2}) + \mathrm{Rb}. \tag{6.6}$$

For 795→780,

$$\mathrm{Rb}(5\mathrm{P}_{1/2}) + \mathrm{Rb} \rightarrow \mathrm{Rb}(5\mathrm{P}_{3/2}) + \mathrm{Rb}. \tag{6.7}$$

As mentioned in the introduction, the avenues of possible exploration on this topic are numerous. For example, with the current investigation, we cannot determine precisely which transition is giving rise to the 795 nm photons. This *can* be done using the etalon filter discussed in Chapter 3.2, and this is presently being explored by other members of our group.

6.4 Conclusion

This short investigation into the generation of background photons confirms that these photons are being generated by a well-understood phenomenon which is encountered in many atomic physics experiments. It goes by many names, and we have chosen to use the term spectral redistribution. Spectral redistribution increases at higher input powers and higher temperatures because it is due to the redistribution of energy during collisions between atoms. To keep this effect to a minimum, it is best to work at low temperatures. However, as we shall see in the next chapter, this also decreases the signal in our experiments, and so a compromise is required.

Chapter 7

Seeded four-wave mixing

Sections 7.1—7.4 of this chapter are based, in part, on the following publication:

D. J. Whiting, R. S. Mathew, J. Keaveney, C. S. Adams, and I. G. Hughes, *Four-wave mixing in a non-degenerate four-level diamond configuration in the hyperfine Paschen–Back regime*, Journal of Modern Optics **65** 713 (2018). [25]
DOI:10.1080/09500340.2017.1377308

7.1 Introduction

Four-wave mixing (4WM[1]) is a nonlinear optical process in which three optical fields are coherently combined to produce a fourth field. It has been used in numerous applications, including for the production of correlated photon-pairs [2] for use in quantum information protocols; for the creation of entangled imaging systems [5]; for creating collective spin excitations [152]; and for observing collective quantum beats [6]. The particular case of seeded 4WM has found use in precision spectroscopic measurements [4]; in displacement measurements in electromechanical cantilevers [153]; and in photon storage [154]. Different configurations of energy levels have been used to generate 4WM, including the double-lambda [155], double-ladder [123], and the diamond [24, 156–158] schemes. In this chapter, we shall experimentally characterise a four-level seeded 4WM system in the diamond configuration and show how the results can be modelled. We shall then use the 4WM signal to extract the corresponding excited state dipole matrix element. Our full characterisation of the seeded 4WM system gives us a good understanding of how the 4WM signal varies with the parameters of the experiment, which becomes crucial in the next chapter where we switch to spontaneous 4WM.

Atomic vapours are especially useful for some of the aforementioned applications as the photons produced are already frequency- and bandwidth-matched to atomic resonances and, for that reason, matched to other elements of an atom-based technology, e.g., atom-based quantum memories [159], quantum repeaters [160], and quantum gates [161]. Experiments with thermal vapours are advantageous in comparison to the experiments with cold-atoms because the former are relatively inexpensive, simpler, and miniaturisable. However, thermal motion leads to the Doppler broadening of absorption lines and multiple path interference occurs, which is often disadvantageous. Interferences occur when $\omega_{\mathrm{D}} \geq \delta$, where ω_{D} is the Doppler width and δ is the separation between lines. ω_{D} is on the order of hundreds of MHz and, for alkali metal D-lines, δ is typically 100 MHz. This means that, under normal laboratory conditions, ω_{D} is indeed greater than δ and interference will occur.

[1]Four-wave mixing is often also abbreviated as FWM in the wider literature, but we prefer 4WM because ‘4’ immediately indicates four whereas ‘F’ does not.

However, the application of a strong magnetic field has been shown to simplify nonlinear atom-light interactions in thermal vapours, resulting in enhanced control of electromagnetically induced transparency [24] and absorption [23], without the need for optical pumping. At a field strength where the Zeeman shift exceeds the internal hyperfine energy intervals of the atom, the hyperfine Paschen–Back (HPB) regime[2] is entered. In the absence of a magnetic field, there is coupling between the total electronic momentum $\mathcal{J}$ and the nuclear spin momentum $\mathcal{I}$ to give the total angular momentum, $\mathcal{F}$ where $\mathcal{F} = \mathcal{I} + \mathcal{J}$. The 'good' quantum numbers to describe the system are therefore F and m_F. In the HPB regime, there is decoupling of $\mathcal{J}$ and $\mathcal{I}$ and so the 'good' quantum numbers for describing states are m_J and m_I. We show here that entering the HPB regime also simplifies 4WM as it does other nonlinear optical phenomena [6, 18–22, 65, 110, 162–164].

The conceptual core of the experiment covered in this chapter is illustrated in Figures 7.1 and 7.2. The seed, pump, and coupling laser beams, i.e., the three optical fields, are coherently combined and the medium, rubidium-87 thermal vapour, generates a fourth signal field when the phase-matching condition is met: $\mathbf{k}_{\text{pump}} + \mathbf{k}_{\text{coupling}} = \mathbf{k}_{\text{seed}} + \mathbf{k}_{\text{signal}}$. On the left of Figure 7.1, we show that in the absence of a magnetic field, the hyperfine splitting of the four levels results in many paths by which the 4WM signal is generated. The electric fields from the different paths mutually interfere, leading to a complicated 4WM signal, as shown in the bottom-left panel of Figure 7.2. In contrast, as shown on the right of the respective figures, when a large magnetic field is applied such that the HPB regime is entered, adjacent levels are separated by more than the Doppler width and a single four-level system can be isolated. We see, in the bottom-right panel, that the 4WM signal thus produced is much simpler and can therefore be quantitatively modelled. We did not note any significant amplitude difference between the two 4WM signals (i.e., with and without the magnet).

In the HPB regime, the four-level system consists of levels $|0\rangle$, $|1\rangle$, $|2\rangle$, and $|3\rangle$ which are explicitly the following states in ^{87}Rb:
$|0\rangle = |5\mathrm{S}_{1/2}, m_J = \frac{1}{2}\rangle$, $|1\rangle = |5\mathrm{P}_{3/2}, m_J = \frac{3}{2}\rangle$, $|2\rangle = |5\mathrm{D}_{3/2}, m_J = \frac{1}{2}\rangle$, and
$|3\rangle = |5\mathrm{P}_{1/2}, m_J = -\frac{1}{2}\rangle$. In all cases, $m_I = \frac{3}{2}$.

[2] See Section 2.6 for a discussion of the HPB regime.

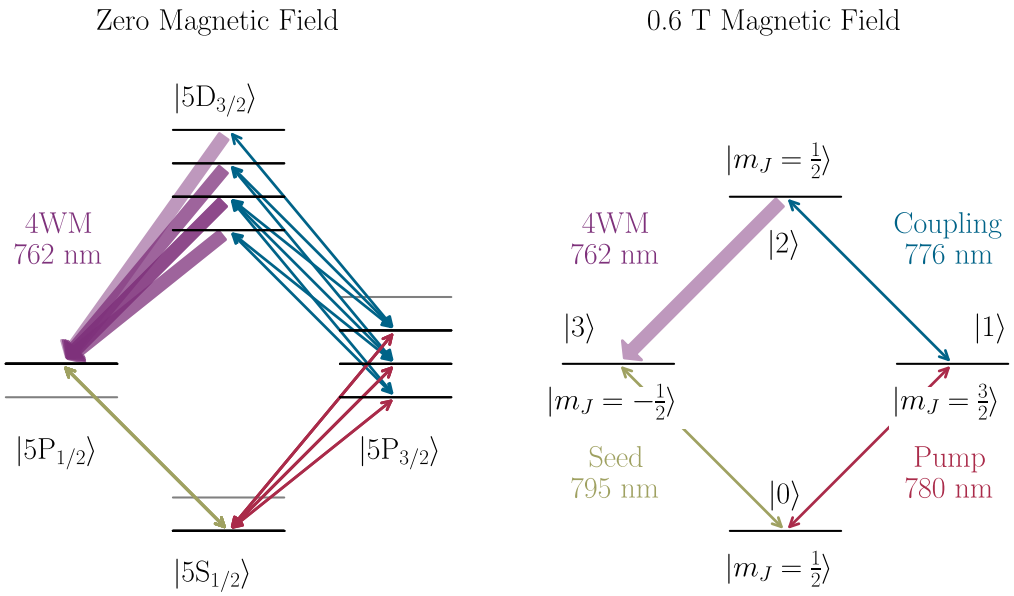

Figure 7.1: Rubidium energy levels in the diamond configuration in the absence (left) and presence (right) of a 0.62 T magnetic field. Three beams (pump, seed & coupling) are added. When the phase-matching condition is fulfilled, the four-wave mixing process (4WM) generates a fourth signal beam. With no magnetic field, there is multiple-path interference; the magnetic field removes this interference. The states labelled $|0\rangle, |1\rangle, |2\rangle$ & $|3\rangle$ are those used in the model.

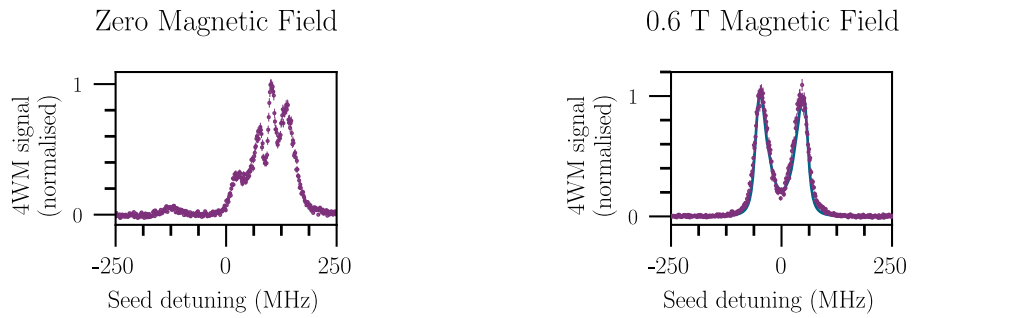

Figure 7.2: Example experimental four-wave mixing spectra in the diamond configuration in the absence (left) and presence (right) of a 0.62 T magnetic field. In the zero-field case, the four-wave mixing spectra are very sensitive to the experimental conditions and thus difficult to model. Application of a magnetic field results in 'textbook' four-wave mixing spectra which can be quantitatively modelled. The right panel shows the model (blue) overlaid on the experiment (purple).

7.2 Theory

To understand how the 4WM signal is produced by the medium we need to consider the polarisation of the medium, $\widetilde{P}(t)$, which depends on the strength of an applied optical field $\tilde{E}(t)$. In linear optics, $\tilde{P}(t) = \epsilon_0 \chi^{(1)} \tilde{E}(t)$ where $\chi^{(1)}$ is the linear electric susceptibility. When an optical field is incident on a medium, a nonlinear response is possible where the polarisation of the medium can develop frequency components that are not present in the incident field. These components then act as sources of new frequency components of the optical field [165]. Four-wave mixing is a nonlinear optical phenomenon; here, the macroscopic optical response can be described by generalising the previous equation by expressing the polarisation as a power series in the field strength [165]:

$$\widetilde{P}(t) = \epsilon_0[\chi^{(1)}\tilde{E}(t) + \chi^{(2)}\tilde{E}^2(t) + \chi^{(3)}\tilde{E}^3(t) + ...],$$

where $\chi^{(2)}$ and $\chi^{(3)}$ are known as the second- and third-order nonlinear optical susceptibilities, respectively. The value of $\chi^{(2)}$ is only nonzero for crystals that do not display inversion symmetry (noncentrosymmetric crystals [165]). Thus for the case of interest to us, which is 4WM in a thermal vapour, it is $\chi^{(3)}\tilde{E}^3(t)$ which is of relevance. In 4WM, because three optical fields $\tilde{E}_1(t), \tilde{E}_2(t)$, and $\tilde{E}_3(t)$ are coherently combined to produce a fourth one, the relevant term becomes: $\chi^{(3)}\tilde{E}^3_{\text{total}}(t)$, where $\tilde{E}_{\text{total}} = \tilde{E}_1(t) + \tilde{E}_2(t) + \tilde{E}_3(t)$.

Our model consists of a system of four levels that interact with three continuous-wave driving fields with wavelengths 780 nm (pump), 776 nm (coupling), 795 nm (seed), with Rabi frequencies Ω_{P}, Ω_{C}, Ω_{S} and detunings Δ_{P}, Δ_{C}, Δ_{S} respectively[3]. We begin by writing the density matrix, $\hat{\rho}$, for the unperturbed atom. We then find the time evolution of the density matrix in response to a Hamiltonian and spontaneous emission by using the master equation in Lindblad form [57, 166]:

$$\frac{\mathrm{d}\hat{\rho}}{\mathrm{d}t} = -\frac{i}{h}[\hat{H}, \hat{\rho}] + \sum_n \frac{1}{2}[2\hat{C}_n\hat{\rho}\hat{C}_n^\dagger - (\hat{\rho}\hat{C}_n^\dagger\hat{C}_n + \hat{C}_n^\dagger\hat{C}_n\hat{\rho})],$$

where $C_n = \sqrt{\gamma_n}\hat{A}_n$, $\hat{A}_n$ are the collapse operators through which the environment

[3]These angular detunings are converted to linear detunings in the figures, i.e., $\Delta/2\pi$.

and system couple and γ_n are the corresponding rates. The total Hamiltonian, $\hat{H} = \hat{H}_0 + \hat{H}_{\text{int}}$, where $\hat{H}_0$ is the unperturbed Hamiltonian and $\hat{H}_{\text{int}}$ is the interaction Hamiltonian. The interaction Hamiltonian $\hat{H}_{\text{int}} = -\hat{d} \cdot \vec{E}$, where $\hat{d}$ is the electric dipole operator and $\vec{E}$ is the electric field vector. We note in passing the relationship between the macroscopic and microscopic descriptions [165]: $\widetilde{P} = N\langle\hat{d}\rangle = N \cdot \text{Tr}(\hat{\rho}\hat{d})$, where N is the atomic number density. We write the interaction Hamiltonian, in the rotating wave approximation, as[4]

$$\hat{H} = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_{\text{P}} & 0 & \Omega_{\text{S}} \\ \Omega_{\text{P}} & -2\Delta_{\text{P}} & \Omega_{\text{C}} & 0 \\ 0 & \Omega_{\text{C}} & -2(\Delta_{\text{P}} + \Delta_{\text{C}}) & 0 \\ \Omega_{\text{S}} & 0 & 0 & -2\Delta_{\text{S}} \end{pmatrix}.$$

Once the appropriate decay rates are included, we can solve the Lindblad master equation to find the steady-state density matrix ρ. There is an emitted electric field which is proportional to the coherence between $|3\rangle$ and $|2\rangle$, ρ_{23}: this is the 4WM signal [165].

Atomic motion is modelled by including the Doppler-shifted detunings for each velocity class v. The total emitted field is the sum over all $\rho_{23}(v)$ weighted by the 1-D Maxwell-Boltzmann distribution, $P(v)\text{d}v = Be^{-Av^2}\text{d}v$, where $A = \frac{m}{2k_{\text{B}}T}$ and $B = \sqrt{A/\pi}$ [167]. It is also necessary to take into account the spatial intensity profiles of the driving fields, i.e., the driving laser beams. These beams have Gaussian intensity profiles in the radial direction and are assumed to be perfectly overlapped. The emitted field is calculated for radial shells where each shell is assumed to be of constant intensity. Because of phase-matching, the emitted field (i.e., the 4WM signal) is in a well-defined direction. In the experimental set-up, the signal is sent to the photon counter through a single-mode fibre. We model the effect of collecting with the fibre by considering the fibre to be a Gaussian spatial filter that acts on the emitted field. Thus the amplitude of the signal is calculated as

$$S \propto \int_0^\infty |\rho_{23}|^2 \exp(-2r^2/w_{\text{D}}^2)\, r\, \text{d}r, \tag{7.1}$$

[4]The Hamiltonian looks slightly different from the theory chapter only because of a relabelling of the states, done to keep this chapter consistent with its associated paper.

where $w_{\mathrm{D}} \approx 90$ μm is the waist that we measure at the output of the optical fibre after the 11 mm collimating lens.

Model parameters

In order to solve for the steady-state density matrix, the correct decay rates must be coded into the computational model. The decay rates of the relevant atomic states are $\Gamma_{30} = 2\pi \times 5.7$ MHz, $\Gamma_{10} = 2\pi \times 6$ MHz, $\Gamma_{20} = 2\pi \times 0.07$ MHz $\Gamma_{21} = 2\pi \times 0.17$ MHz, and $\Gamma_{23} = 2\pi \times 0.43$ MHz [168]. In the model, the decay rates Γ_{20}, Γ_{21}, and Γ_{23} are multiplied by a factor α. This factor includes the contributions from collisional buffer gas broadening on the excited-state transitions [169] and the effect of inhomogeneity in the magnetic field. By using a value of $\alpha = 35$, we obtain a good agreement with the data. This corresponds to an additional broadening of ~20 MHz, which agrees with values which have previously been measured [24].

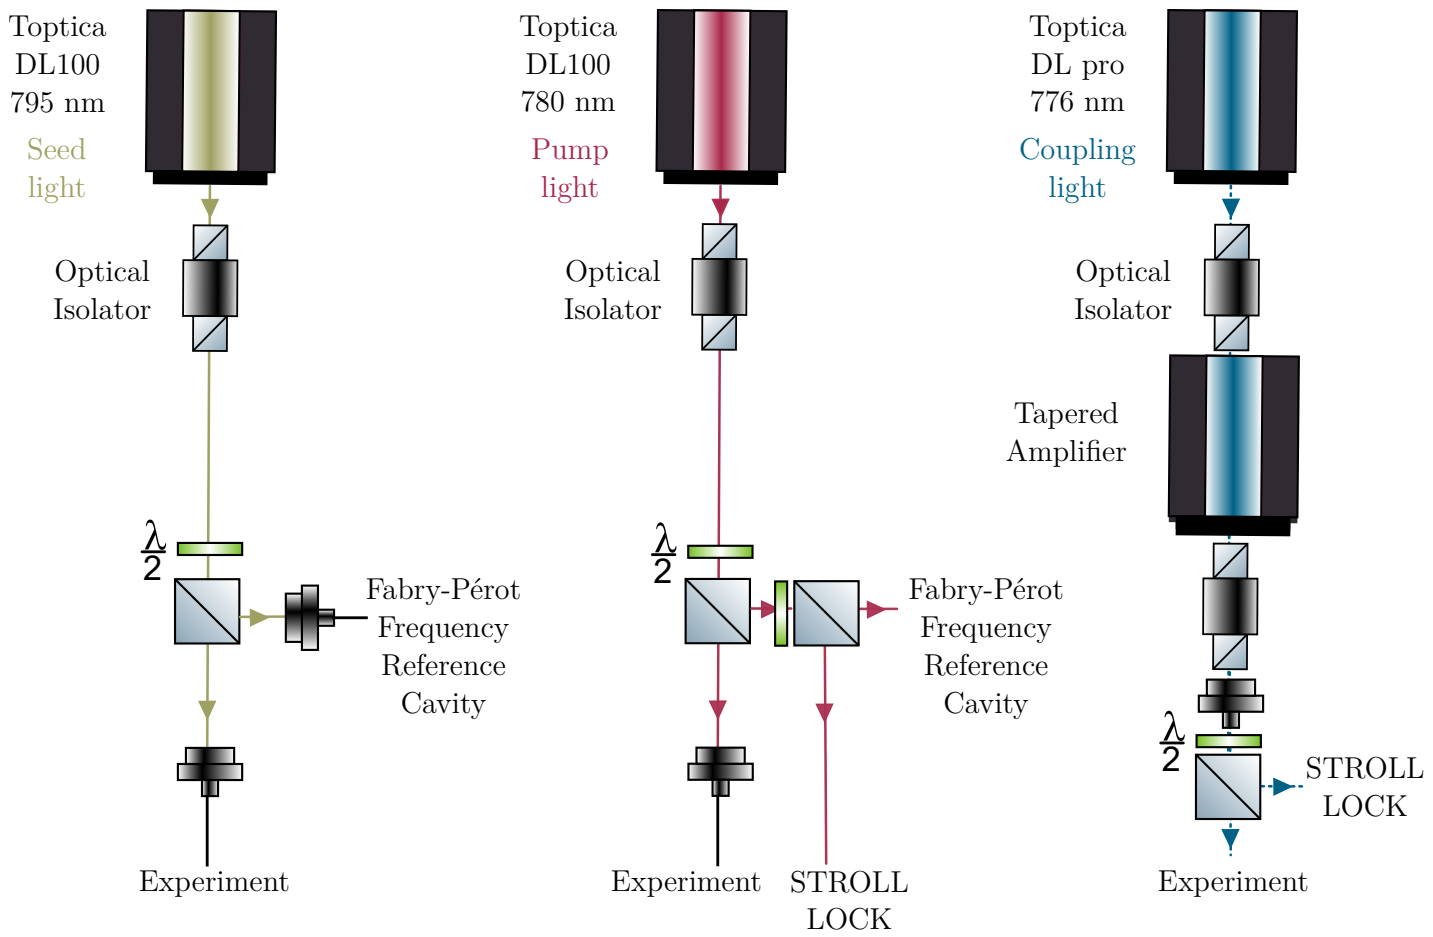

Figure 7.3: Schematic of the routes taken by the lasers in the experiment. Before reaching the experiment, the seed and pump beams are split and sent to the Fabry-Pérot etalon, which is used as a frequency reference. The pump and coupling laser beams are split and sent to the STROLL lock for stabilising the frequencies of both lasers. For a full table layout see Appendix A.

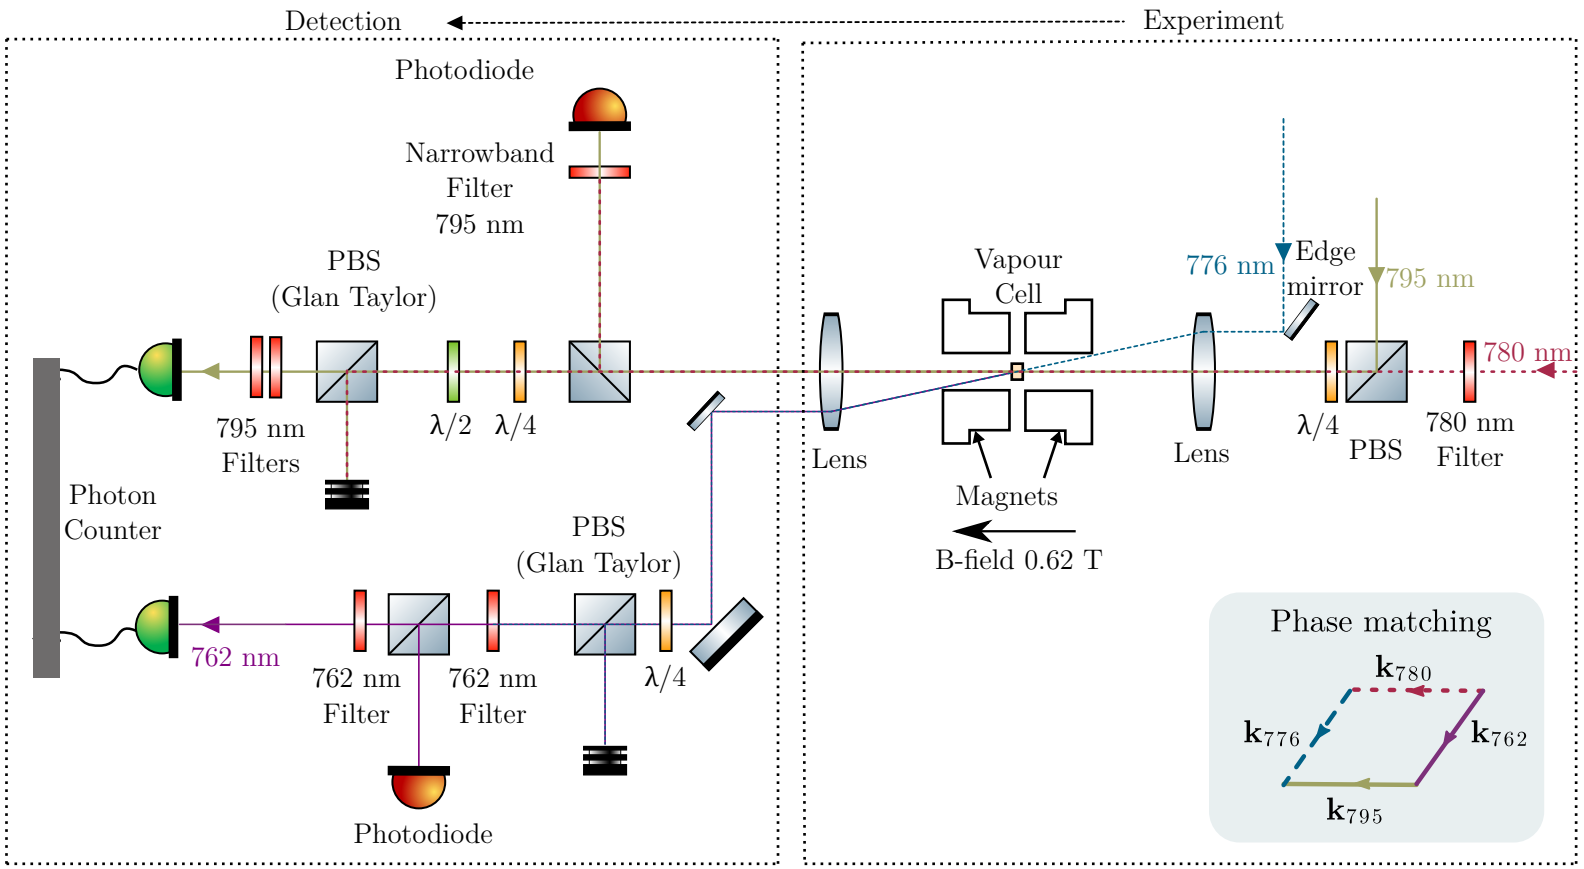

Figure 7.4: Diagram of the experimental set-up, read from right to left. All three beams pass through a heated 2 mm length vapour cell containing 98% rubidium-87. Pump and seed beams are aligned along the 0.62 T magnetic field axis whilst the coupling beam is at a small angle (10 mrad). The quarter-waveplate ($\lambda/4$) sets the pump and seed polarisations to be right-hand circularly polarised and left-hand circularly polarised respectively. The edge mirror is translated up on the diagram during initial alignment. The phase-matching condition is illustrated in the panel.

7.3 Experimental set-up

A schematic of the routes taken by each of the laser beams is shown in Fig. 7.3, and a diagram of the experimental set-up is shown in Fig. 7.4. We first describe the experimental set-up summarily in the précis below and then expand upon the details of the entire set-up in the rest of the section.

Précis

A vapour cell of length 2 mm containing 98% ^{87}Rb, which is housed within a heater, is placed between two NdFeB magnets. These cylindrical magnets have a top-hat profile [75] and allow us to achieve a magnetic field of 0.62 T across the vapour cell, which has a maximum variation of 1 mT across the 2 mm.[5] Adjusting the cell temperature gives us control over the atomic number density, as shown in Fig. 1.1 of [75]. The circularly polarised 795 nm seed beam of 1 μW power drives the σ^- transition between $|0\rangle$ and $|3\rangle$. The circularly polarised 780 nm pump beam, of opposite handedness, of 1 μW power drives the σ^+ transition between $|0\rangle$ and $|1\rangle$. The 776 nm 30 mW coupling beam drives $|1\rangle$ to $|2\rangle$. High intensities[6] are achieved by focussing the beams to a ~60 μm $1/e^2$ radius at the cell centre. The generated 762 nm signal beam ($|2\rangle \rightarrow |3\rangle$) fulfils the phase-matching condition: $\mathbf{k}_{\text{pump}} + \mathbf{k}_{\text{coupling}} = \mathbf{k}_{\text{seed}} + \mathbf{k}_{\text{signal}}$. With the use of an edge mirror on a translation stage, the coupling beam is crossed at a small angle (10 mrad) to the optical axis. This allows the seed/pump beams, which lie along the optical axis, to be physically separated from the coupling/signal beams at the output of the cell so that the signal and seed beams can be detected at different detectors. A photodiode detects the transmitted seed beam and a photon detector detects the four-wave mixing signal, which has a well-specified direction given by phase-matching.

[5] The field profile is shown in Fig. 1 of reference [24].
[6] The equation for converting between power and intensity for a given beam waist is given in Eq. 7.3.

Full experimental set-up

As shown in Fig. 7.3, the seed (795 nm), pump (780 nm) and coupling (776 nm) light are generated by three external cavity diode lasers and coupled into polarisation maintaining fibres before arriving at the experimental cell. For all lasers, polarisation-matching into fibres is performed following the method of [170]. In brief, the method involves using a $\lambda/2$ and $\lambda/4$ waveplate before the fibre input. These waveplates are set at rotations that minimise the polarisation fluctuations at the fibre output, where the *polarisation* fluctuations are deduced by measuring the beam *power* fluctuations after the beam passes through a polarising beam splitter (PBS). These fluctuations can be more quickly observed if the fibre is briefly heated and allowed to cool, although this is not a necessary step.

The pump/780 nm light is emitted by a Toptica DL100 780 nm diode laser. After passing through an optical isolator, it is split into two paths via a $\lambda/2$ waveplate and PBS: one path is fibre coupled to the experimental vapour cell while the other path is split into two again via a PBS. One arm is coupled to a reference Fabry-Pérot etalon which is used as a frequency reference[7], and the other beam is sent to the STROLL locking system discussed in Chapter 5.

En route to the experimental cell, the light is passed through a narrowband 780 nm interference filter (Semrock LL01-780 with a measured extinction of 10^{-6}). This is only essential for the photon counting experiments in the next chapter (Chapter 8) as there is a small but significant number of 795 photons in the pedestal of the 780 spectrum—this must be eliminated before it reaches the detection end of the experiment. The light is *transmitted* through a PBS and then circularly polarised with a $\lambda/4$ waveplate before being focussed by a lens into the experimental cell.

The seed/795 nm light is emitted by a Toptica DL100 795 nm diode laser. As with the pump light, this light passes through another optical isolator and is split with $\lambda/2$ waveplate combined with a PBS. One beam is sent to the aforementioned Fabry-Pérot etalon to create a frequency reference for this laser. We can switch between observing 795 nm or 780 nm transmission peaks with the use of a removable magnetic mirror before the etalon. The rest of the light is sent to the experimental

[7] Section 3.3.

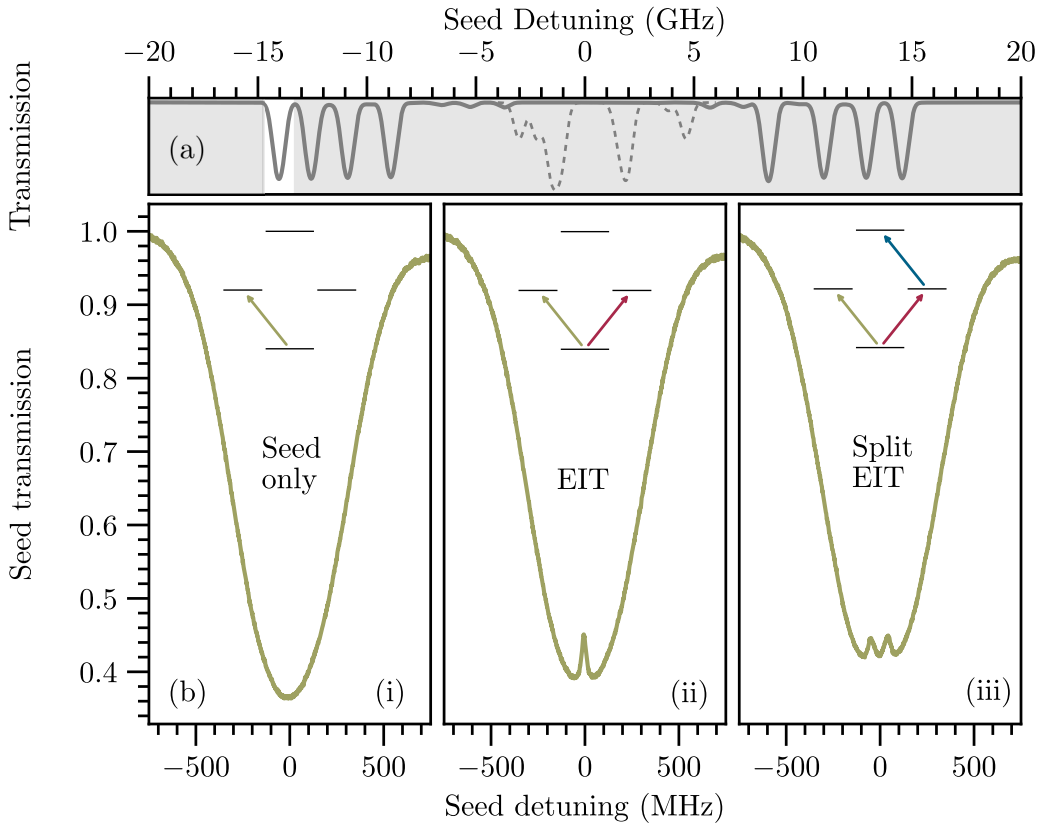

Figure 7.5: (a) Theoretical transmission spectrum of the rubidium D_1 (795 nm) line in the presence (solid grey line) and absence (dotted grey line) of a 0.62 T magnetic field with the vapour cell at a temperature of 80 °C. The unshaded region is the detuning range for the experimental data shown below. (b) Experimental spectra for the transmitted seed light in the presence of (i) only the seed beam (ii) both the seed and pump beams, and (iii) all three beams (seed, pump, and coupling).

cell. The seed light is *reflected* through the same PBS that transmits the pump light and so when it passes through the $\lambda/4$ waveplate it is circularly polarised in the opposite sense to the pump light. It is then focussed by the 20 cm lens into the cell.

The coupling/776 nm light is generated by a Toptica DL Pro 776 nm 100 mW external cavity diode diode laser, passed through an optical isolator and amplified by a tapered amplifier, as fully explained in Chapter 3.1. En route to the cell, it is split by a $\lambda/2$ waveplate and PBS. Some of the light is sent to the STROLL cell; the rest is reflected into the cell, through the 20 cm focusing lens, by an edge mirror that is atop a translation stage.

When setting up from scratch, it can be difficult to obtain the 4WM signal because very precise alignment is required to overlap three focussed beams at the centre of a 2 mm cell, especially as the cell is not visible to the naked eye when it is placed at the centre of the magnet holders. Initially, it is best to perform the alignment procedure without the magnet in place. In the next two subsections, we report the alignment (Section 7.3.1) and realignment (Section 7.3.2) procedure as a detailed set of instructions to aid experimenters working with the same or similar set-up.

7.3.1 Initial alignment

1. Position the cell on the table above an optical table hole.

2. Align the seed beam so that it passes through the cell and simultaneously passes over holes on both the input and output ends of the vapour cell. *This defines the optical axis.*

3. Insert a pair of irises at both the input and output ends of the cell but outside the footprint of the magnet—this will help with realignment in the future.

4. Insert the magnet holder containing the two magnets. The spectrum on the photodiode observing the seed light is shown in Fig. 7.5 (b)(i).

5. Align the pump beam along the optical axis by aligning the beam through

the irises using beam walking techniques. An electromagnetically induced transparency (EIT) signal should appear on the seed spectrum, as shown in Fig. 7.5 (b)(ii).

6. Walk the pump beam so that the amplitude of the EIT feature is maximised.

7. Insert a pair of focusing/defocusing lenses (20 cm) into the beam path at the correct positions.

8. Translate the edge mirror perpendicular to the optical axis so that the centre of the mirror blocks the pump and seed beams. Align the coupling beam through the irises such that the beam is aligned along the optical axis. At this point, all three beams would be perfectly overlapped and copropagating if the edge mirror were not blocking two of the beams.

9. Translate the edge mirror so that the distance between the centre of the lens and the coupling beam spot on the lens is 2 mm. This gives a 10 mrad angle of incidence.[8] Splitting of the EIT feature should be evident, as shown in Fig. 7.5 (b)(iii).

10. Maximise the splitting of the EIT signal by walking the coupling beam. Insert one more pair of irises in the path of the coupling beam to help with future realignment of this beam.

11. The photodiode observing the signal/762 nm light (See Fig. 7.4) should now show a 4WM spectrum, as shown in the right panel of Fig. 7.2.

7.3.2 Realignment

It is important that the previous alignment be performed as well as possible as even slight misalignment leads to some asymmetry in the lineshapes. Due to changes in the laboratory environmental conditions and the mechanical relaxation of optical

[8]This value for the angle of incidence was kept constant throughout all the experiments in this thesis involving the coupling beam. The value was chosen so that the phase mismatch is minimal while still allowing the output beams to be easily separated from the input beams. An interesting future direction for this project would involve checking how the 4WM signal varies with the angle, as more collinear beams may give a better signal.

components, all of the beams require periodic realigning. It is hoped that in the future this realignment can be automated using the Pi Auto-aligner of Chapter 4. Here we outline a possible manual realignment protocol. For all three input beams, begin by maximising the light through their respective irises. It is likely that this will not give the best possible signal. Therefore, it is useful to block the coupling laser and maximise the EIT signal by walking the pump beam. Then unblock the coupling laser and maximise the splitting of the EIT signal by walking the coupling beam.

We found that on those occasions where the 4WM signal was completely lost, it is useful to first detect the blue light (420 nm photons) that arises from the second step of the two-step decay from the upper state to the ground state: $5D_{3/2} \rightarrow 6P_{3/2} \rightarrow 5S_{1/2}$ (See Fig. 1 of [171]). We place a photomultiplier tube (PMT, Hamamatsu H10682-210) with a blue light filter perpendicular to the optical axis so that blue light fluorescence can be detected. Exciting to the upper state requires only the pump and coupling lasers to be switched on (i.e., the seed laser can be blocked). So, maximising blue light detection is one way of ensuring maximal overlap between these two lasers.

In summary, to maximise the overlap of the three beams, first overlap the pump and seed beams by maximising the EIT signal. Then, translate the edge mirror to a position such that the coupling beam crosses at 10 milliradians. Finally, overlap the coupling beam, with the pump beam fixed, by maximising the blue light signal. This procedure results in a maximised 4WM signal.

7.3.3 Detection

There are two detection "arms" to the experiment: in both cases the experimental set-up terminates with the capture of photons on photon detectors. These detectors (Excelitas SPCM-AQRH-14-FC) are connected to a timing device, a SensL HRM-TDC, which precisely records the time of arrival of the photons. Although photon detectors are essential for the experiments done in the next chapter on spontaneous 4WM, they are not necessary here. However, their use gives 4WM spectra with an excellent signal-to-noise ratio. For alignment purposes it is easier to *start* by

looking at both the seed spectrum and the 4WM signal on *photodiodes* (Hamamatsu-C5460) which have a much larger active surface area than photon detectors. A removable mirror on a magnetic mount allows us to switch between observing the four-wave mixing 762 nm/signal light on the photodiode and the photon detector. We calibrate the photon detector and photodiode by detecting a beam of light of known power (1 μW attenuated sufficiently using ND filters).

On the seed/pump (795 nm/780 nm) detection arm, the 795 nm photons are separated from the 780 nm photons by the use of a 795 nm interference filter (IF) (Semrock LL01-808, placed at an angle to allow only 795 nm light through[9], giving a measured extinction of 10^{-6}) and by polarisation filtering with a $\lambda/4$, $\lambda/2$ and Glan-Taylor polariser giving an extra 10^{-4} extinction. A beam splitter on a magnetic mount allows observation of the 795 nm weak probe transmission when required.

On the coupling/signal (776 nm/762 nm) detection arm, the 776 nm light is filtered out using two 780 nm interference filters (Semrock LL01-780 with a measured extinction of 10^{-6}) which are set at an angle so that only 762 nm light passes through. Two filters are required because there are milliwatts of 776 nm power on this arm, compared to the microwatts of 795 nm power on the other arm. Polarisation filtering using a $\lambda/4$, $\lambda/2$ and Glan-Taylor polariser is also performed.

[9]We use a filter designed for 808 nm because at the time of purchase there was no readily available filter specified for 795 nm.

7.4 Results

Here we report the dependence of the 4WM signal on atomic number density, seed power, pump power, coupling detuning, and coupling power. In the next section (Section 7.5), we will show how to use the dependence on coupling power to extract the relevant dipole matrix element. By characterising the seeded 4WM system, we show that it behaves as expected from our model, and we position ourselves to perform the single-photon experiments of the spontaneous 4WM system of the next chapter.

For clarity, we note that the phrase "on resonance" is used to mean that the detunings of all the lasers are zero, i.e., $\Delta_P = \Delta_C = \Delta_S = 0$; this is the only case where the doublet feature of the 4WM signal is perfectly symmetrical. Perfect symmetry is in practice very difficult to achieve as it requires exquisite alignment of the beams, as described in the previous section.

The waists of the three beams are measured by taking images on a CCD camera and fitting a Gaussian curve to each image. The error is the statistical error from fitting many images. This gives waists for the seed, pump and coupling beams of (60 ± 1) μm, (60 ± 1) μm and (63 ± 1) μm respectively.

7.4.1 Dependence on atomic number density

We increase atomic number density by increasing the temperature of the cell and find that the 4WM signal intensity scales quadratically with atomic number density, as shown in Fig. 7.6. This is as expected since the intensity is proportional to the square of the electric field and the electric field scales linearly with the number of atoms because 4WM is a coherent effect. Above a temperature of 80 °C (i.e., number densities above 5×10^{18} m^{-3}), the 4WM signal saturates because the vapour cell becomes optically thick for ground-state transitions.

7.4.2 Dependence on seed power

Increasing the seed power increases the 4WM signal, as shown in Fig. 7.7. We estimate the saturation parameter by inspection of the graph. It becomes apparent

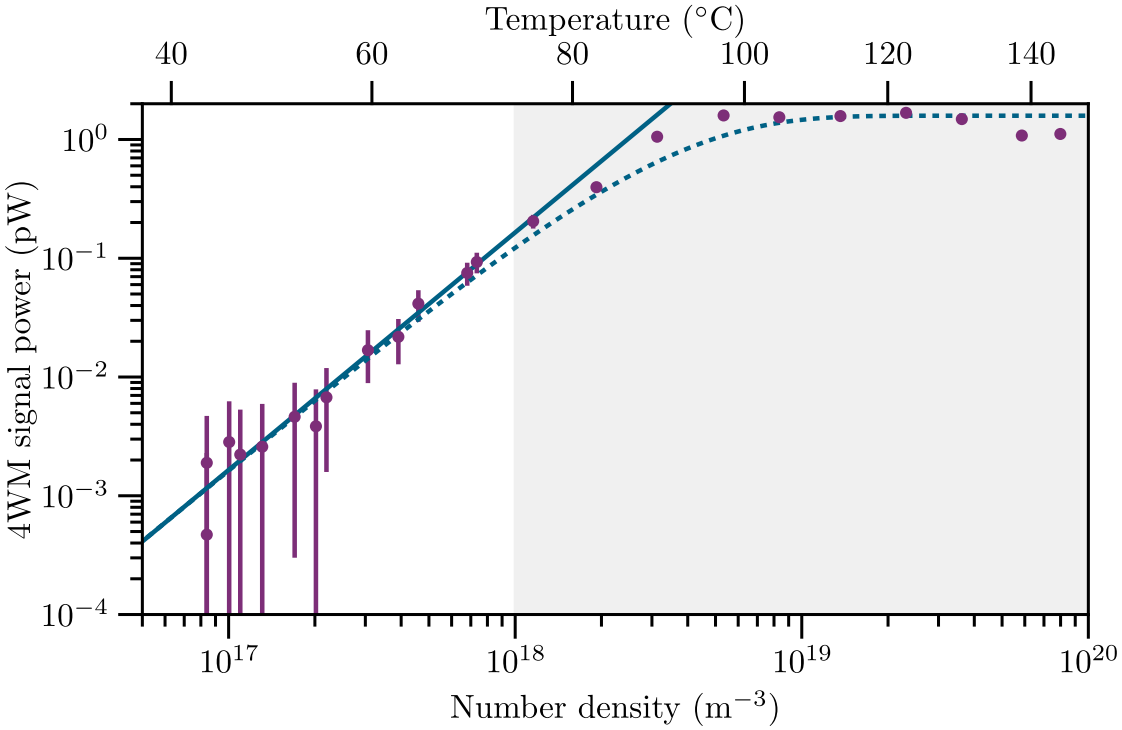

Figure 7.6: The resonant four-wave mixing signal peak power as a function of vapour cell temperature and atomic number density. The signal has an initial quadratic rise (solid blue line) after which it saturates at a temperature of ~90 °C. The unshaded region indicates the points included in the fit for the solid line.

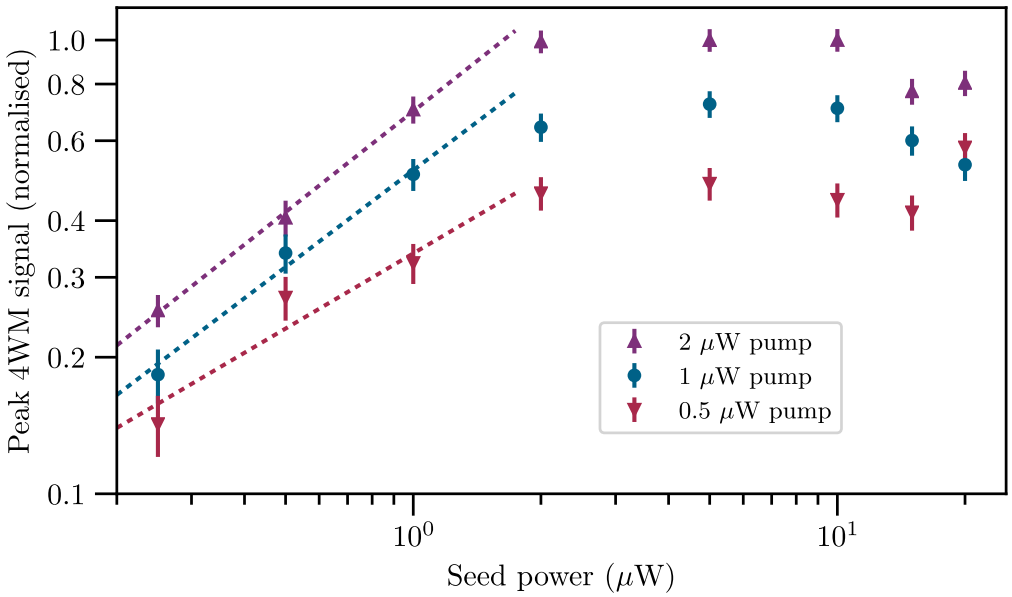

Figure 7.7: The 4WM signal increases with seed power, saturating at $P_{\text{sat}^{795}} = 2$ μW. The three curves correspond to power law fits to the first four points for pump powers 0.5, 1, and 2 μW. If the pump power is increased, the signal increases but $P_{\text{sat}^{795}}$ is unchanged.

at $P_{\text{sat}^{795}} = 2.5$ μW. [10] The signal also increases with pump power but $P_{\text{sat}^{795}}$ remains the same. For the pump powers of 0.5, 1, and 2 μW, the initial increase can be fitted with power laws with exponents 0.6 ± 0.1, 0.7 ± 0.1 and 0.7 ± 0.1 respectively. For seed powers above 10 μW, the signal decreases again, which we attribute to power broadening of the pump transition.

7.4.3 Dependence on pump power

Fig. 7.8 and 7.9 shows that increasing the pump power increases the 4WM signal. Between 0 and 20 μW, the 4WM signal increases linearly. By inspection of the graph, we estimate that saturation occurs at $P_{\text{sat}^{780}} = 40$ μW.[11]

The linear increase in the 4WM signal with pump power can be fitted with

[10] This is a Rabi frequency, $\Omega_{\text{sat}^{795}}$, of 15 ($\times 2\pi$) MHz.
[11] This is a Rabi frequency, $\Omega_{\text{sat}^{780}}$, of 68 ($\times 2\pi$) MHz.

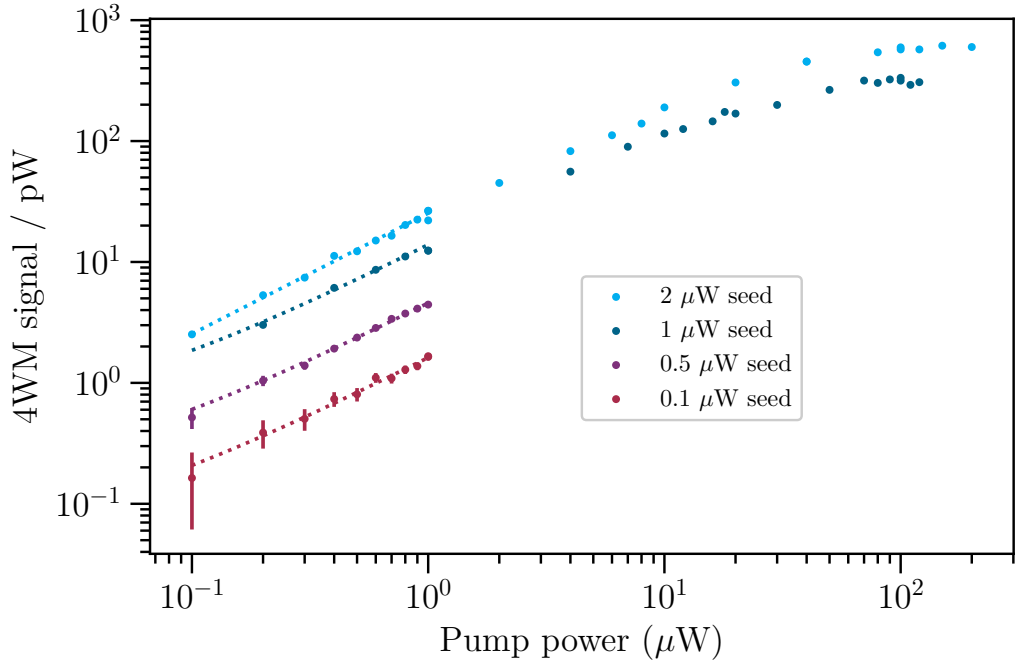

Figure 7.8: The 4WM signal increases linearly with pump power between 0 to 20 μW. The signal saturates at 42 μW. The dotted lines are fits to straight lines.

straight lines of the form $y = mx + c$. At a cell temperature of 114 °C, for seed powers of 0.1 μW, 0.5 μW, 1.0 μW and 2.0 μW, $m = 1.57$ pW/μW, 4.41 pW/μW, 13.38 pW/μW, 25.33 pW/μW and $c = 0.05$ pW, 0.16 pW, 0.5 pW, 0.01 pW respectively. The errors are given by the fitting procedure[12] and are $\Delta m = \pm 0.01$ pW/μW and $\Delta c = \pm 0.01$ pW. We would expect c to equal 0 and the small nonzero value is likely to be an indication of unsubtracted background light.

7.4.4 Dependence on coupling detuning

The strong coupling beam dresses the pump transition, which creates two paths to the 5D state. As we see in the colour map, Figure 7.10 (b), there are two peaks in the 4WM signal corresponding to the resonances with these dressed states. The colour map shows that the theoretical 4WM signal has contributions from many different atomic velocity classes. Above the colour map on panel (a) each

[12]Weighted least-squares minimisation, the lmfit module in python.

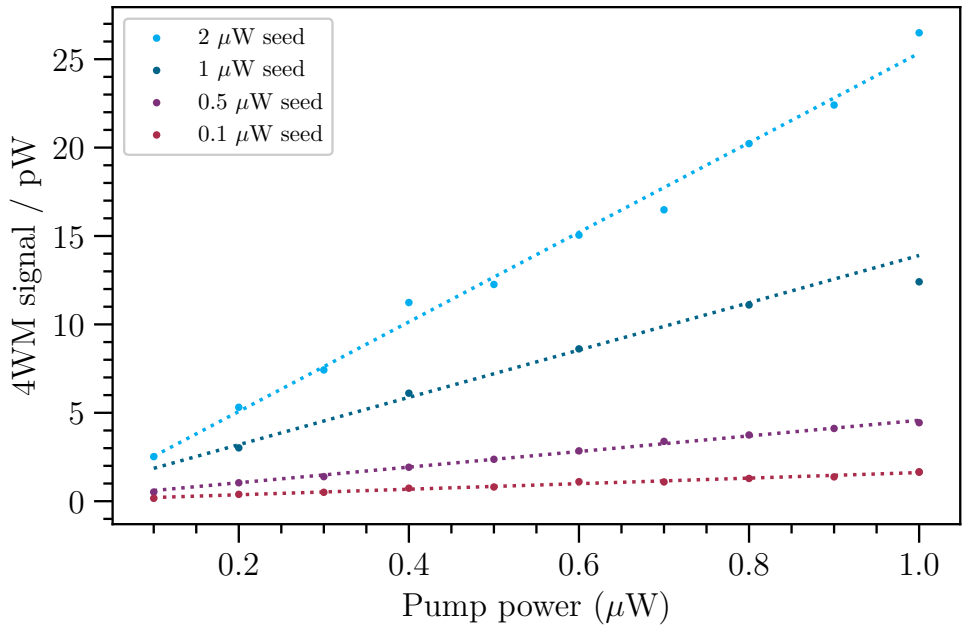

Figure 7.9: Same data as in Figure 7.8 but showing only the fitted region. The dotted lines are fits to straight lines, demonstrating clearly that the 4WM signal increases linearly with pump power. Error bars are too small to be seen.

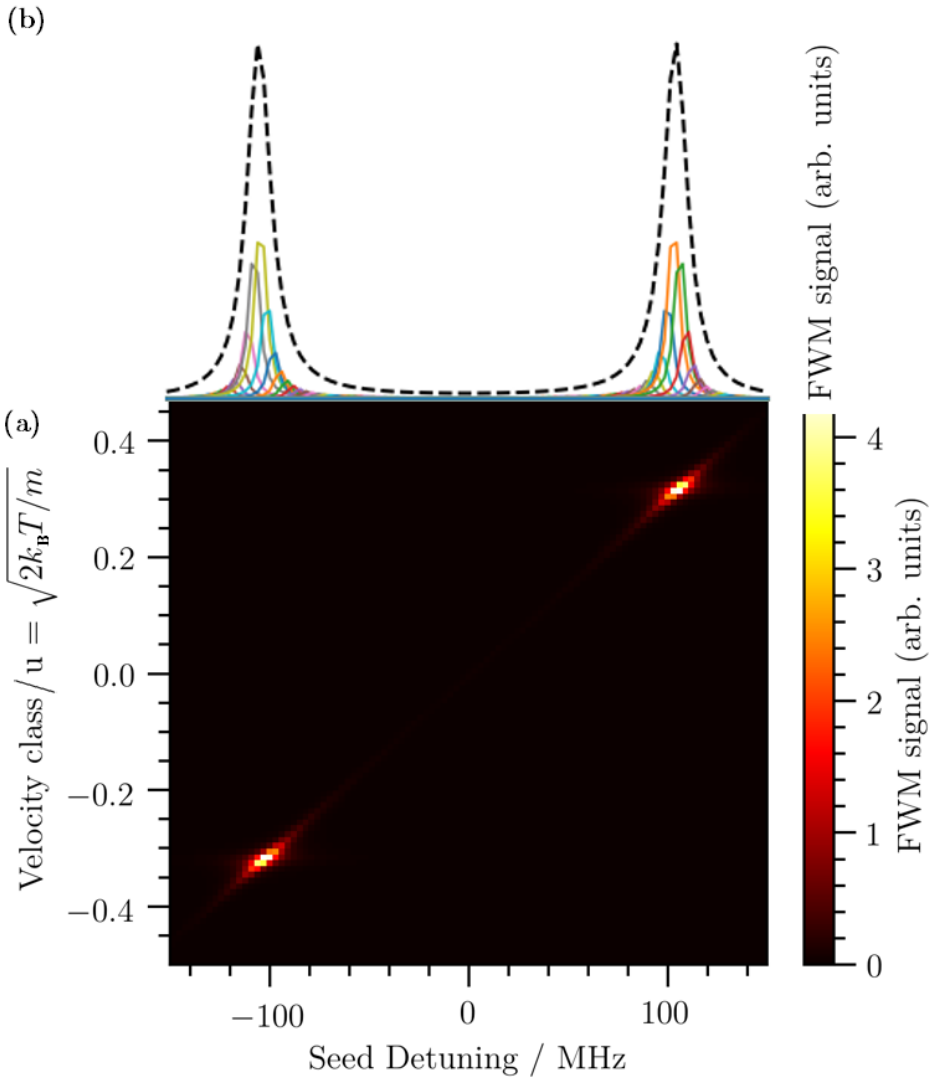

Figure 7.10: (a) The theoretical 4WM signal generated by our computational model has individual contributions from many (but not all) velocity classes. (b) Each of the coloured 4WM signals shown above the colourmap is the contribution from one velocity class as the seed laser is scanned across resonance. The dotted envelope is the sum of these contributions and matches the observed experimental 4WM signal.

individually coloured 4WM signal results from the contribution of one velocity class as the seed laser is scanned across resonance. The dotted envelope is the sum of these contributions.

If we now change the coupling detuning, different velocity classes are involved and the splitting of the peaks in the 4WM signal changes, as does the asymmetry. This is seen in Fig. 7.11 which shows how the splitting of the experimental signal changes as the coupling detuning changes. Only on resonance, where the coupling detuning is zero, are the peaks symmetric.

7.4.5 Dependence on coupling power

Increasing the coupling power leads to a splitting of the dressed-state energy levels. The coupling Rabi frequency, Ω_c, is proportional to the square root of the coupling power and the dressed states split proportionally to $\hbar\Omega_c$ (Autler-Townes splitting). We find that the splitting of the 4WM doublet follows this square-root relationship.

Initially, the absolute value of the 4WM signal increases with coupling power but the signal saturates at a coupling power of 10 mW [25]. Nevertheless, the Autler-Townes splitting, S, continues to increase even above these powers, as can be seen in Fig. 7.12. The asymmetry of the lineshape in the figure is due to imperfect overlap of the three beams in the cell. We can make the features less asymmetric by tweaking of the alignment of all three beams. Nevertheless, it is difficult to achieve and keep perfect symmetry because of exquisite sensitivity to alignment. We find that this does not change the frequency position of the peaks. For fixed beam powers, when all three beams are on resonance with their respective transitions, the splitting between the peaks is minimised. Asymmetry can also be due to the pump or coupling beam being off-resonance. In the case of the data shown, we know this is not the case because any attempt to change these detunings in the experiment led to an increase in splitting. In other words, we set the coupling, pump, and seed powers to fixed values, and then we adjust the pump and coupling detunings so that S is minimised. This ensures that the coupling and pump beams are exactly on resonance.

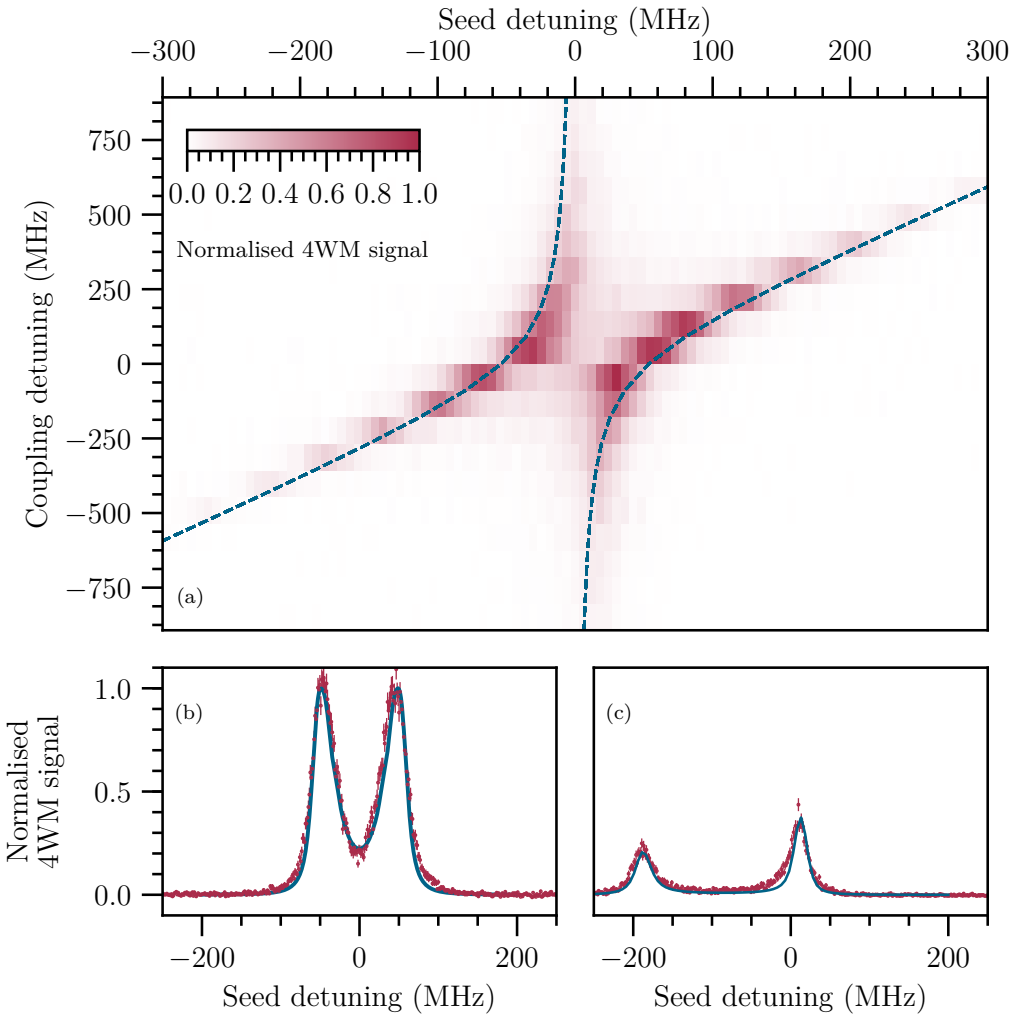

Figure 7.11: (a) The four-wave mixing (4WM) signal as a function of both seed (795 nm) and coupling (776 nm) detuning, where the avoided crossing indicates Autler-Townes splitting. The theoretical prediction of this splitting is displayed as a solid blue line. Panels (b) and (c) show individual 4WM spectra at coupling detunings of 0 and -360 MHz respectively, with theory overlaid.

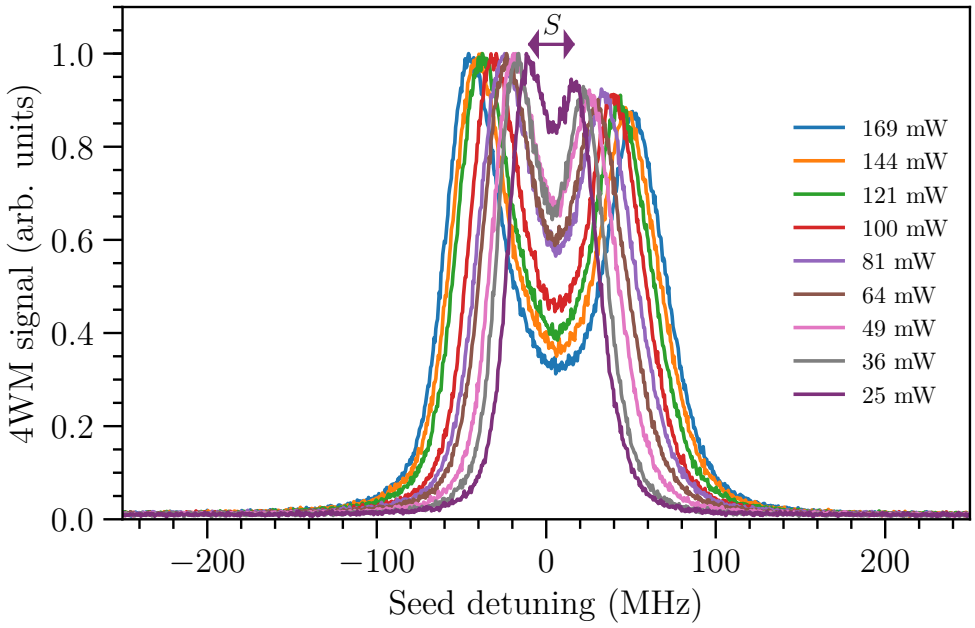

Figure 7.12: Experimental data showing that the splitting, S, of the peaks of the 4WM signal increases with coupling power. The asymmetry of the lineshape is due to a slight imperfection in the overlap of the three beams in the cell.

The separation between the peaks of the doublet, which we call the splitting, S, is proportional to the dipole matrix element of the excited state transition. From careful measurements of S we can extract the dipole matrix element of the excited state transition.

For a two-level system, $S = \Omega$, i.e., the proportionality constant between splitting and Rabi frequency is one [16]. This is not true for a four-level system: when we input a particular value for the coupling Rabi frequency, Ω_c, into the computational model, we find that the splitting of the generated theoretical signal, S is always less than the coupling Rabi frequency such that

$$S = m' \times \Omega_c, \tag{7.2}$$

where $m' < 1$. The exact value of the proportionality constant m' depends upon the power of the seed and pump beams. We ran simulations for different seed and pump powers to determine the correct value of m' for the experiment. In the limit, as these beams are made weaker and weaker, we find $m' = 0.725$. This is shown in Fig. 7.13, where it can be seen that m' is only a constant at very low pump powers. Unfortunately, we cannot make the beams arbitrarily weak in the experiment without losing the 4WM signal, meaning there is some uncertainty in the m' associated with any particular experimental run. The powers of the pump and seed beams used in the experiment were 0.9 ± 0.1 μW, which, as seen in the figure, gives $m' = 0.83^{+0.07}_{-0.02}$.

7.5 Extracting a dipole matrix element

In this section, we show how the splitting of the four-wave mixing signal can be used to extract the relevant dipole matrix element. The transition of interest is that addressed by the 776 nm coupling laser, which has the dipole matrix element $|\langle 5P_{3/2},\ m_J = 3/2|er|5D_{3/2},\ m_{J'} = 1/2\rangle|$. Using Wigner 6-j and 3-j symbols, we can also determine $|\langle 5P_{3/2}||er||5D_{3/2}\rangle|$, $|\langle 5P_{3/2}||er||5D_{5/2}\rangle|$, and $|\langle 5P||er||5D\rangle|$, as explained in the next subsection.

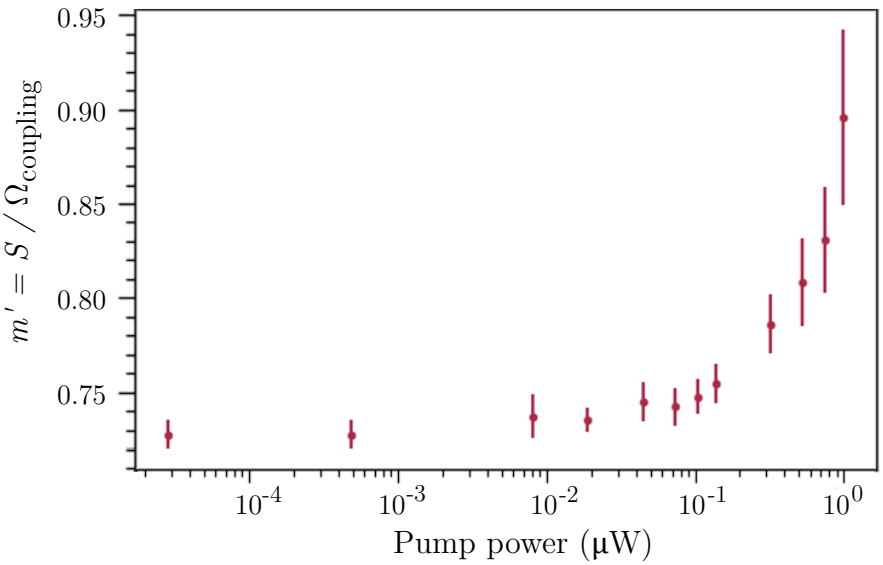

Figure 7.13: Data from simulations. The scaling factor, m', that relates the splitting of the peaks, S, to the coupling Rabi frequency, Ω_{coupling}, is only a constant at very low pump powers ($< 10^{-2}$ μW).

Experimental measurement of dipole matrix elements is of more than pure intellectual interest. The importance of the dipole matrix element becomes clear when Fermi's golden rule is recalled: the transition probability between two states is proportional to the associated matrix element squared [26]. In other words, the strength of a transition is determined by the magnitude of the dipole matrix element, and a transition is "forbidden" when this element goes to zero.

The Rb $|\langle 5P||er||5D\rangle|$ matrix element, which we determine here, is a commonly used transition in atomic physics [172–175]. Even when not addressed directly, it has to be accounted for in those rubidium experiments (e.g., atomic traps) where the degree of the A.C. Stark shift (i.e., the light shift) needs to be established. To determine this shift, the polarisability is calculated from the sum of matrix elements and from the beam detunings.

Safronova (2004) [176] has calculated a value from theory for many atomic transitions, including this one ($|\langle 5P_{3/2}||er||5D_{3/2}\rangle|$). However, it is particularly difficult to measure dipole matrix elements associated with excited state transitions [24]. To

our knowledge, at the time of writing, no experimental value had been established. It had certainly not been established using our method, which we believe to be novel. Because our method relies on the dipole matrix element being directly proportional to the splitting of the 4WM signal, the experiment is relatively simple to perform.

7.5.1 Calculating $|\langle 5P_{3/2},\ m_J = 3/2||er||5D_{3/2},\ m_{J'} = 1/2\rangle|$

We first show how to calculate from theory the expected value of 0.239 ea_0 for $|\langle 5P_{3/2},\ m_J = 3/2||er||5D_{3/2},\ m_J = 1/2\rangle|$, and then we compare against the experimental value. We begin with the fully reduced dipole matrix element $|\langle 5P||er||5D\rangle|$, which from experimental measurements was found to be 2.069 ea_0 [177]. From this we can obtain $|\langle 5P_{3/2}||er||5D_{3/2}\rangle|$ using the Wigner 6-j symbol [74, 178] as so:

$$|\langle J||er||J'\rangle| = \sqrt{(2J+1)(2J'+1)} \begin{Bmatrix} J' & 1 & J \\ L & S & L' \end{Bmatrix} |\langle L||er||L'\rangle|,$$

where $J = 3/2$, $J' = 3/2$, $L = 1$, $L' = 2$ and $S = 1/2$, giving

$$\begin{aligned} |\langle 5P_{3/2}||er||5D_{3/2}\rangle| &= \sqrt{4\times 4} \times \frac{1}{2\sqrt{30}} |\langle 5P||er||5D\rangle| \\ &= \frac{2}{\sqrt{30}}(2.069)\ ea_0 \\ &= 0.755\ ea_0. \end{aligned}$$

We can then obtain the more specific transition matrix element $|\langle 5P_{3/2},\ m_J = 3/2||er||5D_{3/2},\ m_{J'} = 1/2\rangle|$ using the Wigner 3-j symbol [24, 178] as so:

	(ea_0)	ratio
$\lvert\langle 5P\lvert\lvert er\rvert\rvert 5D\rangle\rvert$	2.069	1
$\lvert\langle 5P_{3/2}\lvert\lvert er\rvert\rvert 5D_{3/2}\rangle$	0.755	$\frac{2}{\sqrt{30}}$
$\lvert\langle 5P_{3/2}, m_J = 3/2\lvert\lvert er\rvert\rvert 5D_{3/2}, m_{J'} = 1/2\rangle$	0.239	$\frac{2}{\sqrt{30}} \times \frac{1}{\sqrt{10}}$
$\lvert\langle 5P_{3/2}\lvert\lvert er\rvert\rvert 5D_{5/2}\rangle$	2.266	$\frac{\sqrt{30}}{5}$
$\lvert\langle 5P_{3/2}, m_J = 1/2\lvert\lvert er\rvert\rvert 5D_{5/2}, m_{J'} = 3/2\rangle$	0.717	$\frac{\sqrt{30}}{5} \times \frac{1}{\sqrt{10}}$

Table 7.1: Dipole matrix elements calculated from the measured value in row one, where the measured value is from [177].

$$|\langle J, m_J||er||J', m_{J'}\rangle| = \begin{pmatrix} J & 1 & J' \\ -m_J & q & m_{J'} \end{pmatrix} |\langle J||er||J'\rangle|,$$

where $J = 3/2$, $J' = 3/2$, $m_J = 3/2$, $m_{J'} = 1/2$ and $q = m_J - m_{J'} = 1$, giving

$$\begin{aligned} &|\langle 5P_{3/2},\ m_J = 3/2||er||5D_{3/2},\ m_{J'} = 1/2\rangle| = \frac{1}{\sqrt{10}}|\langle 5P_{3/2}||er||5D_{3/2}\rangle| \\ &= \frac{1}{\sqrt{10}} \times 0.755\ ea_0 \\ &= 0.239\ ea_0. \end{aligned}$$

These results are tabulated in Table 7.1, along with the other matrix elements of interest that we have recalculated here and which can be compared with [177].

7.5.2 Analysing experimental data

In order to obtain the dipole matrix element, we must start with the properties of the strong coupling laser beam. For a Gaussian beam with waist w, we can write the intensity as $I = I_0 e^{-2r^2/w^2}$, where r is the distance from the centre and I_0 is the intensity at the centre. The total power P is the integral of the intensity over the area of the beam in the x-y plane, which gives [73]

$$I_0 = \frac{2P}{\pi w^2}. \tag{7.3}$$

We know that the intensity of a plane EM wave is $I_0 = \frac{1}{2} cn\epsilon_0 E_0^2$, where c is the speed of light, n is the refractive index of the medium, ϵ_0 is the permittivity of free space, and E_0 is the strength of the electric field. Rearranged, this gives

$$E_0 = \sqrt{\frac{2I_0}{cn\epsilon_0}}. \tag{7.4}$$

Lastly, the angular Rabi frequency is

$$\Omega_{\text{ang}} = \frac{dE_0}{\hbar},$$

where d is the strength of the dipole moment given by the appropriate dipole matrix element. The angular Rabi frequency, Ω_{ang}, and the linear Rabi frequency, Ω_{lin} are related as so: $\Omega_{\text{lin}} = \frac{\Omega_{\text{ang}}}{2\pi}$. This gives

$$\Omega_{\text{ang}} = \frac{2d}{\hbar w}\left(\frac{P}{cn\epsilon_0\pi}\right)^{1/2},$$

and

$$\Omega_{\text{lin}} = \frac{d}{\hbar\pi w}\left(\frac{P}{cn\epsilon_0\pi}\right)^{1/2}. \tag{7.5}$$

Since we will be using linear Rabi frequencies throughout for all beams, we shall drop the subscript. We see that Eq. 7.5 is of the form $\Omega = m\sqrt{P}$ and so a plot of linear Rabi frequencies against the square root of the input power should give a

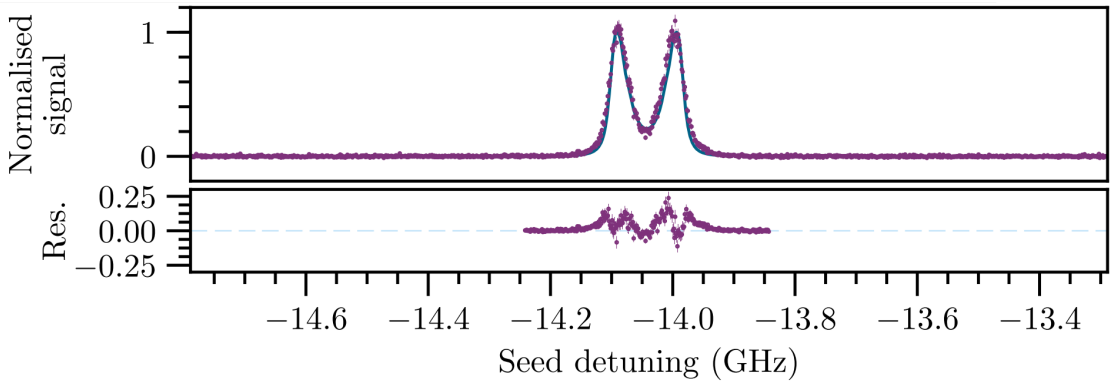

Figure 7.14: Four-wave mixing signal as a function of seed/795 nm beam detuning, with the pump and coupling beams on resonance, and the residuals (experiment minus theory) plotted below. The purple points are experimental data and the blue line is the theoretical signal given by the computational model. Zero detuning is the weighted D1 line centre of rubidium (of natural abundance) in zero magnetic field [132].

straight line going through the origin. From the gradient of this line, $m = \Omega/P^{1/2}$, we can extract the dipole matrix element, d, as so:

$$m \times \hbar\pi w(cn\epsilon_0\pi)^{1/2} = d. \tag{7.6}$$

In the experiment, we measure the splitting S and not the Rabi frequency Ω, and we plot $S = M\sqrt{P}$ and extract M and not the required m. However, from Eq. 7.2, we know these are simply related: in short, we obtain m from M by dividing through by the scaling factor m', i.e., $M/m' = m$.

7.5.3 Estimating the error

Figure 7.14 shows the experimental 4WM signal (purple points) with the theoretical signal generated by the computational model (blue line) overlaid. Here we reanalyse the data for this curve. Measuring the splitting, S, of the peaks on the experimental data gives 95 MHz.

The power measured on the power meter was 25.2 ± 0.1 mW. Before this laser light reaches the atoms it passes through two glass surfaces each with 4% loss.

A multiplicative factor of 0.92 [13] is therefore required, giving an input power of 23.2 mW. We need to further reduce this value to account for the non-uniformity of the power across the Gaussian beam: a factor of 0.74 (equivalent to reducing the seed input Rabi frequency by 14%) gives an input power of 17.12 mW. Using Eq. 7.5, with $w = 63 \pm 2$ μm, this gives the seed Rabi frequency, $\Omega_S = 144 \pm 4$ MHz.

If this Rabi frequency value is inserted into the computational model, the splitting of the peaks on the *theoretical* signal is 106 MHz. So, the experimentally measured value of the splitting (95 MHz) is 0.90 times the value given by the model (106 MHz). This gives an estimate of the error we expect from our dipole matrix element measurement.

7.5.4 Experimental determination of the excited state dipole matrix element

We collected two distinct supersets of data, one with a coupling beam waist of $w = 110$ μm and one with $w = 63$ μm. Each superset contains many datasets which are experimental repeats used to determine the statistical error.

We take 4WM signal spectra (762 nm) for different coupling powers and measure the splitting of the peaks, S, for each spectrum. This doublet is due to Autler-Townes splitting and should equal m' times the Rabi frequency of the coupling laser as given in Eq. 7.2. An example spectrum is shown in Fig. 7.14.

For the first superset of data, where $w = 110$ μm, when the splittings, S, are plotted in Fig. 7.15 and Fig. 7.16, the expected square root dependence is evident. We use the gradient, M, of the gold line of Fig. 7.16 to obtain the dipole matrix element, d, using Eq. 7.6 and obtain the following. The gradient is $(2.83 \pm 0.01) \times 10^8$ s^{-1}W$^{-1/2}$. Dividing by $m' = (0.83 \pm 0.01)$ and multiplying by[14] $\hbar\pi w(cn\epsilon_0\pi)^{1/2}/ea_0$ gives 0.135 ea_0 , which is only 57% of the expected 0.239 ea_0.

[13] i.e., $100 \times (1-0.04)^2$
[14] Recall Eq. 7.6.

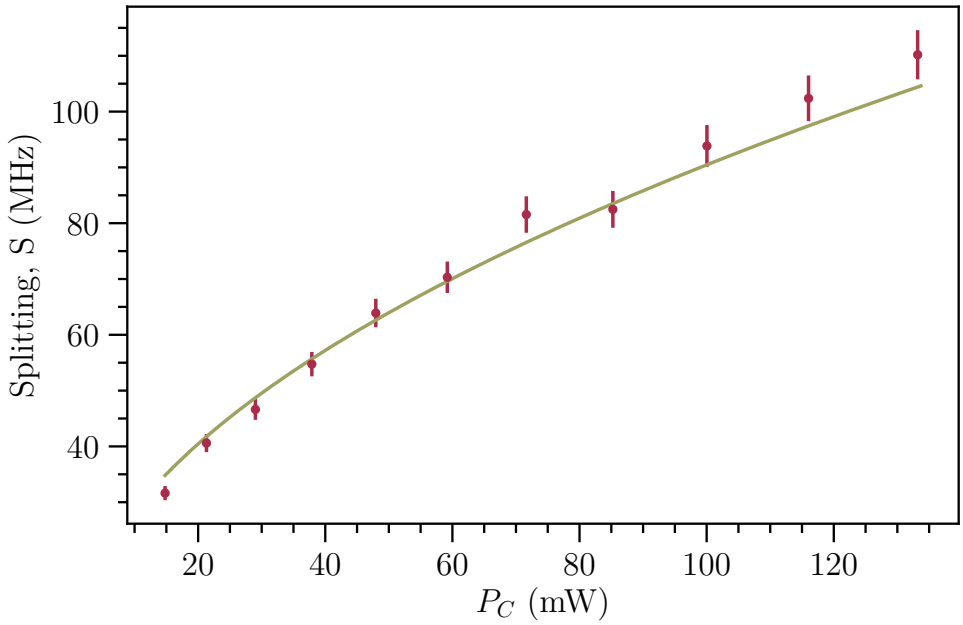

Figure 7.15: Extracted values of the splitting of the peaks, S, for increasing coupling beam power (P_C). The gold line is a least-squares fit to the function $S = M\sqrt{P_C}$. The square root dependence is evident. The beam waist is 110 μm.

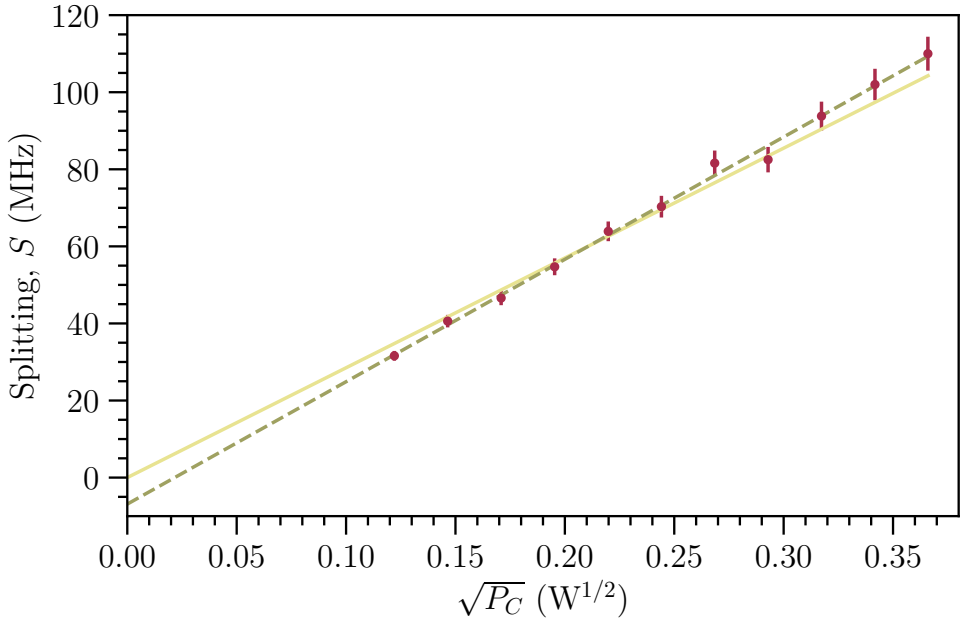

Figure 7.16: Extracted values of the splitting of the peaks, S, for increasing values of the square root of the coupling-beam power ($\sqrt{P_C}$). The gold line is a weighted least-squares fit to the function $S = M\sqrt{P_C}$. The dashed line is a weighted least-squares fit to the function $S = M\sqrt{P_C} + S_0$. This is the same data as Fig 7.15.

If we instead use the line of best fit that does not force the y-intercept, S_0, to equal zero, we obtain a gradient of $(3.18 \pm 0.01) \times 10^8$ $s^{-1}W^{-1/2}$ and a y-intercept of (-7 ± 1) MHz; this gives a final value of 0.150 ea_0 which is closer to the theoretical value but still only 63% of the expected. The reason for the nonzero intercept is likely to be that we did not manage to place the coupling laser exactly on resonance (i.e., the data indicates that it is (-7 ± 1) MHz away from resonance).[15]

For the second superset of data, where $w = 63$ μm, when the experimental data are plotted in Fig. 7.17, again the expected square root dependence is evident. Now when we use the gradient, M, of the gold line of Fig. 7.17 to obtain the dipole matrix element, d, using Eq. 7.6, we obtain a value closer to the theory. The gradient is $(6.32 \pm 0.01) \times 10^8$ $s^{-1}W^{-1/2}$. Dividing by $m' = (0.83 \pm 0.01)$ and multiplying by $\hbar\pi w(cn\epsilon_0\pi)^{1/2}/ea_0$ gives 0.171 ea_0 , which is 72% of the expected 0.239 ea_0. If we now use instead the line of best fit that allows a nonzero y-intercept, we obtain a gradient of $(6.60 \pm 0.01) \times 10^8$ $s^{-1}W^{-1/2}$ and a y-intercept of (-4 ± 1) MHz; this gives a final value of 0.179 ea_0 which is closer to the theoretical value but still only 75% of the expected.

A discussion on errors

We note that our best value for the matrix element is 25% lower than the theoretical value. Here we discuss several possible reasons. We believe that there are two main sources of error, the uncertainty in the measurement of the beam waist and the uncertainty in m'. For the first source, the difficulty lies in knowing the exact position where the three beams overlap within the cell and also placing the CCD element of the camera at exactly that position. Placing the camera is difficult as there is no physical access with the magnets in place. For the first superset of data (where the final value for the matrix element is only 63% of that expected), we attempted to measure the waist using a mirror: We first measured the distance from the first lens to the cell (20.0 cm) and then temporarily inserted a mirror and photographed the reflected beams at the equivalent distance. For the second superset of data, we recorded on the table the position of the cell and then removed

[15] Because of this difficulty with placing the coupling laser exactly on resonance, in the future, we should try to take repeat data for different detunings of the coupling laser.

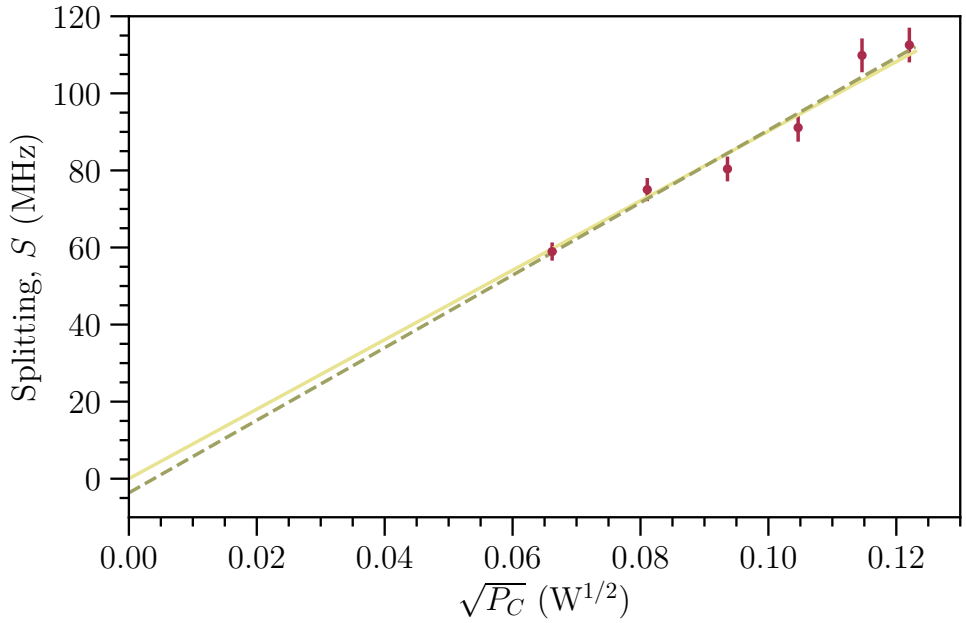

Figure 7.17: Extracted values of the splitting of the peaks, S, for increasing values of the square root of the coupling-beam power ($\sqrt{P_C}$). The gold line is a weighted least-squares fit to the function $S = M\sqrt{P_C}$. The dashed line is a weighted least-squares fit to the function $S = M\sqrt{P_C} + S_0$. The beam waist is 63 μm.

both the cell and magnets and then photographed the beams at that position. This gave a much closer value to the theoretical, only 25% less instead of 37%. But even with this method, there is some uncertainty as to exactly where the CCD element is inside of the camera. So there is still, potentially, ± 3 mm of error in the z-position. We know that the beam width diverges according to $w = w_0 \times \sqrt{1 + (z/z_\mathrm{R})^2}$ where $z_\mathrm{R} = \pi w_0^2/\lambda$, w_0 is the minimum waist, and z_R is the Rayleigh range. Using this equation we find that the error, when propagated, leads to ± 1.3 μm of error in the waist measurement, leading to up to a 5% change in the final value of the dipole matrix element calculation.

The other large source of error is the value of $m' = 0.83$, the proportionality constant that relates S, the splitting, to the Rabi frequency, Ω. Although we know from the computational model that m' is 0.725 in the limit of weak beams, we cannot get a large enough signal-to-noise ratio for such weak beams. This means that, for an

error of ± 0.1 μW, there is up to $^{+0.07}_{-0.02}$ of error on this constant. So comparing the dipole element resulting from using $m' = 0.81$ versus $m' = 0.90$, leads to up to a 20% change in the final value of the dipole matrix element calculation.

In order to perform a more accurate and precise measurement of this dipole matrix element, these two problems need to be solved: the exact value of both m' and the beam waist, w_0, needs to be established.

7.6 Conclusion

In this chapter, we have experimentally studied the nonlinear optical phenomenon of four-wave mixing. We devised a novel method of extracting an excited dipole matrix element but found that, for the level of precision required, we are limited by experimental constraints. However, our model allows us to understand how the four-wave mixing signal varies with different experimental parameters. In particular, we have shown how applying a large magnetic field simplifies the system to allow for agreement between model and experiment. We also characterised our system, which gives us a good understanding of the variables that alter the 4WM signal so that we are well-positioned to study spontaneous four-wave mixing, which is the topic of the next chapter.

Chapter 8

Spontaneous four-wave mixing leading to a heralded single-photon source

We have so far shown that when the phase-matching condition is met, turning on the seed, pump, and coupling lasers results in a *seeded* four-wave mixing (4WM) signal. In this chapter, the seed laser is switched off but, remarkably, the 4WM process still continues as *spontaneous* 4WM because of coupling between the vacuum and the doubly excited state. The 762 nm photon produced on the $|5D_{3/2}\rangle \rightarrow |5P_{1/2}\rangle$ transition and the 795 nm photon produced on $|5P_{1/2}\rangle \rightarrow |5S_{1/2}\rangle$ transition are now a simultaneously produced correlated pair. One member of this pair, the *herald* 762 nm photon, can be used to identify its correlated counterpart, the *signal* 795 nm photon. In this way, we turn our rubidium thermal vapour into a heralded single-photon source in the hyperfine Paschen–Back (HPB) regime.[1]

We begin the chapter by explaining the theory behind herald-signal photon pair production in spontaneous four-wave mixing. The experimental section that follows

[1]Please note the new labelling of the photons. In the wider literature they are often called signal (762 nm) and idler (795 nm). Our labelling clearly indicates that we are using the first 762 nm herald photon to announce the imminent (i.e., within a known time window) arrival of the second 795 nm signal photon. The bottom right of Fig. 8.1 shows the relevant energy levels.

is short as there are only slight modifications to the experimental section of the previous seeded 4WM chapter (Chapter 7). The section that follows details how to transform raw photon counts into the normalised cross-correlation function, $g^{(2)}_{\mathrm{h,s}}(\tau)$. We simulate such a situation computationally, both to gain insight into the function and also to explain the code which analyses experimental data. We then explain the Hanbury Brown–Twiss (HBT) experiment that is required to prove the single-photon nature of our source, and we show the result from a single HBT experiment: the classic $g^{(2)}(\tau)$ dip[2]. From several HBT experiments, we show the important experimental result that a higher $g^{(2)}_{\mathrm{h,s}}(\tau)$ peak gives a deeper $g^{(2)}$ dip. In other words, the greater the maximum correlation between herald and signal photons, $g^{(2)}_{\mathrm{h,s-MAX}}$, the closer the source is to an ideal single-photon source. The rest of the chapter investigates the parameters required to increase this correlation, namely, lower temperature, lower pump and coupling power, greater pump detuning, and use of the etalon lens filter of Chapter 3. In all cases, we find that we can increase the correlation but at the expense of a lower rate of single-photon production.

8.1 Introduction

Sources of single-photons are highly sought-after because photons are excellent quantum information carriers due to the combination of their high speed and long coherence times at room temperature [179–181]. Although photons are the fundamental unit of light, this does not mean that making single-photons is simple. An obvious method would be to simply attenuate a laser beam so that the mean photon number for a single laser pulse is very low. However, photons display Poissonian statistics [54], and there is no guarantee that exactly one photon will be delivered, as opposed to many or none.

Another conceptually simple method is to make use of spontaneous decay in a

[2]To avoid any later confusion, we shall state here that when the term $g^{(2)}(\tau)$ is used without the subscripts, it refers to the $g^{(2)}_{\mathrm{s_1,s_2|h}}(t_{\mathrm{s_1,s_2}})$ of the signal light in the HBT experiment. This is a subtle concept that will be explained in Section 8.6. For now, we note that a $g^{(2)}(\tau)$ plot shows a dip whilst a $g^{(2)}_{\mathrm{h,s}}(\tau)$ plot shows a peak.

single quantum emitter. For example, excite an atom [182] and allow it to decay; when it decays, a single-photon is released. The difficulty with this method lies in the uncertainty as to when and where the photon is released. The temporal uncertainty comes from spontaneous decay not being deterministic: although the decay has a predictable half-life, the exact time of photon release cannot, even in principle, be known beforehand. The spatial uncertainty arises from the spatial distribution of the emitted mode being such that the photon can be released within a very broad angular range. It is thus difficult to capture, especially against the background of other photons. It is possible to improve upon this by placing the quantum emitter in an optical cavity so that photon is preferentially emitted into the direction defined by the axis of the cavity. These sources are known as deterministic single-photon sources because the photons are generated on-demand. Photons produced using deterministic sources include those from single atoms [183, 184], single ions [185, 186], single molecules [187], quantum dots [188, 189] and colour centres in crystals [190, 191].

A different approach is to use a source that creates photon pairs, so that detection of the first member of the pair "heralds" the release of the second member of the pair. Since such sources usually involve many independent emitters, they are probabilistic sources of single-photons. Their advantage is that, due to the phase-matching condition, the photons are emitted into well defined spatial modes. Single-photons have been created in this way in different physical systems including nonlinear optical crystals (via spontaneous parametric downconversion [192]), photonic crystal fibres [193], and atomic ensembles [194], amongst others.

When extended atomic media—thermal vapours [195] and cold-atom ensembles [196]—are used as probabilistic sources, the emitted photons are naturally frequency- and bandwidth-matched to other atomic devices. In particular, the matching to atomic quantum memories allows the probabilistic source to be converted to a deterministic source by allowing the probabilistically generated photons to be stored until they are needed [197]. Spontaneous four-wave mixing has been used to create probabilistic sources in cold atoms [125, 198, 199] and in thermal vapours [2, 123, 152, 200, 201]. In the case of thermal vapours, typically optical pumping is required to isolate the necessary energy levels. As in the case of the seeded 4WM

performed in the last chapter, an alternative method of isolating the atomic levels is to enter the hyperfine Paschen–Back (HPB) regime[3].

In this chapter, we show how spontaneous four-wave mixing in thermal rubidium vapour performed in the HPB regime can be used to create a heralded source of single-photons.

8.2 Theory

We begin by summarising the *seeded* 4WM process using a concept—the *spin-wave*—which has so far not been needed but which will be useful in understanding *spontaneous* 4WM. This section draws heavily from the theoretical model developed by Whiting and Šibalić [6, 131]. Spatiotemporal polarisation oscillations in an atomic medium, called spin-waves, can result from continuous resonant driving of a two-level system. These oscillations of polarisation lead to an EM-field being emitted. In the case of a single resonant driving field, this field is both at the same frequency as the driving field and out of phase with it. It is the destructive interference between these two fields that results in the extinction of the driving field, which we have hitherto described with the absorption coefficient α. In a four-level system that is coupled with three driving fields, the resulting spin-wave oscillates at a sum or difference frequency of the driving fields, and an EM-field is emitted. This generated light is emitted from the uncoupled transition and is the seeded four-wave mixing signal; this whole process is known as seeded four-wave mixing (4WM) [165]. Under the phase-matching condition,

$$\mathbf{k}_{\text{pump}} + \mathbf{k}_{\text{coupling}} = \mathbf{k}_{\text{seed}} + \mathbf{k}_{\text{signal}}, \tag{8.1}$$

the fields emitted by all the atoms add constructively in one direction, and the 4WM signal light is emitted in the phase-matched direction.

Starting with the same atomic transitions as in Chapter 7 (shown again on Fig. 8.1), the seed beam is removed, and there remain only the pump and the coupling beams. Atoms are driven to the doubly excited state and, from there, there is

[3]The HPB regime is described in Chapter 2.

spontaneous decay via the herald and signal channels. As the two lasers are almost co-propagating, the spin-wave imprinted on the vapour has a wavenumber approximately equal to $k_\mathrm{p} + k_\mathrm{c}$. This is approximately twice that of the herald wavenumber, k_h, meaning that there is no phase-matched emission for herald photons. They are thus emitted in all directions. However, in the case of a herald photon that happens to be emitted in the forward direction, the resulting spin-wave has the wavenumber, $k_\mathrm{s} = k_\mathrm{p} + k_\mathrm{c} - k_\mathrm{h}$. Now there is a phase-matched emission for the signal photon in the forward direction. This means that a pair of photons, herald & signal, has been emitted such that detection of the herald photons tells us of the emission of a signal photon in the phase-matched direction. The phase-matching condition is now

$$\mathbf{k}_\mathrm{pump} + \mathbf{k}_\mathrm{coupling} = \mathbf{k}_\mathrm{signal} + \mathbf{k}_\mathrm{herald}\,. \tag{8.2}$$

We can quantify the extent of the correlation between the herald (subscript 'h') and signal (subscript 's') photons using the second-order cross-correlation $g^{(2)}_\mathrm{h,s}$ ([33]),

$$g^{(2)}_\mathrm{h,s}(\tau) = \frac{\langle \hat{E}^\dagger_\mathrm{s}(t+\tau)\hat{E}_\mathrm{s}(t+\tau)\hat{E}^\dagger_\mathrm{h}(t)\hat{E}_\mathrm{h}(t)\rangle}{\langle \hat{E}^\dagger_\mathrm{s}\hat{E}_\mathrm{s}\rangle\langle \hat{E}^\dagger_\mathrm{h}\hat{E}_\mathrm{h}\rangle}. \tag{8.3}$$

This is a measure of the likelihood of detecting a genuine signal photon (as opposed to a 'background' photon) at a time $t+\tau$, given that a herald photon has been detected at time t. The meaning of $g^{(2)}_\mathrm{h,s}$ will be explained more fully in the next few sections.

If we now consider the signal light field alone, measurement of the second-order autocorrelation[4] function allows us to demonstrate nonclassicality of the light. For classical light, $g^{(2)}(0) \geq 1$. In terms of the photon number distribution, the function is given by [33],

$$g^{(2)}(0) = 1 + \frac{(\Delta n)^2 - \langle n\rangle}{\langle n\rangle^2}, \tag{8.4}$$

where $\langle n\rangle$ is the mean photon number and $(\Delta n)^2$ is the photon number variance. A single-photon number state has $\langle n\rangle = 1$ and $(\Delta n)^2 = 0$, giving a $g^{(2)}(0) = 0$. This

[4]All of the 'h' subscript labels in Eq. 8.3 are replaced by 's'.

would be a perfectly pure single-photon source. By comparing how close $g^{(2)}(0)$ is to zero we can quantify how close we are to the ideal. Before explaining how raw data are transformed into the correlation function, we shall first detail the experimental procedure.

8.3 Experimental set-up

Figure 8.1 shows the experimental set-up for spontaneous four-wave mixing. A full table layout is shown in Appendix A. A vapour cell of length 2 mm containing 98% ^{87}Rb, which is housed within a heater, is placed between two NdFeB magnets. These cylindrical magnets have a top-hat profile [75] and allow us to achieve a magnetic field of 0.62 T across the vapour cell. There is a maximum variation of 1 mT across the 2 mm cell.[5]

In order to perform spontaneous 4WM, first the seeded 4WM light must be obtained. Because of the low amount of 4WM light generated, it is important to initially maximise the 4WM light using the variables explored in Chapter 7.[6] The seed beam is then turned off. The circularly polarised 780 nm pump beam drives the σ^+ transition between $|0\rangle$ and $|1\rangle$. The 776 nm coupling beam drives the σ^- transition between $|1\rangle$ to $|2\rangle$. High intensities[7] are achieved by focussing the beams to a ~60 μm $1/e^2$ radius at the cell centre. The generated 762 nm herald ($|2\rangle \rightarrow |3\rangle$) and 795 nm signal ($|3\rangle \rightarrow |0\rangle$) photons fulfil the phase-matching condition: $\mathbf{k}_{\text{pump}} + \mathbf{k}_{\text{coupling}} = \mathbf{k}_{\text{herald}} + \mathbf{k}_{\text{signal}}$. With the use of an edge mirror on a translation stage, the coupling beam is crossed at a small angle (10 mrad) to the optical axis. This allows the signal & pump light, which lie along the optical axis, to be physically separated from the herald & coupling light at the output of the cell. An optical fibre connected to a photon detector (D0) detects the herald beam; this is registered on Channel 0 (CH0) of the timing module. A split optical fibre collects the signal light and acts as a 50:50 beam splitter, sending

[5] The magnetic field is measured with a Hall probe and the value confirmed by fitting with ElecSus. The field profile is shown in Fig. 1 of reference [24].

[6] A high temperature is the most relevant variable for maximising the seeded 4WM light but the temperature will need to be decreased afterwards to obtain strong correlations, as we shall see.

[7] The equation for converting between power and intensity for a given beam waist is given in Eq. 7.3.

the signal photons to two photon detectors (D1 & D2); these are registered on Channel 1 (CH1) and Channel 2 (CH2) of the timing module respectively. The timing module precisely records the arrival of the photons with a resolution of 27 picoseconds. As CH0 is used as the herald channel, it is important that the BNC cable connecting D0 to CH0 is shorter than the cables connecting the other detectors to their respective channels, thus ensuring that the herald photon is registered before the signal photon.

8.4 Normalised cross-correlation, $g^{(2)}_{\mathrm{h,s}}(\tau)$

For any two detectors numbered i and j, we obtain $g^{(2)}_{i,j}$ from the probability of coincidence between the clicks from both detectors ($P_{i,j}$) and the probability of independent clicks, P_i and P_j on each detector, as so[8]:

$$g^{(2)}_{i,j} = \frac{P_{i,j}}{P_i P_j} = \frac{G^{(2)}_{i,j}}{r_i r_j \Delta t_c T}, \tag{8.5}$$

where $G_{i,j}$ is the total number of coincidences in measurement time T, r_i and r_j are the count rates on detectors i and j, and Δt_c is the width of the coincidence time window used for $G^{(2)}_{i,j}$. The term $g^{(2)}_{i,j}$ always refers to the normalised version of $G^{(2)}_{i,j}$. Thus, it can be convenient when analysing raw data from detector clicks to think of $g^{(2)}_{i,j}$ as the normalised double-coincidence (pair) events[9].

[8] This is Eq. 3.4 from [202] or Eq. 7 from [203].

[9] Similarly, $g^{(2)}_{i,j,k}$ can be thought of as normalised triple-coincidence events for the three detectors of Section 8.6.

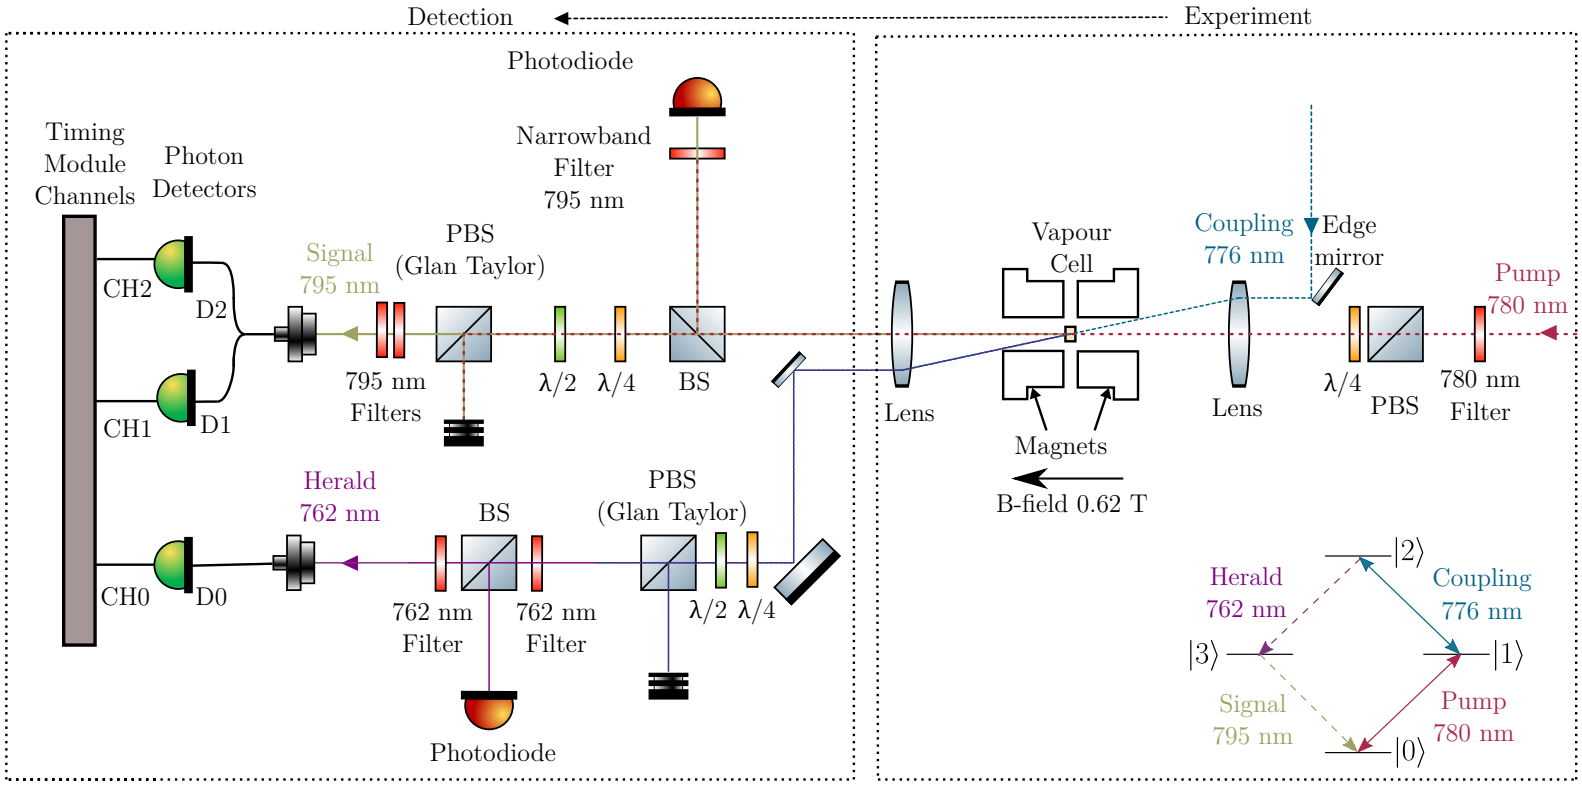

Figure 8.1: Diagram of the experimental set-up, read from right to left. Two beams pass through a heated 2 mm length vapour cell containing 98% rubidium-87. The pump beam is aligned along the 0.62 T magnetic field axis whilst the coupling beam is at a small angle (10 mrad). The quarter-waveplate ($\lambda/4$) sets the pump to be right-hand circularly polarised. The detected 762 nm *herald* photons and 795 nm *signal* photons are produced by spontaneous 4WM at the cell. The two beam-splitters (BS) are easily removable. The differences between this set-up and that of the last chapter are that the seed beam has been switched off and there are now three photon detectors connected to a timing module. The rubidium energy levels used are shown in the bottom right corner. $|0\rangle = |5S_{1/2}, m_J = \frac{1}{2}\rangle$, $|1\rangle = |5P_{3/2}, m_J = \frac{3}{2}\rangle$, $|2\rangle = |5D_{3/2}, m_J = \frac{1}{2}\rangle$, and $|3\rangle = |5P_{1/2}, m_J = -\frac{1}{2}\rangle$.

We shall model the situation where herald photon arrivals are registered on Channel 0 (CH0) and signal photon arrivals are registered on Channel 1 (CH1) of the timing module. The normalised cross-correlation, $g^{(2)}_{\mathrm{h,s}}(\tau) = g^{(2)}_{0,1}(\tau)$, is then given by [125]

$$g^{(2)}_{0,1}(\tau) = \frac{G^{(2)}(\tau)}{r_0 r_1 \Delta\tau T}, \tag{8.6}$$

where $G^{(2)}(\tau)$ is the histogram of coincidence events as a function of time delay τ, T is the total integration time, $\Delta\tau$ is the width of each time bin, and r_0 and r_1 are the rates of the channels 0 and 1 respectively. It is also possible to normalise $G^{(2)}(\tau)$ differently by replacing the denominator by $G^{(2)}(\tau = \infty)$ because $g^{(2)}_{0,1}(\tau \to \infty)$ is one [204]. Experimentally, we do both and find that the two methods agree. This agreement is a useful check that the analysis is correct—especially in the case of the slightly more involved case of the triple-coincidence events of the HBT experiment of Section 8.6.1. The previous equation (Eq. 8.5) allows us to provide a simple interpretation of $g^{(2)}_{\mathrm{h,s}}$. Since the background level is 1 [204], a $g^{(2)}_{\mathrm{h,s}}$ of some value x means that detection of a correlated signal photon is x times more likely than detection of an uncorrelated background photon. There is no theoretical upper limit to $g^{(2)}_{\mathrm{h,s}}$ [204].

To gain further insight into Equation 8.6, we shall first simulate the arrival of photons computationally. This will also explain how we transform the photon detection data taken in the lab to values for $g^{(2)}_{\mathrm{h,s}}(\tau)$. In the following, we will be showing both pseudocode and, on occasion, full python code. Whilst this may seem more detail than a physics thesis warrants, it is our hope that this level of detail will save future researchers time and effort[10].

8.5 Simulating photon arrivals

Photon arrival times obey Poissonian statistics, and so the probability of arrival is given by the cumulative distribution function for the exponential distribu-

[10] Typically, we include such details only when we ourselves would have benefitted if we had previously seen it in this form. In particular, it is hoped that thinking through double-coincidence events will make the case of triple-coincidence events simpler.

tion [16]:

$$\mathrm{P}(t) = 1 - e^{-\lambda t},$$

where the rate parameter, λ, is the average number of arrivals per second. To determine the time between two photons arriving, i.e., the *timegap*, we simply rearrange the above for t, giving

$$t = \frac{-\ln(1 - \mathrm{P}(t))}{\lambda}. \tag{8.7}$$

To transform this[11] into the exact time of each photon arrival (the *timetag*), we need to simply cumulatively sum the timegaps.

Correlated arrivals: Cross-correlation between CH0 and CH1

We can now model herald photons arriving at CH0 and signal photons arriving at CH1. Cumulatively summing the timegaps given by Eq. 8.7 generates a set of timetags for photon arrivals on CH0. To create correlated arrivals on CH1, we could simply duplicate this set. However, to bring our simulation closer to the experimental case, we model the delay associated with electrical signals travelling through BNC cables of different lengths by adding a delay to each CH0 timetag, and this becomes the set of timetags for photon arrivals on CH1. Clearly, in our experiment, we must ensure that the BNC cable from D1 to CH1 is longer than the cable from D0 to CH0.

Because this delay is not fixed but happens within the lifetime of the excited state (with additional broadening due to the atomic motion, see Sec 2.4), we can choose the delay to have a Gaussian distribution[12]. We can also choose the degree of correlation by only adding this delay to some percentage of the timetags from CH0. The rest of the set of CH1 timetags is uncorrelated and generated in the same way as CH0 timetags, i.e., using Eq. 8.7.

We use Eq. 8.6 to build the histogram of coincidence events. It is instructive to

[11] The inbuilt Python function random.expovariate(λ) provides exactly this. However, for the millions of counts we wish to simulate, it is faster to perform array operations using the NumPy module and write Eq. 8.7 as numpy.log(numpy.random.random(λ)) / λ.

[12] Pythonically, this is: delay = μ + (numpy.random.randn(Ch1_length))*σ

consider how this is done computationally as this is exactly how we analyse the data acquired by the photon-counting module during our experiments. We compile all the timetags at which clicks are detected as an ordered array. We also have a corresponding array that lists on which channel each click was detected. We identify the first timetag (t_n) at which a click on CH0 happens. Then, if the next timetag (i.e., t_{n+1}) corresponds to a click on CH1, we subtract the timetags to obtain a timegap. We then proceed sequentially through the array, adding to our list of timegaps every time a click on CH1 follows a click on CH0. We can then bin these timegaps to produce the final histogram.

The method described is summarised by the following pseudocode:

```
for every CH0-click:
  if CH0-click is followed by CH1-click:
    add to histogram timegap between CH0-click and CH1-click
```

It turns out that this is too slow in an actual experiment since there are millions of timetags to iterate through. This also prevents us from saving the raw photon counts in the experiment as 1 GB of data are generated every 10 seconds. This was one of the motivations for simulating the data generation: doing so allows us to test the analysis code without running the experiment. Array operations using the numpy module in Python speed up the code considerably:

```
timetags = time-ordered array of all timetags
channels = time-ordered array of the corresponding channel value (0 or 1)
indices = indices (of above arrays) where a CH0-click is followed CH1-click
timegaps = timetags[indices+1] - timetags[indices]
G = histogram of timegaps
```

The "histogram of timegaps" is the $G^{(2)}(\tau)$ term of Eq. 8.6 and is normalised using that equation.

There is a subtlety here which is easily overlooked: The signal click may not immediately follow the herald click. There may be one or more uncorrelated photons in between both clicks. The meaning of "`followed by`" in "`CH0-click is followed by a CH1-click`" must include these possibilities. This can be done by including an extra loop so that `timegaps = times[indices+1] - times[indices]` becomes `timegaps = times[indices+i] - times[indices]` where `i`, in principle, loops

through every timetag but, in practice, can be truncated at <10 for our time window of interest.

We expect the largest histogram bin to appear at the delay we set. We show in Fig. 8.2 the normalised cross-correlation in the case where we have chosen 50% of the timetags to be correlated, and indeed we see that the histogram peaks at a delay of 20 ns. We also found, as expected, that as the percentage of correlated pairs decreases, the normalised cross-correlation decreases.

So far, we have explained how to transform double-coincidence (pair) data into $g^{(2)}_{\mathrm{h,s}}(\tau)$; next, we shall explain how to turn triple-coincidence data into $g^{(2)}(\tau)$, and show experimental results.

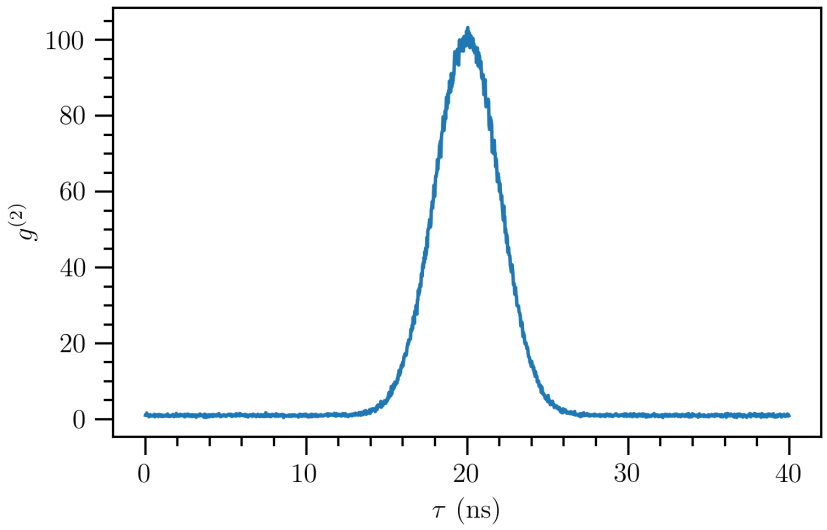

Figure 8.2: The normalised cross-correlation between *simulated* arrivals at channels 0 and 1. The timetags are generated such that 50% of the timetags in CH1 are temporally offset from the timetags in CH0 by a Gaussian distribution of delays with $\mu = 20$ ns and $\sigma = 2$ ns.

8.6 The Hanbury Brown–Twiss (HBT) experiment

To show that the signal light field is a single-photon state, we turn to the Hanbury Brown–Twiss (HBT) experiment. The HBT experiment [205] is straightforward to describe conceptually and harder to implement experimentally. Conceptually, for the case that we are interested in, it is as follows: a stream of single-photons is sent towards a 50:50 beam-splitter (BS)[13]. Detector 1 (D1) clicks when a photon is *transmitted* through the BS while Detector 2 (D2) clicks when a photon is *reflected* by the BS. If only single-photons are passing through the BS, then it is impossible for both D1 and D2 to click simultaneously because either the photon passes through the BS or it is reflected by the BS, i.e., the photon cannot be split. This is exactly the case where $g^{(2)}(0) = 0$ and, in the experiment, the closer this value is to zero[14] the closer we are to an ideal single-photon source. Classically, $g^{(2)}(0) < 1$ is forbidden: obtaining a value in that range is equivalent to a violation of the Cauchy inequality (p. 171 of [204]), which is a measure of the nonclassicality of light.

8.6.1 Obtaining a $g^{(2)}(\tau)$ measurement from experimental data

The double coincidence analysis code is adapted for analysing triple coincidence events as so:

```
for every CH0-click:
  if CH0-click is followed by CH1-click
  (within a set time window)
  and then followed by CH2-click:
    add to histogram timegap between CH1-click and CH2-click
```

In this case, the `if` statement indicates that the third click is counted *only when* the second click happens within a particular coincidence window, T_c, after the first

[13] In our case, the 50:50 beam splitter is a 50:50 optical fibre.

[14] The value of $g^{(2)}(0) < 0.5$ has been used as a cut-off for when a source is emitting single-photons [206].

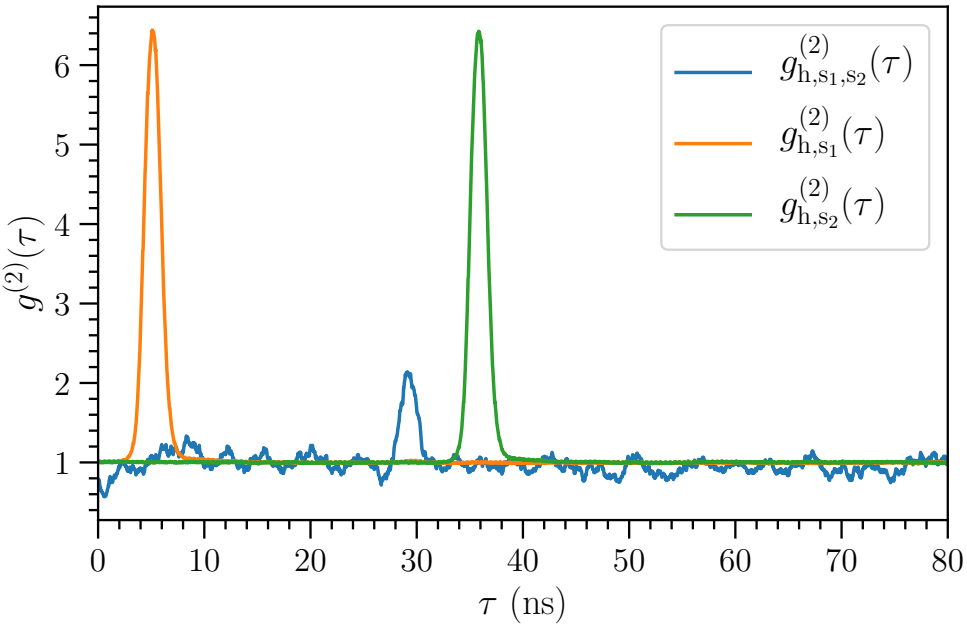

Figure 8.3: The result of a typical experiment: the normalised double coincidence events (orange & green) and the normalised triple coincidence events (blue) are shown after rectangular smoothing. The subscript label 'h' refers to the herald photon that is detected on D0, and the labels 's_1' and 's_2' refer to signal photons that are detected on D1 and D2 respectively.

click[15]. The coincidence window is centred around the peak of the first herald-signal correlation. We chose $T_c = 1.5$ ns as a compromise between a high herald-signal correlation and an adequate number of photon counts. Figure 8.3 shows the results of a typical experiment. The normalised double-coincidence and triple-coincidence histograms after the pythonic processing just described are shown, where we have used a simple rectangular smoothing method[16] to make the peaks associated with each measurement visually obvious. The triple coincidence histogram requires several hours of data-collection—eight hours, in this case, giving 1×10^5 datapoints; in the same time period, each double-coincidence histogram accumulates 1×10^7 datapoints. Thus, the double-coincidence histograms can be created much quicker

[15] We must also account for the difference in the timing between channels due to the differing BNC cable lengths.

[16] Also known as the sliding-average, this simply involves replacing each point with the average of adjacent points.

after collecting only ~20 minutes of data.[17] The figure shows that, as expected, both $g^{(2)}_{\mathrm{h,s_1}}$ and $g^{(2)}_{\mathrm{h,s_2}}$ are almost identical in shape and differ only in the time at which their peaks occurs. The similarity in the shape is expected because the only difference between the two is which arm of the 50:50 fibre the signal photons travel through. The peak height for both, $g^{(2)}_{\mathrm{h,s-MAX}}$, is 6.4; recall that Eq. 8.5 allows us to interpret this as the detection of a correlated signal photon being 6.4 times more likely than detection of an uncorrelated background photon. The full width at half maximum (FWHM) of the peak is 1.8 ± 0.1 ns. As with other Doppler-broadened systems [123], this corresponds to the inverse of the Doppler width of 560 ± 30 MHz (associated with a temperature of $97 \pm 40\,^{\circ}$C.) The time difference in the peaks is simply due to different lengths of the BNC cables connecting detectors to channels. So, in everything that follows, we translate the curves horizontally so that the peaks $g^{(2)}_{\mathrm{h,s_1}}$, $g^{(2)}_{\mathrm{h,s_2}}$, and $g^{(2)}_{\mathrm{h,s_1,s_2}}$ coincide at $\tau = 0$. Note that $g^{(2)}_{\mathrm{h,s_1,s_2}}$, the normalised triple coincidence histogram shown in blue in the figure, is not yet the $g^{(2)}(\tau)$ value for the signal light in the HBT experiment.

To obtain the $g^{(2)}(\tau)$ value for the signal light in the HBT experiment, we must make use of the fact that a single-photon from a heralded single-photon source is defined only with the detection of the herald photon [203] at Detector 0 (D0). We follow the method of Grangier [203] and of Srivathsan[18] [125, 202] to calculate $g^{(2)}(\tau)$ in the HBT experiment:

Recall that, for any two detectors, we obtain $g^{(2)}_{1,2}$ from the probability of coincidence between the clicks from both detectors ($P_{1,2}$) and the probability of independent clicks, P_1 and P_2 on each detector, as so:

$$g^{(2)}_{1,2} = \frac{P_{1,2}}{P_1 P_2}. \tag{8.8}$$

In the case of a heralded source of photons, the signal photon is only defined

[17] If a correlation between the double-coincidence histogram and the triple-coincidence histogram can be established, this obviates the need to always collect enough data for the second histogram as the correlation can be used to infer the properties of the second from the first. This will be done in the next subsection.

[18] I would like to thank Dr. Bharath Srivathsan for providing some clarifying comments on calculating $g^{(2)}(\tau)$ from experimental data.

upon detection of a herald photon [203]. Thus, the probabilities in 8.8 must be conditioned upon detecting a herald photon:

$$g^{(2)}_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}} = \frac{P_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}}}{P_{\mathrm{s}_1|\mathrm{h}} P_{\mathrm{s}_2|\mathrm{h}}}. \tag{8.9}$$

We use the multiplication rule[19] for probability calculus, $P(A \mid B) = P(A, B) \cdot \frac{1}{P(B)}$ to convert the above to:

$$\frac{P_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}}}{P_{\mathrm{s}_1|\mathrm{h}} P_{\mathrm{s}_2|\mathrm{h}}} = \frac{P_{\mathrm{h},\mathrm{s}_1,\mathrm{s}_2} \cdot \frac{1}{P_{\mathrm{h}}}}{(P_{\mathrm{h},\mathrm{s}_1} \cdot \frac{1}{P_{\mathrm{h}}})(P_{\mathrm{h},\mathrm{s}_2} \cdot \frac{1}{P_{\mathrm{h}}})} = \frac{P_{\mathrm{h},\mathrm{s}_1,\mathrm{s}_2} P_{\mathrm{h}}}{P_{\mathrm{h},\mathrm{s}_1} P_{\mathrm{h},\mathrm{s}_2}}. \tag{8.10}$$

Converting to coincidences, this is

$$g^{(2)}_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}} = \frac{G^{(2)}_{\mathrm{s}_1,\mathrm{s}_2,\mathrm{h}} N_{\mathrm{h}}}{G^{(2)}_{\mathrm{h},\mathrm{s}_1} G^{(2)}_{\mathrm{h},\mathrm{s}_2}}, \tag{8.11}$$

where $G^{(2)}_{\mathrm{s}_1,\mathrm{s}_2,\mathrm{h}}$ is the number of triple coincidence events between the herald and the two signal modes; $G^{(2)}_{\mathrm{s}_1,\mathrm{h}}$ and $G^{(2)}_{\mathrm{s}_2,\mathrm{h}}$ are the number of double coincidence (pair) events between herald and each of the signal modes; and N_{h} is the total number of heralding events[20]. For every delay between the two signal modes, $t_{\mathrm{s}_1,\mathrm{s}_2}$, we can write

$$g^{(2)}_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}}(t_{\mathrm{s}_1,\mathrm{s}_2}) = \frac{G^{(2)}_{\mathrm{s}_1,\mathrm{s}_2,\mathrm{h}}(t_{\mathrm{s}_1,\mathrm{s}_2}) N_{\mathrm{h}}}{N(t_{\mathrm{s}_1,\mathrm{s}_2})}. \tag{8.12}$$

The denominator, $N(t_{\mathrm{s}_1,\mathrm{s}_2})$, is *not* simply the product, $G^{(2)}_{\mathrm{h},\mathrm{s}_1}(t_{\mathrm{h,s}}) G^{(2)}_{\mathrm{h},\mathrm{s}_2}(t_{\mathrm{h,s}})$, of the two double coincidence histograms because these histograms are not functions of $t_{\mathrm{s}_1,\mathrm{s}_2}$. To transform $G^{(2)}_{\mathrm{h},\mathrm{s}_1}(t_{\mathrm{h,s}})$ and $G^{(2)}_{\mathrm{h},\mathrm{s}_2}(t_{\mathrm{h,s}})$ into $N(t_{\mathrm{s}_1,\mathrm{s}_2})$ we must integrate over all possible delays of $t_{\mathrm{h,s}}$ within our coincidence window, T_c, that result in a delay

[19] The multiplication rule is often written as $P(A \mid B) = P(A \cap B) \cdot \frac{1}{P(B)}$.

[20] Instead of using N_{h}, it is possible to normalise the function by using the knowledge that $g^{(2)}_{\mathrm{s}_1,\mathrm{s}_2|\mathrm{h}}(t_{\mathrm{s}_1,\mathrm{s}_2})$ must go to 1 for large values of $t_{\mathrm{s}_1,\mathrm{s}_2}$. As we indicated before, this is a helpful check that the analysis is correct.

t_{s_1,s_2}:

$$N(t_{s_1,s_2}) = \int_0^{T_c} G^{(2)}_{h,s_1}(t_{h,s})\, G^{(2)}_{h,s_2}(t_{h,s} + t_{s_1,s_2})\, dt_{h,s}. \tag{8.13}$$

Please see Appendix B for a pythonic implementation of this function.

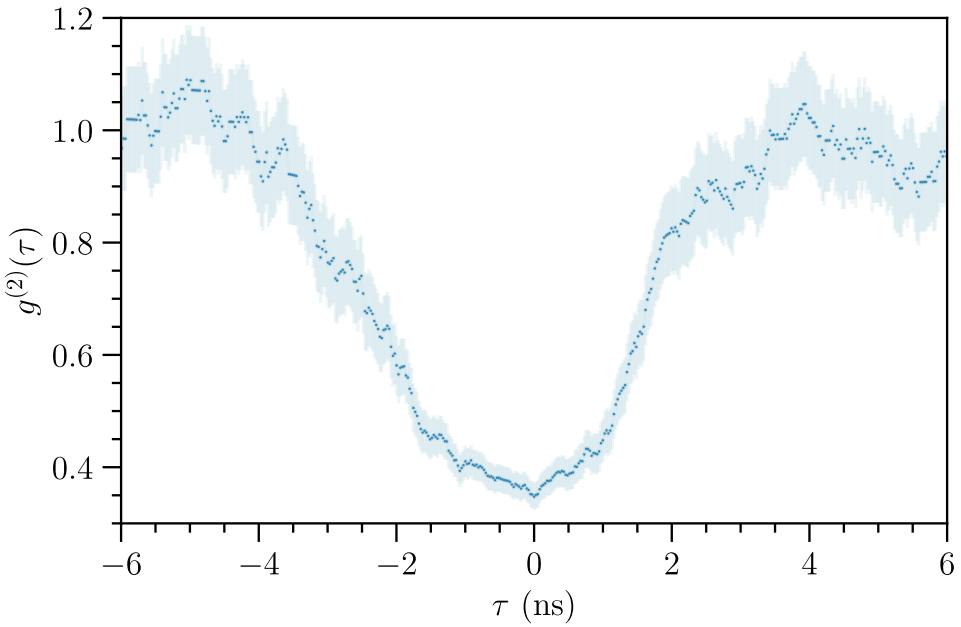

Figure 8.4: Hanbury Brown–Twiss experiment shows a $g^{(2)}(\tau = 0)$ dip below 1.0, demonstrating the presence of signal single-photons. $g^{(2)}(0) = 0.35 \pm 0.02$. Experimental conditions: cell temperature, 114 °C; pump laser power, 55 μW; coupling laser power, 100 mW. The pump laser was detuned from resonance by -2063 MHz, and the coupling laser was detuned from resonance by 1925 MHz. Data have been smoothed using a rectangular smoothing method. Errors (light blue) are calculated by propagating the Poissonian error on the photon counts.

Once this is done, $g^{(2)}_{s_1,s_2|h}(t_{s_1,s_2})$ can be written more simply as $g^{(2)}(\tau)$ and visualised. Figure 8.4 shows the result of a single HBT experiment. The $g^{(2)}$ dip is clearly evident on this curve: $g^{(2)}(0) = 0.35 \pm 0.02$. The error is calculated by propagating the Poissonian error, $\sqrt{N}$, on the photon counts, N. The width of the dip is 3.8 ± 0.1 ns. This is approximately double that of the width of the $g^{(2)}_{h,s}$ peak because of the extra broadening introduced by the $N(t_{s_1,s_2})$ function: Because $g^{(2)}_{h,s_1}$ has a Gaussian lineshape and $g^{(2)}_{h,s_1} \approx g^{(2)}_{h,s_2}$, $N(t_{s_1,s_2})$ is effectively the convolution of $g^{(2)}_{h,s_1}$ with part of itself, resulting in a broadened Gaussian.

8.6.2 Relationship between $g^{(2)}_{h,s}$ and $g^{(2)}$

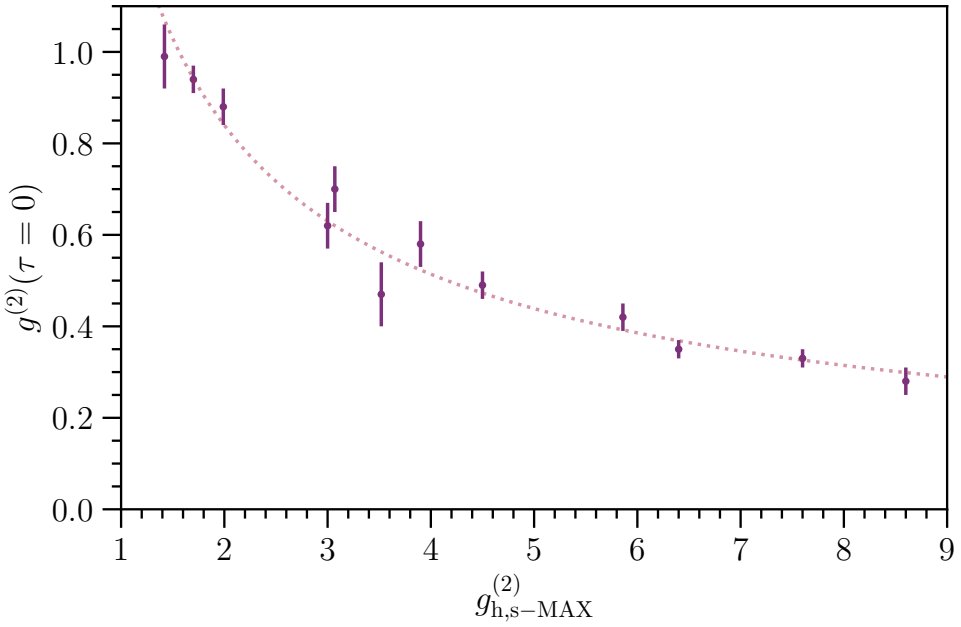

Figure 8.5: The single-photon purity is higher as $g^{(2)}_{h,s-MAX}$ increases. As the correlation between herald and signal increases (at greater detunings of the pump laser), the HBT experiment gives a deeper $g^{(2)}$ dip, i.e., closer to 0. The dotted line is a guide to the eye.

Figure 8.5 shows that the HBT experiment gives a lower $g^{(2)}$ dip as the correlation between the herald and signal pair increases. In other words, we get "purer" single-photons as $g^{(2)}_{h,s-MAX}$ increases. Much of the rest of the chapter focuses on how to improve $g^{(2)}_{h,s-MAX}$. We shall find that there are many ways to do this, albeit at the

expense of the pair rate. The pair rate, or corrected coincidence rate, is calculated by subtracting uncorrelated coincidences from all coincidences: For an arbitrary data-taking time, T, if we only look in a window defined by t_0 & t_1, this is

$$\text{Pair rate } = \frac{1}{T}\left(\sum_{\tau=t_0}^{t_1} G^{(2)}_{\mathrm{h,s}}(\tau)\right) - r_{\mathrm{h}} r_{\mathrm{s}} \left(t_1 - t_0\right). \tag{8.14}$$

The total number of coincidences is the number of data points in the unnormalized double-coincidence histogram, $(\sum_{\tau=t_0}^{t_1} G^{(2)}_{\mathrm{h,s}}(\tau))$. The number of uncorrelated coincidences, $(r_{\mathrm{h}} r_{\mathrm{s}} \left(t_1 - t_0\right))$, is the number of data points underneath the background level. Thus, the ratio between correlated and uncorrelated coincidences varies according to the shape of the $G^{(2)}_{\mathrm{h,s}}(\tau)$ curve.

We generally obtained data with a coupling laser that is frequency stabilised ("locked") to the $5P_{3/2}$–$5D_{3/2}$ transition using STROLL (Chapter 5). This allows us to collect data over a time frame of over 24 hours, which is essential for those regions of the parameter space where the $g^{(2)}_{\mathrm{h,s-MAX}}$ is high but the pair rate is low. When we are not locking the coupling laser, we can choose arbitrary detunings on the transition that this laser couples. Thence it was possible to find a detuning where there is both a higher correlation ($g^{(2)}_{\mathrm{h,s-MAX}} > 7$) and a higher pair rate (>15 Hz) that allows for an HBT experiment to be done within 6-8 hours. If we attempt to further detune, the pair rates are too low for an HBT experiment to be done in a time frame before the laser drifts completely away from the optimal region.

8.7 Towards a more ideal heralded single-photon source

Having now established that higher cross-correlations between herald and signal give deeper $g^{(2)}$ dips, in this section we examine all the experimental parameters giving higher cross-correlations in order to create a purer single-photon source.

The errors on the following data are calculated by obtaining the standard error from repeated measurements [133]. This error dominates over the Poissonian statistics

from counting the photons.[21]

8.7.1 Dependence on temperature

Figure 8.6 shows the variation with temperature of both the maximum herald-signal cross-correlation, $g^{(2)}_{\mathrm{h,s-MAX}}$, and the pair rate. We find that the maximum cross-correlation follows the curve expected from Chapter 6: The value increases with temperature until the value of 85 °C, after which it decreases again. We saw there that the background counts increase linearly with laser input until 85 °C, after which it increases nonlinearly. We also saw, in Chapter 7, that light generated by the seeded 4WM signal increases quadratically with the atomic number density. Thus, just at or below 85 °C would be the ideal temperature to work at if one were interested in maximising $g^{(2)}_{\mathrm{h,s}}$.

However, we also see that the pair rate continues to increases beyond a temperature of 85 °C. It is for this reason that we chose to perform our HBT experiments at higher temperatures. However, having used the higher pair rates to show that a large cross-correlation gives deeper $g^{(2)}$ dips, we would recommend working below 85 °C if possible.[22]

[21] This is partly due to the difficulty in returning to precisely the same laser alignments. We expect the finalised Build 2.0 of the Pi Auto-aligner (Chapter 4) to reduce this difficulty.

[22] To be perfectly clear on this point, we note that we only performed the HBT experiments at higher temperatures because we could not collect enough data at lower temperatures. However, the combination of Fig. 8.6 and Fig. 8.5 tell us that if we could collect enough HBT data at lower temperatures, they would give deeper $g^{(2)}$ dips.

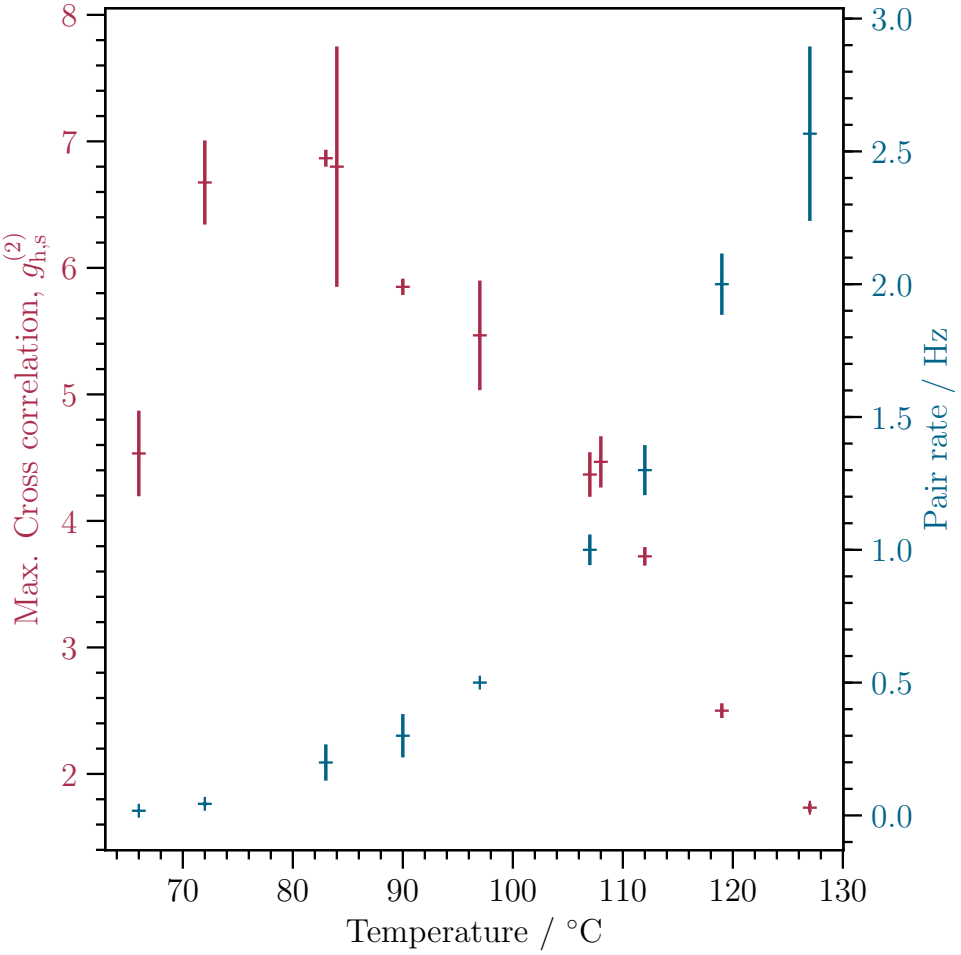

Figure 8.6: The herald-signal pair rate increases as the temperature increases but the maximum cross-correlation decreases. More precisely, with all other parameters fixed, the maximum correlation increases with temperature up to 85 °C but then decreases as the temperature goes above 85 °C.

8.7.2 Dependence on pump detuning

Figure 8.7 shows the variation with pump laser detuning of both the maximum herald-signal cross-correlation and the pair rate. We find that the maximum cross-correlation increases with pump detuning, increasing rapidly at large detunings. This is partly because at larger detuning, there is a decreased population of the intermediate P state from which collisional transfer can create uncorrelated 795 photons, as we saw in Chapter 6. It is also true that when the pump detuning exceeds the Doppler width, there are only a relatively small number of atoms that are two-photon resonant with both lasers. When the detuning is sufficiently large [207], the total coherence due to the off-resonant interaction is stronger than that due to resonant interactions. This means that the relevant coherence decreases more slowly with detuning than the relevant population; again, this increases the correlation. In a now familiar theme, there is a trade-off with pair rate: as the detuning is increased, the rate of pair production decreases while the maximum cross-correlation increases.

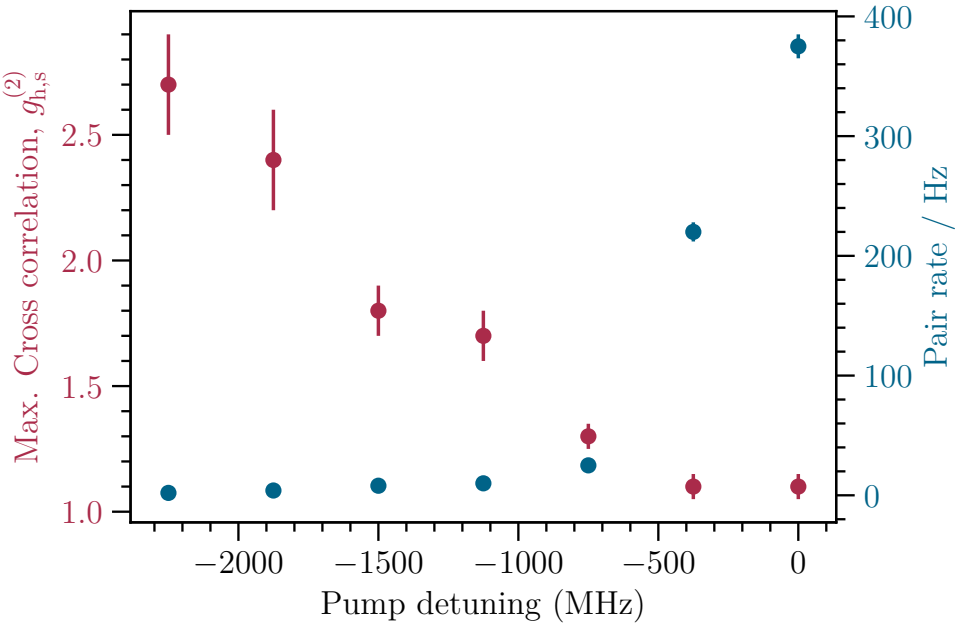

Figure 8.7: The maximum cross-correlation between herald and signal increases with pump detuning while the herald-signal pair rate decreases with pump detuning. The coupling detuning is adjusted so that the two-photon detuning is always equal to zero. The data were taken at a pump laser power of 55 μW, a coupling laser power of 180 mW and a temperature of 114 °C. Error bars are the standard error from repeated measurements.

8.7.3 Dependence on pump power

Figure 8.8 shows that as the pump power increases, the maximum cross-correlation between the herald-signal pairs decreases. This is partly because there is an increased population of the intermediate P state from which collisional transfer can create uncorrelated photons of the same colour as the signal photons, as we saw in Chapter 6. The largest $g^{(2)}_{\mathrm{h,s}}$ that we recorded for our system was $g^{(2)}_{\mathrm{h,s}} = 105$ at the weak pump power of 0.5 μW but, at this power, there is a pair rate that is less than 0.1 Hz.

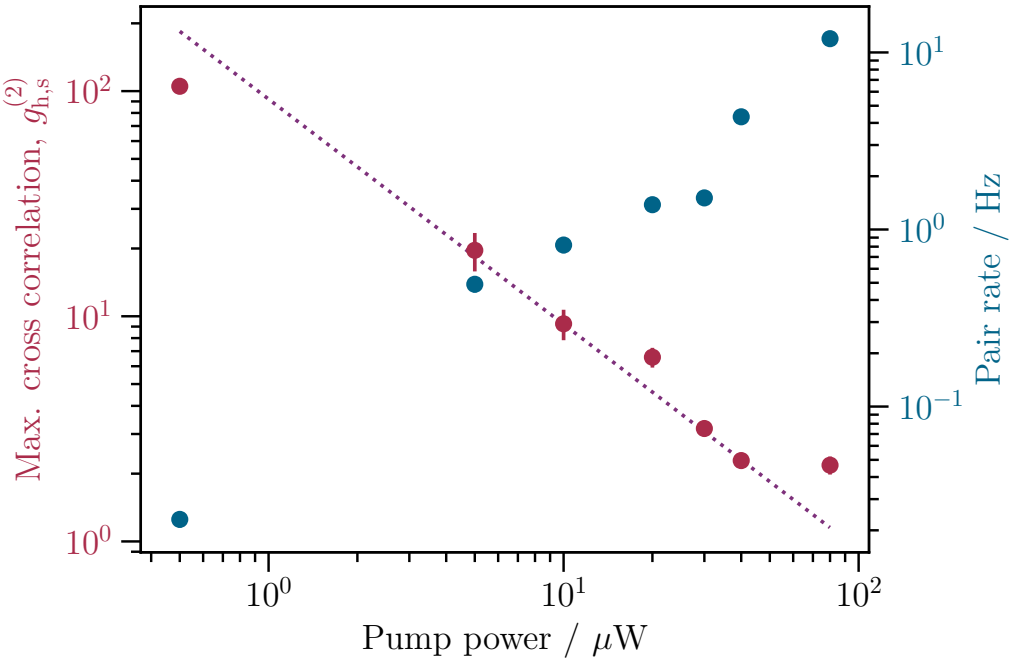

Figure 8.8: The maximum cross-correlation between herald and signal decreases with pump power. The data were taken with pump laser detuned by -1110 MHz from resonance, with a coupling laser power of 30 mW, and the cell at a temperature of 114 °C. Some error bars are too small to be seen. The dotted line is a guide to the eye.

8.7.4 Dependence on coupling power

Figure 8.9 shows the variation with the coupling power of both the maximum herald-signal cross-correlation and the pair rate. Our main motivation for increasing coupling power was to increase the pair rate in order to do an HBT experiment. With increased coupling power there is more population in the doubly-excited state and consequently greater herald-signal pair production. However, as can be seen from the figure, the maximum cross-correlation decreases with increased power. This is partly because atoms may be being de-excited from the D state to the $P_{3/2}$ state, leading to the collisional transfer process releasing uncorrelated background photons. We found that the 30 mW from the coupling laser alone gave only ~1 Hz, which was too low, while the full 330 mW with the TA reduced the correlations too much. We settled on 100 mW as a compromise that gave a ~5 Hz pair rate.

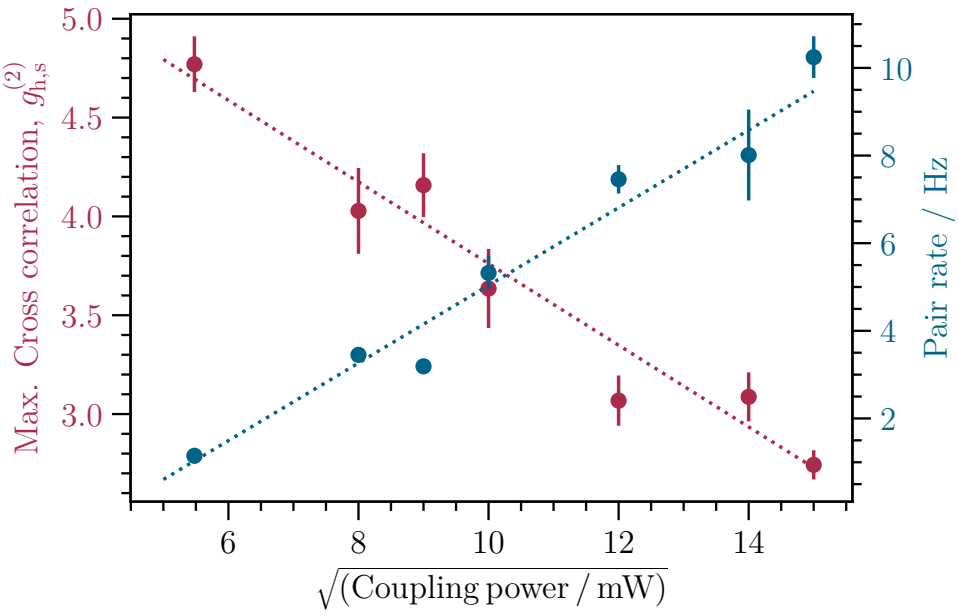

Figure 8.9: The maximum cross-correlation between herald and signal decreases with coupling power. In other words, as the Rabi frequency of the coupling laser increases, the $g^{(2)}_{h,s-MAX}$ decreases. The data were taken at a pump laser power of 55 μW detuned by -1850 MHz, the coupling laser at a frequency that keeps the two lasers two-photon resonant, and a cell temperature of 114 °C. The dotted lines are guides to the eye.

8.7.5 Deepening the $g^{(2)}$ dip using an etalon lens filter

In Chapter 3 we discussed the details of the etalon lens filter that we built for spectral filtration. Here we address two related questions: Is it possible to use the etalon filter to obtain a higher $g^{(2)}_{\mathrm{h,s}}$ peak, and to obtain a deeper $g^{(2)}$ dip?

We found that with enough careful tuning of the etalon central frequency, it is always possible to make the $g^{(2)}_{\mathrm{h,s}}$ peak higher when the 795 nm photons are sent through the etalon. Figure 8.11 shows that, for several detunings, the correlation associated with the light that passes through the etalon is higher than for the light that does not pass through the etalon. Figure 8.10 shows how the ratio between the two peaks in Fig. 8.11 (a) varies with the etalon central frequency. Figure 8.10 also illustrates how the etalon frequency is tuned to find the etalon's transmission peak so that we maximise the transmission of the herald photons: We start by placing the etalon peak at the same detuning as the seeded 4WM signal. Then we plot how the ratio of the $g^{(2)}_{\mathrm{h,s}}$ with and without the etalon varies with etalon frequency (i.e., Fig. 8.10). The peak of the resulting curve gives the transmission peak frequency of the etalon.

Although Fig. 8.5 strongly suggests that whenever we have a large $g^{(2)}_{\mathrm{h,s}}$ peak, we will have a deep $g^{(2)}(\tau = 0)$ dip, the only way to prove this is by doing an HBT experiment. However, when the signal light is sent through the etalon, the pair rate drops by an order of magnitude which meant that enough data could not be collected for the HBT measurement during the time that the lasers stay locked.[23] In the case of the HBT experiment, the slight increase in signal-to-noise ratio as demonstrated by the higher $g^{(2)}_{\mathrm{h,s}}$ is not enough to compensate for the overall decrease in signal. So although the etalon is discarding background photons, it is also discarding signal photons. If the background photons had been sufficiently different in frequency from the signal photons (i.e., greater than the FWHM linewidth of the etalon of 128 MHz), this would not occur. Therefore, we learn that a large number of the background photons must be at a very similar frequency (i.e., different by less than

[23] While, in principle, it would be possible to integrate the results of many experiments, practically this is very difficult for two related reasons: (1) We would still need to integrate over several weeks to get a single result, and (2) it is not feasible to keep experimental conditions the same over that time period.

128 MHz) to the signal photons.

Ideally, we would want to spectrally filter the herald (762 nm) photon as well, but because this photon comes from the decay of an excited state, it is difficult to position the etalon central frequency. Fortunately, it is less important to filter this light as there is no spectral redistribution process generating unwanted light at the herald frequency.

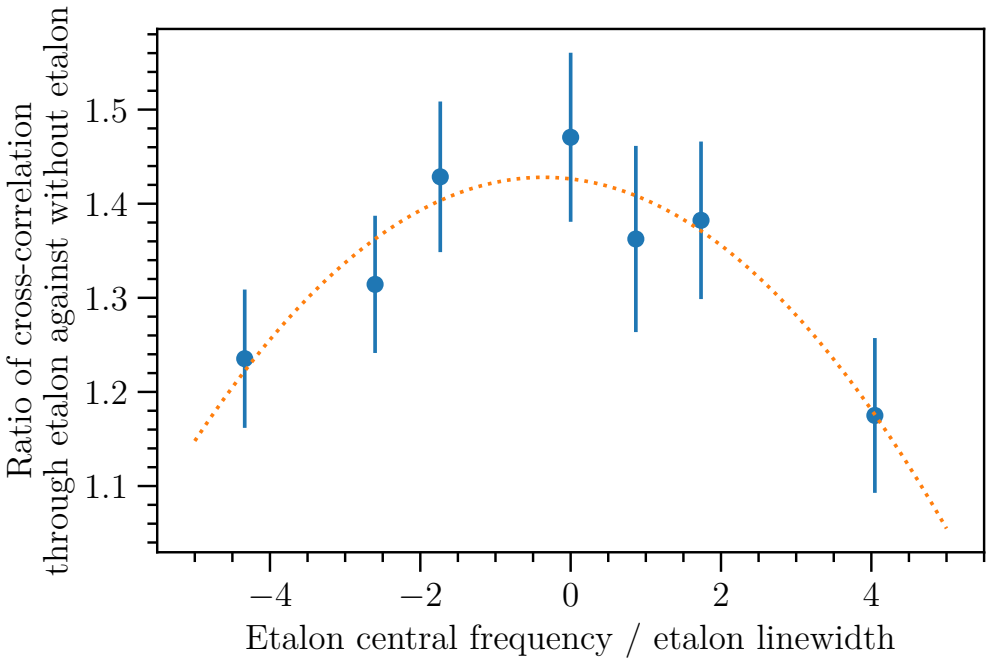

Figure 8.10: Ratio of correlation peaks (between with and without etalon). The x-axis is the etalon central frequency in units of etalon FWHM linewidth; the etalon has a FWHM linewidth of 128 MHz. The etalon central frequency depends on the temperature of the etalon, which is controlled by the voltage on the temperature controller. 1 mV gives a change of 37 MHz.

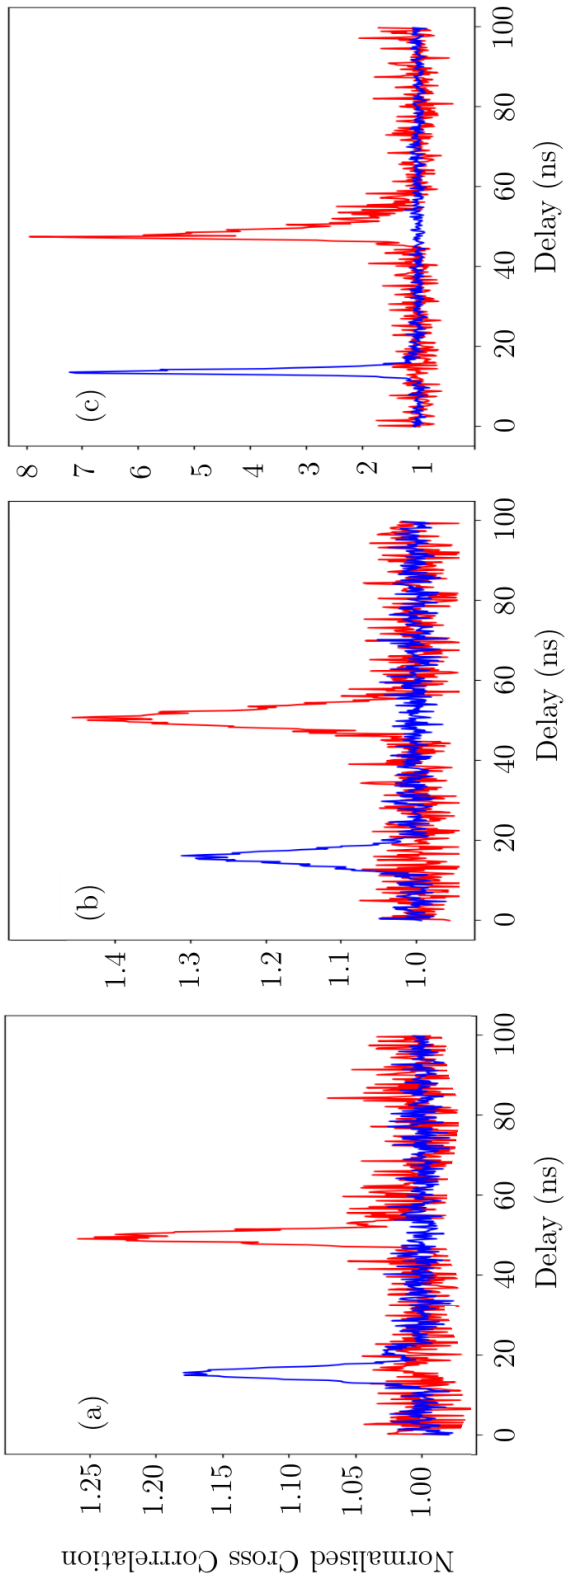

Figure 8.11: Cross-correlation of herald and signal, without (blue) and with (red) the lens etalon filter, at different pump detunings, Δ_p. (a) $\Delta_p = 0$ MHz (b) $\Delta_p = -375$ MHz, and (c) $\Delta_p = -1050$ MHz. The data were taken at a pump laser power of 55 μW, a coupling laser power of 100 mW, and a cell temperature of 114 °C. In all cases, the maximum correlation, $g^{(2)}_{\mathrm{h,s-MAX}}$, is higher with the etalon than without. In each case, we have used a version of Fig. 8.10 to find the correct frequency to set the etalon lens filter at.

8.8 Discussion

For maximum $g^{(2)}_{\mathrm{h,s}}$ values that are ~10, the pair rate is always on the order of ~Hz. At the same time, the raw count rate on the photon counter for the herald channel (which will include background or 'noise' photons) is on the order of ~kHz. This leads to the probability of detecting a signal photon conditional on detecting a herald photon, i.e., the heralding efficiency, of less than 0.1%. This is very low compared to other sources; a comparable, cold-atom, system which uses the same diamond configuration in ^{87}Rb (but not in the HPB regime) has pair rates on the order of kHz and a heralding efficiency of 23% [125].

There are many possible reasons for our low numbers. One is that, as stated in the theory section, there is no phase-matched emission for the herald photons. Only those photons that happen to be emitted in the forward direction will have a phase-matched signal photon emitted. This means that the optical fibre can accidentally collect photons at the colour of the herald photon but those photons do not have a paired signal photon associated with them. Secondly, inefficiencies in the detection of the photons accumulate: the quantum efficiency of the detectors is specified to be 60% at 780 nm; the optical fibre efficiency is estimated to also be 60% and the 50:50 split in the fibre introduces another 50% loss; the losses due to the narrowband interference filters was measured to be 3% for each of the 3 filters; and the loss due to reflection at the vapour cell surfaces was measured to be 8%. In combination, the total efficiency is no more than 10%. Last but not least, there is a problem specific to thermal vapour: for light emitted by atoms from different velocity groups, destructive interference occurs, reducing correlations between photon pairs. This may explain the most significant difference between a cold-atom source and ours.

A comparable thermal vapour system is the double-ladder configuration in thermal ^{87}Rb where the $5S_{1/2}$, $5P_{3/2}$, and $5D_{3/2}$ terms [123] are used, resulting in 780 nm herald and 776 nm signal photons. In the double-ladder configuration, detuned 780 nm and 776 nm laser beams are used to generate, via 4WM, resonant 780 nm and 776 nm pairs. Compared to our system, they report a lower $g^{(2)}(0) = 0.037(3)$ with a higher pair rate of ~1 kHz. However, we found a higher peak $g^{(2)}_{\mathrm{h,s}}$ value of

105 compared to their 85. We note, however, that this was done at a very low pump power, giving a small pair rate and they may not have investigated this region of the parameter space. There are two main differences between our system and that system. First, ours is in the HPB regime; that system is not. We would not expect this to degrade the performance of our source because, as we have seen, both in this thesis and the wider literature [6, 18–25], the HPB regime is used precisely because it leads to a simplified, more controllable experiment. Indeed, use of the HPB regime would remove the extra beam that system requires for performing optical pumping to prepare the initial state of the atoms. Secondly, that system uses a double-ladder configuration whilst ours is a diamond configuration. We expect the root of the difference to lie there. Thus, a logical step for the future direction of our project is to repeat the experiments in this chapter in a double-ladder configuration in the HPB regime.[24] Before ending, we should note that the double-ladder must be done in the counter-propagating geometry. It is possible that this alone contributes to some improvements in the important metrics because a broad range of velocity groups is involved in generating the photon pairs. Therefore, before changing to the double-ladder configuration, it would be worthwhile trying to change the direction of the beams in our system from a co-propagating to a counter-propagating geometry. For this we require greater angular access to our cell; a new cell holder and heater are being built for this purpose.

We note that there is no single overarching figure-of-merit used in the literature on single-photons [179], and the metric(s) of importance depends on the application for which the single-photons will be used. Perhaps, if our source were combined with a quantum memory, as suggested in the introduction [197], the low pair rate may not be an issue because the single-photons generated slowly by our source could be stored by the quantum memory until they are needed. We will note here two advantages of our system. Compared to other configurations, the diamond configuration in rubidium vapour gives rise to pairs of photons that differ in wavelength by nanometres and so are easily separated from each other.

[24] Of course, if then all our metrics are still worse, then we must seriously consider that the physics of the HPB regime is responsible though at present we have no reason to suspect this.

Compared to cold-atoms, use of a thermal vapour means that none of the elaborate equipment and mechanisms of cooling are required in order to build our single-photon source.

8.9 Conclusion

In this chapter, we showed that we have built a heralded single-photon source in thermal rubidium vapour, as demonstrated by the Hanbury Brown–Twiss experiment giving a $g^{(2)}$ dip of $g^{(2)}(0) = 0.35 \pm 0.02$. We have shown that maximising the cross-correlation, $g^{(2)}_{\mathrm{h,s}}$, between herald and signal photons leads to a deeper $g^{(2)}$ dip, and have explored all the experimental parameters that increase this cross-correlation. Furthermore, we have shown that our bespoke etalon lens filter also increases $g^{(2)}_{\mathrm{h,s}}$ by spectrally filtering out uncorrelated photons but at the cost of a lower pair rate. In the next chapter, we will note possible avenues of future research and conclude the thesis.

Chapter 9

Summary and Outlook

We conclude the thesis by summarising each chapter, noting important links amongst chapters, and describing some possible further avenues of research.

Chapter 2: Theory: Atom-light interactions. We introduced the basic atomic physics underlying the work in this thesis. Starting from a two-level system coupled by a seed (probe) beam, we wrote the associated density matrix and Hamiltonian. Computationally solving the Lindblad equation gave the steady-state solution, which allowed us to determine the electric susceptibility, χ, of the medium. From this, we showed how the absorptive and dispersive properties of the system vary with the seed frequency. We then showed how these properties vary as extra energy levels and couplings between these levels were introduced. Specifically, we noted how the addition of an extra energy level coupled by a strong pump laser leads to an electromagnetically induced transparency (EIT) window in the seed transmission. Addition of a further level coupled by a coupling laser led to the splitting of the EIT feature. These EIT features appear experimentally in Chapter 7. We ended by explaining the hyperfine Paschen–Back (HPB) regime and how the rubidium energy level structure changes there. We noted that the simplification of energy levels in the HPB regime allows us to create the effective four-level system required for Chapters 7 & 8.

Chapter 3: Experimental techniques and apparatus. [50] We detailed the method by which the output power of a tapered amplifier (TA) can be maximised, finding that careful mode-matching, at both the input and output ends of the TA, was required. We described the design of a bespoke etalon lens filter that we commissioned to aid in filtering out unwanted light in our spontaneous four-wave mixing (4WM) experiments of Chapter 8. We also described, in brief, the method of using a separate optical cavity to calibrate the relative frequency spectrum of our lasers.

Chapter 4: Machine learning in quantum optics experiments. [49] We noted that a perennial problem in laser physics experiments is the need to continually align and realign lasers. We demonstrated that this process can be automated with the use of inexpensive components, controlled by a machine learning algorithm. With Build 1.0 of the "Raspberry Pi Auto-aligner", we found that, after hysteresis correction, the device could align weak beams detectable by a photodiode. However, the device was not as reliable in aligning photons detectable only on a photon-counter. Thus, a more fine-tuned device (Build 2.0) was designed and built.

When Build 2.0 of this device has been fully tested, we expect to be able to set as its cost function the negative of the $g^{(2)}_{\mathrm{h,s}}$ peak value: This should allow for a more reliable way of finding the best alignment to obtain the purest photons. Once this first alignment is performed, it is expected that the machine will be used for automatic realignment—it can simply be set to run every morning before experiments are undertaken. At the time of writing, one of our colleagues outside the Quantum Light & Matter group[1] has indicated his intention to implement a copy of our device in his laboratory for a use which we had not originally envisioned, indicating to us that there a large number of uses for our device that are yet to be discovered.

Chapter 5: Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime. [48] We demonstrated a novel technique to simultaneously lock two laser frequencies to a two-photon transition ($5S_{1/2} \rightarrow 5P_{3/2} \rightarrow 5D_{5/2}$) in ^{87}Rb. We showed data demonstrating

[1] Dr A Hindmarch, Durham University, personal communication.

a frequency instability of less than 6 MHz for a time scale of 10 hours. In practice, the lasers can stay locked for much longer: We regularly locked for over 24 hours to perform the Hanbury Brown–Twiss experiments of Chapter 8 despite, in that chapter, locking the second laser to the weaker $5P_{3/2} \rightarrow 5D_{3/2}$ transition.

The present STROLL scheme works well for our purposes. We are able to detune on the first step, Δ_{pump}, of the two-photon transition by varying the magnetic field strength. The lock automatically adjusts the detuning of the second laser, Δ_{coupling}, so that the $\Delta_{\text{pump}} + \Delta_{\text{coupling}} = 0$. This is an advantage when one always wants to lock to the two-photon frequency, but it does not allow the freedom to arbitrarily detune the coupling laser. This limitation could be circumvented by incorporating an acousto-optic modulator into the setup.

Chapter 6: Spectral redistribution. We described the phenomenon of spectral redistribution where 780 nm light is generated by the input of 795 nm light into the rubidium vapour cell and vice versa. This phenomenon is due to the collisional transfer of states between the $5P_{1/2}$ and $5P_{3/2}$ levels. We found that this transfer increases, as expected, with both the power of the input beam and the cell temperature (and hence the atomic number density). There were two distinct temperature regimes, one where the generated photon count rates increase linearly as a function of atomic number density and a second regime where it increases nonlinearly.

Using an etalon lens filter will allow for the determination of which hyperfine transitions give rise to which photons. Capturing photons emerging at 90° to the laser beam path will allow for a comparison of the relative rates of 780 nm and 795 nm photon production from atoms interacting with a single laser beam.

Chapter 7: Seeded four-wave mixing. [25] We experimentally studied the nonlinear optical phenomenon of seeded four-wave mixing in a diamond configuration in rubidium and showed that a simple model allows us to understand the variation of the 4WM signal light with experimental parameters. It was possible to work with this simple model because—as explained in Chapter 2—the application of a high magnetic field allowed us to isolate a four-level system. Thus, when we characterised the system, there was good agreement between theory and data. We

also devised a novel method of extracting the excited state dipole matrix element $|\langle 5P_{3/2},\ m_J = 3/2|er|5D_{3/2},\ m_{J'} = 1/2\rangle|$ by measuring how the splitting (S) of the peaks of the 4WM signal increases as the coupling beam power (P) increases. From the gradient, M, of the graph of $S = M\sqrt{P}$, the dipole matrix element can be determined. Although a measurement was made that was close to the theoretical figure, the precision was limited by experimental constraints.

Chapter 8: Spontaneous four-wave mixing leading to a heralded single photon source. The thesis culminated in the demonstration of how the nonlinear optical phenomenon of spontaneous 4WM in rubidium vapour can be used as a source of single-photons. We explained the computational techniques used to investigate photon statistics. We demonstrated the single-photon nature of the rubidium vapour source by performing a Hanbury Brown–Twiss (HBT) experiment that gave a $g^{(2)}(\tau = 0)$ of 0.35 ± 0.02. We found the greater the cross-correlation between herald and signal photon, the closer $g^{(2)}(\tau = 0)$ is to zero (i.e., "purer" photons). We explored all the experimental parameters that increase this correlation, including the use of the etalon lens filter to remove background photons. We found that, in all cases, there is a compromise involved: the purer the single-photons produced, the fewer of them there are.

We investigated our particular excitation configuration (diamond) and beam geometry (co-propagating) exhaustively. We would expect that changing the *geometry* to counter-propagating would give more seeded 4WM signal light as more velocity classes of atoms would be excited. It is not clear that this necessarily means a higher-quality single-photon source, however. The counter-propagating geometry requires a greater angular difference between the beams and, in preparation for this change in geometry, a new cell holder is under commission to enable greater optical access. In the wider community, a different 4WM *configuration*—double-ladder [123] instead of diamond—has been used to create a single-photon source. The future of this project will involve attempting this configuration in the HPB regime, as it will be simpler to understand and model within the effective four-level atomic system. The opportunities for continuing research into nonlinear optical phenomena in the hyperfine Paschen–Back regime have by no means been exhausted.

Appendix A

Full Optical Table Layout

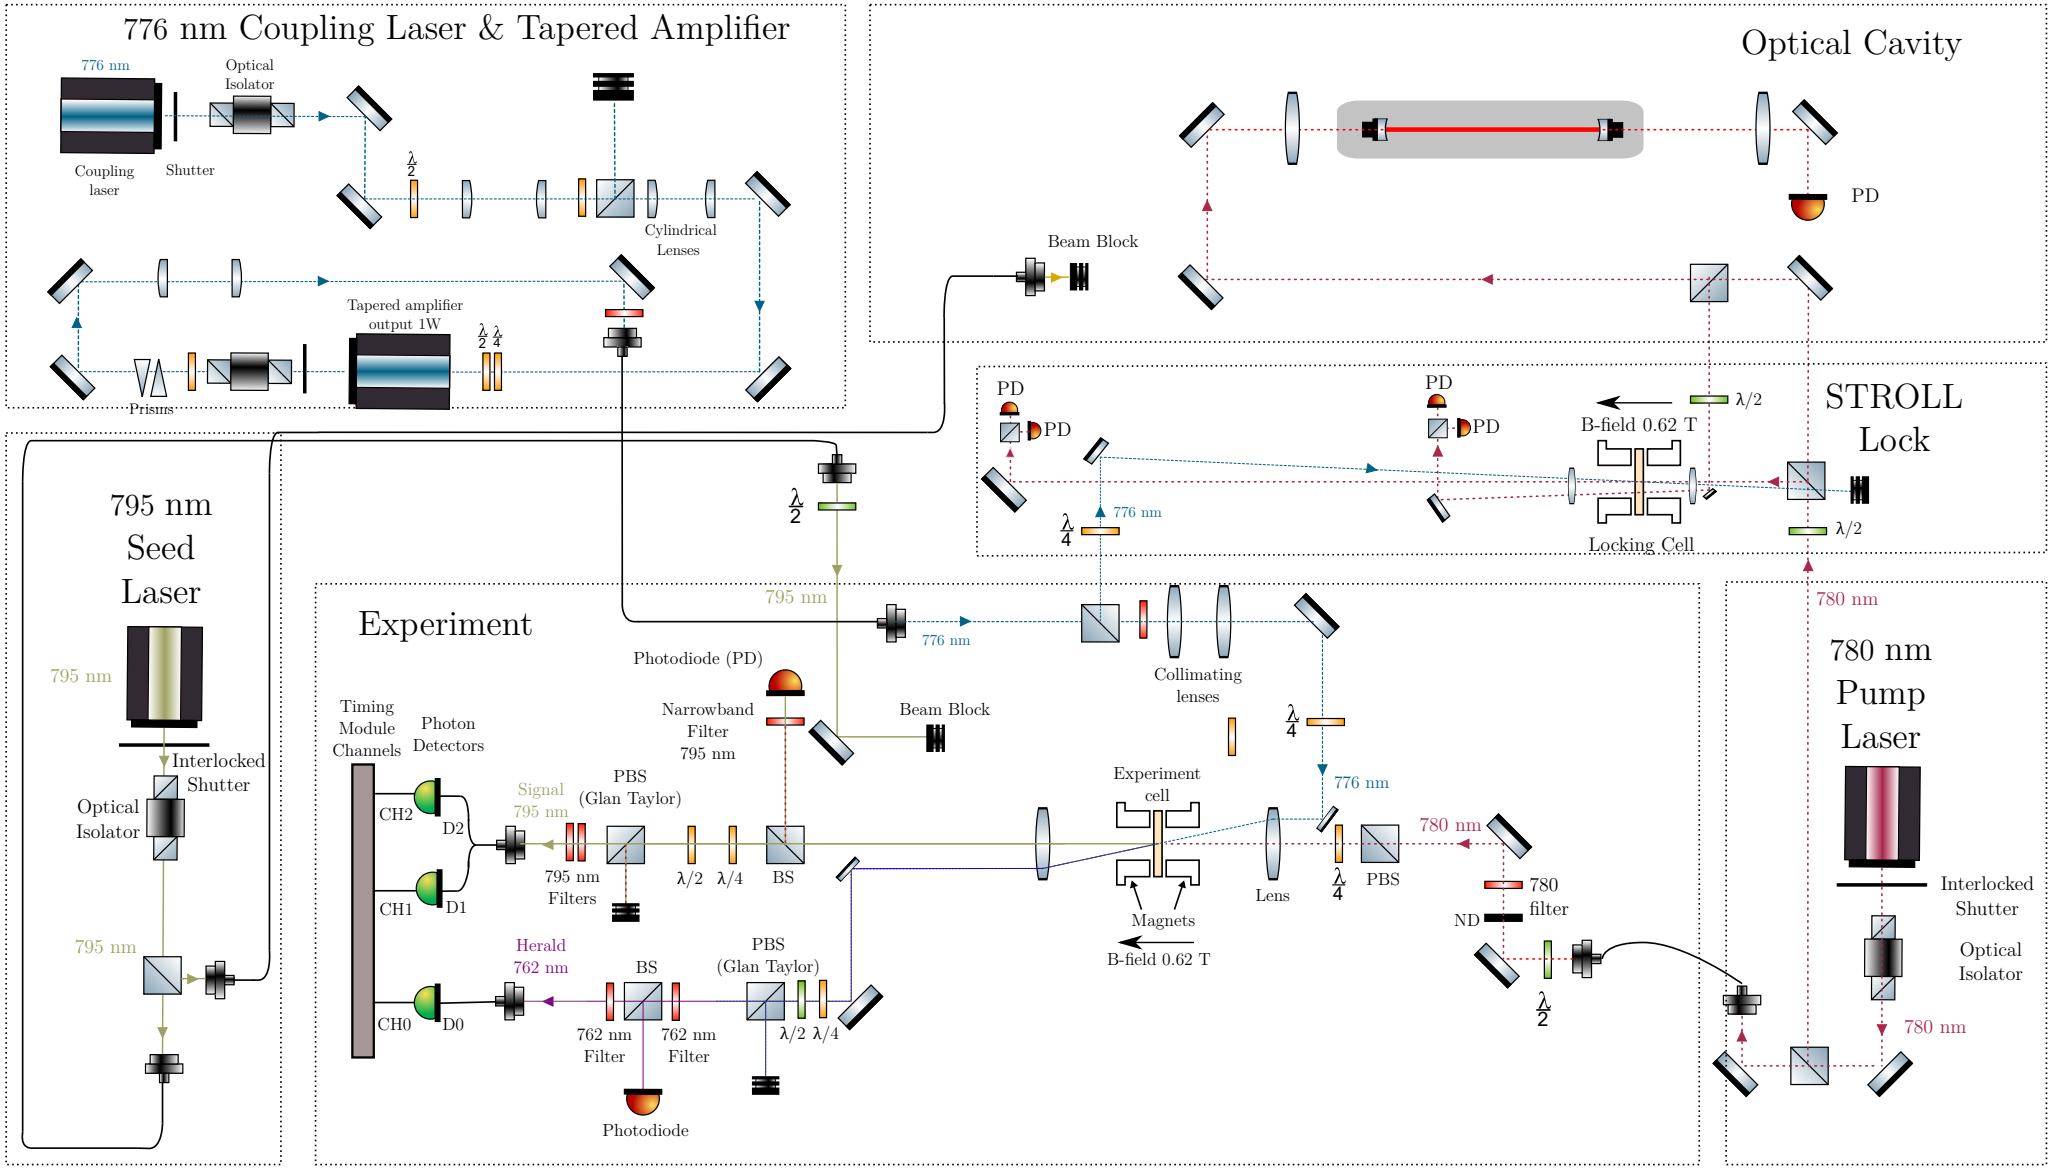

Figure A.1: Full optical table layout for spontaneous four-wave mixing.

Appendix B

Python Code

We show code for calculating Eq. 8.13 from Chapter 8, which is displayed again below for the reader's convenience.

$$N(t_{s_1,s_2}) = \int_0^{T_c} G^{(2)}_{h,s_1}(t_{h,s})\, G^{(2)}_{h,s_2}(t_{h,s} + t_{s_1,s_2})\, dt_{h,s}. \tag{8.13}$$

The code is shown written with loops as this is the clearer implementation. Using the cross-correlation function in the numpy module (numpy.correlate) would be faster but less instructive.

```
def Nplus(x,y):
  Y = np.pad(y, (0,len(y)), 'constant')
  Z = []
  for j in range(len(x)):
    z = 0
    for i in np.arange(len(x)):
      z = z + (x[i]*Y[i+j])
    Z.append(z)
  return np.array(Z)
```

Bibliography

[1] E. Curie. *Madame Curie: A Biography*. Hachette Books (1937). [Cited on page xiv.]

[2] R. T. Willis, F. E. Becerra, L. A. Orozco, and S. L. Rolston. *Photon statistics and polarization correlations at telecommunications wavelengths from a warm atomic ensemble*. Opt. Express **19**, 14632 (2011). [Cited on pages 1, 67, 99, and 136.]

[3] G. Walker, A. S. Arnold, and S. Franke-Arnold. *Trans-Spectral Orbital Angular Momentum Transfer via Four-Wave Mixing in Rb Vapor*. Phys. Rev. Lett. **108**, 243601 (2012). [Cited on page 1.]

[4] P. H. S. Ribeiro, C. Schwob, A. Maître, and C. Fabre. *Sub-shot-noise high-sensitivity spectroscopy with optical parametric oscillator twin beams*. Opt. Lett. **22**, 1893 (1997). [Cited on pages 1 and 99.]

[5] V. Boyer, A. M. Marino, R. C. Pooser, and P. D. Lett. *Entangled Images from Four-Wave Mixing*. Science (80-.). **321**, 544 (2008). [Cited on pages 1 and 99.]

[6] D. J. Whiting, N. Šibalić, J. Keaveney, C. S. Adams, and I. G. Hughes. *Single-Photon Interference due to Motion in an Atomic Collective Excitation*. Phys. Rev. Lett. **118**, 253601 (2017). [Cited on pages 1, 2, 66, 76, 99, 100, 137, and 163.]

[7] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt. *Optical atomic clocks*. Rev. Mod. Phys. **87**, 637 (2015). [Cited on page 1.]

[8] N. Hinkley, J. A. Sherman, N. B. Phillips, M. Schioppo, N. D. Lemke, K. Beloy, M. Pizzocaro, C. W. Oates, and A. D. Ludlow. *An atomic clock with* 10^{-18} *instability*. Science (80-.). **341**, 1215 (2013). [Cited on page 1.]

[9] L. A. Downes, A. R. MacKellar, D. J. Whiting, C. Bourgenot, C. S. Adams, and K. J. Weatherill. *Full-Field Terahertz Imaging at Kilohertz Frame Rates Using Atomic Vapor*. Phys. Rev. X **10**, 11027 (2020). [Cited on page 1.]

[10] C. L. Degen, F. Reinhard, and P. Cappellaro. *Quantum sensing*. Rev. Mod. Phys. **89**, 035002 (2017). [Cited on page 1.]

[11] J. Kitching, S. Knappe, and E. A. Donley. *Atomic sensors - A review* (2011). [Cited on page 1.]

[12] M. A. Zentile, J. Keaveney, R. S. Mathew, D. J. Whiting, C. S. Adams, and I. G. Hughes. *Optimization of atomic Faraday filters in the presence of homogeneous line broadening*. J. Phys. B At. Mol. Opt. Phys. **48**, 185001 (2015). [Cited on page 1.]

[13] W. D. Phillips. *Laser cooling and trapping of neutral atoms*. Rev. Mod. Phys. **70**, 721 (1998). [Cited on page 1.]

[14] M. Pellaton, C. Affolderbach, Y. Pétremand, N. de Rooij, and G. Mileti. *Study of laser-pumped double-resonance clock signals using a microfabricated cell*. Phys. Scr. **T149**, 14013 (2012). [Cited on page 1.]

[15] W. C. Griffith, R. Jimenez-Martinez, V. Shah, S. Knappe, and J. Kitching. *Miniature atomic magnetometer integrated with flux concentrators*. Appl. Phys. Lett. **94**, 23502 (2009). [Cited on page 1.]

[16] C. J. Foot. *Atomic physics*. Oxford master series in physics. Oxford University Press (2005). [Cited on pages 2, 10, 11, 12, 13, 123, and 143.]

[17] D. D. Budker and M. Romalis. *Optical magnetometry*. Nat. Phys. **3**, 227 (2007). [Cited on page 2.]

[18] B. A. Olsen, B. Patton, Y.-Y. Jau, and W. Happer. *Optical pumping and spectroscopy of Cs vapor at high magnetic field*. Phys. Rev. A **84**, 63410 (2011). [Cited on pages 2, 66, 100, and 163.]

[19] A. Sargsyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, A. Papoyan, D. Sarkisyan, and M. Auzinsh. *Hyperfine Paschen–Back regime in alkali metal atoms: consistency of two theoretical considerations and experiment*. J. Opt. Soc. Am. B **31**, 1046 (2014). [Cited on pages 2, 66, 100, and 163.]

[20] A. Sargsyan, A. Tonoyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, and D. Sarkisyan. *Complete hyperfine Paschen-Back regime at relatively small magnetic fields realized in potassium nano-cell*. Europhys. Lett. **110**, 23001 (2015). [Cited on pages 2, 66, 100, and 163.]

[21] A. Sargsyan, A. Tonoyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, and D. Sarkisyan. *Atomic transitions of Rb, D2 line in strong magnetic fields: Hyperfine Paschen–Back regime.* Opt. Commun. **334**, 208 (2015). [Cited on pages 2, 66, 100, and 163.]

[22] A. Sargsyan, E. Klinger, A. Tonoyan, C. Leroy, and D. Sarkisyan. *Hyperfine Paschen–Back regime of potassium D 2 line observed by Doppler-free spectroscopy.* J. Phys. B At. Mol. Opt. Phys. **51**, 145001 (2018). [Cited on pages 2, 66, 100, and 163.]

[23] D. J. Whiting, E. Bimbard, J. Keaveney, M. A. Zentile, C. S. Adams, and I. G. Hughes. *Electromagnetically induced absorption in a non-degenerate three-level ladder system.* Opt. Lett. **40**, 4289 (2015). [Cited on pages 2, 66, 100, and 163.]

[24] D. J. Whiting, J. Keaveney, C. S. Adams, and I. G. Hughes. *Direct measurement of excited-state dipole matrix elements using electromagnetically induced transparency in the hyperfine Paschen-Back regime.* Phys. Rev. A - At. Mol. Opt. Phys. **93**, 43854 (2016). [Cited on pages 2, 66, 71, 97, 99, 100, 104, 107, 124, 125, 139, and 163.]

[25] D. J. Whiting, R. S. Mathew, J. Keaveney, C. S. Adams, and I. G. Hughes. *Four-wave mixing in a non-degenerate four-level diamond configuration in the hyperfine Paschen–Back regime.* J. Mod. Opt. **65**, 713 (2018). [Cited on pages 2, 7, 66, 98, 120, 163, and 167.]

[26] H. Paul. *Introduction to quantum theory*, volume 9780521876. Cambridge University Press (2008). [Cited on pages 3 and 124.]

[27] T. Young. *II. The Bakerian Lecture. On the mechanism of the eye.* Philos. Trans. R. Soc. London pp. 23–88 (1801). [Cited on page 3.]

[28] A. Fresnel. *Premier mémoire sur la diffraction de la lumière.* Mémoire Adress. à l'Académie des Sci. le **15**, 1866 (1815). [Cited on page 3.]

[29] G. I. Taylor. *Interference fringes with feeble light.* In *Proc. Camb. Philol. Soc.*, volume 15, pp. 114–115 (1909). [Cited on pages 3 and 4.]

[30] R. S. Aspden, M. J. Padgett, and G. C. Spalding. *Video recording true single-photon double-slit interference.* Am. J. Phys. **84**, 671 (2016). [Cited on page 3.]

[31] M. Planck. *On an Improvement of Wien's Equation for the Spectrum.* Verh. Deut. Phys. Ges **2**, 202 (1900). [Cited on page 4.]

[32] M. Planck. *On the law of distribution of energy in the normal spectrum.* Ann. Phys. **4**, 1 (1901). [Cited on page 4.]

[33] R. Loudon. *The Quantum Theory of Light.* OUP Oxford (2000). [Cited on pages 4 and 138.]

[34] P. Ball. *Beyond Weird.* Random House (2018). [Cited on page 4.]

[35] M. Kumar. *Quantum: Einstein, Bohr and the Great Debate About the Nature of Reality.* Icon Books Limited (2008). [Cited on page 4.]

[36] M. Planck and F. Gaynor. *Scientific Autobiography, and Other Papers: With a Memorial Addresses on Max Planck, by Max Von Laue. Tr. from German by Frank Gaynor.* Philosophical Library (1949). [Cited on page 4.]

[37] A. Einstein. *On the Electrodynamics of Moving Bodies (Translation).* Ann. Phys. **7**, 891 (1905). [Cited on page 4.]

[38] G. N. Lewis. *The conservation of photons.* Nature **118**, 874 (1926). [Cited on page 4.]

[39] L. T. Troland. *On the measurement of visual stimulation intensities.* J. Exp. Psychol. **2**, 1 (1917). [Cited on page 4.]

[40] W. E. Lamb. *Anti-photon.* Appl. Phys. B **84**, 77 (1995). [Cited on page 5.]

[41] P. A. M. Dirac. *The Quantum Theory of the Emission and Absorption of Radiation.* P. Roy. Soc. A **114**, 243 (1927). [Cited on page 4.]

[42] I. A. Walmsley. *Light: A Very Short Introduction.* Very short introductions. Oxford University Press (2015). [Cited on page 5.]

[43] T. H. Maiman. *Stimulated Optical Radiation in Ruby.* Nature **187**, 493 (1960). [Cited on page 5.]

[44] A. L. Schawlow and C. H. Townes. *Infrared and Optical Masers.* Phys. Rev. **112**, 1940 (1958). [Cited on page 5.]

[45] R. J. Glauber. *Coherent and incoherent states of the radiation field.* Phys. Rev. **131**, 2766 (1963). [Cited on page 5.]

[46] H. J. Kimble, M. Dagenais, and L. Mandel. *Photon antibunching in resonance fluorescence.* Phys. Rev. Lett. **39**, 691 (1977). [Cited on page 5.]

[47] R. H. Brown and R. Q. Twiss. *A test of a new type of stellar interferometer on Sirius.* Nature **178**, 1046 (1956). [Cited on page 5.]

[48] R. S. Mathew, F. Ponciano-Ojeda, J. Keaveney, D. J. Whiting, and I. G. Hughes. *Simultaneous two-photon resonant optical laser locking (STROLLing) in the hyperfine Paschen–Back regime.* Opt. Lett. **43**, 4204 (2018). [Cited on pages 7 and 166.]

[49] R. S. Mathew, R. O'Donnell, D. Pizzey, and I. G. Hughes. *The Raspberry Pi auto-aligner: Machine learning for automated alignment of laser beams.* Rev. Sci. Instrum. **92**, 015117 (2021). [Cited on pages 7, 42, 48, 63, and 166.]

[50] C. R. Higgins, D. Pizzey, R. S. Mathew, and I. G. Hughes. *Atomic line versus lens cavity filters: a comparison of their merits.* OSA Contin. **3**, 961 (2020). [Cited on pages 7, 30, 38, and 166.]

[51] B. H. Bransden and C. J. Joachain. *Quantum mechanics.* Pearson Education, Harlow, 2nd edition (1979). [Cited on pages 9 and 24.]

[52] S. Rand and S. C. Rand. *Lectures on Light: Nonlinear and Quantum Optics Using the Density Matrix.* OUP Oxford (2010). [Cited on page 10.]

[53] K. Blum. *Density Matrix Theory and Applications.* Springer Series on Atomic, Optical, and Plasma Physics. Springer Berlin Heidelberg (2012). [Cited on page 11.]

[54] A. M. Fox, M. Fox, D. Fox, and O. U. Press. *Quantum Optics: An Introduction.* Oxford Master Series in Physics. OUP Oxford (2006). [Cited on pages 13 and 135.]

[55] D. Manzano. *A short introduction to the Lindblad master equation.* AIP Adv. **10**, 025106 (2020). [Cited on page 13.]

[56] H. P. Breuer, P. Breuer, F. Petruccione, and S. Petruccione. *The Theory of Open Quantum Systems.* Oxford University Press (2002). [Cited on page 13.]

[57] J. R. Johansson, P. D. Nation, and F. Nori. *QuTiP 2: An open-source Python framework for the dynamics of open quantum systems.* Comput. Phys. Commun. **184**, 1234 (2013). [Cited on pages 13, 15, and 102.]

[58] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-Photon Interactions.* Wiley (1998). [Cited on pages 13, 17, and 70.]

[59] J. Gea-Banacloche, Y. Q. Li, S. Z. Jin, and M. Xiao. *Electromagnetically induced transparency in ladder-type inhomogeneously broadened media: Theory and experiment.* Phys. Rev. A **51**, 576 (1995). [Cited on page 14.]

[60] B. E. Sherlock and I. G. Hughes. *How weak is a weak probe in laser spectroscopy?* Am. J. Phys. **77**, 111 (2009). [Cited on page 14.]

[61] A. M. Fox and D. Fox. *Optical Properties of Solids*. Oxford master series in condensed matter physics. Oxford University Press (2001). [Cited on page 15.]

[62] S. J. Blundell and K. M. Blundell. *Concepts in Thermal Physics*. OUP Oxford (2010). [Cited on page 20.]

[63] P. Zeeman. *The Effect of Magnetisation on the Nature of Light Emitted by a Substance*. Nature **55**, 347 (1897). [Cited on pages 22 and 24.]

[64] D. A. Steck. *Rubidium 87 D Line Data. http://steck.us/alkalidata*. [Cited on page 23.]

[65] L. Weller, K. S. Kleinbach, M. A. Zentile, S. Knappe, C. S. Adams, and I. G. Hughes. *Absolute absorption and dispersion of a rubidium vapour in the hyperfine Paschen–Back regime*. J. Phys. B At. Mol. Opt. Phys. **45**, 215005 (2012). [Cited on pages 23, 24, and 100.]

[66] M. A. Zentile, J. Keaveney, L. Weller, D. J. Whiting, C. S. Adams, and I. G. Hughes. *ElecSus: A program to calculate the electric susceptibility of an atomic ensemble*. Comput. Phys. Commun. **189**, 162 (2015). [Cited on pages 24, 71, 73, 87, and 90.]

[67] J. Keaveney, C. S. Adams, and I. G. Hughes. *ElecSus: Extension to arbitrary geometry magneto-optics*. Comput. Phys. Commun. **224**, 311 (2017). [Cited on pages 24 and 73.]

[68] G. K. Woodgate. *Elementary Atomic Structure*. Oxford science publications. Clarendon Press (1980). [Cited on page 25.]

[69] P. Palittapongarnpim, A. MacRae, and A. I. Lvovsky. *Note: A monolithic filter cavity for experiments in quantum optics*. Rev. Sci. Instrum. **83**, 2010 (2012). [Cited on page 37.]

[70] G. Ghosh. *Sellmeier Coefficients and Chromatic Dispersions for Some Tellurite Glasses*. J. Am. Ceram. Soc. **78**, 2828 (1995). [Cited on page 38.]

[71] *TIE-19 Temperature Coefficient of the Refractive Index*. Technical report, SCHOTT (2016). [Cited on pages 38 and 39.]

[72] Andrew John MacRae. *An Atomic Source of Quantum Light*. Ph.D. thesis, University of Calgary (2012). [Cited on page 39.]

[73] C. S. Adams and I. G. Hughes. *Optics f2f: From Fourier to Fresnel*. OUP Oxford (2018). [Cited on pages 39 and 127.]

[74] L. Weller. *Absolute Absorption and Dispersion in a Thermal Rb Vapour at High Densities and High Magnetic Fields*. Ph.D. thesis, Durham University (2013). [Cited on pages 40 and 125.]

[75] M. A. Zentile. *Applications of the Faraday Effect in Hot Atomic Vapours*. Ph.D. thesis, Durham University (2015). [Cited on pages 40, 71, 83, 107, and 139.]

[76] P. B. Wigley, P. J. Everitt, A. Van Den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. Mcdonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush. *Fast machine-learning online optimization of ultra-cold-atom experiments*. Sci. Rep. **6**, 1 (2016). [Cited on pages 43 and 46.]

[77] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. Adaptive Computation and Machine Learning series. MIT Press (2012). [Cited on page 44.]

[78] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. Adaptive Computation and Machine Learning series. MIT Press (2016). [Cited on pages 44 and 46.]

[79] R. S. Judson and H. Rabitz. *Teaching lasers to control molecules*. Phys. Rev. Lett. **68**, 1500 (1992). [Cited on pages 44, 46, and 64.]

[80] J. Nathan Kutz, X. Fu, and S. Brunton. *Self-tuning fiber lasers: Machine learning applied to optical systems*. Nonlinear Photonics, NP 2014 pp. 1–2 (2014). [Cited on pages 44 and 64.]

[81] U. Andral, J. Buguet, R. S. Fodil, F. Amrani, F. Billard, E. Hertz, and P. Grelu. *Toward an autosetting mode-locked fiber laser cavity*. J. Opt. Soc. Am. B **33**, 825 (2016). [Cited on page 44.]

[82] T. Baumeister, S. L. Brunton, and J. N. Kutz. *Deep learning and model predictive control for self-tuning mode-locked lasers*. J. Opt. Soc. Am. B **35**, 617 (2018). [Cited on page 44.]

[83] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, and Others. *Fast machine-learning online optimization of ultra-cold-atom experiments*. Sci. Rep. **6**, 1 (2016). [Cited on page 44.]

[84] Y. Zhang, C. Wu, Y. Song, K. Si, Y. Zheng, L. Hu, J. Chen, L. Tang, and W. Gong. *Machine learning based adaptive optics for doughnut-shaped beam*. Opt. Express **27**, 16871 (2019). [Cited on page 44.]

[85] A. A. Melnikov, H. P. Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel. *Active learning machine learns to create new quantum experiments*. Proc. Natl. Acad. Sci. **115**, 1221 (2018). [Cited on page 44.]

[86] T. Baumeister, S. L. Brunton, and J. Nathan Kutz. *Deep learning and model predictive control for self-tuning mode-locked lasers*. J. Opt. Soc. Am. B **35**, 617 (2018). [Cited on pages 44 and 46.]

[87] B. Adams. *Nonlinear Optics, Quantum Optics, and Ultrafast Phenomena with X-Rays: Physics with X-Ray Free-Electron Lasers*. Springer US (2012). [Cited on page 45.]

[88] B. Burger, P. Maffettone, V. Gusev, C. Aitchison, Y. Bai, W. Xiaoyan, R. S. Sprick, and A. I. Cooper. *A Mobile Robotic Researcher*. Nature **583** (2020). [Cited on page 45.]

[89] D. Deutsch. *The Beginning of Infinity: Explanations that Transform the World*. Penguin Books (2012). [Cited on page 45.]

[90] B. J. Pearson, J. L. White, T. C. Weinacht, and P. H. Bucksbaum. *Coherent control using adaptive learning algorithms*. Phys. Rev. A **63**, 063412 (2001). [Cited on pages 46 and 64.]

[91] D. E. Goldberg. *Genetic Algorithms*. Pearson Education (2006). [Cited on page 46.]

[92] M. A. Nielsen. *Neural networks and deep learning*, volume 2018. Determination press, San Francisco, CA, USA: (2015). [Cited on page 46.]

[93] D. R. Burt and C. E. Rasmussen. *Convergence of Sparse Variational Inference in Gaussian Processes Regression Mark van der Wilk*. J. Mach. Learn. Res. **21**, 1 (2020). [Cited on page 46.]

[94] R. J. H. O. Donnell. *Machine Learning for Optimisation of Laser Physics Experiments*. Master's thesis, Durham University (2020). [Cited on pages 46 and 63.]

[95] R. Dawkins. *In defence of selfish genes*. Philosophy **56**, 556 (1981). [Cited on page 55.]

[96] S. Kobtsev, S. Kandrushin, and A. Potekhin. *Long-term frequency stabilization of a continuous-wave tunable laser with the help of a precision wavelengthmeter.* Appl. Opt. **46**, 5840 (2007). [Cited on page 66.]

[97] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. *Laser phase and frequency stabilization using an optical resonator.* Appl. Phys. B **31**, 97 (1983). [Cited on page 66.]

[98] T. Day, E. K. Gustafson, and R. L. Byer. *Sub-Hertz Relative Frequency Stabilization of Two-Diode Laser-Pumped Nd: YAG Lasers Locked to a Fabry-Perot Interferometer.* IEEE J. Quantum Electron. **28**, 1106 (1992). [Cited on page 66.]

[99] T. Uehara, K. Tsuji, K. Hagiwara, and N. Onodera. *Optical beat-note frequency stabilization between two lasers using a radio frequency interferometer in the gigahertz frequency band.* Opt. Eng. **53**, 124109 (2014). [Cited on page 66.]

[100] G. C. Bjorklund. *Frequency-modulation spectroscopy: a new method for measuring weak absorptions and dispersions.* Opt. Lett. **5**, 15 (1980). [Cited on page 66.]

[101] J. H. Shirley. *Modulation transfer processes in optical heterodyne saturation spectroscopy.* Opt. Lett. **7**, 537 (1982). [Cited on page 66.]

[102] D. J. McCarron, S. A. King, and S. L. Cornish. *Modulation transfer spectroscopy in atomic rubidium.* Meas. Sci. Technol. **19**, 105601 (2008). [Cited on page 66.]

[103] K. L. Corwin, Z. T. Lu, C. F. Hand, R. J. Epstein, and C. E. Wieman. *Frequency-stabilized diode laser with the Zeeman shift in an atomic vapor.* Appl. Opt. **37**, 3295 (1998). [Cited on page 67.]

[104] T. Petelski, M. Fattori, G. Lamporesi, J. Stuhler, and G. M. Tino. *Doppler-free spectroscopy using magnetically induced dichroism of atomic vapor: A new scheme for laser frequency locking.* Eur. Phys. J. D **22**, 279 (2003). [Cited on page 67.]

[105] M. L. Harris, S. L. Cornish, A. Tripathi, and I. G. Hughes. *Optimization of sub-Doppler DAVLL on the rubidium D2 line.* J. Phys. B At. Mol. Opt. Phys. **41**, 085401 (2008). [Cited on page 67.]

[106] F. E. Becerra, R. T. Willis, S. L. Rolston, and L. a. Orozco. *Two-photon dichroic atomic vapor laser lock using electromagnetically induced transparency and absorption.* J. Opt. Soc. Am. B **26**, 1315 (2009). [Cited on page 67.]

[107] C. P. Pearman, C. S. Adams, S. G. Cox, P. F. Griffin, D. A. Smith, and I. G. Hughes. *Polarization spectroscopy of a closed atomic transition: Applications to laser frequency locking.* J. Phys. B At. Mol. Opt. Phys. **35**, 5141 (2002). [Cited on page 67.]

[108] G. R. Hanes and C. E. Dahlstrom. *Iodine hyperfine structure observed in saturated absorption at 633 nm.* Appl. Phys. Lett. **14**, 362 (1969). [Cited on page 67.]

[109] G. T. Purves, G. Jundt, C. S. Adams, and I. G. Hughes. *Refractive index measurements by probe-beam deflection.* Eur. Phys. J. D - At. Mol. Opt. Phys. **29**, 433 (2004). [Cited on page 67.]

[110] M. A. Zentile, R. Andrews, L. Weller, S. Knappe, C. S. Adams, and I. G. Hughes. *The hyperfine Paschen–Back Faraday effect.* J. Phys. B At. Mol. Opt. Phys. **47**, 75005 (2014). [Cited on pages 67, 69, 70, 76, and 100.]

[111] D. J. Reed, N. Šibalić, D. J. Whiting, J. M. Kondo, C. S. Adams, and K. J. Weatherill. *Low-drift Zeeman shifted atomic frequency reference.* arXiv p. 1804.07928 (2018). [Cited on page 67.]

[112] A. Sargsyan, A. Tonoyan, R. Mirzoyan, D. Sarkisyan, A. M. Wojciechowski, A. Stabrawa, and W. Gawlik. *Saturated-absorption spectroscopy revisited: atomic transitions in strong magnetic fields (>20 mT) with a micrometer-thin cell.* Opt. Lett. **39**, 2270 (2014). [Cited on page 67.]

[113] P. M. T. Barboza, G. G. Nascimento, M. O. Araújo, C. M. da Silva, H. L. D. d. S. Cavalcante, M. Oriá, M. Chevrollier, and T. P. de Silans. *Stabilization of a laser on a large-detuned atomic-reference frequency by resonant interferometry.* J. Phys. B At. Mol. Opt. Phys. **49**, 085401 (2016). [Cited on page 67.]

[114] J. Keaveney, W. J. Hamlyn, C. S. Adams, and I. G. Hughes. *A single-mode external cavity diode laser using an intra-cavity atomic Faraday filter with short-term linewidth <400 kHz and long-term stability of <1 MHz.* Rev. Sci. Instrum. **87**, 95111 (2016). [Cited on page 67.]

[115] A. Y. Kalatskiy, A. E. Afanasiev, P. N. Melentiev, and V. I. Balykin. *Frequency stabilization of a diode laser on the 5 $P \rightarrow$ 5 D transition of the Rb atom.* Laser Phys. **27**, 055703 (2017). [Cited on page 67.]

[116] R. P. Abel, A. K. Mohapatra, M. G. Bason, J. D. Pritchard, K. J. Weatherill, U. Raitzsch, and C. S. Adams. *Laser frequency stabilization to excited state transitions using electromagnetically induced transparency in a cascade system.* Appl. Phys. Lett. **94**, 71107 (2009). [Cited on page 67.]

[117] S. C. Bell, D. M. Heywood, J. D. White, J. D. Close, and R. E. Scholten. *Laser frequency offset locking using electromagnetically induced transparency.* Appl. Phys. Lett. **90**, 1 (2007). [Cited on page 67.]

[118] A. Sargsyan, A. V. Papoyan, D. Sarkisyan, and A. Weis. *Efficient technique for measuring laser frequency stability.* Eur. Phys. J. Appl. Phys. **48**, 20701 (2009). [Cited on page 67.]

[119] M. Parniak, A. Leszczyński, and W. Wasilewski. *Magneto-optical polarization rotation in a ladder-type atomic system for tunable offset locking.* Appl. Phys. Lett. **108**, 161103 (2016). [Cited on page 67.]

[120] C. Carr, C. S. Adams, and K. J. Weatherill. *Polarization spectroscopy of an excited state transition.* Opt. Lett. **37**, 118 (2012). [Cited on page 67.]

[121] H.-R. Noh and H. S. Moon. *Discrimination of one-photon and two-photon coherence parts in electromagnetically induced transparency for a ladder-type three-level atomic system.* Opt. Express **19**, 11128 (2011). [Cited on page 67.]

[122] H.-R. Noh and H. S. Moon. *Three-photon coherence in a ladder-type atomic system.* Phys. Rev. A **92**, 013807 (2015). [Cited on page 67.]

[123] Y.-S. Lee, S. M. Lee, H. Kim, and H. S. Moon. *Highly bright photon-pair generation in Doppler-broadened ladder-type atomic system.* Opt. Express **24**, 28083 (2016). [Cited on pages 67, 99, 136, 148, 162, and 168.]

[124] Y.-S. Lee, S. M. Lee, H. Kim, and H. S. Moon. *Single-photon superradiant beating from a Doppler-broadened ladder-type atomic ensemble.* Phys. Rev. A **96**, 63832 (2017). [Cited on page 67.]

[125] B. Srivathsan, G. K. Gulati, B. Chng, G. Maslennikov, D. Matsukevich, and C. Kurtsiefer. *Narrow band source of transform-limited photon pairs via four-wave mixing in a cold atomic ensemble.* Phys. Rev. Lett. **111**, 123602 (2013). [Cited on pages 67, 136, 142, 148, and 162.]

[126] B. Huber, A. Kölle, and T. Pfau. *Motion-induced signal revival in pulsed Rydberg four-wave mixing beyond the frozen-gas limit.* Phys. Rev. A **90**, 53806 (2014). [Cited on page 67.]

[127] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. *Electromagnetically induced transparency: Optics in coherent media.* Rev. Mod. Phys. **77**, 633 (2005). [Cited on page 70.]

[128] H. S. Moon and H.-R. Noh. *Polarization dependence of double-resonance optical pumping and electromagnetically induced transparency in the 5S(1/2)-5P(3/2)-5D(5/2) transition of Rb-87 atoms.* Phys. Rev. A **84**, 033821 (2011). [Cited on page 70.]

[129] D. Budker, D. Kimball, S. Rochester, and V. Yashchuk. *Nonlinear Magneto-optics and Reduced Group Velocity of Light in Atomic Vapor with Slow Ground State Relaxation.* Phys. Rev. Lett. **83**, 1767 (1999). [Cited on page 70.]

[130] J. Keaveney. *Automated translating beam profiler for in situ laser beam spot-size and focal position measurements.* Rev. Sci. Instrum. **89**, 035114 (2018). [Cited on page 70.]

[131] D. J. Whiting. *Nonlinear Optics in a Thermal Rb Vapour at High Magnetic Fields.* Ph.D. thesis, Durham University (2017). [Cited on pages 71 and 137.]

[132] P. Siddons, C. S. Adams, C. Ge, and I. G. Hughes. *Absolute absorption on rubidium D lines: comparison between theory and experiment.* J. Phys. B At. Mol. Opt. Phys. **41**, 155004 (2008). [Cited on pages 72 and 128.]

[133] I. G. Hughes and T. P. A. Hase. *Measurements and their uncertainties: a practical guide to modern error analysis.* Oxford University Press, Oxford (2010). [Cited on pages 73, 91, and 152.]

[134] W. J. Riley and D. A. Howe. *Handbook of Frequency Stability Analysis.* Spec. Publ. (NIST SP) - 1065 (2008). [Cited on page 74.]

[135] D. Allan. *Statistics of atomic frequency standards.* Proc. IEEE **54**, 221 (1966). [Cited on page 74.]

[136] R. W. Wood and F. L. Mohler. *Resonance Radiation of Sodium Vapor Excited by one of the D Lines.* Phys. Rev. **11**, 70 (1918). [Cited on page 80.]

[137] J. F. Sell, M. A. Gearba, B. M. Patterson, T. Genda, B. Naumann, and R. J. Knize. *Enhancement of Rb fine-structure transfer in 4He due to three-body collisions.* Opt. Lett. **35**, 2146 (2010). [Cited on pages 80 and 81.]

[138] M. D. Rotondaro and G. P. Perram. *Role of rotational-energy defect in collisional transfer between the* $5^2P_{1/2,3/2}$ *levels in rubidium.* Phys. Rev. A **57**, 4045 (1998). [Cited on page 80.]

[139] M. A. Gearba, J. H. Wells, P. H. Rich, J. M. Wesemann, L. A. Zimmerman, B. M. Patterson, R. J. Knize, and J. F. Sell. *Collisional excitation transfer and quenching in* Rb($5P$)*-methane mixtures.* Phys. Rev. A **99**, 022706 (2019). [Cited on pages 80 and 81.]

[140] J. F. Sell, M. A. Gearba, B. M. Patterson, D. Byrne, G. Jemo, T. C. Lilly, R. Meeter, and R. J. Knize. *Collisional excitation transfer between Rb(5P) states in 50–3000 Torr of 4 He.* J. Phys. B At. Mol. Opt. Phys. **45**, 055202 (2012). [Cited on page 80.]

[141] A. Omont, E. W. Smith, and J. Cooper. *Redistribution of Resonance Radiation. I. The Effect of Collisions.* Astrophys. J. **175**, 185 (1972). [Cited on page 81.]

[142] J. L. Carlsten and A. Szoke. *Collisional redistribution of near-resonant scattered light in Sr vapour.* J. Phys. B At. Mol. Phys. **9**, L231 (1976). [Cited on page 81.]

[143] A. Gallagher and D. E. Pritchard. *Exoergic collisions of cold Na*-Na.* Phys. Rev. Lett. **63**, 957 (1989). [Cited on page 81.]

[144] R. J. Beach, W. F. Krupke, V. K. Kanz, S. A. Payne, M. A. Dubinskii, and L. D. Merkle. *End-pumped continuous-wave alkali vapor lasers: experiment, model, and power scaling.* J. Opt. Soc. Am. B **21**, 2151 (2004). [Cited on page 81.]

[145] H. Cai, Y. Wang, J. Han, G. An, W. Zhang, L. Xue, H. Wang, J. Zhou, M. Gao, and Z. Jiang. *Reviews of a Diode-Pumped Alkali Laser (DPAL): a potential high powered light source.* In *Selected Pap. from Conf. Photoelectron. Technol. Comm. Chinese Soc. Astronaut. 2014, Part I*, pp. 421 – 433 (2015). [Cited on page 81.]

[146] G. A. Pitz, D. M. Stalnaker, E. M. Guild, B. Q. Oliker, P. J. Moran, S. W. Townsend, and D. A. Hostutler. *Advancements in flowing diode pumped alkali lasers.* High Energy/Average Power Lasers Intense Beam Appl. IX **9729**, 972902 (2016). [Cited on page 81.]

[147] E. Lewis. *Collisional relaxation of atomic excited states, line broadening and interatomic interactions.* Phys. Rep. **58**, 1 (1980). [Cited on page 81.]

[148] M. D. Rotondaro and G. P. Perram. *Collision-induced transitions between the Zeeman-split (J,m) levels of Rb(5P).* Phys. Rev. A **58**, 2023 (1998). [Cited on page 81.]

[149] C. B. Alcock, V. P. Itkin, and M. K. Horrigan. *Vapour pressure equations for the metallic elements: 298-2500k.* Can. Metall. Q. **23**, 309 (1984). [Cited on page 87.]

[150] M. Tetenbaum, J. F. Babelot, and J. Magill. *An assessment of the melting, boiling, and critical point data of the alkali metals.* Pure Appl. Chem. **57**, 1407 (1985). [Cited on page 87.]

[151] J. H. Lienhard. *A Heat Transfer Textbook: Fifth Edition.* Dover Books on Engineering. Dover Publications (2019). [Cited on page 90.]

[152] A. Macrae, T. Brannan, R. Achal, and A. I. Lvovsky. *Tomography of a high-purity narrowband photon from a transient atomic collective excitation.* Phys. Rev. Lett. **109**, 33601 (2012). [Cited on pages 99 and 136.]

[153] R. C. Pooser and B. Lawrie. *Ultrasensitive measurement of microcantilever displacement below the shot-noise limit.* Optica **2**, 393 (2015). [Cited on page 99.]

[154] F. Ripka, Y.-h. Chen, R. Löw, and T. Pfau. *Rydberg polaritons in a thermal vapor.* Phys. Rev. A **93**, 053429 (2016). [Cited on page 99.]

[155] H. Kang, G. Hernandez, J. Zhang, and Y. Zhu. *Backward four-wave mixing in a four-level medium with electromagnetically induced transparency.* J. Opt. Soc. Am. B **23**, 718 (2006). [Cited on page 99.]

[156] F. Becerra, R. Willis, S. Rolston, and L. Orozco. *Nondegenerate four-wave mixing in rubidium vapor: The diamond configuration.* Phys. Rev. A **78**, 13834 (2008). [Cited on page 99.]

[157] R. T. Willis. *Photon pair production from a hot atomic ensemble in the diamond configuration.* Ph.D. thesis, University of Maryland (2009). [Cited on page 99.]

[158] M. M. M. Parniak and W. Wasilewski. *Interference and nonlinear properties of four-wave-mixing resonances in thermal vapor: Analytical results and experimental verification.* Phys. Rev. A - At. Mol. Opt. Phys. **91**, 23418 (2015). [Cited on page 99.]

[159] K. Heshami, D. G. England, P. C. Humphreys, P. J. Bustard, V. M. Acosta, J. Nunn, and B. J. Sussman. *Quantum memories: emerging applications and recent advances.* J. Mod. Opt. **63**, 2005 (2016). [Cited on page 99.]

[160] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin. *Quantum repeaters based on atomic ensembles and linear optics*. Rev. Mod. Phys. **83**, 33 (2011). [Cited on page 99.]

[161] D. Paredes-Barato and C. S. Adams. *All-Optical Quantum Information Processing Using Rydberg Gates*. Phys. Rev. Lett. **112**, 040501 (2014). [Cited on page 99.]

[162] L. Weller, T. Dalton, P. Siddons, C. S. Adams, and I. G. Hughes. *Measuring the Stokes parameters for light transmitted by a high-density rubidium vapour in large magnetic fields*. J. Phys. B At. Mol. Opt. Phys. **45**, 055001 (2012). [Cited on page 100.]

[163] A. Sargsyan, E. Klinger, G. Hakhumyan, A. Tonoyan, A. Papoyan, C. Leroy, and D. Sarkisyan. *Decoupling of hyperfine structure of Cs D1 line in strong magnetic field studied by selective reflection from a nanocell*. J. Opt. Soc. Am. B **34**, 776 (2017). [Cited on page 100.]

[164] A. Sargsyan, A. Papoyan, I. G. Hughes, C. S. Adams, and D. Sarkisyan. *Selective reflection from an Rb layer with a thickness below $\lambda/12$ and applications*. Opt. Lett. **42**, 1476 (2017). [Cited on page 100.]

[165] R. W. Boyd. *Nonlinear Optics*. Elsevier Science, USA, 3rd edition (2008). [Cited on pages 102, 103, and 137.]

[166] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-photon interactions: basic processes and applications*. Wiley-Interscience publication. J. Wiley (1992). [Cited on page 102.]

[167] A. M. Guenault. *Statistical Physics*. Student Physics Series. Springer Netherlands (2007). [Cited on page 103.]

[168] O. S. Heavens. *Radiative Transition Probabilities of the Lower Excited States of the Alkali Metals*. J. Opt. Soc. Am. **51**, 1058 (1961). [Cited on page 104.]

[169] A. Sargsyan, D. Sarkisyan, U. Krohn, J. Keaveney, and C. Adams. *Effect of buffer gas on an electromagnetically induced transparency in a ladder system using thermal rubidium vapor*. Phys. Rev. A **82**, 045806 (2010). [Cited on page 104.]

[170] P. Arora, A. Agarwal, and A. S. Gupta. *Simple alignment technique for polarisation maintaining fibres*. Rev. Sci. Instrum. **82**, 125103 (2011). [Cited on page 108.]

[171] O. López-hernández, S. Hernández-gómez, F. S. Ponciano-ojeda, F. Ramírez-martínez, and J. Jiménez-mier. *A laser spectroscopy system with combined absorption, polarization rotation and fluorescence detection to study two photon transitions in atomic rubidium.* Rev. Mex. Trastor. Aliment. **13**, 543 (2015). [Cited on page 112.]

[172] R. R. Moseley, S. Shepherd, D. J. Fulton, B. D. Sinclair, and M. H. Dunn. *Two-photon effects in continuous-wave electromagnetically-induced transparency.* Opt. Commun. **119**, 61 (1995). [Cited on page 124.]

[173] Y. Q. Li, S. Z. Jin, and M. Xiao. *Observation of an electromagnetically induced change of absorption in multilevel rubidium atoms.* Phys. Rev. A **51**, 1754 (1995). [Cited on page 124.]

[174] M. Xiao, Y.-Q. Li, S.-Z. Jin, and J. Gea-Banacloche. *Measurement of Dispersive Properties of Electromagnetically Induced Transparency in Rubidium Atoms.* Phys. Rev. Lett. **74**, 666 (1995). [Cited on page 124.]

[175] D. J. Maas, C. W. Rella, P. Antoine, E. S. Toma, and L. D. Noordam. *Population transfer via adiabatic passage in the rubidium quantum ladder system.* Phys. Rev. A - At. Mol. Opt. Phys. **59**, 1374 (1999). [Cited on page 124.]

[176] M. S. Safronova, C. J. Williams, and C. W. Clark. *Relativistic many-body calculations of electric-dipole matrix elements, lifetimes and polarizabilities in rubidium.* Phys. Rev. A **69**, 22509 (2003). [Cited on page 124.]

[177] D. J. Whiting, J. Keaveney, C. S. Adams, and I. G. Hughes. *Erratum: Direct measurement of excited-state dipole matrix elements using electromagnetically induced transparency in the hyperfine Paschen-Back regime.* Phys. Rev. A **97**, 49902 (2018). [Cited on pages 125 and 126.]

[178] M. Weissbluth. *Atoms and Molecules.* Academic Press, London (1978). [Cited on page 125.]

[179] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov. *Invited Review Article: Single-photon sources and detectors.* Rev. Sci. Instrum. **82**, 071101 (2011). [Cited on pages 135 and 163.]

[180] B. Lounis and M. Orrit. *Single-photon sources.* Reports Prog. Phys. **68**, 1129 (2005). [Cited on page 135.]

[181] J. L. O'Brien, A. Furusawa, and J. Vučković. *Photonic quantum technologies.* Nat. Photonics **3**, 687 (2009). [Cited on page 135.]

[182] O. R. Frisch. *Take a photon....* Contemp. Phys. **7**, 45 (1965). [Cited on page 136.]

[183] B. Darquié, M. P. A. Jones, J. Dingjan, J. Beugnon, S. Bergamini, Y. Sortais, G. Messin, A. Browaeys, and P. Grangier. *Controlled single-photon emission from a single trapped two-level atom.* Science (80-.). **309**, 454 (2005). [Cited on page 136.]

[184] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. *Single-atom single-photon quantum interface.* Science (80-.). **317**, 488 (2007). [Cited on page 136.]

[185] J. Höffges, H. Baldauf, T. Eichler, S. Helmfrid, and H. Walther. *Heterodyne measurement of the fluorescent radiation of a single trapped ion.* Opt. Commun. **133**, 170 (1997). [Cited on page 136.]

[186] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. *Continuous generation of single photons with controlled waveform in an ion-trap cavity system.* Nature **431**, 1075 (2004). [Cited on page 136.]

[187] Y. L. A. Rezus, S. G. Walt, R. Lettow, A. Renn, G. Zumofen, S. Götzinger, and V. Sandoghdar. *Single-Photon Spectroscopy of a Single Molecule.* Phys. Rev. Lett. **108**, 93601 (2012). [Cited on page 136.]

[188] E. Moreau, I. Robert, J. M. Gérard, I. Abram, L. Manin, and V. Thierry-Mieg. *Single-mode solid-state single photon source based on isolated quantum dots in pillar microcavities.* Appl. Phys. Lett. **79**, 2865 (2001). [Cited on page 136.]

[189] R. Keil, M. Zopf, Y. Chen, B. Hoefer, J. Zhang, F. Ding, and O. G. Schmidt. *Solid-state ensemble of highly entangled photon sources at rubidium atomic transitions* pp. 1–16 (2016). [Cited on page 136.]

[190] E. Wu, J. R. Rabeau, G. Roger, F. Treussart, H. Zeng, P. Grangier, S. Prawer, and J.-F. Roch. *Room temperature triggered single-photon source in the near infrared.* New J. Phys. **9**, 434 (2007). [Cited on page 136.]

[191] J. M. Smith, S. A. Meynell, A. C. Bleszynski Jayich, and J. Meijer. *Colour centre generation in diamond for quantum technologies* (2019). [Cited on page 136.]

[192] E. Waks, E. Diamanti, and Y. Yamamoto. *Generation of photon number states.* New J. Phys. **8**, 4 (2006). [Cited on page 136.]

[193] J. G. Rarity, J. Fulconis, J. Duligall, W. J. Wadsworth, and P. S. J. Russell. *Photonic crystal fiber source of correlated photon pairs.* Opt. Express **13**, 534 (2005). [Cited on page 136.]

[194] M. M. Müller, A. Kölle, R. Löw, T. Pfau, T. Calarco, and S. Montangero. *Room-temperature Rydberg single-photon source.* Phys. Rev. A **87**, 53412 (2013). [Cited on page 136.]

[195] F. Ripka, H. Kübler, R. Löw, and T. Pfau. *A room-temperature single-photon source based on strongly interacting Rydberg atoms.* Science (80-.). **362**, 446 (2018). [Cited on page 136.]

[196] C. W. Chou, S. V. Polyakov, A. Kuzmich, and H. J. Kimble. *Single-Photon Generation from Stored Excitation in an Atomic Ensemble.* Phys. Rev. Lett. **92**, 213601 (2004). [Cited on page 136.]

[197] J. Nunn, N. K. Langford, W. S. Kolthammer, T. F. M. Champion, M. R. Sprague, P. S. Michelberger, X. M. Jin, D. G. England, and I. A. Walmsley. *Enhancing multiphoton rates with quantum memories.* Phys. Rev. Lett. **110**, 133601 (2013). [Cited on pages 136 and 163.]

[198] S. Du, P. Kolchin, C. Belthangady, G. Y. Yin, and S. E. Harris. *Subnatural Linewidth Biphotons with Controllable Temporal Length.* Phys. Rev. Lett. **100**, 183603 (2008). [Cited on page 136.]

[199] L. Zhao, X. Guo, C. Liu, Y. Sun, M. M. T. Loy, and S. Du. *Photon pairs with coherence time exceeding 1 μs.* Optica **1**, 84 (2014). [Cited on page 136.]

[200] D.-S. Ding, Z.-Y. Zhou, B.-S. Shi, X.-B. Zou, and G.-C. Guo. *Generation of non-classical correlated photon pairs via a ladder-type atomic configuration: theory and experiment.* Opt. Express **20**, 11433 (2012). [Cited on page 136.]

[201] C. Shu, P. Chen, T. K. A. Chow, L. Zhu, Y. Xiao, M. M. T. Loy, and S. Du. *Subnatural-linewidth biphotons from a Doppler-broadened hot atomic vapor cell.* Nat. Commun. **7**, 1 (2016). [Cited on page 136.]

[202] B. Srivathsan. *Heralded Single Photons for Efficient Interaction With Single Atoms.* Ph.D. thesis, National University of Singapore (2015). [Cited on pages 140 and 148.]

[203] P. Grangier, G. Roger, and A. Aspect. *Experimental evidence for a photon anticorrelation effect on a beam splitter: A new light on single-photon interferences.* EPL **1**, 173 (1986). [Cited on pages 140, 148, and 149.]

[204] C. Gerry and P. Knight. *Introductory quantum optics* (2005). [Cited on pages 142 and 146.]

[205] R. H. Brown and R. Q. Twiss. *Correlation between photons in two coherent beams of light.* Nature **177**, 27 (1956). [Cited on page 146.]

[206] P. Grünwald. *Effective second-order correlation function and single-photon detection*. New J. Phys. **21**, 093003 (2019). [Cited on page 146.]

[207] I. G. Hughes. *Velocity selection in a Doppler-broadened ensemble of atoms interacting with a monochromatic laser beam* pp. 1–12 (2017). [Cited on page 155.]