Wisniewski-Barker, Emma (2015) *Slow light in ruby: a study in spatial and temporal domains*. PhD thesis.

http://theses.gla.ac.uk/6591/

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Glasgow Theses Service
http://theses.gla.ac.uk/
theses@gla.ac.uk

Slow Light in Ruby:
a Study in Spatial and Temporal Domains

by

Emma Wisniewski-Barker

Submitted in fulfilment of the requirements for the
Degree of Doctor of Philosophy

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

August 2015

Abstract

Slow light is the study of the dramatic change in the velocity of light as it travels through certain media. This thesis focuses on slowing caused by transmitting light through a ruby crystal. When ruby experiences spatial or temporal modulation from a laser beam, the velocity of the light is greatly changed from its speed in a vacuum. The underlying mechanism for slow light in ruby is not fully understood and is, therefore, the subject of much debate.

In this thesis, I examine many experimental parameters and their effects on slow light in ruby. First, I investigate the delay of images with both bright and dark regions through a spatial modulation of the ruby. I then turn to a temporal modulation of the ruby to answer the question of whether light can be delayed beyond the input pulse, the answer to which has the potential to differentiate between two proposed models for the mechanism that causes slow light in ruby.

I return to the spatial domain to study the effect of spatial intensity distribution on the slow-light effect in ruby. I show that beams carrying orbital angular momentum rotate by an amount determined by the spatial feature of the beam profile. I present experimental evidence supporting a complicated mechanism of slow light in ruby, which informs the ongoing debate on the cause of slow light in ruby and provides direction for applications dependent upon the preservation of complex patterns in slow-light media.

Contents

Abstract **i**

List of Figures **v**

List of Tables **xii**

Acknowledgements **xiii**

Declaration of Authorship **xiv**

Publications **xv**

Foreward **xviii**

1 Slow Light Introduction **1**
- 1.1 Overview . . . 1
  - 1.1.1 Refractive Index . . . 1
  - 1.1.2 Phase Velocity . . . 2
  - 1.1.3 Group Velocity . . . 3
  - 1.1.4 Transit Time and Types of Slow Light . . . 3
- 1.2 Early Research and Observation . . . 4
  - 1.2.1 Slow Light . . . 4
  - 1.2.2 Stopped Light . . . 5
  - 1.2.3 Fast Light . . . 6
- 1.3 Structural Slow Light . . . 9
  - 1.3.1 Optical Fibers . . . 9
  - 1.3.2 Photonic Crystals . . . 9
  - 1.3.3 Other Types of Structural Slow Light . . . 10
- 1.4 Material Slow Light . . . 11
  - 1.4.1 The Kramers-Kronig Relation . . . 11
  - 1.4.2 Optical Interactions with Atoms . . . 12
  - 1.4.3 Electromagnetically Induced Transparency . . . 13
  - 1.4.4 Stimulated Brillouin and Raman Scattering . . . 14
  - 1.4.5 Coherent Population Oscillations . . . 15
- 1.5 Applications . . . 17
  - 1.5.1 Optical Delay Lines and Optical Processing . . . 17

1.5.2 Interferometry . . . 18
1.5.3 Enhanced Photon Drag . . . 19
1.6 Conclusions . . . 21

2 Slow Darkness 22
2.1 Introduction . . . 22
2.2 Initial Image Rotation . . . 25
2.3 Experimental Motivation and Methods . . . 26
2.4 Results . . . 30
2.5 Discussion . . . 33
2.6 Alternate Methods . . . 33
2.7 Rotating Complicated Patterns . . . 37
2.8 Fluorescence . . . 39
2.9 Absorption Study . . . 41
2.10 Conclusions . . . 43

3 Delaying Energy Beyond the Input Pulse 45
3.1 Introduction . . . 45
3.2 Slow Light in the Time Domain . . . 46
3.3 Methods and Materials . . . 48
3.4 Results . . . 49
3.5 Pulse Distortion and Harmonic Components . . . 52
3.6 Conclusions . . . 53

4 Orbital Angular Momentum Introduction 54
4.1 Historical Introduction . . . 54
4.2 Creating Beams with Orbital Angular Momentum . . . 58
4.3 Micro-Manipulation Through the Use of OAM . . . 63
4.4 Beam Transformations . . . 66
4.5 Measuring Beams with Orbital Angular Momentum . . . 69
4.6 Orbital Angular Momentum in Classical Imaging . . . 71
4.7 Orbital Angular Momentum in Nonlinear and Quantum Optics . . . 72
4.8 Conclusions . . . 74

5 Rotating Orbital Angular Momentum 76
5.1 Introduction . . . 76
5.2 Magnetic and Mechanical Faraday Effects . . . 77
5.3 Creating Orbital Angular Momentum States . . . 77
5.4 Photon Drag and Slow Light . . . 79
5.5 Methods and Materials . . . 79
5.6 Results and Discussion . . . 82
5.7 Summary and Conclusions . . . 88

6 Conclusions and Future Work 89

A Slow Light Interferometer **93**

B Additional Rotated Patterns **99**

Bibliography **102**

List of Figures

1.1 Light cone showing the region in XY space, as time progresses on the vertical axis, that can be affected by an event at the red point position in spacetime. 8

1.2 Fibre Bragg grating, waveguide Bragg grating, and planar photonic crystal waveguides. Dark and light regions indicate positions of high and low refractive indices in the first two cases. In the planar photonic crystal, regions of low refractive indices arise from the air holes created in the crystal. 10

1.3 (*left*) A narrow absorption band, leads to a large dispersion in n_ϕ, which in turn gives rise to a very large n_g. (*right*) A wider absorption band causes a much smaller dispersion in n_ϕ, and therefore a small n_g. 12

1.4 Models of the energy levels in (a) two- and (b) three-level systems. The coupling beam is resonant between levels $|2\rangle$ and $|3\rangle$. The probe beam is resonant between levels $|1\rangle$ and $|3\rangle$. 13

1.5 Models of the (*top*) energy levels and (*bottom*) absorption profiles with (a) the coupling beam off and (b) the coupling beam on. . . . 14

2.1 Distortion of an incident pulse passing through a slow-light medium. Part a shows a single pulse passing through the slow-light medium. (a) Optical bleaching and (c) large group index effects are indistinguishable for small delay of a single pulse, typical of most experimental conditions. To distinguish the two cases, one can pass a pulse with an imprinted zero through the slow-light medium. Optical bleaching and pulse reshaping would never cause the position of the dark region to shift, although it may become brighter due to fluorescence from the saturable absorber (b), an effect that is in agreement with [1]. However, large group index would cause the position of the dark region to be maintained but shifted (d). By imprinting an intensity zero in the pulse, I can determine which mechanism is causing the apparent slowing of the light through the ruby, even with a small delay. 23

2.2 Image of the elliptical beam while the ruby is spinning clockwise at 1 Hz, 5 Hz, and 10 Hz. 25

2.3 Largest recorded image drag between (a) anticlockwise and (b) clockwise rotations of the ruby was measured to be ≈ 55°. 26

2.4 Experimental procedure. A beam of 3 W, 532 nm light is bisected by a coverslip with a thickness chosen so as to create a π-radian phase shift between the two halves of the beam. The entire beam then passes through a quarter waveplate and two cylindrical telescoping lenses before being focused on a ruby window. The ruby is spun to $\approx \pm 20$ rotations per second (rps), and the light leaving the ruby is imaged onto a camera. 27

2.5 (a) Modelled intensity profile of an elliptical beam. (b) Unwrapped annular intensity distributions of a measured around the origin as indicated by the black circle on a, where r is the radial distance from the origin in a. (c) Modelled intensity profile of an elliptical beam with a black region through the centre. (d) Unwrapped annular intensity distributions of c measured around the origin as indicated by the black circle on c, where r is the radial distance from the origin in c. 28

2.6 Images of the beam profile when a microscope coverslip is inserted halfway into the beam. The coverslip can be tilted so (a, b) there is no phase shift or (c, d) there is a π-radian phase shift between the two halves of the beam. Figures a and c are taken before the ruby and use a shear plate to show the interference pattern for the coverslip set to two different thicknesses. These images show that tilting the coverslip does create a π-radian phase shift between the two halves. Figures b and d show the corresponding intensity profiles after the ruby. Note the creation of an intensity null in d. . 30

2.7 Images of the beam (a) without the phase discontinuity and (b) with the phase discontinuity when the ruby is spinning from -19 rps to +19 rps. The rotation of the elliptical beam increases with speed to a maximum rotation angle of $\approx 5°$ between clockwise and anticlockwise rotations for both the case without and with the discontinuity. 31

2.8 Two images overlaid to show the rotation of the black region. One image is artificially coloured blue and the other red. Overlapping regions of the two images are shown in white. (b) The two images overlaid while the ruby was stopped in both cases. (a) The two images overlaid while the ruby is spinning anticlockwise for the red image and clockwise for the blue image. (c) The two images overlaid while the ruby is spinning in the opposite direction from a. 32

2.9 Intensity profiles. A strip of the beam profile at the radius with the greatest photon drag is unwrapped in *a* when the ruby is spinning anticlockwise (*blue*), spinning clockwise (*pink*), and removed (*olive*). The intensity of each of those strips is plotted in *b*. The blue and pink curves show the position of the region of darkness while spinning anticlockwise and clockwise, respectively, to be outside the position of the dark region when the ruby is removed (and therefore not spinning). One observation in this figure is the movement of the dark line causing there to be light in the clockwise and anticlockwise cases where there had been darkness in the stopped case, which could be explained by fluorescence from the saturable absorber. However, the region of darkness is preserved through the rotation to a position that had previously been bright. 34

2.10 Intensity as a function of azimuthal position. The stopped trace (*olive*) has an intensity dip at approximately 270°, whereas in the spinning case (*blue*), the intensity dip has moved to approximately 265° and the intensity at 270° has increased. 35

2.11 Images of the beam with (*top row*) the wire and (*bottom row*) the coverslip when the ruby is spinning at -19 rps, stopped and spinning at +19 rps. The dotted white line shows the position of the black line when the ruby is stopped as a reference. 36

2.12 Image of the elliptical beam while the coverslip creates a π-radian phase difference between the two halves of the beam and the ruby spins anticlockwise. Fringes can be seen on the left and right sides of the image, caused by interference from multiple reflections of the light in the coverslip. 37

2.13 Rotation of pattern made by the corner of a glass microscope coverslip when the ruby is (*b*) still and rotating (*a*) anticlockwise and (*c*) clockwise. 38

2.14 Rotation of pattern made by a metal mask when the ruby is rotating (*a*) anticlockwise and (*b*) clockwise. Red vertical lines are added for reference. 38

2.15 Images in the shape of UG are shown for the (*a*) green coherent light, (*b*-*d*) red fluorescent light, and (*e*-*g*) simulated fluorescence. The z-position of the focus lens was purposefully misaligned, such that the coherent light was not slowed. Therefore, the transmitted green coherent light (shown in *a*) was unchanged by the rotation of the ruby. Images *b* and *e* show the red fluorescence while the ruby was not rotating. Images *c* and *d* show the fluorescence at medium and high speeds, respectively, which gave increasing amounts of blurring. For processing, the image in *e* was changed to an intensity image, rather than a colour one. The intensity image in *e* was artificially rotated by a LabVIEW program, giving more image rotation and blurring with increasing numbers of iterations in *f* and *g*. 40

2.16 Percentage of transmitted intensity plotted against incident intensity for two orientations of the ruby. Yellow squares show the percentage of transmitted intensity when the ruby is turned to the orientation with the weakest interaction between the laser light and the ruby crystal. Blue circles show the transmission percentage for the orientation with the strongest interaction between the laser light and the ruby crystal. 43

3.1 Signals from reference and ruby arms are shown on an oscilloscope. The frequency of the sinusoidal modulation of the signals decreases from *a* to *c*, and the ruby signal becomes more delayed as the modulation frequency decreases. In *c*, the two signals are vertically displaced from each other by a small amount for clearer viewing of the relative delay between them. 48

3.2 Intense 532 nm laser light was focussed onto an optical chopper (OC) and then recollimated by 160 mm focal-length lenses. The light was then split by a beamsplitter (BS), sending some of the light to be focussed by a 60 mm focal-length lens onto photodiode (PD1). The remaining light was focused onto the front surface of a 90-mm-long ruby rod with a 50 mm focal-length spherical lens. The light transmitted through the ruby was measured by a second photodiode (PD2) after a spatial filter comprised of a 60 mm focal-length lens and a 400 μm pinhole (PH). Dichroic bandpass filters (DF) were placed before both photodiodes. 49

3.3 Time evolution of the reference and transmitted signals shown on semi-log plots. (*a*) When the ruby was removed, the transmitted (beige, dotted) and reference (green, solid) signals have the same intensity profiles. (*b*) When the ruby was added, the overall intensity of the transmitted signal (beige, dotted) decreases, although the intensity at the trailing edge of the pulse increases above that of the reference signal (green, solid). The region of interest in *b* is shown in detail in *c*. For times greater than that marked by the dashed vertical line in *c*, the intensity of the transmitted signal is greater than that of the reference pulse. Dark signals taken with the laser off are shown for the reference (pink, large dashes) and transmitted (blue, dash dots) arms in *a*, *b*, and *c*. 50

3.4 In an ideal slow-light medium, a square-top pulse would simply be delayed (shifted) in time, as shown in the insert in *d*. In contrast, the results I observe are shown for a square wave with a modulation of (*a*) 7 Hz, (*b*) 16 Hz, and (*c*) 28 Hz. Unlike Fig. 3.3, all pulse shapes in this figure have been normalised to have the same peak intensity, to illustrate the difference in pulse shapes between the reference and transmitted signals. Part *d* shows all three traces overlaid temporally, where the falling edge of each reference pulse is aligned. 51

3.5 Trace of square wave pulse with thin “black line” embedded in the brightest part of the signal. Energy is delayed into the black line in the same way it is delayed into the trailing edge of the pulse, as shown in the regions circled in red. 51
3.6 Time delay plotted against the frequency of the Fourier component for square-wave signals of different modulation frequencies. The delays of the Fourier components from the 7 Hz-modulated square wave are marked with large filled diamonds. The trace of a corresponding 7 Hz-modulated pulse is shown in Fig. 3.4*a*. Delays of Fourier components from square waves modulated at different frequencies all follow the same curve. Because different frequency components experience different time delays, the square wave is distorted upon transmission through ruby, resulting in the transmitted traces in Fig. 3.4. 53

4.1 Helical phase fronts interfered with plane waves result in spiral fringes, shown here for $l = 0$ through $l = +3$. 57
4.2 (*left*) $\pi/2$ converters switch between HG and LG modes, whereas (*right*) π converters switch modes to their mirror images. 59
4.3 The effects of (*top*) wave plates and (*bottom*) mode converters. Quarter-wave plates convert from linear to circular polarisation in an analogous way to how $\pi/2$ mode converters convert HG modes to LG mods. Half-wave plates reverse the direction of the polarisation analogously to how π mode converters (or dove prisms) switch the direction of the OAM. 59
4.4 A plane wave gains a helical phase front when shone through a spiral phase plate. 60
4.5 Plane waves gain a helical phase front when shone through a forked diffraction grating, resulting in 1st and -1st order diffracted beams. The 1st order diffracted beam is shown here with $l = +1$. 61
4.6 A spiral phase plate illuminated by incoherent light produces an incoherent vortex in the far field. 63
4.7 Different rotating conditions for spin and orbital angular momentum. (*a*) SAM makes the particle spin about its own axis, whereas (*b*) OAM orbits the particle about the beam axis. (*c*) For a beam much smaller than the particle, SAM and OAM will both make the particle spin in the same way. (*d*) For a large beam, OAM will make the particle orbit the beam. 64
4.8 Equivalent dragging of (*top*) polarisation and (*bottom*) transmitted image through a rotating medium. 68
4.9 (*left*) Poincaré sphere for polarisation and (*right*) the analogous Bloch sphere for OAM. 68

4.10 A vortex beam with a helical phase structure is incident upon a forked diffraction grating. When the OAM value of the vortex beam and the diffraction grating are l and $-l$ respectively, the resulting beam will be Gaussian and will successfully couple into a single mode fibre (SMF). An example of this process is shown here for a vortex beam of $l = +1$. 70

5.1 Petal patterns created by a superposition of $\pm\ell$ beams. The second and fourth lines have a phase shift, i, between the $+\ell$ and $-\ell$ beams, causing a 45 degree rotation of the petal pattern. 78

5.2 532 nm light passes through two spherical lenses to be expanded before arriving at the spatial light modulator (SLM). The beam is then focused onto the front face of a ruby window, which is spun about its axis by a motor. The light is imaged from the back of the ruby onto a screen and then captured by a camera. Petal patterns are made from superpositions of LG beams with different l values while the ruby window spins at ±19 Hz. Patterns shown in the inset are (from left to right) 2 petals from $l = \pm 1$; 3 petals from $l = (+1, -2)$; 4 petals from $l = \pm 2$; and 5 petals from $l = (+2, -3)$. 80

5.3 Image of pattern with 4 petals, shown rotating (a) anticlockwise and (b) clockwise. A red X is overlaid on the images to show the orientation of the pattern. (c) Images in a and b are added to show the angle between the orientation of the images while rotating in different directions. 82

5.4 Normalised peak intensity of petal patterns with $N = 2$ through 8 petals with constant input laser power. 83

5.5 Rotation angle of petal patterns with $N = 2$, 4, 6, and 8 petals as a function of input laser power while rotated at a constant speed. . 84

5.6 Rotation angle of petal patterns with $N = 2$ through 5 petals as a function of rotational speed while held at constant peak intensity. Patterns with different N saturate at different rotational speeds. Error bars represent the standard deviations of independent data runs. 85

5.7 Images of N =2 through 5 for unsaturated (low rotational frequency, first column) and saturated (high rotational frequency, second column) modulation frequencies with the ruby spinning anticlockwise (ACW) and clockwise (CW). 86

5.8 Fraction of a petal rotated by $N = 2$ through 5 petals as a function of rotational speed while held at constant peak intensity. Error bars represent the standard deviations of independent data runs. 87

A.1 A 3 W beam of 532 nm light passes in both directions around a Sagnac interferometer. The light is focused on a ruby window, which moves longitudinally along one arm of the interferometer. Light passing clockwise through the interferometer spends less time in the ruby than light passing counterclockwise through the interferometer. The interference pattern of the recombined light is imaged onto a camera. Red (633 nm) laser light follows the same path as the 532 nm light and is viewed on the same camera. 94

A.2 Fringe positions (*a*) while the ruby was stationary, (*b*) while the ruby was accelerating, and (*c*) while the ruby was moving at its highest velocity. Green fringes are seen to shift by approximately 0.7 fringes, while red fringes do not. Red and green reference lines are added in all three images at the same position and orientation to show the relative movement of the fringes. 98

B.1 Rotation of pattern of an E made by a metal mask. Images shown were taken when the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. 99

B.2 Rotation of pattern made by a metal mask. Images shown were taken when the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. 99

B.3 Rotation of a square pattern made by a metal mask. Images shown were taken when the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. 100

B.4 Rotation of a pattern of horizontal lines made by a metal mask. Images shown were taken when the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. 100

B.5 Rotation of a pattern of vertical lines made by a metal mask. Images shown were taken when the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. 101

List of Tables

2.1 Ruby Absorption, Orientation A 42
2.2 Ruby Absorption, Orientation B 42

Acknowledgements

I would first like to thank my thesis advisor, Prof. Miles J. Padgett, for all of his advice, guidance, patience, and support over the entire course of my PhD. He has helped me think critically about my work and take a step back to examine the big picture. His energy and enthusiasm for optics have inspired me over the past three and a half years, constantly reminding me of the best parts of research.

Second, I want to express my sincere gratitude to Dr. Graham Gibson. He has been such an important part of my experience that I do not know how other labs function without him. I have gained so much hands-on knowledge and experience from working with him, which has been essential to the completion of my research and to my development as a physicist.

I would also like to thank Prof. Robert W. Boyd for his invaluable help. His insights into and perspective on my research greatly enhanced my comprehension of complicated optical phenomena. Additionally, my work was strongly impacted by the opportunity to visit Prof. Boyd's research groups in Rochester and Ottawa, for which I am extremely grateful.

For all their help with interpreting experimental results, I would like to thank Dr. Sonja Franke-Arnold and Dr. Paul Narum. Additionally, I would like to express my thanks to Prof. Zhimin Shi for his deep understanding of my work and his experimental and interpretation insights.

Additionally, I am so grateful for the wonderful support from our group secretary, Lucy Murray, who always went above and beyond to assist with navigating the paperwork, planning, and organization that accompanies PhD research.

From all the research assistance and help with theoretical problems to all the cups of tea, Scottish word translations, cake, and support, I am unendingly grateful to all of my fellow PhD students and postdocs in the Optics group.

I would also to thank my parents, who have always believed in me and fully supported me. A special thank you to Dr. Nicola Mancell for always knowing what I needed to hear and to Rebecca Harper for always being there for me.

And finally, to all of my family and friends who have supported me through the past few years, both on this continent and across the ocean, I could not have done it without you.

Declaration of Authorship

I, EMMA WISNIEWSKI-BARKER, declare that this thesis titled, ‘Slow Light in Ruby: a Study in Spatial and Temporal Domains’ and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: ______________________________________

Date: ____

Publications

1. E. Wisniewski-Barker, G. Gibson, S. Franke-Arnold, Z. Shi, R. W. Boyd, and M. J. Padgett, “Evidence of slow-light effects from rotary drag of structured beams,” *New Journal of Physics* **15** (2013).

2. E. Wisniewski-Barker, G. Gibson, S. Franke-Arnold, Z. Shi, R. W. Boyd, and M. J. Padgett, “Reply to comment on ‘Evidence of slow-light effects from rotary drag of structured beams,’ ” *New Journal of Physics* **16** (2014).

3. E. Wisniewski-Barker, G. Gibson, S. Franke-Arnold, R. W. Boyd, and M. J. Padgett, “Mechanical Faraday effect for orbital angular momentum-carrying beams,” *Optics Express* **22** (2014).

4. E. Wisniewski-Barker, G. Gibson, S. Franke-Arnold, Z. Shi, P. Narum, R. W. Boyd, and M. J. Padgett, “Experimental investigation of the transient dynamics of slow light in ruby,” *New Journal of Physics* **16** (2014).

5. E. Wisniewski-Barker and M. J. Padgett. (2015). “Orbital Angular Momentum” in David L. Andrews. (Ed.), *Photonics: Scientific Foundations, Technology and Applications, Volume 1* (321-340). John Wiley & Sons, Inc.

“What is wanted is not the will to believe, but the will to find out, which is the exact opposite.”

Bertrand Russell

Dedicated to the family I was born with
as well as the family I developed along the way

Foreward

The work contained in this thesis comes from research conducted during my doctoral studies at the University of Glasgow. I conducted the experimental work explained within with assistance and advice from Dr. Graham Gibson and my advisor Prof. Miles J. Padgett. Experimental planning and interpretation of results were a collaborative effort between several researchers in the Optics Group at the University of Glasgow as well as research groups at collaborating institutions. Please note that I am the first author on all of the publications listed below.

Chapter 2 is based on an expanded version of a paper titled "Evidence of slow-light effects from rotary drag of structured beams" published in *New Journal of Physics* **15** 8 (2013), written with G. Gibson, S. Franke-Arnold, Z. Shi, R. W. Boyd, and M. J. Padgett. The experimental work was conducted by myself and G. Gibson, with assistance from Z. Shi and M. J. Padgett. Chapter 2 also includes work published in "Reply to comment on 'Evidence of slow-light effects from rotary drag of structured beams' " published in *New Journal of Physics* **16** 3 (2014), written with G. Gibson, S. Franke-Arnold, Z. Shi, R. W. Boyd, and M. J. Padgett and work published in "Observation of slowed light through a ruby window" *Proc. SPIE 8636* Advances in Slow and Fast Light VI (2013).

Chapter 3 expands on the paper titled "Experimental investigation of the transient dynamics of slow light in ruby" published in *New Journal of Physics* **16** 12 (2014), written with G. Gibson, S. Franke-Arnold, Z. Shi, P. Narum, R. W. Boyd, and M. J. Padgett. I conducted the experimental work in Glasgow with assistance from G. Gibson, and all authors contributed to the exchange of ideas and interpretation of the results.

Chapter 4 contains an edited version of a book chapter that I wrote with assistance from M. J. Padgett. A version of this chapter was published in *Photonics: Scientific Foundations, Technology and Applications, Volume 1*, edited by David L. Andrews and published by John Wiley & Sons, Inc (2015).

Chapter 5 contains an extended version of a research paper titled "Mechanical Faraday effect for orbital angular momentum-carrying beams" that was published in *Optics Express* **22** 10 (2014), written with G. Gibson, S. Franke-Arnold, R. W. Boyd, and M. J. Padgett. I conducted the experimental research in Glasgow with assistance from G. Gibson and with experimental planning and data interpretation by all authors.

Chapter 1

Slow Light Introduction

Slow light is the study of the dramatic change in the velocity of light as it travels through certain media. In this chapter, I present the important concepts, methods for creating, and applications of slow light.

1.1 Overview

1.1.1 Refractive Index

When light passes from a vacuum into a medium, it slows down. Different materials slow down the velocity of light by different amounts. The refractive index of a material is the measure of how much slower light travels through that material than it would through a vacuum. This commonly-discussed refractive index is usually denoted n. This n is often referred to as the phase refractive index, and for clarity, it will henceforth be called n_ϕ. The need for clarity arises with the realisation that there is also a group refractive index, n_g, often simply called the group index. The group index determines how quickly the peak or envelope of a light wave travels through a medium, whereas the phase refractive index is the measure of how fast wavefronts of constant phase, also called phasefronts, travel.

The group delay, τ_g measures how long a pulse is delayed by a medium. It is relatively simple to measure by comparing the arrival time of reference and delayed pulses, and it is used to calculate n_g from the simple relation $\tau_g = Ln_g/c$, where L is the length of the material and c is the speed of light in a vacuum. The group delay τ_g is a straightforward time delay, but often a more useful measurement to report is the fractional delay. Dividing τ_g by the width of the pulse, usually measured as the full width at half of the maximum height (FWHM), results in a delay proportional to the pulse width.

1.1.2 Phase Velocity

The phase velocity, v_ϕ, is understood to be the velocity of points of constant phase in monochromatic light. For a plane wave

$$E(z,t) = Ae^{i\phi} + C, \tag{1.1}$$

A and C are constants, and $\phi = kz - \omega t$ is the phase of the light, where $k = n_\phi \omega / c$, ω is the frequency, z is the propagation distance, and t is time. Because n_ϕ is wavelength-dependent, k is rewritten as

$$k(\omega) = \frac{n_\phi(\omega)\omega}{c}, \tag{1.2}$$

which accounts for dispersion. For the pulse to travel without distortion, the change in the z-component must correspond to the change in time, $k(\omega)\Delta z = \omega \Delta t$. Dividing this through by Δz gives

$$k(\omega) = \frac{\omega \Delta t}{\Delta z} = \frac{\omega}{v_\phi}. \tag{1.3}$$

Therefore, solving for the phase velocity in Eq. (1.3) and combining it with Eq. (1.2) to eliminate $k(\omega)$ results in the definition of the phase velocity being

$$v_\phi = \frac{\omega}{k(\omega)} = \frac{c}{n_\phi(\omega)}. \tag{1.4}$$

1.1.3 Group Velocity

A pulse can be composed of many individual monochromatic waves with varying frequencies. Each monochromatic wave will have a phase velocity, as described above. The sum of the phases must be constant, $\phi = kz - \omega t = \frac{n_\phi(\omega)\omega}{c} z - \omega t$, for all values of z for the pulse to travel without distortion. In other words, $\delta\phi/\delta\omega = 0$.

The group velocity, v_g, describes the velocity with which the envelope of a light wave travels through a medium. It can be derived from the group refractive index through

$$v_g = \frac{c}{n_g}, \tag{1.5}$$

where by definition

$$n_g = n_\phi(\omega) + \omega \frac{\delta n_\phi(\omega)}{\delta\omega}, \tag{1.6}$$

which arises from setting the derivative of ϕ with respect to ω equal to zero. The group velocity can be used to calculate z, the distance traveled by the pulse, by simply multiplying it with time elapsed, $z = v_g t$. A thorough derivation of these equations is found in [2].

1.1.4 Transit Time and Types of Slow Light

The study of slow light includes the study of any case where the effective velocity of light greatly varies from the speed of light in a vacuum. The effective velocity can be faster than c, much slower than c, or even negative, making the time it takes to pass through the medium vary greatly from the time it would take to travel that distance in a vacuum.

The amount of time it takes for light to pass through a medium is the transit time, $T = L/v_g$, where L is the length of the medium through which the light is passing. When the n_g of a material is very large, it is called a slow-light material, meaning the transit time of light through that material is very large. Because $n_g >> n_\phi$, $v_g << v_\phi$ and $v_g << c$. The study of slow light also includes the cases where $v_g > c$, $v_g = 0$, and even $v_g < 0$. When $v_g > c$, called fast light, the transit time through the medium is very small, because n_g of the medium is between 0 and 1. Although v_g is greater than the speed of light, there is no superluminal information transfer, which means that fast light follows causality and special relativity, as will be discussed in Section 1.2.3.

Backwards light occurs when both n_g and v_g are less than 0, resulting in a negative transit time. In this situation, the peak of the pulse appears to leave the material before entering it, with the pulse peak traveling backwards through the medium. However, the energy and information carried by the pulse always move forwards. This mechanism occurs through pulse reshaping that amplifies the leading edge and attenuates the original peak of the pulse. In stopped light, v_g is zero, as its name suggests. This is often achieved by manipulating the medium through which the light pulse is passing to temporarily store the pulse before retrieving it. All of these mechanisms for dramatically changing the speed of light will be discussed in more detail below.

1.2 Early Research and Observation

1.2.1 Slow Light

The experimental study of slow light really began to take off in the late 1960's and early 1970's [3–6]. A subset of these studies focussed on weak Gaussian pulses passing through an amplifying or absorbing medium. Garrett and McCumber [7] proposed a theory in 1970 that a weak Gaussian pulse experiences a negligible distortion to the pulse width when passing through many exponential absorption or gain lengths. Garrett and McCumber also proposed that the pulse peak travels at v_g. In 1982, Chu and Wong experimentally tested Garrett and McCumber's theory with picosecond laser pulses in a GaP:N crystal, where the frequency was tuned through the absorption resonance. Chu and Wong measured positive as well as negative values of v_g without significant pulse distortion [8]. Group velocity is a robust concept, with the peak of the pulse travelling at v_g, even when $v_g > c$

or $v_g < 0$. Katz and Alfano [9] noted that the pulse shapes are distorted through the measurement process, although this does not negate the importance and persistence of the concept of group velocity. A follow-up numerical simulation study in 1985 [10] showed that the pulses were experiencing ringing in addition to the pulse compression reported by Katz and Alfano. The same numerical simulation study also examined fast light using absorption resonance in millimetre waves in a molecular absorber.

Atomic vapours were very commonly used in some of the early slow light experiments [11–14]. Xiao *et al.* used hot rubidium vapour in a Doppler-free electromagnetically induced transparency (EIT) experiment [11]. In this experiment, $v_g = c/13.2$ m/s was indirectly determined by measuring the phase shift of the wave passing through the rubidium vapour. The same year, Kasapi *et al.* used nanosecond pulses passing through a lead vapour cell with a coupling field under EIT conditions that showed a further reduction of the group velocity to $v_g = c/165$ m/s [12]. A few years later, Kash *et al.* used hot rubidium atoms that were driven coherently by a narrow EIT resonance [14]. In order to achieve the narrow EIT resonance, a group of researchers slowed light to an incredible 17 m/s by using EIT in a laser-cooled sodium atomic vapour at 450 nK [13]. This experiment was the first to slow light so dramatically, and it ignited the scientific community's interest in and curiosity about slow light.

1.2.2 Stopped Light

Hau and colleagues expanded the work that allowed them to slow light to 17 m/s in order to completely stop the light for a short time [15]. They achieved this through control of the coupling field in their EIT experiment, by recognising that the probe beam can only propagate when the coupling beam is also on (see section 1.5 for more details on EIT). Their method involved keeping the coupling beam on while the probe pulse enters the medium and quickly switching off the coupling beam before the probe pulse can leave the medium. When the coupling beam is turned back on, the probe pulse exits the medium. This is made possible by the fact that the phase and amplitude information from the probe beam is stored coherently in the medium, so that when the coupling beam turns on again, the coupling field scatters coherently from the medium to recreate the probe pulse. It has been claimed just to be a semantic difference whether the light is being stopped or being turned into material degree of freedom, and then that degree of freedom is turned back into an optical field when the coupling beam is turned on again [2]. This stopped light can be thought of as an optical memory, albeit one where the memories decay with time, due to the eventual decoherence experienced by EIT media [16].

1.2.3 Fast Light

Fast light can be separated into two categories, that with $v_g > c$ and that with $v_g < 0$. The first is where the light appears to travel through the material faster than it would through a vacuum. The second is also known as backwards light, and when the group velocity is negative, the pulse appears to leave the material before entering it.

In 1914, Brillouin wrote a paper discussing the distinction between group velocity and front velocity, as well as discussing the implications of that distinction on special relativity [17]. However, it wasn't until the 1960's that much work was done researching fast light. Icsevgi and Lamb, apparently unaware of the work done by Brillouin, theorised that intense laser pulses in a laser amplifier would create fast light [18], a theory that is consistent with the work of Brillouin. Their work expanded Brillouin's to also include nonlinear optical media. They described Brillouin's front velocity as having compact support, which is to say the amplitude of the pulse is nonzero over a finite range of times. This velocity cannot be greater than the speed of light. On the other hand, they described Brillouin's group velocity as having infinite support, that is the pulse is nonzero for all times and can propagate with an effective velocity (specifically the group velocity, v_g) greater than the speed of light. In agreement with Brillouin, they stated that there is no violation of causality where fast light is concerned, as the pulse exists at all times.

The pulse has leading and trailing edges that extend, theoretically, infinitely, and all parts of the pulse contain the information needed to reconstruct the pulse shape, which allows for the pulse to be reconstructed beyond the medium before the peak of the pulse has entered the medium. Basov *et al.* [19] used optical saturation of an amplifier, which is made of a group of inverted atoms, to experimentally create fast light in 1966, where the high intensity of the pulse caused a nonlinear effect.

In the subsequent decades, many advances have been made in the area of fast light research. Chu and Wong [8] observed a large anomalous dispersion in 1982 that was accompanied, however, by a large absorption. In an attempt to reduce the absorption but retain the fast light, Bolda *et al.* [20] used a nearby gain line to create a region with a negative group velocity. Instead of tuning near one gain line, Steinberg and Chiao [21] suggested tuning between a pair of Raman gain lines in order to create a transparency, similar to EIT, although the sign of the dispersion in this case is opposite of that in EIT, which is why the group velocity is negative. Wang *et al.* experimentally realised Steinberg and Chiao's theory, resulting in a $v_g = -c/310$, where the pulse was advanced by 62 ns with only negligible amounts of pulse reshaping [22]. An alternate method of reducing absorption while achieving superluminal propagation was discussed in 1996 by Chiao *et al.* who proposed using the collective emissions from a system of inverted two-level atoms [23], in a method similar to superfluorescence, a pulse of the stored energy from atoms that initially undergo spontaneous emission before becoming coupled in phase by the radiation field [24].

In 1981, Sherman and Oughstun wrote an algorithm for the propagation of a short pulse through a dispersive system with loss [25]. Diener showed in 1996 [26] that the part of the pulse traveling superluminally could be mathematically calculated by extending the extreme leading edge of the pulse within the future "light cone." When an observer experiences an event happening at a position in space and time, the future light cone describes every position in space that could be affected by that event as time progresses forwards (or backwards for the past light cone). When considering a simplified two dimensional space, the region potentially affected by the event at any position in time is a circle with a radius that grows at a rate

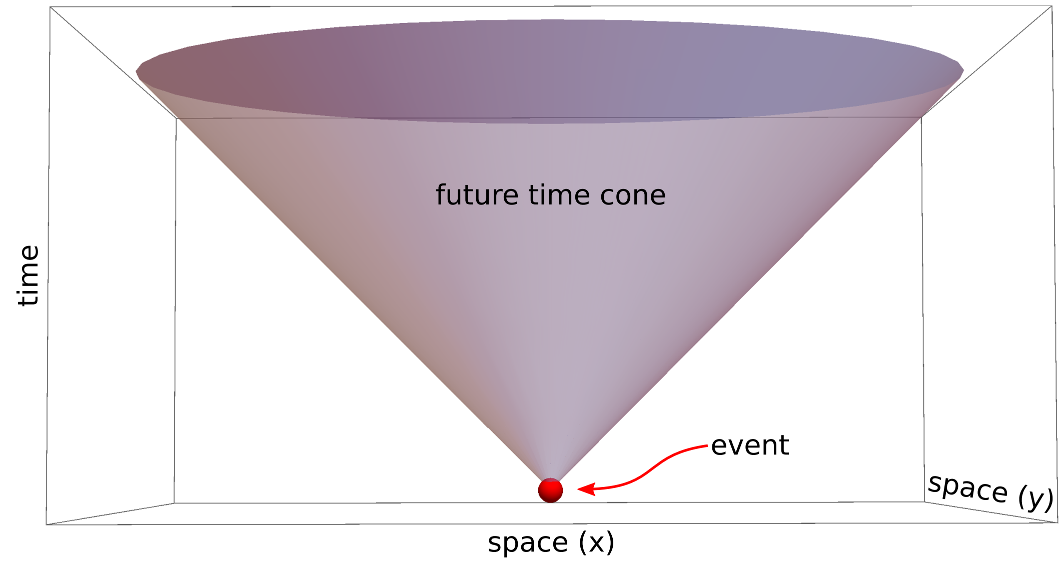

FIGURE 1.1: Light cone showing the region in XY space, as time progresses on the vertical axis, that can be affected by an event at the red point position in spacetime.

proportional to the speed of light in a vacuum, as shown in Fig. 1.1. Anything outside the light cone cannot be affected by the event, as the information from the event cannot travel to that position quickly enough. However, because the signal is analytically continued without deviating from the signal within the light cone, the superluminal light does not contain new information, and therefore causality is not broken. The group velocity is once again shown to not carry information, which has an upper limit of the speed of light in a vacuum.

Akulshin *et al.* used a two-level atomic system for electromagnetically induced absorption, EIA, [27] (as opposed to transparency, in the more common EIT). Through EIA, they created very large anomalous dispersion, and from the measured dispersion, they inferred a v_g of $-c/23000$, although this negative group velocity was accompanied by a large absorption.

In 2000, Segev *et al.* posited that a superluminal signal would have a very low signal-to-noise ratio [28]. The following year, Kuzmich *et al.* [29] defined the “signal velocity” to be a measure of the signal-to-noise velocity, allowing for useful descriptions of pulse propagation involving only a few photons.

1.3 Structural Slow Light

Slow-light systems can generally be broken down into two main categories: material slow-light systems, presented in Section 1.4 and structural slow-light systems. Structural slow light systems depend upon the spatially-varying (often periodic) refractive index changes in the system. A review of periodic structures for creating slow light can be found in [30], and a couple of the most common ones are presented below.

1.3.1 Optical Fibers

Optical fibres can contain a grating with a refractive index change in one direction (one dimensional structural slow light systems). The refractive index changes usually occur in a sinusoidal manner. Erbium-doped optical fibres have been used for both slow and superluminal light propagation [31, 32]. Optical fibres are especially useful for understanding light propagation through a material with a negative group velocity, as the intensity of the pulse can be measured at multiple positions through the fibre. By measuring throughout the length of the fibre, the researchers could prove that the peak of the pulse travels backwards through the fibre, although the energy always moves forward [31].

1.3.2 Photonic Crystals

Photonic crystals depend on the repetitive (periodic) changes in the refractive index of a dielectric or other highly nonlinear material. When the refractive index changes along two directions, it is called a planar photonic crystal waveguide. Photonic crystals can, however, be made in one, two, or three dimensions. Examples of photonic crystals in one and two dimensions are shown in Fig. 1.2. A common example of a periodic structure for slow light is a dielectric material with air holes in a rectangular lattice pattern. As light passes from the dielectric to the air and back, it experiences the periodic modulation of the refractive index required for a

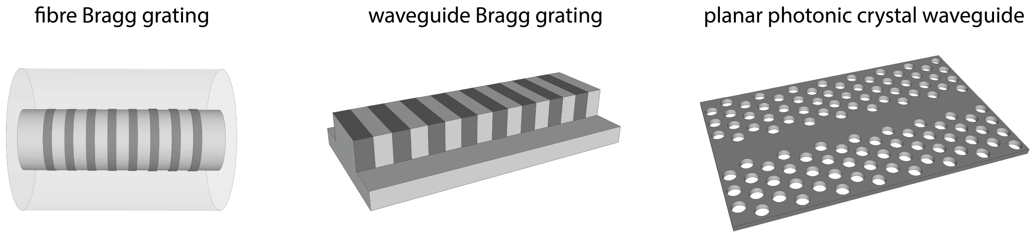

Figure 1.2: Fibre Bragg grating, waveguide Bragg grating, and planar photonic crystal waveguides. Dark and light regions indicate positions of high and low refractive indices in the first two cases. In the planar photonic crystal, regions of low refractive indices arise from the air holes created in the crystal.

structural slow-light system. Materials like this permit or prohibit light of certain wavelengths depending upon the spacing of the air holes, with groups of forbidden wavelengths known as photonic band gaps (PBG).

Photonic crystals are often used in slow-light systems [33–36] and are one type of resonator that can be used to make coupled-resonator optical waveguides (CROWs), also often called coupled-cavity waveguides (CCWs). In coupled resonators, the light couples into the first resonator with a high quality factor (Q factor), which is a measure of how much energy is stored versus how much is lost. A resonator with a high Q factor has very low damping. From the first resonator, the light couples weakly into a second high-Q resonator. It continues coupling weakly into the subsequent resonators until the light has passed through the system. However, because it resonates at each resonator, the group velocity is very low. Aside from CROWs in photonic crystals, CROWs can also be created in evanescently coupled Fabry-Pérot resonators [37] and coupled ring resonators [38].

1.3.3 Other Types of Structural Slow Light

Other types of structural slow light systems include metamaterials (engineered substances that are made of structures smaller than the wavelength of the light they are designed to slow) [39] and semiconductor nanostructures [40].

1.4 Material Slow Light

Material slow-light systems are those that function because of the frequency-dependent refractive index of the material through which the light propagates [41], as opposed to the spatially-varying refractive index that causes structural slow light. Many material slow-light systems rely on materials with strong non-linear properties, and some of these systems and their underlying principles are described below.

1.4.1 The Kramers-Kronig Relation

The Kramers-Kronig relation describes the link between the intensity absorption coefficient, α, and the phase index:

$$n_\phi(\omega) = 1 + \frac{c}{\pi} P \int_0^{+\infty} \frac{\alpha(f)}{f^2 - \omega^2} df, \tag{1.7}$$

where P is the Cauchy principal value, ω is the frequency of the light, and c is the speed of light in a vacuum [42]. The phase index is calculated over all frequencies, f. In the case of slow light, a narrow absorption band is created from hole burning at or near the resonant response of a nonlinear medium. The narrow absorption band gives rise to a large dispersion in $n_\phi(\omega)$. Using $n_\phi(\omega)$ to calculate the group index, n_g, from Eq. (1.6) leads to a large n_g when the dispersion in $n_\phi(\omega)$ is large. These relations are shown in Fig. 1.3. In turn, for large n_g, the simple definition of $v_g = c/n_g$ gives a small group velocity, v_g, which is the basis of material slow light.

As a pulse enters a material, its power density and peak electric field remain constant, and it experiences spatial compression [43]. The intensity of the pulse remains constant, as the increase in energy density, u, upon compression is counteracted by the decrease in velocity, as $I = uv_g$, where I is the intensity of the pulse [2].

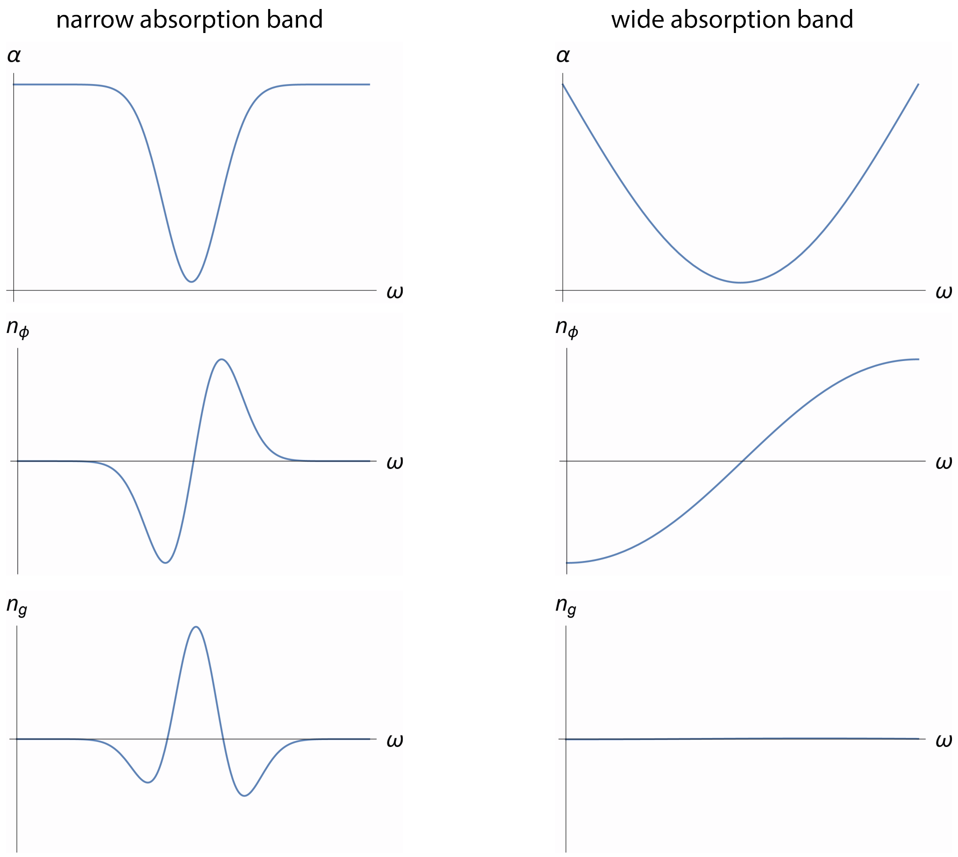

FIGURE 1.3: (*left*) A narrow absorption band, leads to a large dispersion in n_ϕ, which in turn gives rise to a very large n_g. (*right*) A wider absorption band causes a much smaller dispersion in n_ϕ, and therefore a small n_g.

1.4.2 Optical Interactions with Atoms

Optical interactions with atoms are often described using either a two- or a three-level model. In a two-level model, the ground state ($|1\rangle$) and excited state ($|2\rangle$) are separated by the photon energy $E = \hbar\omega$, where ω is the frequency of the light and $\hbar$ is Planck's constant divided by 2π. Therefore, light of frequency ω is resonant with the transition from $|1\rangle$ to $|2\rangle$, and the possibility of a transition to any other energy level is negligibly small and can be ignored.

A three-state model includes two light fields of different frequencies, one that is resonant to the transition from $|1\rangle$ to $|3\rangle$ and the other that is resonant to the $|2\rangle$ to $|3\rangle$ transition. Often dipole selection rules eliminate the possibility of a transition from $|1\rangle$ to $|2\rangle$.

Figure 1.4 shows (a) two-level and (b) three-level models.

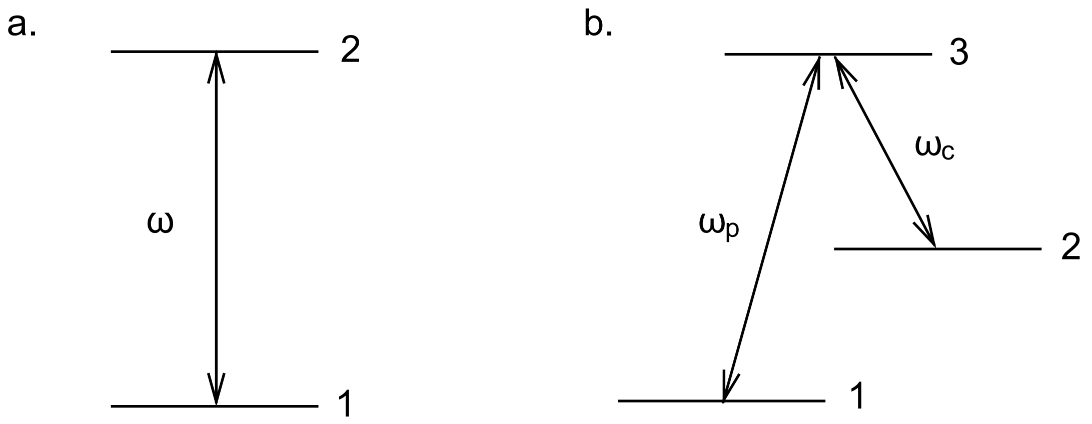

FIGURE 1.4: Models of the energy levels in (a) two- and (b) three-level systems. The coupling beam is resonant between levels $|2\rangle$ and $|3\rangle$. The probe beam is resonant between levels $|1\rangle$ and $|3\rangle$.

1.4.3 Electromagnetically Induced Transparency

Electromagnetically induced transparency occurs when a spectral hole is created in a region of absorption. A coupling beam (DC field) is used to prepare a medium such that there is a separation of the upper state ($|3\rangle$) into hyperfine states ($|3\rangle + \epsilon$ and $|3\rangle - \epsilon$). When a probe beam resonant with the upper state $|3\rangle$ is added into the system, it is unable to couple with the hyperfine states, so the probe beam experiences extremely low absorption at the frequency of the pump beam, creating an absorption spectrum that varies widely across a short frequency region. Both the states with couplings and the absorption profiles are shown in Fig. 1.5. This sharp absorption profile leads to large dispersion and therefore a large group index through the Kramers-Kronig relation, as described earlier. EIT also results in an increase in the material's optical nonlinearity that occurs alongside, but is not caused by, slow light.

EIT needs to be conducted in an environment that preserves quantum coherence, as EIT depends on quantum interferences. This means the experiments are conducted with materials at cryogenic temperatures or vapours at low pressure. Despite the complexity of the experimental requirement for EIT, it has been a popular method for creating slow light, due to the extreme degree of slowing possible. Hau *et al.* performed EIT with Bose-Einstein condensates to slow light to 17 meters per second in an experiment that greatly increased the scientific community's

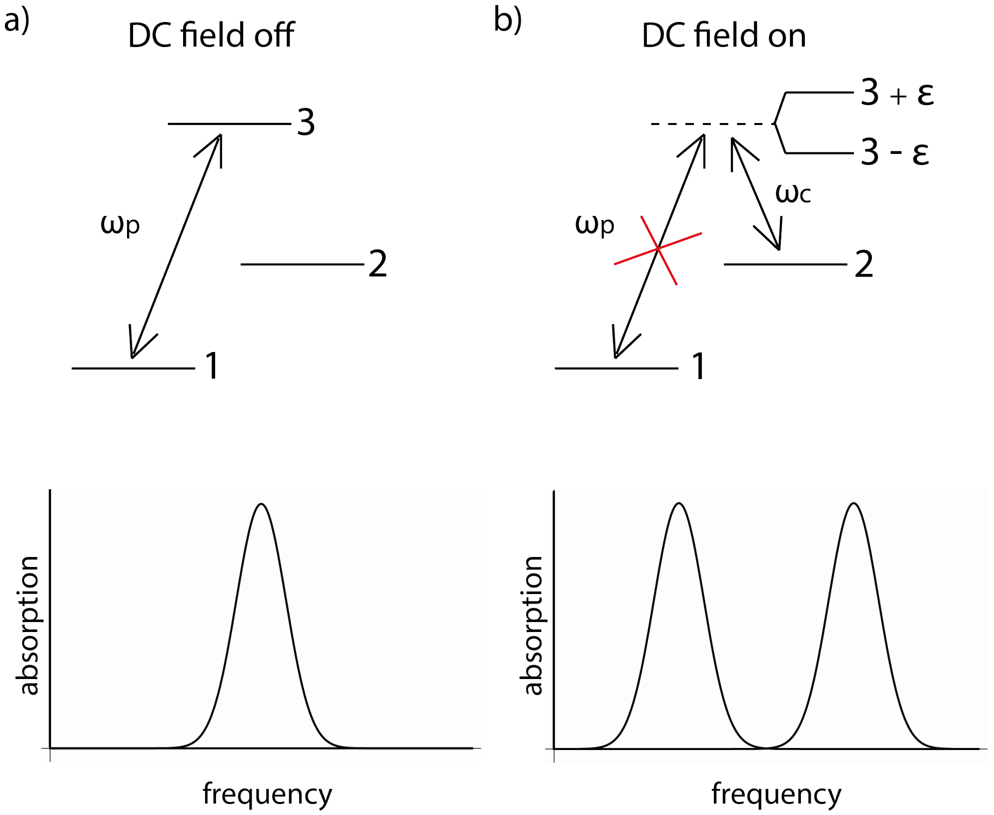
a)
DC field off
3
ωp
2
1
b)
DC field on
3 + ε
3 - ε
ωc
ωp
2
1
absorption
frequency
absorption
frequency

Figure 1.5: Models of the (*top*) energy levels and (*bottom*) absorption profiles with (*a*) the coupling beam off and (*b*) the coupling beam on.

interest in slow light research, as mentioned earlier [13].

1.4.4 Stimulated Brillouin and Raman Scattering

Stimulated Brillouin scattering (SBS) and stimulated Raman scattering (SRS) occur when pump light of frequency ω is scattered off of a vibrational wave with frequency ω_v (see [44] for a review of SBS and SRS). A field is created at the Stokes frequency, ω_S, corresponding to the difference between the pump and vibrational wave frequencies:

$$\omega_S = \omega - \omega_v. \tag{1.8}$$

When a probe field is added at the Stokes frequency, it will gain energy from the scattered field (at the same energy). That gain line converts the material in the region to a slow-light medium through the Kramers-Kronig relation of a strongly dispersive material over a short frequency range [16].

Stimulated Brillouin scattering occurs in a medium that changes in density when an optical field is applied to it. This change in density, an electrostrictive effect, results in a change in the refractive index. The beating of the pump and probe fields travels through the material, causing a moving modulation of the density as a pressure or acoustic wave. The probe wave scatters off of this acoustic wave at the probe frequency. The enhanced probe field then feeds back to increase the acoustic wave, creating a positive feedback system.

Stimulated Raman scattering occurs when the difference in frequency between the pump and probe fields causes a beating that excites molecular vibrations at frequency ω_v. Energy from the pump field is scattered at the Stokes frequency, ω_S, which in turn causes a gain in the probe field. Meanwhile, in the presence of the pump beam, the excess of molecules in the ground state can also be excited by light at a frequency called the anti-Stokes frequency, $\omega_a = \omega + \omega_v$. Unlike SBS, SRS can be seen in any system that can be excited vibrationally, including atoms and crystals, as well as molecules. For SBS, the vibrational wave has frequencies up to the gigahertz region, whereas SRS has frequencies raining up to the terahertz region [16].

Many experiments have been conducted where SBS [45–47] and SRS [48] have been used to create slow (or fast) light.

1.4.5 Coherent Population Oscillations

Coherent population oscillations (CPO) is another method of creating material slow light [49–57]. In coherent population oscillations, the atoms in a saturable absorber are excited by intense, modulated laser light. This coherent light excites the atoms to the upper state. The atoms coherently de-excite, creating an oscillation of the atoms between the upper and lower states. CPO can be modelled as a two-level system, with a coupling (pump) beam (ω_c) and probe beam (frequency of $\omega_p = \omega_c + \delta$) at the same transition. The two beams at this transition causes a cycling of the population between the excited and ground states at the beat frequency of $\omega_p - \omega_c = \delta$. The hole in the absorption is centred at a frequency of ω_c and is very narrow, leading to a large change in absorption over a short frequency range, which gives the large dispersion and therefore high group refractive index necessary for slow light. CPO uses narrow linewidths, although the linewidth is limited to $1/T_1$ of its atomic transition [16].

CPO can be conducted in crystals [58, 59] and many other experimental setups [55, 60].

While CPO usually relies on coupling and probe lasers, in 2003 Bigelow *et al.* showed that a single intense beam could be used to created slow light, without the need for separate coupling and probe lasers [58]. A benefit of this method is that one does not need to laser frequency lock, as the "pump" and "probe" lasers are the same, so even if the frequency drifts slightly, the medium is prepared at the exact same frequency as the probing light.

Coherent population oscillations in ruby has some drawbacks, when compared to other slow light methods, including the high power needed to create the slow-light effect. Ruby also has a very high absorption, as shown later in Tables 2.1 and 2.2. This method becomes impractical for many low-power situations, including single-photon experiments. However, there are also many benefits to CPO in ruby, including a relatively simple setup. No large vacuum or cooling system is needed; this experiment can be done in a simple tabletop setup. CPO in ruby also does not require any temperature or frequency stabilisation, as mentioned earlier.

Observations of CPO-based slow light in ruby has been the subject of much debate [61–64]. Several researchers [1, 63, 65, 66] have claimed that a CPO model of n_g-based slow light is not required to explain the results presented by Bigelow *et al.* [58]. Instead, Zapasskii *et al.* [65] claim that saturable absorption from the ruby could explain the observed effects. As I have discovered, the effects of slow light in ruby are not fully explained by either CPO or saturable absorption. Both my observations of the effects of slow light in ruby and the two previously-considered models will be discussed in detail in Chapters 2 and 3.

CPO-based slow light has been observed in alexandrite as well [59]. Whereas ruby only experiences decreases in absorption over narrow frequency bandwidths, in alexandrite there are short frequency bandwidths over which the material experiences greatly increased absorption as well as frequency bandwidths over which it experiences greatly decreased absorption. These increased absorption areas allow for superluminal propagation velocities, fast light, as well as greatly decreased propagation velocities, slow light.

1.5 Applications

Slow light has already begun to have applications in many areas. I describe a few of the important applications below.

1.5.1 Optical Delay Lines and Optical Processing

One application for slow light is in optical delay lines [67, 68]. For example, instead of needing to reflect a signal back and forth across the lab to delay it the necessary amount of time (a distance that could be quite large) one could instead pass the light through a slow-light system to achieve the same time delay of the signal. This would especially be useful in an optical communication system that might need to be quite compact.

Tuneable optical delay lines that can store and buffer light pulses could aid in compensation from jitter caused by random processes, such as vibrations, nonlinearities or changes in temperature of the transmission medium. Other uses for optical delay lines include optical coherence tomography (OCT), ultrafast pulse metrology, and other optical signal processing [69, 70].

Optical systems are currently used to send signals. However, when optical signals reach their destination, they have to be converted to electronic signals in order to be processed, as current technology cannot store the signals in a manner necessary for processing. After the signals are processed electronically, they are converted back to optical signals and are sent off again. However, there are two benefits to be gained from conducting the entire process optically, with optical buffers and routers. Firstly, removing the conversion to electronic signals and back (called OEO conversion) would reduce the required energy, making the system more energy-efficient. Secondly, a problem faced by the current system is that the electronic processing can only handle one packet of information at a time. When two packets of information reach a port at the same time, one is dropped while the other is processed. The dropped packet then has to be retransmitted, a process that increases the network latency [16]. If this process were occurring in an all-optical system with slow-light capacities, one packet could simply be stored until the first one was processed. This would eliminate the need for packets to be retransmitted, increasing the speed of the network.

1.5.2 Interferometry

An interferometer compares two path lengths by combining the output from the two paths. The fringe pattern measured from the superposition of the beams traveling along each path gives information about the relative phases of the two beams, and thus the optical path length difference (OPD). A change in frequency changes the measured OPD, and adding a slow-light medium into one arm of an interferometer will magnify the change in path length by a factor of the group index [71].

In recent years, there has been a growing interest in the use of slow-light materials to increase the spectral sensitivity of interferometers [71–73]. The spectral sensitivity of an interferometer is equal to the change in phase between the arms of the interferometer per change in input frequency and is proportional to Ln_g, where L is the length of the slow-light material. Adding a fast-light material to an interferometer would broaden the cavity linewidth and decrease its sensitivity to frequency changes, and conversely, a slow-light material in an interferometer would narrow the cavity linewidth and increase its sensitivity [16]. Due to the potential for increased sensitivity, slow light has been suggested for use in Fourier transform interferometry [74] and interferometric rotation sensing [75, 76]. Using ruby as a self-pumped, slow-light material is suggested to increase the sensitivity of a Sagnac interferometer by up to a factor of one million, due to its large n_g. By moving a ruby window longitudinally along the path of the propagating light, I have begun to show signs of heightened sensitivity to path length in a Sagnac interferometer. See Appendix A for more details.

1.5.3 Enhanced Photon Drag

When light passes through a moving material, it is dragged by that material. This was shown for the longitudinal [77] and transverse [78] motion of the material over 100 and 50 years ago, respectively. Measurement of the rotation of the polarisation state by the moving material [79] followed a few years after the observation the transverse photon drag measurement. However, detecting the rotational dragging of an image proved much more difficult. The polarisation of light is unaffected by vertical and horizontal displacements of the light, making it easier to measure under conditions where the position of the beam is not stable, such as when it is passing through a rotating medium. A small amount of image rotation, on the other hand, becomes extremely difficult to measure when the position of the beam is not stable.

The degree of rotation expected when a laser beam is passing through a rotating material can be calculated through an equation originally proposed by Fermi in 1923 [80]. Fermi's equation calculated dragging of plane waves without dispersion. That equation was later refined by Player to include dispersion [81], resulting in the definition of the degree of rotation per unit length of the material, δ, as:

$$\delta = \left(n_g - \frac{1}{n_\phi} \right) \frac{\Omega}{c}, \tag{1.9}$$

where n_g and n_ϕ are the group and phase refractive indices as defined earlier, Ω is the rotational frequency of the material, and c is the speed of light in a vacuum. Multiplying Eq. (1.9) by L, the length of the material through which the light travels, results in a value for the total angle by which the light is rotated, $\Delta\theta_{image}$, as:

$$\Delta\theta_{image} = \left(n_g - \frac{1}{n_\phi} \right) \frac{\Omega L}{c}. \tag{1.10}$$

Passing an image through a 30 cm-long BK7 glass bar with a refractive index of $n_\phi = 1.5168$ that is spinning with a rotation speed on the order of a few hundred meters per second, the calculated angle or rotation would still be less than a microradian. Such a small rotation becomes difficult to detect, especially when there is image wobble introduced by the rotation of the medium and the machinery noise.

The longer the light spends in the material, the larger of an angle it can be dragged through before exiting the material. To enhance the amount of dragging experienced by the light, one could change a few experimental factors. To determine which factors could be manipulated to increase the rotation angle, one need only refer to Eq. (1.10). First, one could use a longer material. However, the length of the material would need to be increased by a factor of 10^4 to get a rotation angle on the order of 1 degree. Trying to rotate a glass bar 3000 meters long would prove impractical, experimentally. Similarly, one could increase the rotation speed. However, a similar problem would arise, with the increase in rotary speed needed to result in an observable dragging angle outside practical experimental conditions.

As it quickly becomes apparent that neither the length nor the rotary speed of the material could feasibly be increased enough to enhance the photon drag to an observable level, attention shifted to the refractive indices of the material. If one could greatly enhance the first factor in Eq. (1.10), the angle of rotation could be enhanced to a large enough degree to be observeable.

It was recognised by Franke-Arnold *et al.* that a slow-light material like ruby would have a group refractive index large enough to cause a dragging angle that could be observed directly [61]. The group refractive index of ruby depends on many factors, including the intensity and the modulation rate of the incident light. Through careful control of those factors, Franke-Arnold *et al.* were able to observe photon drag by eye, with an angle of rotation of approximately 5 degrees. Because the group index of ruby can be so large, the experiment was conducted with a ruby that was only 100 mm long and rotated with speeds only up to ±30 Hz [61].

By observing the rotation angle of the incident beam to be approximately 5 degrees at a rotation speed of 30 Hz, an effective group index can be calculated as approximately 1.4×10^6, which gives an effective group velocity of approximately 200 m/s.

1.6 Conclusions

In recent years, slow light has become the focus of much research due to its potential for changing the effective velocity of light. This ability leads to many potential applications in areas of optical delay lines, optical processing, interferometry, and enhanced photon drag, some of which have already begun to be realised. While some methods of creating slow light are well known, others are still in the process of being understood. The great potential of slow-light systems makes this an area of rich possibility for researchers.

Chapter 2

Slow Darkness

2.1 Introduction

Slow light is a general term applied to systems that exhibit a greatly reduced group velocity for the propagation of optical signals [13, 82]. Underlying mechanisms that can be used to create slow light range from traditional optical delay lines to structural or material optical resonances, as outlined in Chapter 1 and described in [41, 83–85]. Common to many of these mechanisms are narrow-frequency optical features, which give rise to extremely high dispersion and hence a large group index of the material [86]. For example, one experimentally realisable class of systems is that in which intense laser light induces a coherent oscillation of the ground state population of a dopant in a solid-state medium of the type often encountered in laser amplifiers [31, 32, 66, 87]. As the intensity increases, the absorption of the medium saturates, forming a narrow spectral feature that, via the Kramers-Kronig relation, gives a modification of the group index [82]. In materials with long upper-state lifetimes, such as ruby or alexandrite, the velocity of a transmitted laser pulse can be reduced by a factor of one million, to only hundreds of meters per second [58, 59]. While other slow-light mechanisms require both pump and probe beams, Bigelow *et al.* realised in 2003 that the same effect could be observed

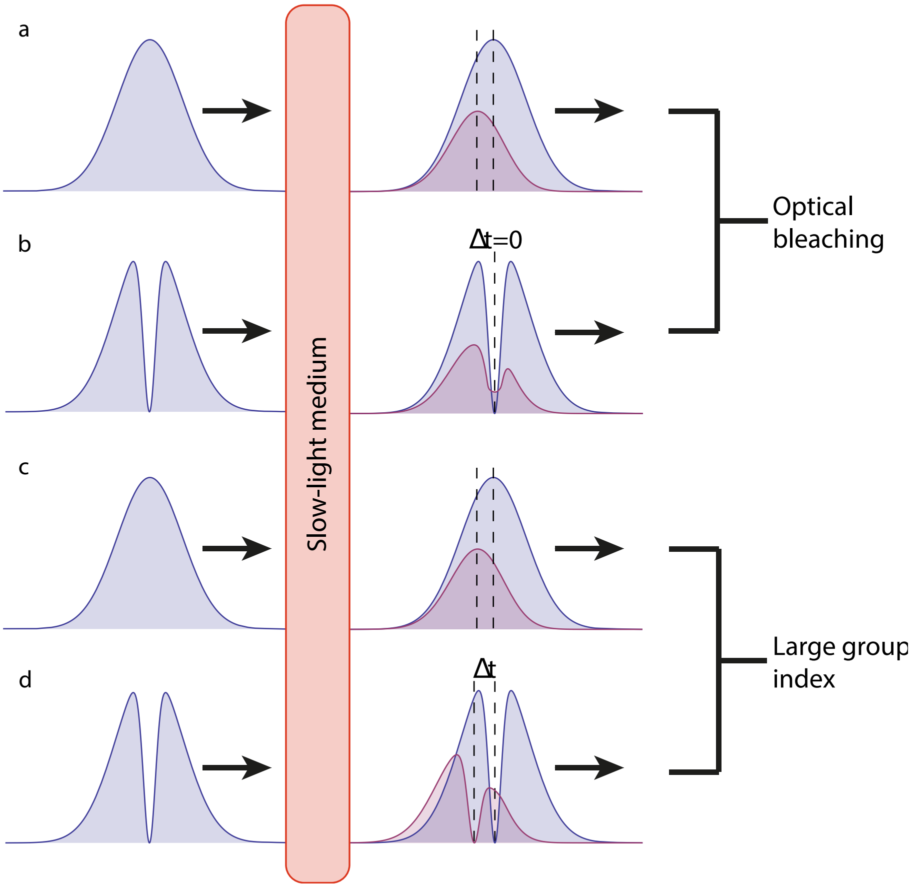

Figure 2.1: Distortion of an incident pulse passing through a slow-light medium. Part *a* shows a single pulse passing through the slow-light medium. (*a*) Optical bleaching and (*c*) large group index effects are indistinguishable for small delay of a single pulse, typical of most experimental conditions. To distinguish the two cases, one can pass a pulse with an imprinted zero through the slow-light medium. Optical bleaching and pulse reshaping would never cause the position of the dark region to shift, although it may become brighter due to fluorescence from the saturable absorber (*b*), an effect that is in agreement with [1]. However, large group index would cause the position of the dark region to be maintained but shifted (*d*). By imprinting an intensity zero in the pulse, I can determine which mechanism is causing the apparent slowing of the light through the ruby, even with a small delay.

when illuminating ruby with a single, high-intensity laser. This simple approach is called self-pumped slow light.

Broadly speaking, within a saturable absorber, one may identify two distinct mechanisms of pulse delay. The first possible mechanism arises from pulse reshaping through optical bleaching. In the case of optical bleaching, the bright incoming pulse becomes a weaker outgoing pulse, with the peak position being dependent on the differing absorption experienced by the leading and trailing edges of the pulse. As the medium absorbs energy from the leading edge of the peak, it becomes saturated, resulting in a decreased absorption of the trailing edge of the pulse. Reference [1] correctly identified that the resulting distortion of the pulse leads to a time delay of the peak as well as a change in the amplitude of the pulse. Such a time delay may be interpreted as evidence of slow light. The second possible mechanism of the pulse delay arises from the rapid change of the phase index n_ϕ near the vicinity of a narrow absorption feature through the Kramers-Kronig relation. Since the width of the coherent population oscillation is of the order of kHz in ruby [82], the induced high dispersion can give rise to extremely large group indices of the order of 10^6. In both possible causes, the extent of the saturation, and hence of the slow-light effect, depends strongly upon the incident intensity [58, 87]. When the input is a single pulse and with a small delay, as expected for most experimental conditions, the two effects are practically indistinguishable from each other, as shown schematically in Fig. 2.1 *a* and *c*. If, however, the input beam includes an intensity zero imprinted in the centre of the pulse, the two cases can be distinguished. If the slow-light effect arises purely from broadband optical bleaching, the region of zero intensity may become brighter from fluorescence but will not be shifted (see Fig. 2.1 *b*). Additionally, with fluorescence removed, the region of zero intensity could not become brighter from optical bleaching alone. In contrast to optical bleaching, if the slow-light effect stems from a narrow absorption and the resulting dispersion and high group index via the Kramers-Kronig relation as described in Chapter 1, the light will be delayed such that the beam retains its structure, shifting both the high-intensity and dark regions. The difference between the two cases can be seen by comparing Fig. 2.1 *b* and *d*.

As described in Chapter 1, it was recently shown that slow-light effects could be observed in the spatial domain [61]. The transverse movement of an optical medium is known to laterally displace a transmitted light beam [78]. Although predicted many years ago, such photon drag effects are usually very small and difficult to observe. In addition to *lateral* displacement, a spinning medium was also shown to cause a *rotation* of the polarisation state [79], termed the mechanical

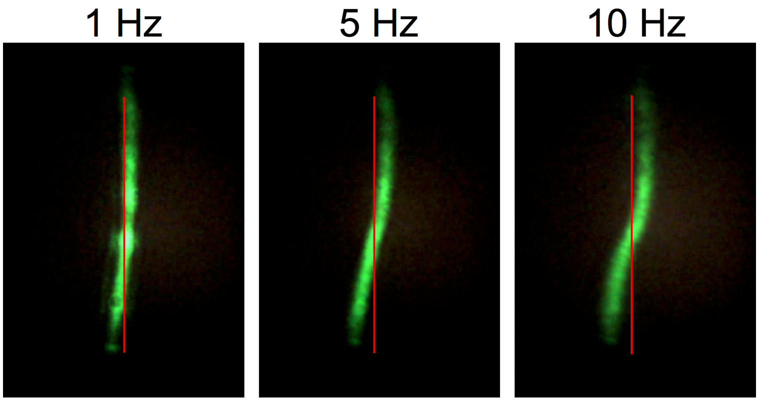

Figure 2.2: Image of the elliptical beam while the ruby is spinning clockwise at 1 Hz, 5 Hz, and 10 Hz.

Faraday effect [88]. Through the equivalence of spin and orbital angular momentum, the rotation of the medium affects both the polarisation and the transmitted image [89, 90], which is discussed in greater detail in Chapter 5. This image rotation, $\Delta\theta_{image}$, is given by Eq. (1.10). Since n_g is of order unity for ordinary materials, an image rotation caused by a spinning medium is on the order of a few microradians. It was shown that the increase in propagation time of light in a slow-light medium leads to a large increase in the rotation [61]. Specifically, it was observed that transmission through a spinning ruby window would rotate an elliptical beam profile by many degrees, a demonstration of slow light-enhanced rotary photon drag. However, in that experiment, the simple rotation of a bright line could not distinguish between optical bleaching and an increase in the group index arising from the dispersion, as shown in Fig. 2.1 *a* and *c*.

2.2 Initial Image Rotation

Before investigating whether dark regions of a beam are rotated, I first looked at the intensity and speed dependence of an elliptical beam formed from an elongated gaussian laser beam. An elliptical beam with a very elongated shape is rotated at three different speeds to show the effect of rotation speed on the angle of dragging experienced by the image (see Fig. 2.2). The image is rotated in an S-shape, due to the varying intensity across the length of the image.

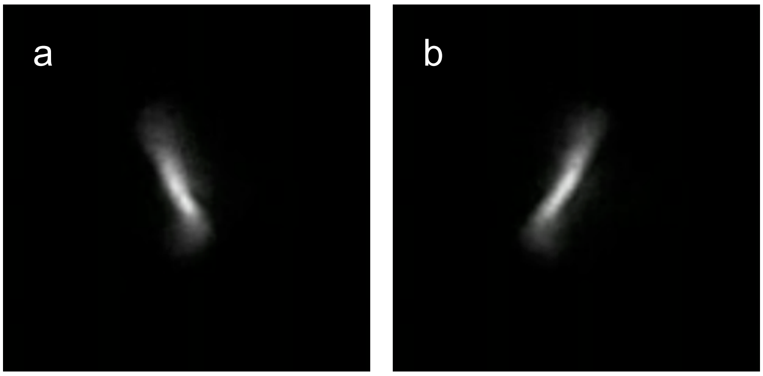

Figure 2.3: Largest recorded image drag between (*a*) anticlockwise and (*b*) clockwise rotations of the ruby was measured to be ≈ 55°.

I also recorded the largest measured rotary drag between clockwise and anticlockwise rotations of the ruby for a simple elliptical image as ≈ 55°, as shown in Fig. 2.3. Due to the strong absorption of the incident beam illuminating the ruby, these rotations, while very large, do not separate the effects of the two proposed models.

2.3 Experimental Motivation and Methods

To determine whether the rotation of the elliptical beam profile [61] arises from optical bleaching or an increase in n_g that arises from the dispersion, I study an intensity null in the centre of a high-intensity pulse. This experiment was motivated by the realisation that it is easy to create an intensity null in the spatial domain. To create this null, I introduced a π-radian phase step into the Gaussian beam cross section, which creates a black line across the beam in a close approximation of an $\text{HG}_{0,1}$ mode. As this beam is a close approximation to a solution to the paraxial wave equation, it is stable upon propagation. Given that the slow-light effect, and thus rotary photon drag, is dependent on self-pumping at high intensity, the question then posed is whether a spinning medium can rotate the orientation of a structured beam containing both bright and dark regions.

As shown in Fig. 2.4, I used a 3 W beam from a Laser Quantum Opus solid state, diode-pumped laser. The initial beam diameter was approximately 2 mm. The

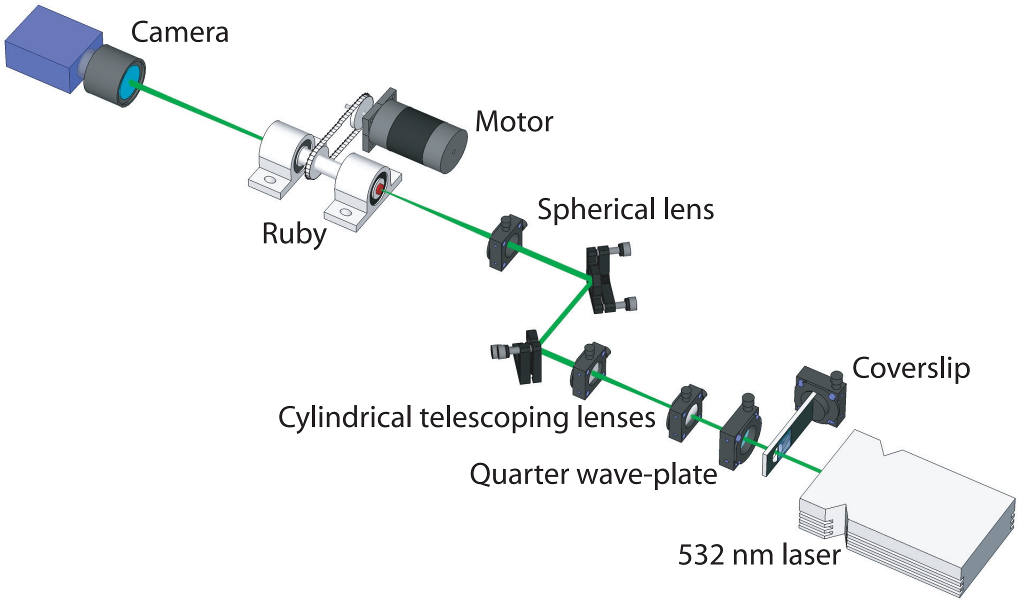

FIGURE 2.4: Experimental procedure. A beam of 3 W, 532 nm light is bisected by a coverslip with a thickness chosen so as to create a π-radian phase shift between the two halves of the beam. The entire beam then passes through a quarter waveplate and two cylindrical telescoping lenses before being focused on a ruby window. The ruby is spun to $\approx \pm 20$ rotations per second (rps), and the light leaving the ruby is imaged onto a camera.

linearly polarised 532 nm light was passed through a quarter waveplate to create circularly polarised light. Using circularly polarised light ensures rotational symmetry, which is important as ruby is birefringent. This beam was then expanded through 40 mm and 60 mm focal-length cylindrical lenses, resulting in a collimated, elliptical beam. The beam entered the ruby as an elongated spot, as modelled in Fig. 2.5*a*. Figure 2.5*b* shows the unwrapped annular intensity distributions at different radial distances from the centre of the modelled elliptical beam in *a*. A spherical lens with a focal length of 60 mm focused the elliptical beam onto the front face of a 6 mm-thick standard laser ruby crystal window, along the window's rotational axis. When pumped with 532 nm light, the ruby acts as a slow-light medium, and thus as the light passes through the ruby window, its group velocity is effectively slowed to tens of meters per second.

A Parker Hannifin, SY-series stepper motor with a built-in ViX250IM controller has encoders that allow us to accurately control the rotation rate of the ruby. A

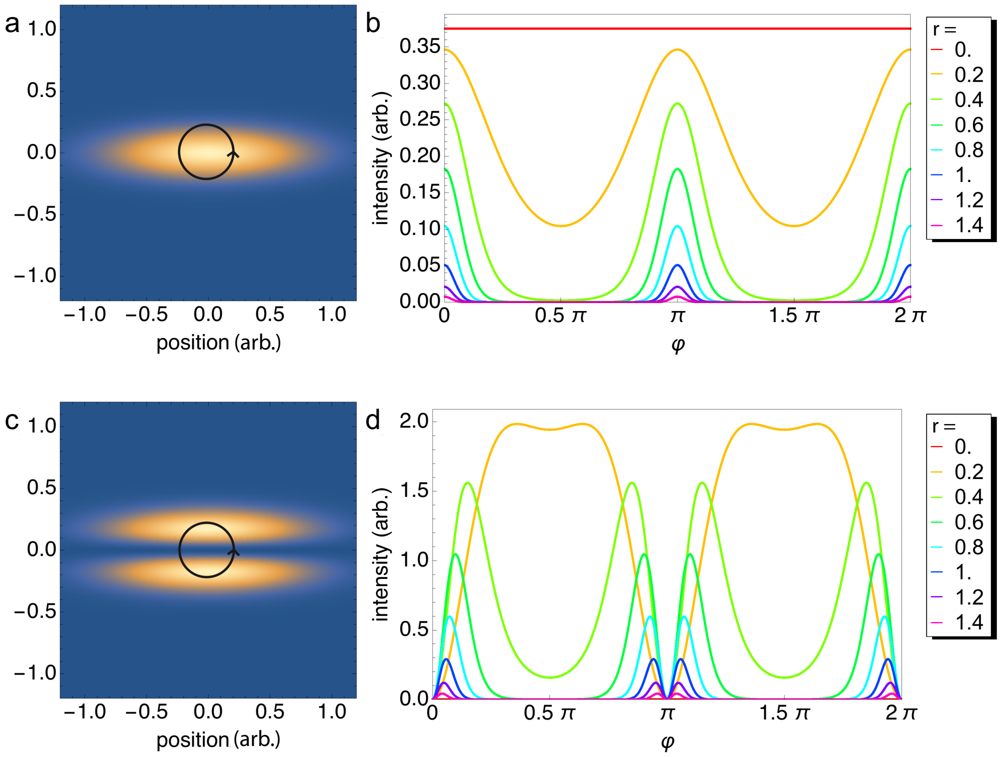

FIGURE 2.5: (*a*) Modelled intensity profile of an elliptical beam. (*b*) Unwrapped annular intensity distributions of *a* measured around the origin as indicated by the black circle on *a*, where r is the radial distance from the origin in *a*. (*c*) Modelled intensity profile of an elliptical beam with a black region through the centre. (*d*) Unwrapped annular intensity distributions of *c* measured around the origin as indicated by the black circle on *c*, where r is the radial distance from the origin in *c*.

toothed drive belt and pulley couple the motor to the ruby mount in a 1:1 gear arrangement. The motor has encoders that spin the ruby window about its axis up to $\approx \pm 20$ rps. Software control of both rotation rate and direction was straightforward to implement using a standard desktop PC computer interface and a National Instruments LabVIEW Virtual Instrument. The magnitude of the image rotation is dependent upon the angular frequency at which the ruby spins [61]. Below ±10 rps, there is a near-linear dependence on angular frequency, as described in Eq. (1.10). Above ±10 rps, the dependence levels out asymptotically with little noticeable increase in effect as the angular speed is increased, as a consequence of the 20 ms upper-state lifetime of the ruby and the corresponding relaxation time of 3 ms, which limit its response [61]. The beam was imaged from the back face of the ruby onto a Dalsa Genie CCD array and recorded for subsequent analysis using another National Instruments LabVIEW Virtual Instrument.

Alignment was critical in this experiment, with the sensitivity of the slow-light effect being so great that any slight misalignments of the telescoping lens or the final focusing lens before the ruby would completely eliminate the effect. The intensity of the light illuminating the ruby was also extremely important in determining the group velocity, and therefore the rotation angle, of the light transmitted through the ruby, as discussed by Franke-Arnold *et al.* [61] and later in Chapter 5. Another experimental difficulty arose from the flatness of the faces of the ruby crystal. Because the light was passing through a rotating medium, if the faces of the ruby deviated at all from parallel to each other (and perpendicular to the direction of the propagation of the laser beam) the resulting image would travel in a circle as the ruby spun. Any small rotary deflection made the image very hard to capture and even harder to analyse. Aligning a laser ruby crystal that had end flatness below manufacturing specifications would not guarantee a still image.

In order to introduce the black line into the elliptical beam, I mounted a glass coverslip onto a tilt stage and inserted it halfway into the output beam of the laser, as shown in Fig. 2.4. The tilt stage allowed us to adjust the orientation and therefore the thickness of the glass through which the light is travelling. At the correct tilt angle, the difference in path lengths caused a π-radian phase discontinuity between the two halves of the beam, as shown in Fig. 2.6. This phase discontinuity created a line of darkness, with a contrast limited only by scattering and crosstalk between pixels, in an otherwise high-intensity beam. Figure 2.5*c* shows the modelled beam profile once the coverslip is tilted to the correct angle to give a π-radian phase discontinuity between the two halves of the beam. Figure 2.5*d* shows the unwrapped annular intensity profile for the beam with the discontinuity down the centre. I estimated that the intensity of the incident beam exceeded the saturation intensity or the ruby by a factor of 4 or 5. By contrast, as measured from the camera image, the dark line (which is ideally of zero intensity but may be subject to experimental imperfection) is at least an order of magnitude darker than this.

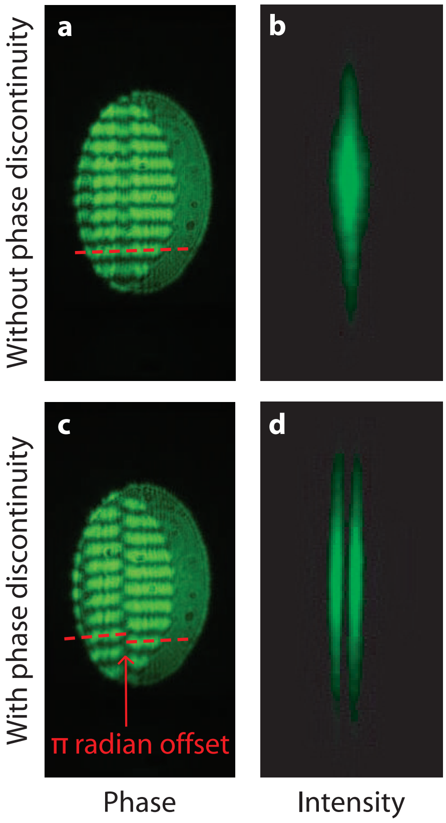

FIGURE 2.6: Images of the beam profile when a microscope coverslip is inserted halfway into the beam. The coverslip can be tilted so (*a*, *b*) there is no phase shift or (*c*, *d*) there is a π-radian phase shift between the two halves of the beam. Figures *a* and *c* are taken before the ruby and use a shear plate to show the interference pattern for the coverslip set to two different thicknesses. These images show that tilting the coverslip does create a π-radian phase shift between the two halves. Figures *b* and *d* show the corresponding intensity profiles after the ruby. Note the creation of an intensity null in *d*.

2.4 Results

Images of the light beam transmitted by the ruby were recorded, and representative images from the results are shown in Fig. 2.7. As the direction of the rotation of the ruby window was changed, one can see that the central portion of the elliptical beam became rotated about 5.0 ± 0.2 degrees, preserving the structure such that the dark region was also rotated through a similar angle. From Eq. (1.10), a rotation of 5.0 ± 0.2 degrees implies an effective n_g of $3.6 \times 10^7 \pm 0.9$, which corresponds to a group velocity of $v_g = 8.3 \pm 0.3$ m/s. It is also evident from Fig. 2.7 that the rotation angle is radially dependent, specifically that positions of

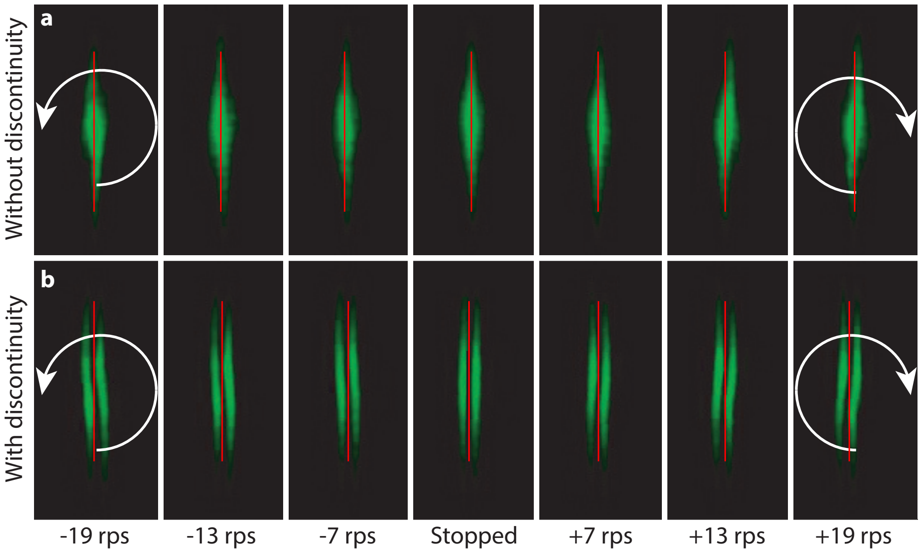

FIGURE 2.7: Images of the beam (*a*) without the phase discontinuity and (*b*) with the phase discontinuity when the ruby is spinning from -19 rps to +19 rps. The rotation of the elliptical beam increases with speed to a maximum rotation angle of ≈ 5° between clockwise and anticlockwise rotations for both the case without and with the discontinuity.

the beam at larger radii are rotated less. As the Gaussian intensity of the beam becomes elongated in the vertical direction, the positions of the beam at larger radii have less intensity. The intensity dependence of the group index results in an S-shape of both the transmitted beam and the phase discontinuity.

In order to better observe the motion of the black region, it is useful to overlay two images taken while the ruby is spinning in opposite directions. Figure 2.8*b* shows two images of the beam while the ruby was stopped, overlaid on top of each other. One set of images was artificially coloured blue, and the other set of images was artificially coloured red. Where the two beams overlap, the image is white. As is seen in Fig. 2.8*b*, while stopped, the two beams are almost entirely overlapped, leaving a mostly white image. However in Figs. 2.8*a* and *c*, the positions of the black lines in either direction can be seen. In *a*, the blue image was taken while the ruby was spinning clockwise and the red while the ruby was spinning anticlockwise. The positions of the black lines for the two images can be seen, and they are spatially separated, with a region of white between them, meaning the black line while rotating in the clockwise direction is entirely beyond the position

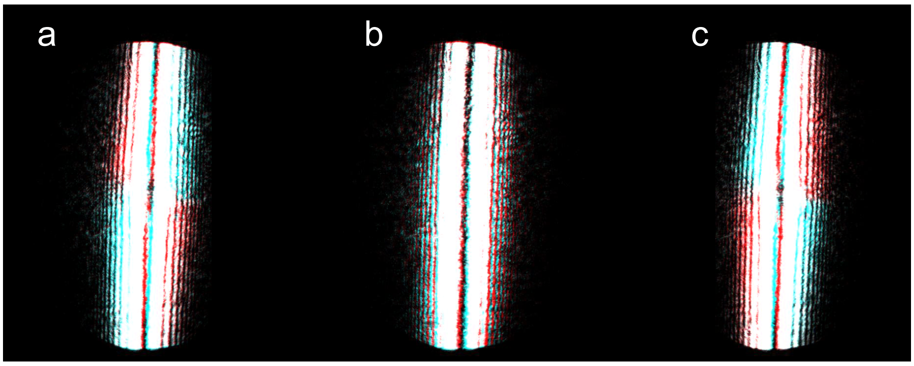

Figure 2.8: Two images overlaid to show the rotation of the black region. One image is artificially coloured blue and the other red. Overlapping regions of the two images are shown in white. (*b*) The two images overlaid while the ruby was stopped in both cases. (*a*) The two images overlaid while the ruby is spinning anticlockwise for the red image and clockwise for the blue image. (*c*) The two images overlaid while the ruby is spinning in the opposite direction from *a*.

of the black line while rotating anticlockwise. Depending on the focusing before and imaging after the ruby, the shape of the image can change quite dramatically, as seen by comparing Figs. 2.7 and 2.8. Optimising both allows for the observation of the largest image rotation.

To investigate the rotation, I analysed intensity-averaged images of the beam when the ruby was removed, rotating clockwise, and rotating anticlockwise to eliminate the effects of vibrations from the motor. For each case, I extracted an annulus at the radius with the most rotation. I then performed a polar to Cartesian coordinate transformation to unwrap the annulus into a straight strip (Fig. 2.9*a*). The intensities for each x-position of the strip were summed, and the resulting intensity distribution was plotted (Fig. 2.9*b*). The dips around 90 and 270 degrees are the locations of the phase discontinuity for the three cases. The enlarged regions of Fig. 2.9*b* show an increase in intensity in previously dark regions and a reduction in intensity in previously light regions. The movement of the dark region outside of the position where it is located when the ruby is removed implies that the slow-light phenomenon cannot be described by a simple model of optical bleaching and pulse reshaping. Upon transmission through the ruby, much of the initial intensity is absorbed. Differential absorption across the beam when the ruby is spinning in each direction could potentially appear as if the image were being rotated, without any actual image drag. When the ruby is removed, the intensities measured across the beam are the highest possible intensities at each location. Therefore, an increase in intensity at any position upon transmission through a spinning ruby as compared to the intensities with the ruby removed could be attributed to a slow light-enhanced photon drag effect.

For clarity, Fig. 2.10 shows only the removed intensity and the anticlockwise spinning traces. The enlarged region, centred on 270°, shows the increase in intensity (shown with red dashed arrow) when the ruby is spinning (*blue*) as compared to when the ruby is removed (*olive*). The increase in intensity allows us to separate the two possible causes of slow light in ruby, as an optical bleaching (saturable absorption alone) effect would never increase the intensity at any location when fluorescence is removed.

2.5 Discussion

I have found that passing an image through a spinning, self-pumped, slow-light medium rotated the bright and dark regions by the same amount and that light in the output cross section was azimuthally displaced to positions where the input beam was dark. Such displacements cannot be explained solely by optical bleaching, since such effects cannot lead to a shift in a region of darkness. As such, rather than optical bleaching, it appears that narrow band absorption, and the associated change in group refractive index, are responsible for the observed slowing. These results further the understanding of the mechanisms that cause slow light and pave the way for applications dependent on the preservation of complex patterns in slow-light media.

2.6 Alternate Methods

I also investigated the simpler possibility of creating the black line through the use of a thin wire placed through the centre of the beam. A comparison was performed

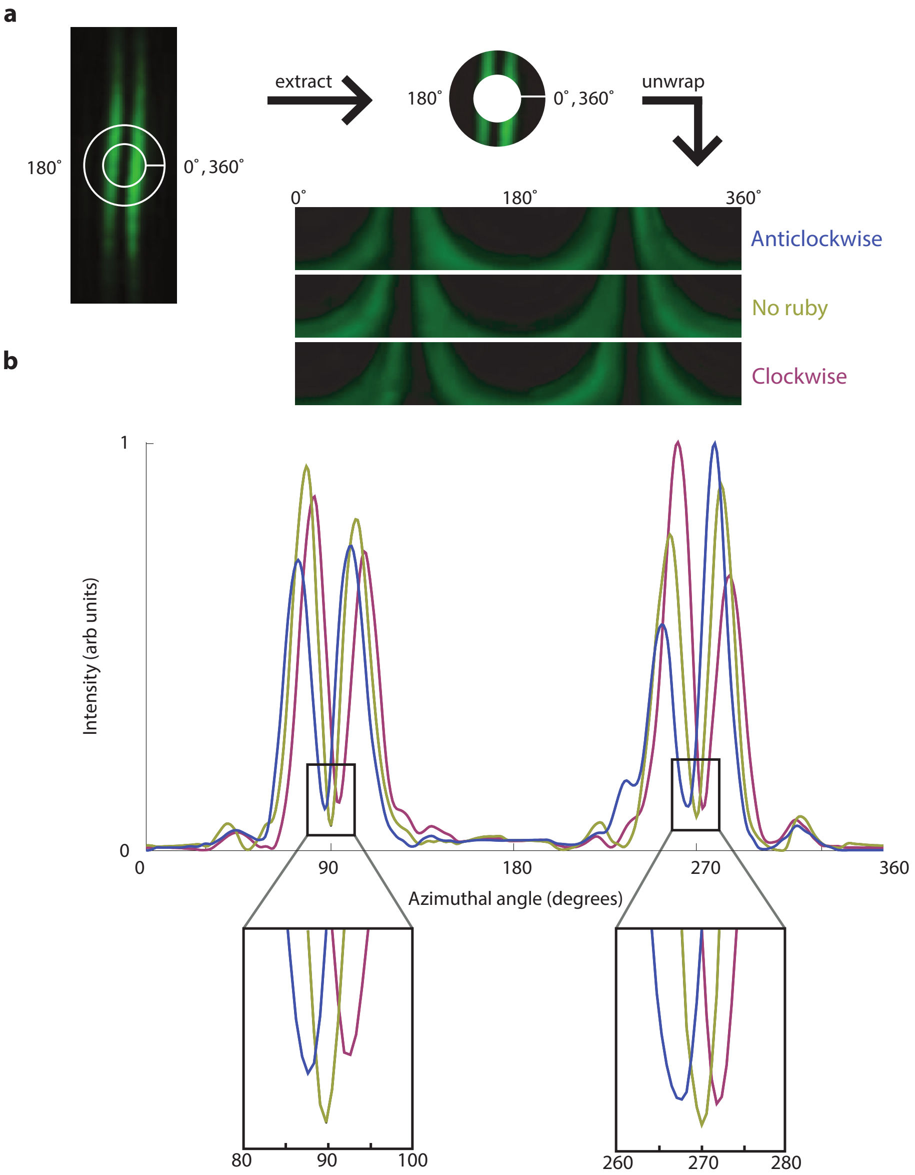

FIGURE 2.9: Intensity profiles. A strip of the beam profile at the radius with the greatest photon drag is unwrapped in *a* when the ruby is spinning anticlockwise (*blue*), spinning clockwise (*pink*), and removed (*olive*). The intensity of each of those strips is plotted in *b*. The blue and pink curves show the position of the region of darkness while spinning anticlockwise and clockwise, respectively, to be outside the position of the dark region when the ruby is removed (and therefore not spinning). One observation in this figure is the movement of the dark line causing there to be light in the clockwise and anticlockwise cases where there had been darkness in the stopped case, which could be explained by fluorescence from the saturable absorber. However, the region of darkness is preserved through the rotation to a position that had previously been bright.

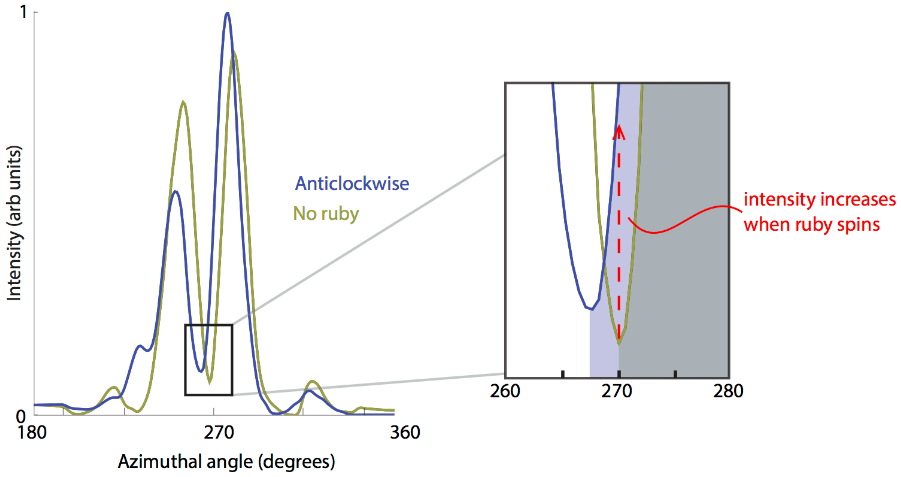

FIGURE 2.10: Intensity as a function of azimuthal position. The stopped trace (*olive*) has an intensity dip at approximately 270°, whereas in the spinning case (*blue*), the intensity dip has moved to approximately 265° and the intensity at 270° has increased.

between this method and the method using a coverslip as described above, both in round beams.

Figure 2.11 shows a comparison between rotation of the black region when created by the wire and the coverslip, both in a round beam. As discussed later in chapter 5, the shape of the bright beam (and specifically the amount of intensity modulation the material experiences) has a large effect on the degree of rotation of an image. The round beam, being almost entirely bright with very little modulation (caused by alternating regions of darkness and light) only at the region of the black line, results in a very small angle of rotation. However, evidence of rotation is observable in both cases.

There are two reasons that I ultimately settled on using the glass coverslip to create a discontinuity for the black line. First, the wire simply blocks the light, casting a simple shadow. Shadows aren't necessarily stable upon propagation. In other words, just because light is blocked by an object, creating a dark region in the shape of that object directly behind it, the region will not necessarily remain dark in all z-positions (where the beam is traveling along the z-axis). Diffraction from the wire causes changing light patterns from the near-field to the far-field, resulting in intensities (and phases) that differ when examined at different z-positions. On

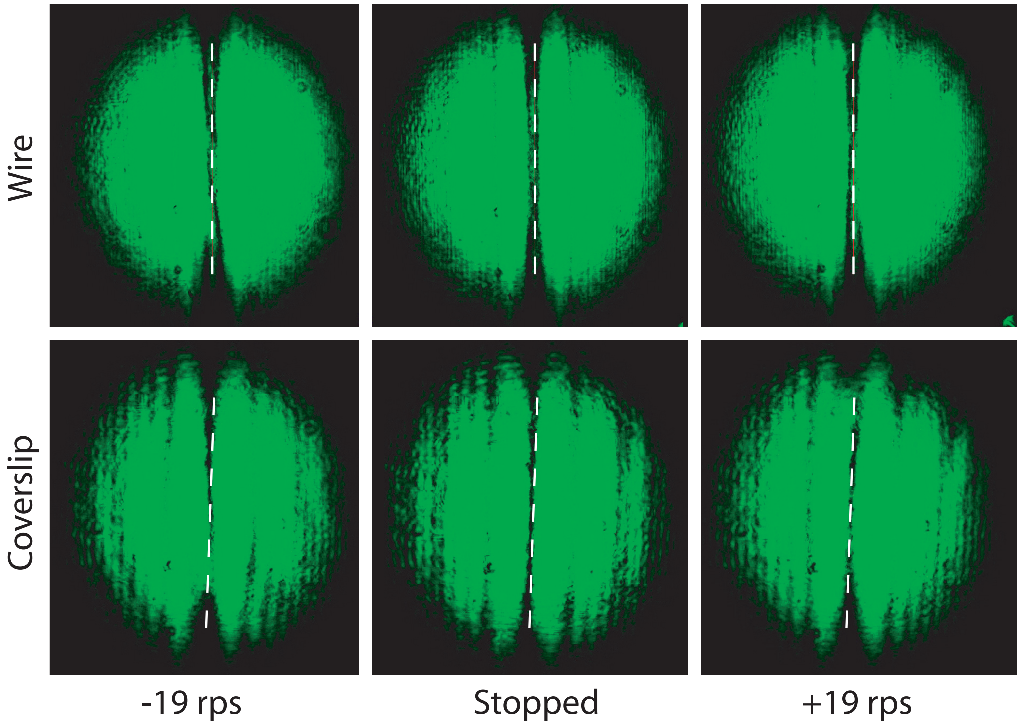

FIGURE 2.11: Images of the beam with (*top row*) the wire and (*bottom row*) the coverslip when the ruby is spinning at -19 rps, stopped and spinning at +19 rps. The dotted white line shows the position of the black line when the ruby is stopped as a reference.

the other hand, the beam shape created using a coverslip to generate a phase discontinuity is stable upon propagation, as it approximates an $HG_{0,1}$ mode, as described earlier.

The second reason that I chose the coverslip over the wire has to do with the width of the dark region. As seen in Fig. 2.11, when the ruby is stopped in the wire case, the dark shadow from the wire is very wide. In the rotating cases, although some rotation of the dark line can be seen, the rotation arises more from a brightening (filling-in) of the previously dark regions, but the region of darkness does not move outside of its original position. The combination of these results being taken without a dichroic filter and with only the wire causing a shadow raises doubts that this is actual energy transport. Conversely, the dark region caused by the phase discontinuity from the coverslip is significantly thinner, and the dark region when the ruby is rotating appears to be outside of the dark region when the ruby is stopped.

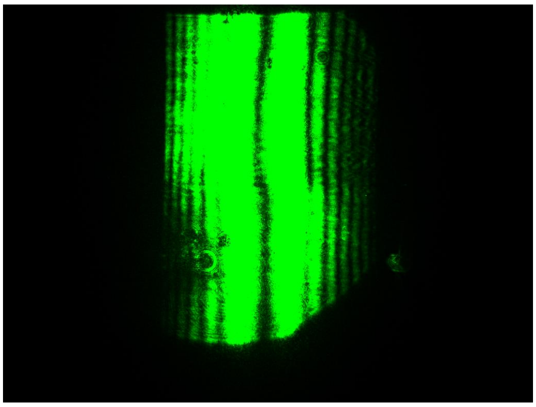

Figure 2.12: Image of the elliptical beam while the coverslip creates a π-radian phase difference between the two halves of the beam and the ruby spins anticlockwise. Fringes can be seen on the left and right sides of the image, caused by interference from multiple reflections of the light in the coverslip.

One drawback to using the coverslip is that the thickness of the glass coverslip caused some interference fringes that were not present when the wire was used. However, when the back face of the ruby was imaged correctly, this problem became negligible. An example of interference fringes from the coverslip can be seen in Fig. 2.12. When imaged more carefully, as in Fig. 2.7, the fringes were not visible, and the image was much cleaner.

2.7 Rotating Complicated Patterns

In addition to the very careful experiment focused on differentiating between the two models of slow light, I also studied the rotation of more complicated patterns. I did this two ways. The first method used the glass coverslip. However, instead of putting one edge of the coverslip halfway through the beam, I placed the corner of the coverslip into the beam. I then tilted the coverslip to the correct thickness to achieve the π-radian phase discontinuity along both edges of the coverslip. Figure 2.13 shows the resulting images while the ruby was (a) rotating anticlockwise, (b) stopped, and (c) rotating clockwise with the corner of the coverslip inserted into the beam. The main challenge with this idea was trying to get both edges of the coverslip aligned for the correct π-radian phase difference. Because the

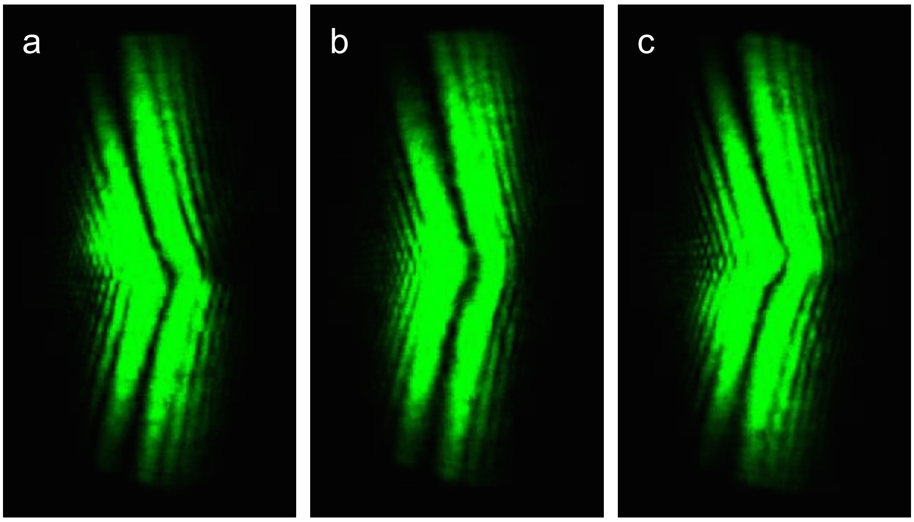

Figure 2.13: Rotation of pattern made by the corner of a glass microscope coverslip when the ruby is (b) still and rotating (a) anticlockwise and (c) clockwise.

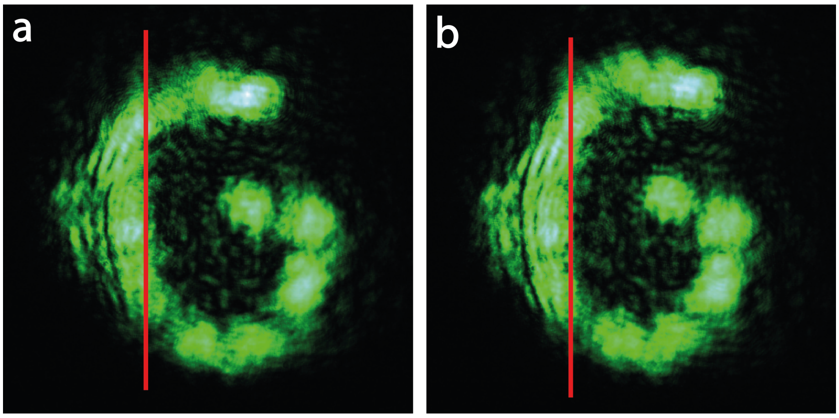

Figure 2.14: Rotation of pattern made by a metal mask when the ruby is rotating (a) anticlockwise and (b) clockwise. Red vertical lines are added for reference.

two edges were connected as part of the same coverslip, they could not be independently adjusted. Very careful alignment resulted in the results presented in Fig. 2.13, but even then I could not ascertain that both sides had the correct phase discontinuity. As can be seen in Fig. 2.13, in a the top edge of the figure seems to rotate more, whereas in c the bottom edge seems to rotate more. This seems to imply that the two different edges are slightly misaligned with each other, giving an uneven image rotation as a result.

In order to observe the reaction of even more complex patterns to slow light-enhanced rotation, I used simple metal masks to block parts of the beam in order to illuminate the rotating ruby. Figure 2.14 shows one such pattern. A clear rotation can be seen between an image taken while the ruby was rotating (*a*) anticlockwise and (*b*) clockwise. Due to the sharp edges from the mask, the image had ragged edges, and (as with the wire-based black line in the elliptical beam) the shape could not be used to differentiate between the two slow-light models. Other examples of patterns created with metal masks can be found in Appendix B, with varying degrees of detectable rotations.

2.8 Fluorescence

I investigated the possible effects of fluorescence on this experiment by purposefully misaligning the focusing lens immediately before the ruby. This misalignment meant that necessary focusing and power to achieve the slow-light effect were not met. Therefore, the rotation of the coherent green light was too small to be measured. With an unrotated green coherent image, I removed all filtering and instead measured the fluorescent light from the ruby. This was easily distinguished from the coherent transmitted light because the coherent transmitted light is green, whereas the fluorescence from the ruby is red. I recorded the red fluorescent light and immediately saw a large rotation and blurring. This makes intuitive sense, as ruby atoms that are excited by the intense green laser light will release light for an amount of time dependent on the upper-state lifetime of the ruby. As ruby has a long upper-state lifetime, the atoms continue to fluoresce as they rotate away from the initially illuminated position, decreasing in intensity as the fluorescence subsides. This results in a blurring of the initial image.

I also applied a simple model to the data to test my understanding of this effect. I rotated an image of the pattern recorded when the ruby was stationary by a series of small angles. I then summed all rotated frames, with each weighted by an exponentially decaying amplitude, i.e. images that were rotated farther had less intensity. The degree of blur in the resulting image could be adjusted by changing the rate of the exponential intensity drop off. As can be seen from

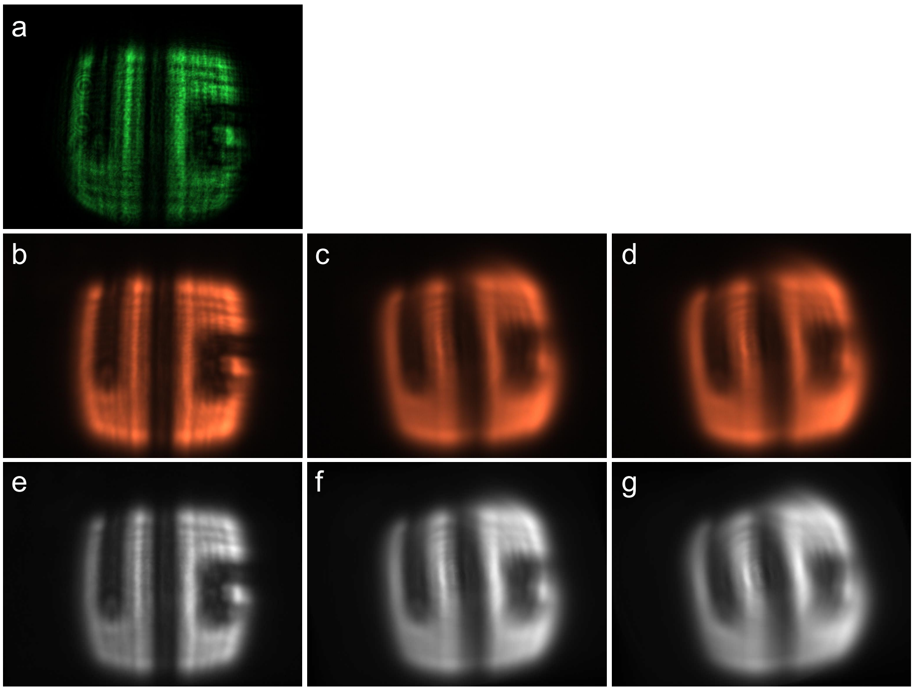

FIGURE 2.15: Images in the shape of UG are shown for the (*a*) green coherent light, (*b*-*d*) red fluorescent light, and (*e*-*g*) simulated fluorescence. The z-position of the focus lens was purposefully misaligned, such that the coherent light was not slowed. Therefore, the transmitted green coherent light (shown in *a*) was unchanged by the rotation of the ruby. Images *b* and *e* show the red fluorescence while the ruby was not rotating. Images *c* and *d* show the fluorescence at medium and high speeds, respectively, which gave increasing amounts of blurring. For processing, the image in *e* was changed to an intensity image, rather than a colour one. The intensity image in *e* was artificially rotated by a LabVIEW program, giving more image rotation and blurring with increasing numbers of iterations in *f* and *g*.

Fig. 2.15, I was able to reproduce the observed blurring effect with a decay time of approximately 4 ms, which corresponds to the upper state lifetime of the ruby. When the fluorescence was removed by a dichroic filter and the coherent green light was observed, there was no blurring and no discernible rotation of the image.

The comparison between the observed blurring, the model blurring, and the green (non-fluorescent) image is shown in Fig. 2.15.

These results reinforced my understanding of how important it was to remove the fluorescence when using rotary photon drag to learn about the nature of slow light in this experiment.

2.9 Absorption Study

Another factor that was manipulated to maximise the slow-light effect was the orientation of the ruby. Ruby is birefringent, leading to different interactions when the light is polarised parallel to the two orientations of the ruby. There was a larger slow-light effect when the light was polarised parallel to one of the two orientations, which I will call orientation A. When the light is polarised parallel to orientation A, the slow-light effect is greatest, but the absorption is also largest. When the light is polarised perpendicular to orientation A (which I will call orientation B), the slow-light effect and absorption are both minimised. I took a more careful look at the absorption to learn more about the saturation of the ruby.

The absorption of the ruby was measured at the two orientations, A and B. The absorptions listed below were measured when transmitting 532 nm light through the 90 mm-long ruby. The pre-ruby measurements were taken immediately before the lens that focused the beam onto the ruby, and the post-ruby measurements were taken immediately following the ruby. Both measurements were taken when the beam diameter was approximately 2 mm. The strong interaction orientation (orientation A) is designated as the orientation of the ruby at which the light experienced the greatest absorption and the greatest slowing (the slowest v_g), as stated before. Similarly, the weak interaction orientation (orientation B) is at 90° to orientation A and is the position with the smallest absorption and the least slowing (the fastest v_g).

The relationship between the incident and transmitted light is roughly linear with a positive slope for both orientations of the ruby, i.e. as a greater intensity is incident on the ruby, more intensity is transmitted in both orientations. However, while the percentage of transmitted light in orientation B (weak interaction) is mostly constant with a very slight upward trend with increasing intensity, the percentage of transmitted light in orientation A (strong interaction) has a downward trend with increasing incident power (see Fig. 2.16).

Table 2.1: Ruby Absorption, Orientation A

Pre Ruby (W)	Post Ruby (W)	Percent Transmitted
0.504	4.82E-06	9.56E-04
0.7483	6.99E-06	9.34E-04
1.01	8.95E-06	8.86E-04
1.246	1.06E-05	8.53E-04
1.504	1.25E-05	8.32E-04
1.749	1.41E-05	8.04E-04
2.001	1.57E-05	7.83E-04
2.505	1.91E-05	7.60E-04
3.008	2.21E-05	7.33E-04
3.502	2.49E-05	7.10E-04
4.001	2.76E-05	6.91E-04
4.506	3.05E-05	6.77E-04
4.999	3.32E-05	6.64E-04

Table 2.2: Ruby Absorption, Orientation B

Pre Ruby (W)	Post Ruby (W)	Percent Transmitted
0.5	9.81E-06	1.96E-03
0.752	1.49E-05	1.98E-03
1.001	1.98E-05	1.98E-03
1.249	2.47E-05	1.98E-03
1.497	2.95E-05	1.97E-03
1.755	3.45E-05	1.97E-03
2.003	3.98E-05	1.99E-03
2.508	4.95E-05	1.97E-03
2.997	5.91E-05	1.97E-03
3.504	6.96E-05	1.99E-03
4.007	7.97E-05	1.99E-03
4.504	8.93E-05	1.98E-03
4.999	9.98E-05	2.00E-03

In the strong interaction case, the higher the incident intensity, the less (relative) transmission there is. Although this is on a very small scale, a clear downward trend can be observed. As the amount of light absorbed does not level off in this power range, this implies that the ruby is not saturated. If the ruby were saturated, the percentage of incident light that was transmitted would level off or even increase, as the ruby would be unable to absorb additional incident intensity.

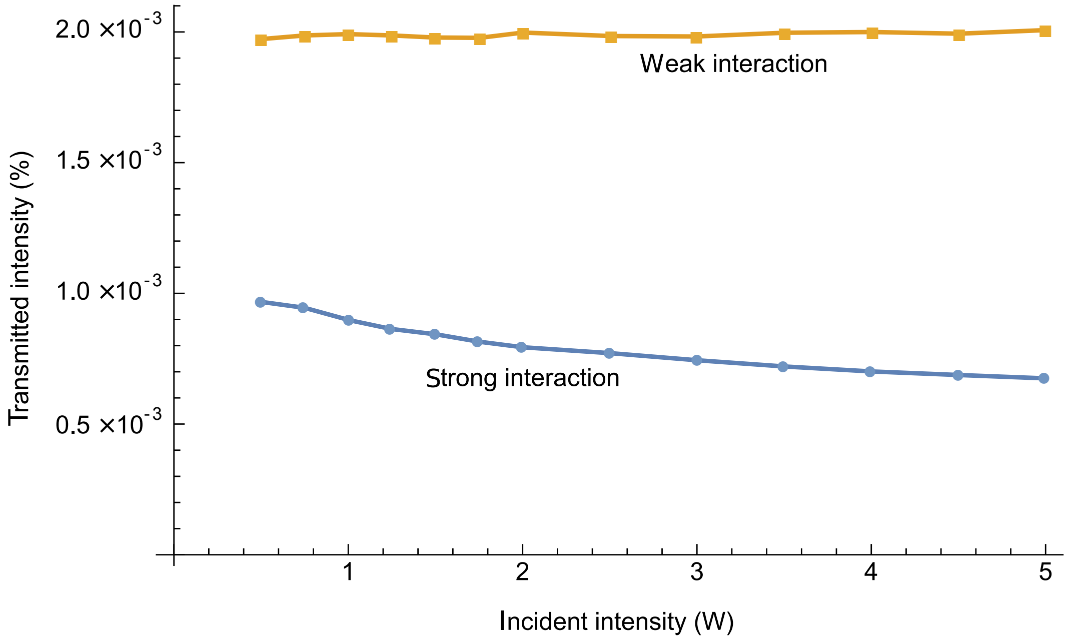

FIGURE 2.16: Percentage of transmitted intensity plotted against incident intensity for two orientations of the ruby. Yellow squares show the percentage of transmitted intensity when the ruby is turned to the orientation with the weakest interaction between the laser light and the ruby crystal. Blue circles show the transmission percentage for the orientation with the strongest interaction between the laser light and the ruby crystal.

2.10 Conclusions

Through careful examination of slow-light enhanced rotary photon drag, I have observed photon drag of an image containing both bright and dark regions, where bright and dark regions rotated by the same amount. Light was azimuthally dragged to regions that had previously been dark. The observed movement of the bright and dark regions could not be explained by saturable absorption alone, since such effects cannot lead to an increase in energy without fluorescence. The observed effects could, however, be explained by narrow band absorption and a change in group refractive index.

Whether created by either the shadow of a wire or a π-radian phase discontinuity, the dark region shows signs of movement. However, the phase discontinuity provides clearer results, as explained previously. I also found that the slow-light effect is stronger when the ruby is illuminated by an elliptical beam rather than a round one, because the elliptical beam causes a stronger modulation of the ruby. The effect of beam shape on the slow-light effect is studied in more detail in Chapter 5.

I have shown the importance of eliminating fluorescence when studying slow light in ruby, because the illumination from fluorescence can mask the actual slow-light effect. With fluorescence present, any observed dragging could not confidently be ascribed to the effects of slow light in ruby. I confirmed the ability of fluorescence to mask the desired effect with a simple model. When fluorescence is removed and other experimental parameters are carefully controlled, the measured rotation can be studied to learn about the effects of slow light in ruby.

Due to power limitations of the laser used in this experiment, the ruby was not fully saturated. This leaves open the question of whether a more powerful laser that could fully saturate the ruby would increase the slow-light effect and therefore the rotation angle. Future work should be done in this area to fully understand the exact relationship between laser power and the slow-light effect in ruby.

Chapter 3

Delaying Energy Beyond the Input Pulse

3.1 Introduction

When a pulsed light beam propagates through ruby, it is delayed by a slow-light mechanism. The debate about whether light is slowed in ruby through a simple saturable absorption mechanism or a more complicated mechanism involving the strongly varying absorption over a short frequency range [1, 62–65] has been discussed in Chapters 1 and 2. Chapter 2 focused on the use of photon drag to infer the effective group velocity while a light propagates through ruby. However, in this chapter, I simplify the experimental procedure to directly investigate the slowing of light. To do this, I move away from the spatial domain, wherein the modulation of the ruby occurs through the rotary movement of the ruby itself. In the spatial domain, a patterned beam is incident on the ruby. Whenever a position on the ruby rotates into and then out of a bright area of the stationary beam, the ruby atoms at that position experience an intensity modulation. Leaving the spatial domain, I move into a temporal modulation, where the intensity of the laser beam is modulated and the ruby is stationary. By directly measuring the arrival time of a pulsed signal, one can remove the experimental dependence on rotary photon drag (which was used to infer the effective group velocity of the light, as described in Chapter 2) and distinguish between the two main proposed mechanisms with less ambiguity. To do so, I investigate the trailing edge of a square-wave pulsed laser beam propagating through ruby. The observation of the creation of a tail at the trailing edge of the pulse provides evidence for a complicated model of slow light in ruby that requires more than pulse reshaping.

3.2 Slow Light in the Time Domain

As discussed previously, one possible mechanism for creating slow light that has lead to some debate is coherent population oscillation (CPO). One example of CPO occurs when intense green laser light propagates through ruby, causing the atoms in the ruby to be excited coherently. The CPO method requires only a single laser beam [58], as opposed to more complicated methods with co- or counter-propagating pump and probe lasers [57].

A debate has grown around different proposed mechanisms by which light is slowed in ruby [1, 61–65, 87]. Following the initial claims of CPO in ruby in 2003 [58], an alternate explanation for the observations in [58] was proposed in 2006 [1, 65], where the apparent slowing of an intense pulse of laser light was explained by the pulse reshaping brought about by a saturable absorber. Early demonstrations of slow light in ruby (see [58, 61]) could not differentiate between the two proposed mechanisms. I conducted a study in the spatial domain in an attempt to differentiate between the two mechanisms, as described in Chapter 2. In that work, I introduced a line of darkness into a bright image that was slowed, and hence azimuthally displaced, in its propagation through a rotating ruby rod. A Comment was published stating that the results in that study were simply due to saturable absorption [63]. Although we disagreed about the mechanism causing the observations in Chapter 2, the authors of that Comment, my collaborators, and I all agreed that the observation of a pronounced tail on the trailing edge of the transmitted pulse cannot be explained solely by the effects of a time-varying absorber acting upon the incident pulse. The Comment recognised that the region of darkness could be produced cleanly in the time domain through the use of a chopper to completely modulate the intensity of the pulse. In their attempt to conduct that experiment quickly, Kozlov *et al.* did not observe slow light [63]. I base this experiment on Kozlov *et al.*'s experimental setup, although slight differences in my setup allowed me to achieve different results from those they presented.

My work [91] demonstrated that the delaying of transmitted laser intensity into the trailing edge of a pulse (tail) could only be caused by a temporal slowing of light, not by a time-varying (saturable) absorption. In the absence of fluorescence, an absorber can only decrease the intensity of light present at a given moment in time. Therefore, detecting more intensity in the tail, as compared to the tail of a reference pulse, provides strong evidence that the pulse delay in ruby is caused by a mechanism more complicated than that of time-varying (saturable) absorption alone. The different delays of individual Fourier components of the pulse signal explain the pulse distortion that occurs upon transmission through the medium and must be accounted for by any model that attempts to describe the effects of slow light in ruby.

I first reproduced the results presented by Bigelow *et al.* in 2003 [58]. By comparing the sinusoidal signal in a reference arm to the identical signal passed through a slow-light medium, in this case ruby, I can measure the delay of the signal due to the slow-light medium. As was determined by Bigelow *et al.* , the slow-light effect is strongest when the signal is modulated at lower frequencies. Increasing the frequency of the sinusoidally-modulated signal decreases the slowing experienced by the signal. This is shown in Fig. 3.1, where slow-light effect increases as the the frequency decreases. The two signals become increasingly separated when one is delayed by a greater amount.

After verifying the expected behaviour of the sinusoidally-modulated signal, I turned to a square wave-modulated signal, where the sharp turn-off is used to investigate whether energy can be delayed beyond the input pulse by transmission through ruby.

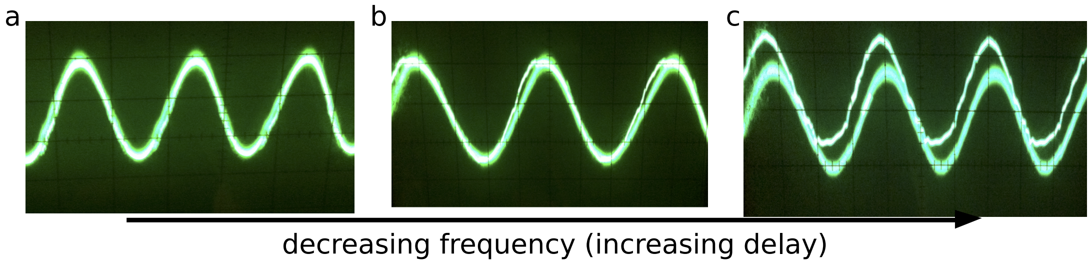

FIGURE 3.1: Signals from reference and ruby arms are shown on an oscilloscope. The frequency of the sinusoidal modulation of the signals decreases from *a* to *c*, and the ruby signal becomes more delayed as the modulation frequency decreases. In *c*, the two signals are vertically displaced from each other by a small amount for clearer viewing of the relative delay between them.

3.3 Methods and Materials

As shown in Fig. 3.2, I produced 4 W of intense 532 nm laser light within a single longitudinal mode. The initial beam diameter was approximately 4 mm. This collimated laser beam was focussed onto a mechanical chopper by a 160 mm focal-length lens. Rotation of the mechanical chopper caused a square-wave intensity modulation of the laser beam. The modulated laser beam was recollimated with a second 160 mm focal-length lens before passing through a beamsplitter. Half of the incident light was reflected onto a 60 mm focal-length lens that focussed the light through a dichroic bandpass filter onto a photodiode. The signal measured with this photodiode is designated as the reference signal. The remaining laser light was transmitted through the beamsplitter and a spherical lens with a focal length of 50 mm, which focused the beam onto the front face of a 90 mm-long standard laser ruby crystal rod. The optical axis was collinear to the rod's z axis. A 60 mm focal-length lens focussed the light through a dichroic bandpass filter and a 400 μm pinhole onto a second photodiode. The bandpass filter and pinhole ensured that only 532 nm light is measured, eliminating virtually all incoherent fluorescent light from the ruby. The intensity signals from both photodiodes were collected using a high speed data acquisition device that is controlled by a National Instruments LabVIEW Virtual Instrument, allowing for easy measurement of the intensities of the two signals.

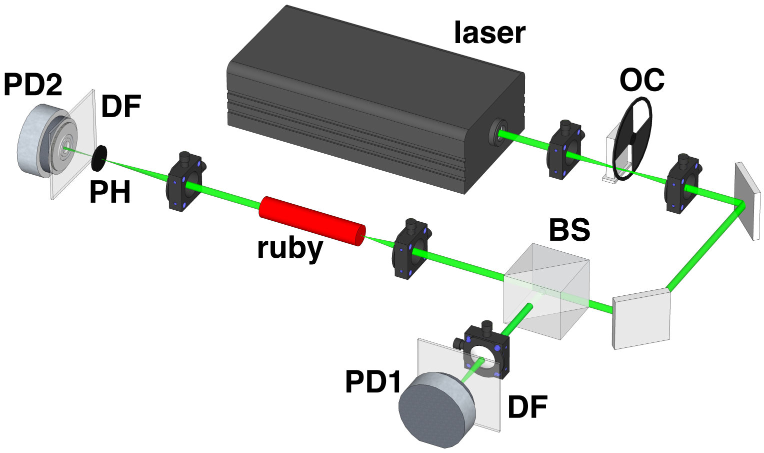

Figure 3.2: Intense 532 nm laser light was focussed onto an optical chopper (OC) and then recollimated by 160 mm focal-length lenses. The light was then split by a beamsplitter (BS), sending some of the light to be focussed by a 60 mm focal-length lens onto photodiode (PD1). The remaining light was focused onto the front surface of a 90-mm-long ruby rod with a 50 mm focal-length spherical lens. The light transmitted through the ruby was measured by a second photodiode (PD2) after a spatial filter comprised of a 60 mm focal-length lens and a 400 µm pinhole (PH). Dichroic bandpass filters (DF) were placed before both photodiodes.

3.4 Results

Figure 3.3 shows the reference and transmitted signals for measurements taken (a) without or (b) with the ruby rod in place. Both a and b are plotted on identical semi-log scales and show signals averaged to reduce noise. The change in intensity of the transmitted (beige) signal from a to b is solely due to the addition of the ruby. The region of particular interest is immediately following the bright to dark transition and is shown enlarged in c. When the ruby is added into the system, some of the energy is delayed, increasing the intensity of the trailing edge of the pulse (tail) of the transmitted signal, as highlighted by Fig. 3.3c. This increase in the intensity in, and hence delay of optical energy into, the tail of the pulse, indicates that the slow-light mechanism in ruby is more complicated than described by a simple time-dependent (saturable) absorber.

I investigated various potential systematic errors. I use identical equipment in both the reference and ruby arms, including detectors, amplifiers, and gain settings. I tested all equipment in both data collection arms, and two data acquisition devices

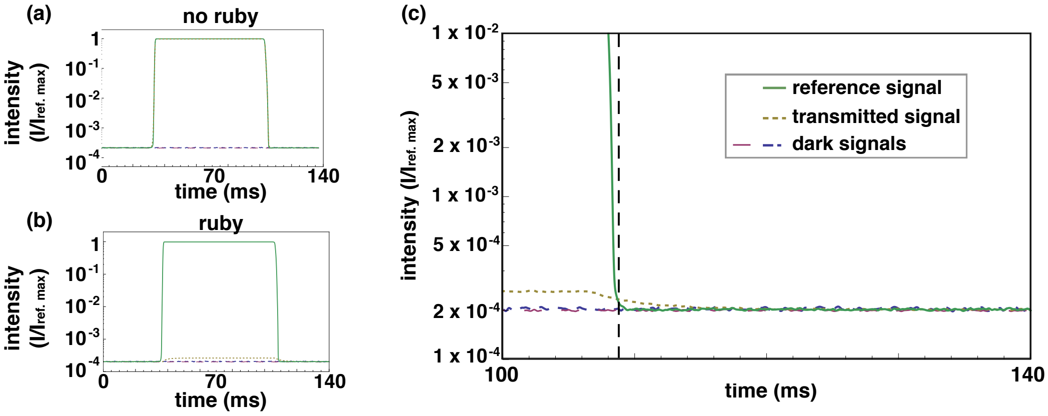

FIGURE 3.3: Time evolution of the reference and transmitted signals shown on semi-log plots. (*a*) When the ruby was removed, the transmitted (beige, dotted) and reference (green, solid) signals have the same intensity profiles. (*b*) When the ruby was added, the overall intensity of the transmitted signal (beige, dotted) decreases, although the intensity at the trailing edge of the pulse increases above that of the reference signal (green, solid). The region of interest in *b* is shown in detail in *c*. For times greater than that marked by the dashed vertical line in *c*, the intensity of the transmitted signal is greater than that of the reference pulse. Dark signals taken with the laser off are shown for the reference (pink, large dashes) and transmitted (blue, dash dots) arms in *a*, *b*, and *c*.

were used, with multiple channels tested on each. For all of these variables, I observed the same trend; adding the ruby delays energy from the pulse, causing an increase in intensity in the tail of the pulse. Other experimental parameters were also investigated. I replaced the 90-mm-long ruby with a 6-mm-long ruby, which also showed an increase in the energy in the tail of the ruby pulse, albeit to a smaller degree. The existence of a pronounced tail of the ruby pulse is observable for both linearly and circularly polarised light.

After investigating the falling edge of the square wave pulse, I investigated how the system would react to a very narrow dark region in the centre of the high-intensity pulse, similar to the black line down the centre of the pulse in Chapter 2. As seen in Fig. 3.5, energy was delayed into the narrow dark region in the same manner as it was delayed at the falling edge of the main pulse.

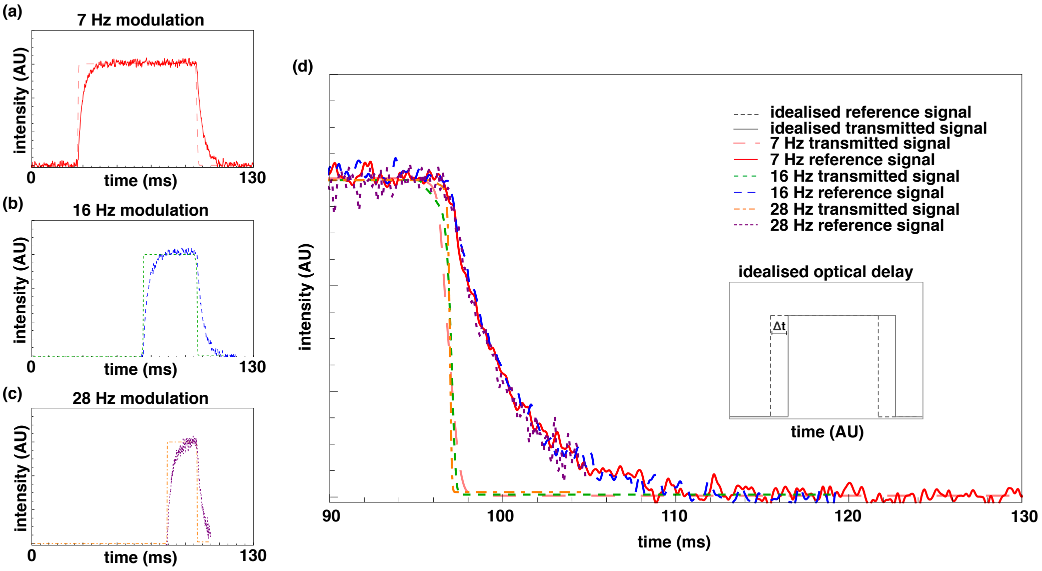

FIGURE 3.4: In an ideal slow-light medium, a square-top pulse would simply be delayed (shifted) in time, as shown in the insert in *d*. In contrast, the results I observe are shown for a square wave with a modulation of (*a*) 7 Hz, (*b*) 16 Hz, and (*c*) 28 Hz. Unlike Fig. 3.3, all pulse shapes in this figure have been normalised to have the same peak intensity, to illustrate the difference in pulse shapes between the reference and transmitted signals. Part *d* shows all three traces overlaid temporally, where the falling edge of each reference pulse is aligned.

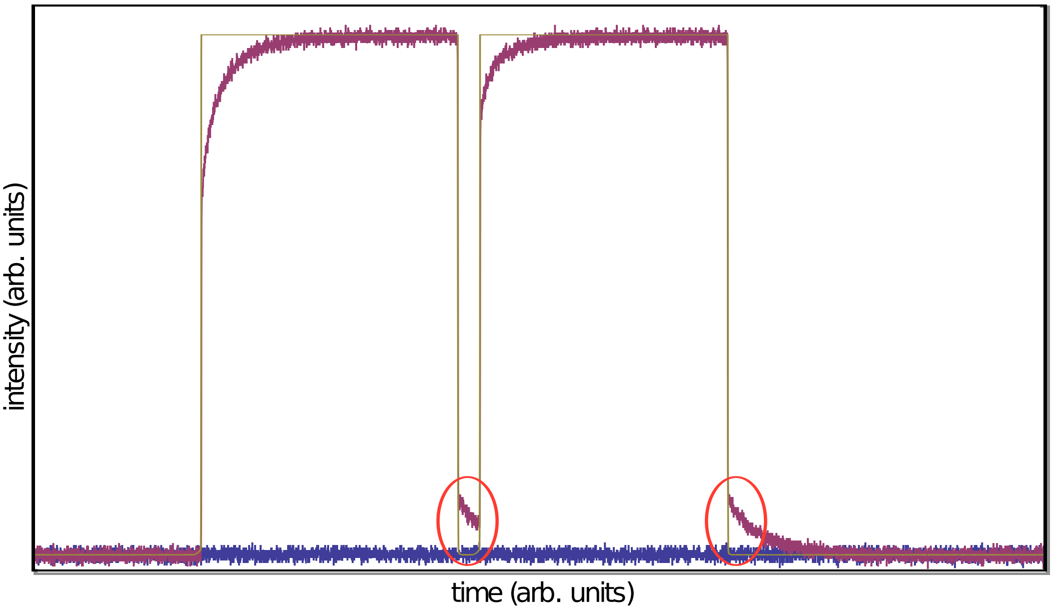

FIGURE 3.5: Trace of square wave pulse with thin "black line" embedded in the brightest part of the signal. Energy is delayed into the black line in the same way it is delayed into the trailing edge of the pulse, as shown in the regions circled in red.

3.5 Pulse Distortion and Harmonic Components

One might expect the delayed pulse to look exactly like the reference pulse with a simple shift in time, Δt, as illustrated by the insert in Fig. 3.4*d*. However, the time delay of a signal depends on the Fourier components of which the intensity signal is comprised. As stated previously, Bigelow *et al.* [58] reported that laser beams modulated with sine waves of different frequencies have different time delays through ruby. More specifically, the higher the frequency of the sine wave, the smaller the time delay the signal experienced, a fact that I verified (see Fig. 3.1). By taking the Fourier transform of the square-wave intensity signal, one can see the sinusoidal waves of many different frequencies that form the Fourier components of this signal. Measuring the difference in phase of the Fourier components of the reference and ruby signals allows us to observe the time delays as a function of frequency of the Fourier components. As shown in Fig. 3.6, the time delay of the Fourier components follow the same exponential decay trend with increasing component frequency, regardless of the modulation frequency of the square wave. Lower frequencies are delayed more than higher frequencies, which makes it evident that the square wave will be distorted when delayed by a slow-light medium. Figure 3.4 depicts the observed shape of the delayed square wave, as measured under the same conditions as the data presented in Figs. 3.3 and 3.6.

Figure 3.6 shows that the delays of individual Fourier components are independent of the modulation frequency of the signal. In other words, the shape of the tail should be independent of the modulation frequency of the pulse. Representative pulse traces taken at 7 Hz, 16 Hz, and 28 Hz modulations are shown in Fig. 3.4*a*-*c*. Figure 3.4*d* overlays the three traces so that the end of the traces coincide temporally. As can be seen in Fig. 3.4*d*, although the three pulses have different modulation frequencies, the traces have tails with the same shape, which is consistent with the delays of Fourier components all following the same trend in Fig. 3.6.

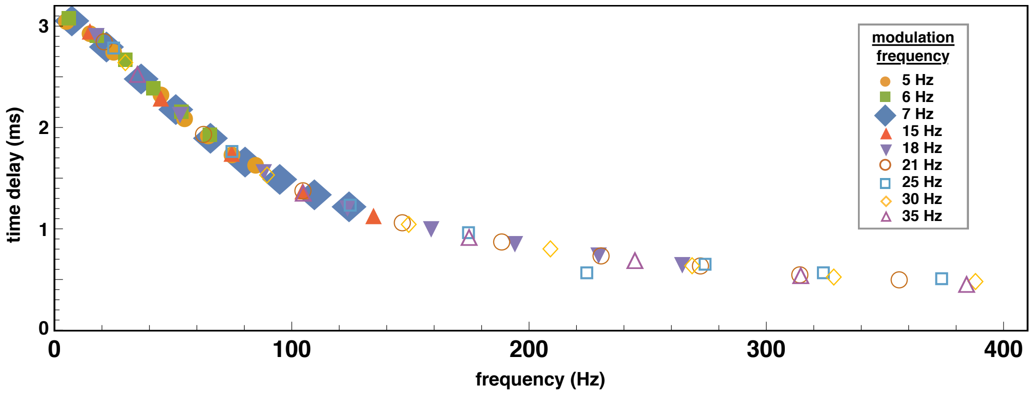

FIGURE 3.6: Time delay plotted against the frequency of the Fourier component for square-wave signals of different modulation frequencies. The delays of the Fourier components from the 7 Hz-modulated square wave are marked with large filled diamonds. The trace of a corresponding 7 Hz-modulated pulse is shown in Fig. 3.4*a*. Delays of Fourier components from square waves modulated at different frequencies all follow the same curve. Because different frequency components experience different time delays, the square wave is distorted upon transmission through ruby, resulting in the transmitted traces in Fig. 3.4.

3.6 Conclusions

I fit the tail of both the ruby and the reference pulses to exponential decays to find the decay time of each. The reference tail has a decay time of approximately $\tau = 0.1$ ms, probably arising from the finite bandwidth of the detector and the associated electronics. By contrast, when the ruby is in place, the tail has a decay time of approximately $\tau = 3.0$ ms. This increased decay time resembles the upper state lifetime of the trap level in ruby, which is approximately 3.4 ms at room temperature [92].

Through careful control of the experimental parameters, I have shown the existence of a pronounced tail on the trailing edge of the transmitted signal, due to the light pulse being slowed as it propagates through the ruby, which is not compatible with a simple model of pulse delay in a time-varying (saturable) absorber. Instead, this experimental evidence supports a more complicated model of slow light in ruby that results in a delay of the transmitted optical energy and a distortion of the pulse shape, as individual Fourier components of the signals are delayed by different amounts.

Chapter 4

Orbital Angular Momentum Introduction

4.1 Historical Introduction

For more than a century, scientists have known that light carries momentum. In 1909, Poynting stated that circularly polarised light should have angular momentum of $\sigma\hbar$ per photon, where $\hbar$ is Plank's constant (h) divided by 2π and $\sigma = \pm 1$ for right- and left-circularly polarised light respectively [93]. This angular momentum became known as spin angular momentum (SAM). Poynting reasoned that the effects of SAM were very small and would be extremely difficult to detect experimentally, but in 1936 Beth performed an experiment in which he could observe the physical effects of SAM. Beth suspended a waveplate from a quartz fiber, and as the circularly polarised light passed through the waveplate, the handedness of the light was switched. The quartz fiber gave a low friction suspension so that the angular momentum transferred from the light to the waveplate by the reversal of the handedness of the circular polarisation caused the waveplate to deflect [94].

While the physical effects of angular momentum can be difficult (although not impossible) to see, it is simple to determine the interactions between light and atoms. The angular momentum of an atom is changed by $\pm\hbar$ when it emits a circularly polarised photon [95], however a higher order transition, such as a quadrupole transition, results in an angular momentum transfer of a higher integer multiple of $\hbar$. Circularly polarised light, and thus SAM, alone cannot account for this; therefore another type of angular momentum is required. In 1932, Darwin (the grandson of the famous evolutionary scientist) described how a photon emitted a short distance from the centre of mass of an atom would provide the additional angular momentum needed for higher order transitions [96]. This angular momentum eventually become known as orbital angular momentum (OAM).

In 1992, Allen *et al.* released a seminal article that described OAM as a natural property of light beams with helical phase fronts and noted that these beams could be generated quite easily in the laboratory [95]. Orbital angular momentum-carrying beams often take the form of (but are not limited to) Laguerre-Gaussian (LG) beams, with amplitude distributions, LG_{pl} described as

$$\begin{aligned} LG_{pl} = \sqrt{\frac{2p!}{\pi(p+|l|)!}}\frac{1}{w(z)}\left[\frac{r\sqrt{2}}{w(z)}\right]^{|l|} \exp\left[\frac{-r^2}{w^2(z)}\right] L_p^{|l|}\left(\frac{2r^2}{w^2(z)}\right)\exp\left[il\theta\right] \\ \exp\left[\frac{ik_0 r^2 z}{2(z^2+z_R^2)}\right]\exp\left[-i(2p+|l|+1)\tan^{-1}\left(\frac{z}{z_R}\right)\right], \end{aligned} \quad (4.1)$$

where r is the radius, $w(z) = w(0)[(z^2+z_R^2)/z_R^2]^{1/2}$ with $w(0)$ as the beam waist, z_R as the Rayleigh range, and $(2p+|l|+1)\tan^{-1}(z/z_R)$ as the Gouy phase [97]. $L_p^{|l|}$ is a Laguerre polynomial that is a result of

$$L_p^{|l|}(x) = (-1)^{|l|}\frac{d^{|l|}}{dx^{|l|}}L_{p+|l|}(x), \quad (4.2)$$

where l is the azimuthal index resulting in OAM of $L = l\hbar$ per photon, as seen in Fig. 4.1, and p is the number of radial nodes in the intensity distribution. The simplest OAM beam has a phase in the transverse plane of $\phi(\theta) = l\theta$, where θ is the angular coordinate and l is any integer. Allen *et al.* derived that all light beams with helical phasefronts described by $e^{il\theta}$ have an OAM of $l\hbar$ per photon [95]. The skew angle of the Poynting vector is $\beta = l/k_0 r$. This leads to an azimuthal component of linear momentum of

$$p = \frac{\hbar k_0 l \lambda}{2\pi r} \tag{4.3}$$

per photon, where λ is the wavelength of light and $k_0 = 1/\lambda$ is the wavenumber [98]. Multiplying the linear momentum by the radius vector, gives an angular momentum of $l\hbar$ per photon [99]. In a similar geometrical argument to that presented above, SAM can be shown to be $\hbar$ per photon. The circular path of the circumference λ has a radius of $\lambda/2\pi$, and a linear momentum of $\hbar k_0$ acting around a circle of this circumference results in an angular momentum of $\hbar$.

One ubiquitous feature of beams with helical phase fronts is that they have a phase singularity running down their centre, which is an intensity null surrounded by a $2\pi l$ phase change [100]. The mathematical similarity between helically phased beams and superfluid vortices led to these phase singularities being named optical vortices [101]. These optical vortices do not carry any angular momentum themselves, as they are points of zero energy; all angular momentum in the beam comes from the field surrounding the singularity.

In the 1970s, there were two important developments in the study of phase singularities. First, Nye and Berry investigated phase singularities in ultrasound waves in order to model radio wave echoes off arctic ice sheets [102]. Second, Vaughan and Willets created a phase singularity in a laser beam by combining high-order Hermite-Gaussian modes [103], which are the solutions the paraxial wave equation in Cartesian coordinates. The handedness of these singularity-carrying beams was studied by breaking the cylindrical symmetry of an output beam of a laser [104] and through the recognition of a bistable helical phase [105, 106]. These developments culminated in the realisation of the importance of Laguerre-Gaussian modes, which are the solutions of the paraxial wave equation in cylindrical coordinates, as described by D'Alessandro and Oppo [107]. Helical phase fronts are not unique to Laguerre-Gaussian beams; Bessel beams [108], Mathieu beams [109], and Ince-Gaussian beams [110] also have helical phase fronts. All beams with

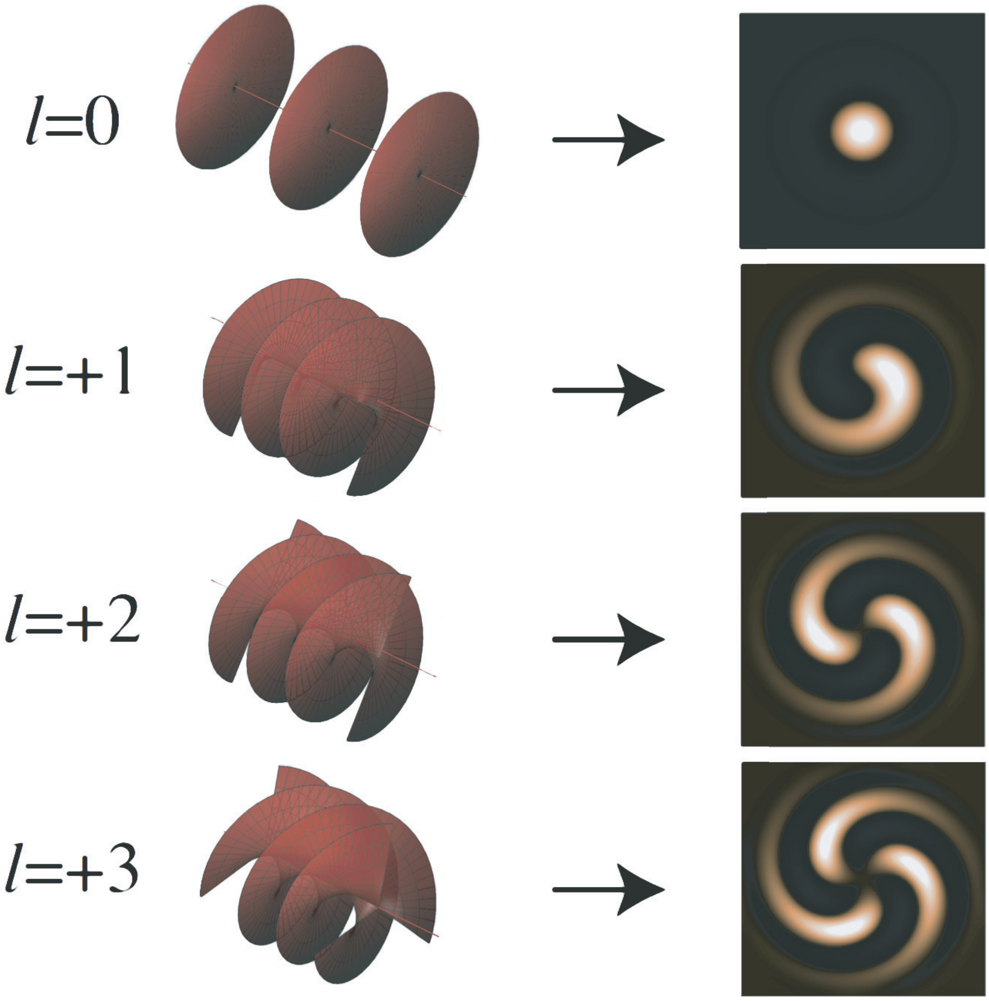

Figure 4.1: Helical phase fronts interfered with plane waves result in spiral fringes, shown here for $l = 0$ through $l = +3$.

helical phase fronts give rise to spiral interference fringes when interfered with a plane wave [111–113], shown in Fig. 4.1.

Orbital angular momentum can lead to quite simple descriptions of situations that are rotationally symmetric or contain singularities. One instance where OAM is useful in understanding a broad concept is in the angular uncertainty relationship. Similar to the more widely recognised Heisenberg uncertainty relationship, an uncertainty relationship for angular momentum can be described quite simply. If a beam did not have an angular restriction from an aperture, the full cyclic nature of the azimuthal phase could be observed and would result in a single-valued OAM state with no uncertainty in its value. However, light restricted to a certain azimuthal range by an aperture is described by a spread of OAM states. The relation between the width of the aperture and the spread of OAM values has been studied extensively [114–116] and is an example of an uncertainty relationship. For apertures of width $\Delta\theta$, the spread in OAM states, ΔL, is described by the angular uncertainty relationship:

$$\Delta\theta\Delta L \geq \frac{1}{2}\hbar|1 - 2\pi P(\theta)|, \tag{4.4}$$

where $P(\theta)$ is the angular probability density of the boundary of the angular range and takes the periodicity of the angle into account [117]. For a narrow Gaussian aperture, this relationship simplifies to $\Delta\theta\Delta L > \hbar/2$.

4.2 Creating Beams with Orbital Angular Momentum

There are many ways to create beams that have orbital angular momentum. One of the earliest methods of creating OAM-carrying beams is to use a pair of cylindrical lenses that transform a Hermite-Gaussian beam to a Laguerre-Gaussian beam. This kind of mode transformation was the method used by Allen *et al.* in 1992 [95]. Hermite-Gaussian and Laguerre-Gaussian modes are both complete sets, which means that a mode of either set can be described as a superposition of modes of the other. As the beam is focused in the cylindrical lens pair, it acquires an additional Gouy phase determined by the orientation and mode indices of the HG mode used, which will cause the HG mode to rephase into a specific LG mode [118]. There are two varieties of cylindrical lens mode converters, both of which are shown in Fig. 4.2; a $\pi/2$ converter will transform HG modes into LG modes when there's a beam waist midway between the cylindrical lenses, whereas a π converter will transform any mode of collimated light into its own mirror image. In this way, a π converter is optically equivalent to a Dove prism [119]. These $\pi/2$ and π mode-converters are mathematically analogous to polarisation-converting quarter-wave plates and half-wave plates, respectively. Figure 4.3 shows the effects of $\pi/2$ and π mode converters and their similarities to quarter- and half-wave plates.

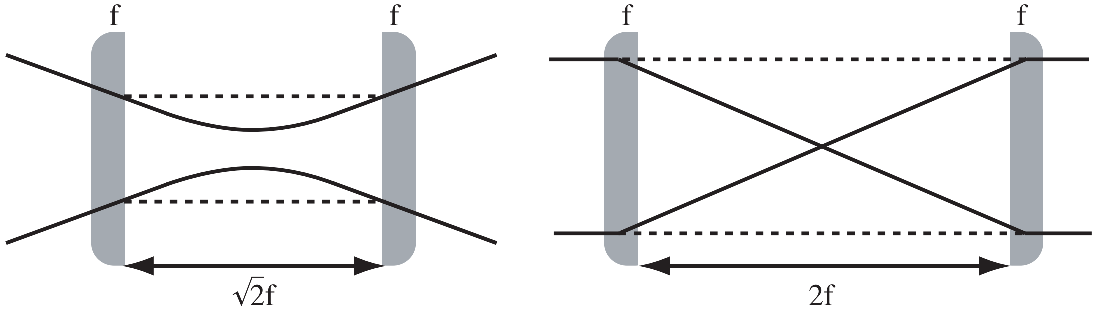

FIGURE 4.2: (*left*) $\pi/2$ converters switch between HG and LG modes, whereas (*right*) π converters switch modes to their mirror images.

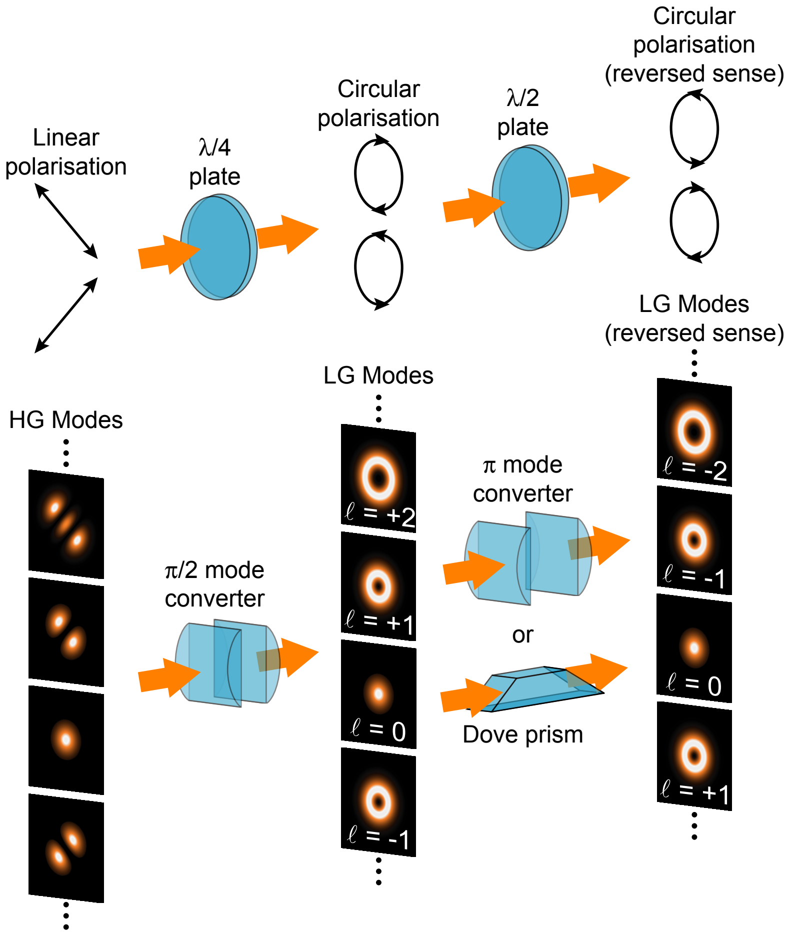
Linear polarisation
λ/4 plate
Circular polarisation
λ/2 plate
Circular polarisation (reversed sense)
HG Modes
π/2 mode converter
LG Modes
ℓ = +2
ℓ = +1
ℓ = 0
ℓ = -1
π mode converter
or
Dove prism
LG Modes (reversed sense)
ℓ = -2
ℓ = -1
ℓ = 0
ℓ = +1

FIGURE 4.3: The effects of (*top*) wave plates and (*bottom*) mode converters. Quarter-wave plates convert from linear to circular polarisation in an analogous way to how $\pi/2$ mode converters convert HG modes to LG mods. Half-wave plates reverse the direction of the polarisation analogously to how π mode converters (or dove prisms) switch the direction of the OAM.

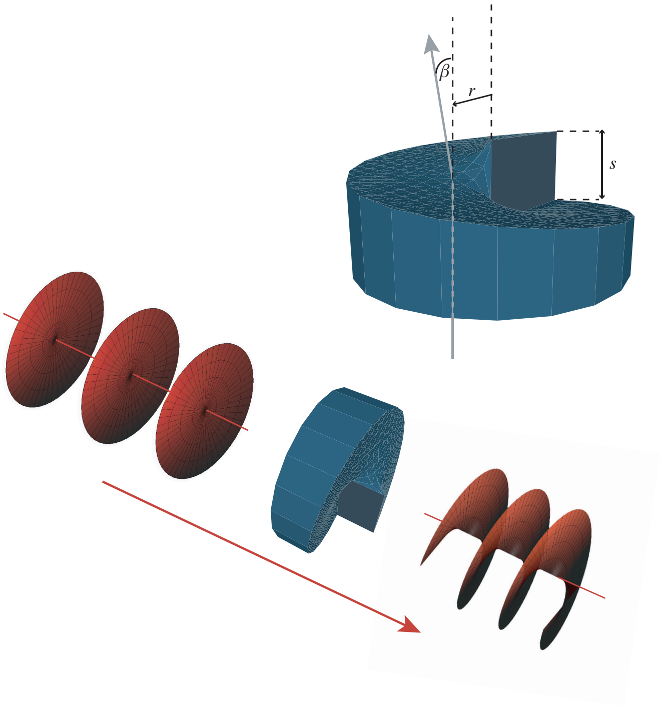

FIGURE 4.4: A plane wave gains a helical phase front when shone through a spiral phase plate.

Instead of using cylindrical lenses to convert between HG and LG modes, spiral phase plates can be used to change plane waves into helically phased waves. Spiral phase plates are transmissive plates that increase in thickness with azimuthal angle with a discontinuity step of one wavelength, such that an incoming plane wave exits the plate with a helical phase front. An example of this is shown in Fig. 4.4. The required step height, s, of a spiral phase plate for light with a wavelength of λ is

$$s = \frac{l\lambda}{(n_\phi - 1)}, \tag{4.5}$$

where n_ϕ is the refractive index of the medium from which the plate is made. Applying Snell's law leads to the angular change in the transmitted ray, or skew angle, being $\beta = l/k_0 r$, as stated previously.

The required level of precision of the spiral phase plate's thickness is very difficult to obtain for light beams, and so to circumvent these difficulties, Beijersbergen *et al.* placed a spiral phase plate in a fluid bath. The temperature of the bath was tuned to achieve the proper index of refraction and thus the difference in thickness, or step height, for a certain wavelength [120]. More recently, micro

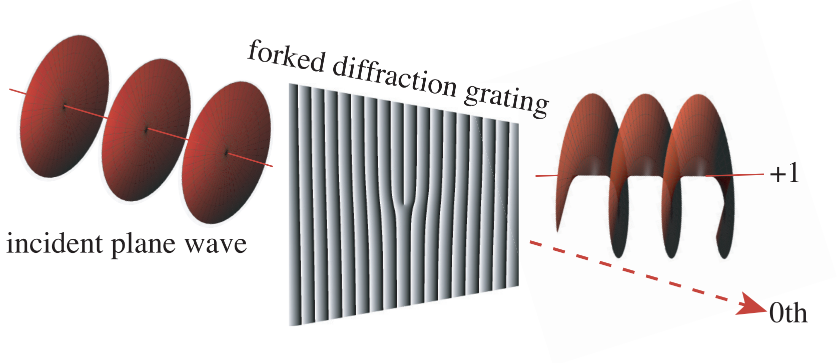

Figure 4.5: Plane waves gain a helical phase front when shone through a forked diffraction grating, resulting in 1st and -1st order diffracted beams. The 1st order diffracted beam is shown here with $l = +1$.

fabrication techniques have allowed for direct construction of millimeter wave [121] and optical spiral phase plates [122, 123]. However, individual physical spiral phase plates have to be constructed for each wavelength of light being investigated, due to the relation of the wavelength to the necessary thickness of the spiral phase plate.

Holographic alternatives to physical spiral phase plates are created through the use of diffractive optical elements (DOE's). Holograms can be used to shape a beam's intensity or phase. The advantage of phase-only holograms is that nearly all the incident light can be directed with the desired phase profile into the first diffracted order. By adding the desired phase profile modulo 2π to a linear phase ramp, any phase profile can be generated in the first diffracted order, shown in Fig. 4.5. In the case of a helical beam, this hologram is a diffraction grating with an l-pronged dislocation at the centre [124–126], seen in Fig. 4.5. Spatial light modulators (SLMs) have been developed and refined over the past decade and have greatly increased the ease with which holograms can be implemented. Spatial light modulators use liquid crystal controlled by a video interface on a computer to act as holograms that can be easily changed in real time. SLMs have become widely used due to the ease of creating and changing the displayed hologram using grayscale images, the circumvention of the manufacturing problems associated with making spiral phase plates, and their ability to work for a wide range of wavelengths within the visible spectrum.

As an alternative to diffractive optics, OAM-carrying beams can be made with Q-plates, which are birefringent liquid crystal plates with spatially varying optical axes[127]. Q-plates have an azimuthal spatial dependence and thus are able to transform between spin and orbital angular momenta by taking a circularly polarised beam and adding an azimuthally-dependent phase that gives the transmitted beam a helical phase. This process also makes Q-plates useful for creating entanglement between spin and orbital angular momentum [128].

Phase singularities can also occur naturally, without the use of spiral phase plates, DOE's, or Q-plates. Simply interfering three or more plane wave components of similar intensity results in many vortices in any field cross-section [129, 130]. A common example of this is laser speckle, where each black speck is an optical vortex [131].

While the previously-described methods of obtaining helically phased beams all use spatially coherent light, spatial and temporal coherence are not fundamental requirements for creating OAM. The description of a helical phasefront implies a coherence of phase across the beam, however, it is possible to illuminate a spiral phase plate or forked diffraction grating with a spatially incoherent source. If light rays are projected through the centre of a spiral phase plate at small radii, as shown in Fig. 4.6, incoherent vortices are generated in the far field; their time-averaged intensity is not zero [132]. Therefore, a vortex beam can be used to determine spatial coherence. A vortex beam with zero on-axis intensity implies spatial coherence, whereas nonzero on-axis intensity indicates spatial incoherence.

One can see by comparing the energy, momentum, and angular momentum of a photon that linear momentum per unit energy is inversely proportional to phase velocity and that angular momentum is inversely proportional to frequency. These relationships are not exclusive to light waves but do in fact apply to all forms of waves. It follows, then, that orbital angular momentum is not only a property of light. Early rotational Doppler shift studies were conducted at millimeter wave frequencies, as experiments at these wavelengths were easier to align [114, 121]. Radio waves carrying OAM have also been proposed for applications in astronomy

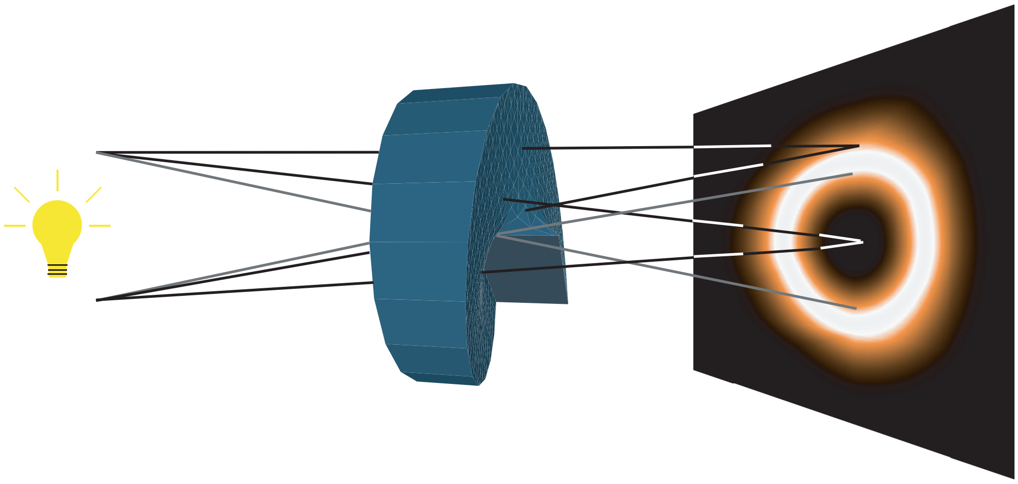

Figure 4.6: A spiral phase plate illuminated by incoherent light produces an incoherent vortex in the far field.

and radio frequency communications, using arrays of antennae for detection [133]. Coherent arrays of detectors or emitters act on longer wavelengths much like a spatial light modulator does for optical frequencies. A cyclical particle accelerator along with the correct diffractive optics results in an X-ray vortex, which could be used for determining molecular structure in x-ray diffraction [134]. Orbital angular momentum can also be created in electron beams. Putting a spiral phase plate into an electron microscope can increase edge detection in low-contrast imaging [135], which has led to a new holographic reconstruction technique in a transmission electron microscope (TEM) [136].

4.3 Micro-Manipulation Through the Use of OAM

Orbital angular momentum has been observed in several experiments and has even become a method of manipulating physical objects. One way of using the orbital angular momentum of light to manipulate particles was discovered through the development of optical tweezers. Optical tweezers use highly focused laser radiation to trap and move micron-sized dielectric particles [137], which are often made of silica or polystyrene. Introducing OAM-carrying beams with the use of

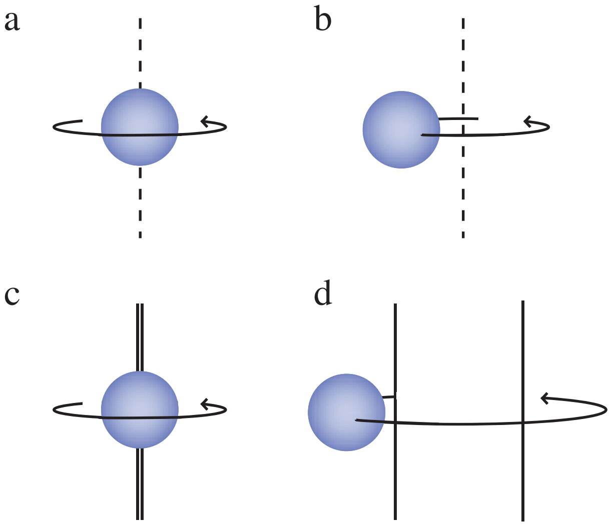

Figure 4.7: Different rotating conditions for spin and orbital angular momentum. (*a*) SAM makes the particle spin about its own axis, whereas (*b*) OAM orbits the particle about the beam axis. (*c*) For a beam much smaller than the particle, SAM and OAM will both make the particle spin in the same way. (*d*) For a large beam, OAM will make the particle orbit the beam.

SLMs allows us to expand the functionality of optical tweezers, enabling them to act as optical spanners [138]. Optical spanners work in a very similar way to optical tweezers, with the added cability to rotate the trapped object. For example, for a beam with $l = 1$ and $\sigma = \pm 1$, when the sign of the SAM is positive, the particle will spin. When the sign of the SAM is negative, the particle will stop [139]. If the centre of mass of the particle is on the beam axis, then SAM and OAM will both spin the particle about its axis, although they can be in the same or in opposite directions. If the centre of mass is off-axis, SAM and OAM will act differently, with the SAM spinning the particle about its own axis and the OAM spinning the particle in an orbit [140]. Spin angular momentum is intrinsic, as its rotation is independent of the choice of axis used to calculate it. Orbital angular momentum is said to be quasi-intrinsic when the centre of mass is on-axis, because in that case it spins about both its axis and the beam axis, which happen to be the same. However, when the centre of mass is off-axis, it is said to be extrinsic, or dependent upon the axis about which it is calculated [140]. These situations are illustrated in Fig. 4.7.

As the phase of an OAM-carrying beam changes with the spatial position in the beam, mechanical effects of SAM and OAM depend on the size of the beam with respect to the size of the particle. If the beam size is larger than the particle size, then a beam carrying SAM will make the particle spin about its own axis and one carrying OAM will make the particle rotate around the beam. If the beam size is much smaller than the particle size, SAM and OAM will both make the particle spin the same way, as a rotation around a beam of that size does not move the particle and instead just spins it as if the beam were the rotation axis, also shown in Fig. 4.7.

When light with SAM and/or OAM is spinning a particle of radius r, the light can exert a maximum torque of the order $\hbar k_0 r$ on the particle[141]. This torque can be used to create micromachines for the manipulation of particles. The sign of the total angular momentum (spin angular momentum and orbital angular momentum combined) determines the direction of the particle's rotation [142]. One benefit of using angular momentum as a micromachine is that the "axle" is optical and hence will never wear out. One of the most common type of micromachines is the micropump, which can be made by an array of beams with OAM that causes a circulation of microparticles. The fluid flows through the array and carries other particles with it [143]. Another way to design a micropump using angular momentum is to transfer SAM to two birefringent particles to spin them in opposite directions. This leads to a flow along the channel between them[144].

Spin angular momentum and orbital angular momentum interact with cold atoms in distinct ways. Circular polarisation and SAM play a role in atomic selection rules, whereas OAM does not, as a helical phase front is locally indistinguishable from an inclined phase front. However, absorption of a photon from a plane wave by a gas will lead to isotropic spontaneous emission with the atoms or molecules recoiling away from the incident light. For a helically phased beam with a force acting at any distance from the beam axis, there will be a torque continually acting on the centre of mass of the atom such that the repeated recoil will guide the atom in a spiral [145]. As the atoms travel around the beam axis, this torque is seen as a manifestation of the rotational Doppler shift [116], described later. The annular intensity of a helically phased beam with the region of darkness in the centre allows for blue-detuned atom trapping, where there is a force attracting the particles to the centre of the beam [146]. This blue-detuned atom trapping acts in a similar fashion to optical tweezers for low index particles.

4.4 Beam Transformations

Spin and orbital angular momentum are usually separable, as they are typically independent when traveling through a transparent, homogeneous medium with no spatially dependent optical properties. However, transfer of spin angular momentum to orbital angular momentum can happen in such a medium if a Gaussian beam [147] or a circularly polarised vortex beam [148] is tightly focused through a high numerical-aperture lens. This occurs in the optical spanners described above. Although SAM and OAM can interchange in this way, the mechanisms for transferring SAM and OAM are not the same. Spin angular momentum affects particles that are birefringent. Birefringent materials change circular polarisation to linear polarisation, causing the SAM to change and therefore the birefringent particle to spin [149]. Birefringence does not affect helical phase fronts and thus does not interact with the orbital angular momentum at all. Astigmatism, however, does change the helical phase front, which means that it affects the orbital angular momentum of the beam, while leaving polarisation and spin angular momentum unchanged.

There are two basic types of transformations. The first type of transformation is a coordinate transformation, which is equivalent to a rotation of the state about the optical axis and results in no change in angular momentum. For polarisation, this type of rotation about the optical axis could arise from optical activity, Faraday-type effects, or a phase difference between the circularly polarised states. Optically active materials do not alter the OAM state of the transmitted light [150, 151]. Image rotation is the equivalent transformation for an OAM-carrying beam to that of optical activity for an SAM-carrying beam[152]. This second type of transformation is a rephasing of linear states. This rephasing of the linear states does change the angular momentum, which for polarisation comes from birefringence. For orbital angular momentum, the rotation of the state manifests as a rotation about the optical axis of the mode or the superposition of modes. This rotation corresponds to a changing in phase between positive and negative OAM states. The equivalence between the rotation of polarisation and the rotation of an image has been used recently to investigate photon drag [61].

Photon drag was studied in depth experimentally by Jones in the 1970s. He called it Fresnel or aether drag, and he showed that light passing through a moving medium would be dragged either longitudinally or transversely by the medium [78]. In addition, he showed that the polarisation of light through a spinning disk could be rotated by a small angle, $\Delta\theta_{pol}$, which is defined as

$$\Delta\theta_{pol} = \left(n_g - \frac{1}{n_\phi}\right)\frac{\Omega L}{c}, \tag{4.6}$$

where Ω is the angular velocity of the medium, L is its length, and n_g and n_ϕ are, respectively, the group and phase refractive indices of the light in the medium [79]. Recent experimental work by Franke-Arnold *et al.* has shown that a spinning disk drags the image through the same rotation angle as the polarisation [61],

$$\Delta\theta_{pol} = \Delta\theta_{image}, \tag{4.7}$$

where $\Delta\theta_{image}$ is defined in Eq. (1.10) In other words, photon drag has the same effect on spin and orbital angular momenta. The equivalent rotation of polarisation and image is shown in Fig. 4.8.

As mentioned above, Hermite-Gaussian modes and Laguerre-Gaussian modes are both complete sets, and any mode in one set can be made by a superposition of modes in the other set. Polarisation is a useful analogy for the transformation between HG and LG modes. A Poincaré sphere shows how right- and left-circularly polarised states can be combined to form linearly-polarised states, as seen in Fig. 4.9. Any two-dimensional space can be described by an analogous Bloch sphere [153, 154]. A generalised two-state system in the OAM basis, $|\mathbf{a}\rangle$,

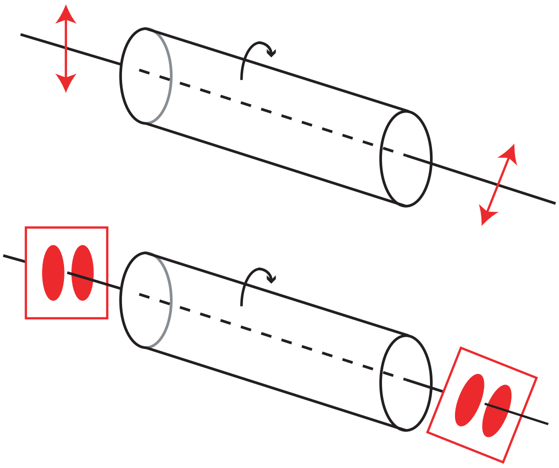

FIGURE 4.8: Equivalent dragging of (*top*) polarisation and (*bottom*) transmitted image through a rotating medium.

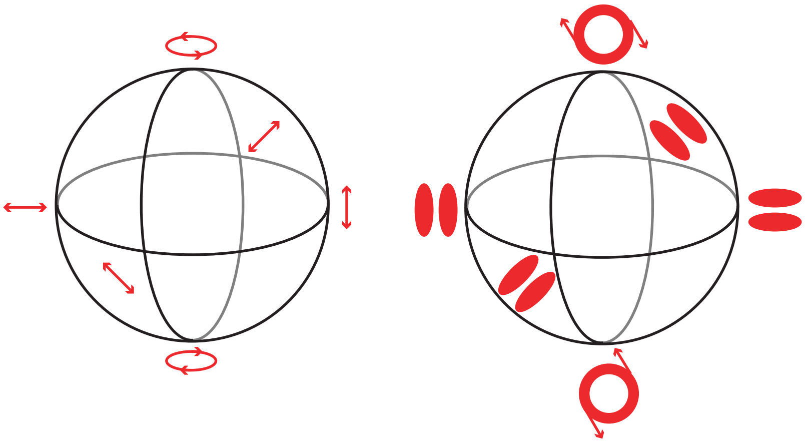

FIGURE 4.9: (*left*) Poincaré sphere for polarisation and (*right*) the analogous Bloch sphere for OAM.

can be written:

$$|\mathbf{a}\rangle = \cos\left(\frac{\theta_a}{2}\right)|l\rangle + e^{i\phi_a}\sin\left(\frac{\theta_a}{2}\right)|-l\rangle, \tag{4.8}$$

where $\mathbf{a} = (\sin(\theta_a)\cos(\phi_a), \sin(\theta_a)\sin(\phi_a), \cos(\theta_a))$ is a vector with longitude $0 \leq \phi_a \leq 2\pi$ and latitude $0 \leq \theta_a \leq \pi$. On the Bloch sphere, the longitudinal position, ϕ, gives the orientation of the mode superposition. Figure 4.9 shows both the Poincaré sphere for polarisation and the analogous Bloch sphere for orbital angular momentum.

It is well known that light emitted from a moving source undergoes a frequency shift, a concept known as the Doppler shift. This effect was seen by Garetz in the 1970s for rotation between a detector and a source of circularly polarised light. Garetz showed this effect by spinning one half-wave plate with respect to another half-wave plate, showing that the polarisation state was rotated at twice the rotational speed of the waveplate, and measuring the corresponding frequency shift [155, 156]. There is a similar effect for a helically phased beam. The single rotation of a helically phased beam changes the field by l cycles and the frequency by $\Delta\omega = l\Omega$[157], and this can be easily implemented by rotating Dove prisms or cylindrical lenses [114]. When light has both spin and orbital angular momentum, a single rotation of the beam about the beam axis changes the phase by $J = (l + \sigma)$ cycles and the frequency by $\Delta\omega = J\Omega$ [158], where J is the total angular momentum of light. The beam cross-section has a J-fold rotation symmetry. Spin angular momentum and orbital angular momentum act similarly when being rotated about the beam axis, so the total angular momentum is the important parameter.

4.5 Measuring Beams with Orbital Angular Momentum

I have discussed how to create a beam having nonzero orbital angular momentum, how to observe the physical effects of OAM, how to manipulate OAM, and how to transform between different modes and between SAM and OAM. I have yet to discuss how to measure OAM. Sometimes it is not enough to simply know that a beam has a helical phase front; there are situations where one wants to quantify the OAM. There are many ways to measure OAM. First is through the use of forked diffraction gratings. As I have discussed previously, to generate an OAM-carrying beam, one can use the collimated output of a single mode fiber to illuminate a forked diffraction grating and get a vortex beam in the far field [124], as shown in Fig 4.5. If instead, one wants to measure a beam with OAM, this process can be reversed. A vortex beam incident on a forked grating will result in a fundamental Gaussian beam only if the OAM of the incoming beam is opposite to that of the diffraction grating. By scanning the diffraction grating through all possible OAM

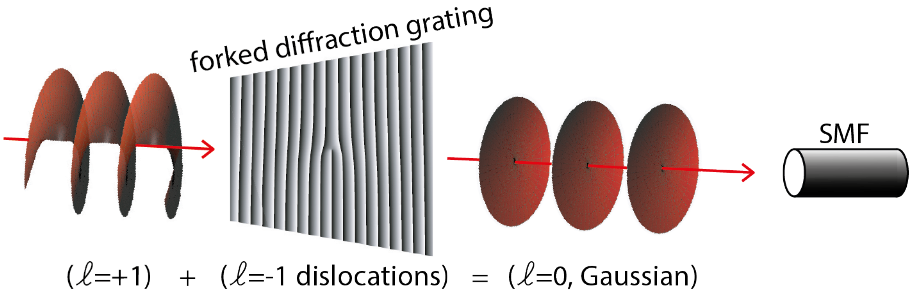
forked diffraction grating
SMF
(ℓ=+1) + (ℓ=-1 dislocations) = (ℓ=0, Gaussian)

FIGURE 4.10: A vortex beam with a helical phase structure is incident upon a forked diffraction grating. When the OAM value of the vortex beam and the diffraction grating are l and $-l$ respectively, the resulting beam will be Gaussian and will successfully couple into a single mode fibre (SMF). An example of this process is shown here for a vortex beam of $l = +1$.

states and detecting which state corresponds to a detected photon after a single mode fiber, the mode of the incident photon can be deduced [159]. An example of this where the vortex beam carries OAM of $l = +1$ is shown in Fig. 4.10. This method works with low levels of light down to single photons. However, when working with single photons, the efficiency of the detection system is only $1/N$, where N is the number of states to be assessed, and so for effective detection of the OAM mode, many photons must be prepared in the same state.

Another method for measuring OAM is to interfere the unknown helically-phased beam with a plane wave. The resulting fringes are spiral in nature, and the number of forked fringes matches the l of the beam [111], shown in Fig. 4.1. Interference can be used to determine OAM for a diffraction pattern from any aperture as well. For single and double slits, there is a displacement due to the helical phase front that is perpendicular to the slit direction and leads to a sharp bend in the otherwise straight fringes. The direction of the bend depends on the handedness of the helicity [159, 160]. More complicated diffraction patterns can also be used to measure OAM. One interesting case is when an LG beam is diffracted through a triangular aperture. The result is an array of $(l+1)(l+2)/2$ spots in a triangular configuration, set at a right angle to the aperture's orientation, with the sign of l determining the direction of the orientation [161]. All of these interferometric methods allow for the measure of both the magnitude and the sign of the OAM.

However, these approaches require many photons, and thus cannot measure the l of single photons.

The rotational frequency shift that was discussed previously can also be used to measure orbital angular momentum, where each value of l results in a frequency sideband [158]. However, in low light it is difficult to measure both the frequency and the beam rotation. If a fixed rotation is used, a static, l-dependent phase shift is created [162]. Dove prisms can be used to make an image rotate, and this has been built into Mach-Zehnder [163] and Sagnac [164] interferometers. This rotational frequency shift method can be used to measure the orbital angular momentum or the total angular momentum [165]. This method requires $N - 1$ interferometers to measure N states, which limits this technique to measuring a small number of states.

A recently-developed method for measuring OAM is through the use of mode sorters. The azimuthal position of the input beam is transformed into a transverse position in the output beam [166, 167]. This takes a helically phased beam and turns it into a transverse phase gradient. The reformatting introduces a phase aberration, but it can be corrected in the output plane. Spatial light modulators have been used to create a system that both reformats the image and also corrects the phase. A lens then focuses the input modes to separate lateral positions, where they can be detected using a detector array [168]. This method is less technically demanding than interferometric approaches and can measure the l of single photons [169], and although not perfect, it is more efficient than other common techniques.

4.6 Orbital Angular Momentum in Classical Imaging

Orbital angular momentum measurement methods can be used in imaging by placing a forked hologram in the Fourier plane of the object [126, 170]. This technique is especially useful for contrast enhancement in optical microscopy. When the phase mask corresponds to a spiral phase plate, the point spread function of the microscope is changed into a helically phased ring, and the phase edges appear as bright lines, a process that is referred to as unidirectional edge enhancement. This technique has been used at the single-photon level in ghost imaging [171], which I will describe later. Varying the technique slightly results in spiral interferometry, where spiral fringes replace the common circular fringes. This allows for the recovery of information that is usually lost by the up/down degeneracy of traditional interferometric techniques. Beyond forked holograms, the same technique can be used in phase contrast and dark field microscopy, or even in combinations of different techniques simultaneously. A spatial light modulator can be used instead of having separate phase masks for each imaging modality[172].

Apart from being used in a microscope, orbital angular momentum can also be exploited in telescopes for the detection of off-axis light. The phase discontinuity caused by a spiral phase plate placed within a telescope blocks on-axis light and allows for fainter off-axis light to be detected [173]. This has proven interesting to astronomers, who hope to attenuate bright stars in order to look for nearby planets whose light would otherwise be eclipsed by the star [174, 175].

4.7 Orbital Angular Momentum in Nonlinear and Quantum Optics

Optical vortices in nonlinear materials are often studied within a Kerr medium. A Kerr medium is one in which the refractive index changes with the intensity of the incident light [176], a third-order nonlinear interaction. If the increase in intensity increases the refractive index, the transmitted beam will self-focus. Conversely, if the increase in intensity decreases the refractive index, the transmitted beam will self-defocus as it passes through the medium. In this situation, a beam carrying OAM, which has an annular intensity, will create a path of high refractive-index material that can guide a second beam through the Kerr medium [177].

Another category of nonlinear processes is second-order nonlinear interactions, where two low-frequency waves, ω_1 and ω_2, exchange energy with ω_3, a higher frequency wave [176]. Energy conservation requires the frequencies to be related by $\omega_1 + \omega_2 = \omega_3$ and momentum conservation requires the momenta to be related by $\mathbf{k_1}+\mathbf{k_2} = \mathbf{k_3}$, where $\mathbf{k_1}$, $\mathbf{k_2}$, and $\mathbf{k_3}$ are the momentum vectors of the three waves and $|\mathbf{k}| = \omega n(\omega)/c$. The conservation of momentum is much more complicated than energy conservation, because the refractive index, $n(w)$ is a function of the frequency, w. The various approaches to modify $n(\omega)$ include temperature and angle tuning of the crystal, and these lead to different phase matching conditions. For example, type I phase matching is where the medium is such that the refractive indices of all three waves are the same.

A specific example of second-order nonlinear interactions is parametric down conversion, where the energy from one high-frequency wave is transferred to two lower-frequency waves. This process is only partially constrained in that the frequencies of ω_1 and ω_2 can vary, so long as they sum to give ω_3, the frequency of the starting wave [178]. The two emitted beams are each spatially incoherent, whilst being phase conjugate with respect to each other, leading to quantum entanglement of the spatial modes [163, 179]. The measurement of one photon in a particular basis collapses the two-photon state and, because of entanglement, immediately gives knowledge about the same basis for the other photon, regardless of the photons' spatial separation.

Orbital angular momentum can be used for the encoding of both classical and quantum information [161]. OAM is potentially useful in quantum cryptography, as it gives access to an unbounded number of states, in contrast to the two states associated with SAM, thus greatly increasing the amount of information that can be transmitted on each photon [180, 181]. Other quantum phenomena explored with OAM-carrying light include hyperentanglement, multiple variables including OAM are entangled simultaneously [182], and hybrid entanglement, where there is entanglement between OAM and a different variable [183, 184].

Orbital angular momentum can be used in various tests of quantum mechanics. As discussed earlier, a Poincaré sphere equivalent exists for OAM, and by drawing direct analogy with polarisation, it is possible to formulate a Bell-type argument for OAM that discounts the local hidden-variable theory[185]. Violation of the Bell inequality has been shown for the case where the phase filter is nonlocal with respect to the object [171]. In a different type of test of quantum mechanics using wedge-shaped apertures and forked holograms, it has also been shown that the EPR paradox also applies to OAM and angular position [186].

With possibilities for carrying an unbounded state space, OAM presents itself as an extremely useful tool to increase information capacity in quantum information. However, a drawback to using OAM is that the decomposition of an OAM-carrying beam into its eigenstates depends on both its transverse and angular alignment. Because it is an extended beam, the phase fronts and the OAM spectrum are sensitive to atmospheric turbulence [186–188]. Despite these sensitivities, OAM multiplexing has recently been shown to result in very high data rates [189].

4.8 Conclusions

Orbital angular momentum has helped scientists think differently about light and has led to a number of discoveries both on the microscopic and macroscopic level. It has proven useful for applications in micromanipulation, imaging, and communication systems. Knowledge and study of OAM has increased with the development of spatial light modulators, which have allowed increased control over both the phase and intensity of light. Orbital angular momentum has also allowed for new demonstrations of quantum entanglement, and the multi-dimensional OAM basis set has led to higher numbers of states for encoding of information. A question that arises is whether information encoded in an OAM basis set would be stable under transformations. After establishing that both bright and dark regions of an image can experience slow-light effects (see Chapter 2), I investigate how an OAM-carrying beam would experience slow-light effects and what parameters would affect the amount by which OAM-carrying beams were slowed.

Chapter 5

Rotating Orbital Angular Momentum

5.1 Introduction

Much of the work in this thesis has involved the rotation of images. Another path to understanding how an image is rotated is by considering orbital angular momentum (OAM). When linearly polarised light is transmitted through a spinning window, the plane of polarisation is rotated. This rotation arises through a phase change that is applied to the circularly polarised states corresponding to the spin angular momentum (SAM). Here I show an analogous effect for OAM, where a differential phase between the positive and negative modes ($\pm\ell$) is observed as a rotation of the transmitted image. For normal materials, this rotation is on the order of a micro radian, but by using a slow-light medium, I show a rotation of a few degrees. I also note that, within the bounds of these experimental parameters, this rotation angle does not exceed the scale of the spatial features in the beam profile.

5.2 Magnetic and Mechanical Faraday Effects

The magnetic Faraday effect is a well-known phenomenon that occurs when a light beam passes through a medium subject to a magnetic field, and it has been studied for over a century [88, 190, 191]. The component of the magnetic field in the direction of the light propagation causes a rotation of the polarisation state of the light. This magneto-optical effect has a mechanical analogy, known as the mechanical Faraday effect, where the rotation of a medium causes a rotation of the polarisation state of the light passing through it, as discussed in previous chapters and in [89, 90, 152]. Polarisation rotation through the mechanical Faraday effect was first observed by Jones in 1976 [79]. In the mechanical Faraday effect, the rotation of linear polarisation arises from a phase shift between the right- and left-circularly polarised light that is superimposed to create linearly polarised light. Polarisation can be represented visually on a Poincaré sphere, as seen in Fig. 4.9a, and is associated with a spin angular momentum of $\pm\hbar$ per photon.

5.3 Creating Orbital Angular Momentum States

As discussed in Chapter 4, in addition to SAM, light carries another form of angular momentum, which is known as orbital angular momentum and has been studied in detail for the last 20 years [95, 176]. Beams carrying OAM, such as Laguerre-Guassian (LG) beams, have a Poynting vector with an azimuthal component that corresponds to phase fronts with a helical structure [95], rather than plane waves. Helical phase fronts can be described by $\exp(i\ell\phi)$ and carry OAM of $\pm\ell\hbar$ per photon, where ℓ is an integer and ϕ is the angular coordinate.

The superposition of right- and left-handed circularly polarised light creates linearly polarised light with an orientation determined by the relative phase between the two beams. In a similar way, the superposition of two LG beams with $\pm\ell$ gives rise to a petal pattern with 2ℓ petals. Interfering positive ℓ and negative $(\ell+1)$ LG beams results in $2\ell+1$ petals, allowing for the creation of odd petal patterns

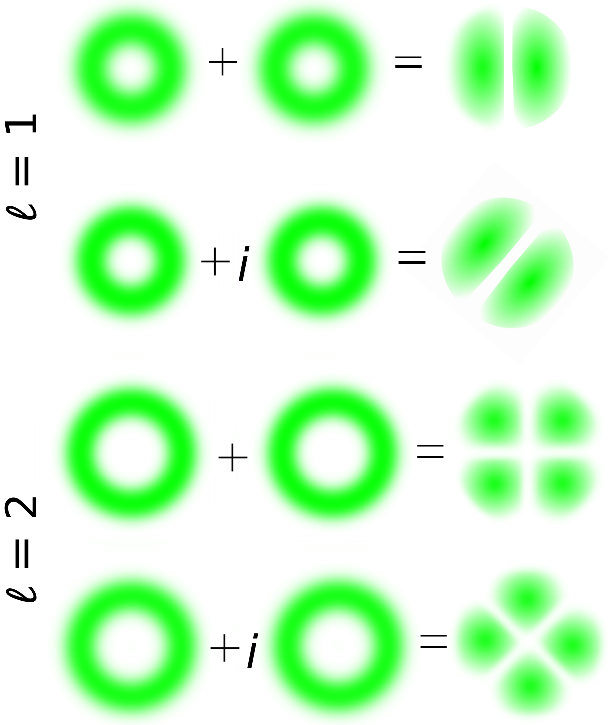

FIGURE 5.1: Petal patterns created by a superposition of $\pm\ell$ beams. The second and fourth lines have a phase shift, i, between the $+\ell$ and $-\ell$ beams, causing a 45 degree rotation of the petal pattern.

as well as even ones. A change in the relative phase between the LG beams causes a rotation of the petal pattern. Analogous to the Poincaré sphere for polarisation is a Bloch sphere, which can be constructed for OAM using LG beams with $\pm\ell$ as the polar states [153]. A Bloch sphere for LG beams with $\ell = \pm 1$ is shown in Fig. 4.9*b*. In Fig. 4.9, the states lying along the equator are Hermite-Gaussian (HG) states with orientations dependent on the relative phase of the $\pm\ell$ polar states. Figure 5.1 shows the superposition of $\ell = \pm 1$ and $\ell = \pm 2$ resulting in two and four petals respectively, with phase shifts being introduced to show the rotation of the petal pattern in the second and fourth cases. The mechanical Faraday effect, which is related to SAM and rotates the polarisation state of the light, therefore also rotates the image and is related to OAM.

5.4 Photon Drag and Slow Light

As discussed earlier, when light enters a translating medium, it is dragged by that medium. By applying well-known formulas for transverse drag in a rotating medium, Padgett *et al.* [90] predicted the angle through which an image should be rotated and found that the polarisation state and the image are dragged through the same angle, which implies some equivalence between SAM and OAM.

Making this reasonable but unproven assumption that SAM and OAM are subject to the same phase change, one can see that the image is dragged by the moving medium through an angle $\Delta\theta_{image}$, as defined in Eq. (1.10). [81, 90]. In most media, the speed of light is much larger than $\Omega L n_g$; the transit time of light through the medium ($T = L/v_g$) is then very short, which leads to a small angle of dragging. To increase the angle of rotation, one can increase the transit time of the light through the medium. As there is a physical limit to the experimentally realistic length and rotational speed of the medium, the practical method for increasing transit time is to use a slow-light medium. When the light passes through a slow-light medium, the transit time increases such that there is a large angle of dragging [58, 59, 61]. The rotation of polarisation was measured in 1976 by Jones [79]. However image rotation proved much harder to detect, with the first observation of image rotation occurring 35 years later by using the enhancing effect of a rotating slow-light medium [61], as discussed in Chapter 1. Slow light is a complicated phenomenon whose mechanism has been the subject of much debate. Whereas previous research has investigated questions about general rotational drag [62, 114, 192], this chapter investigates the rotation of specific superpositions of LG beams by a rotating medium.

5.5 Methods and Materials

In order to examine the effect of a rotating medium on an OAM-carrying beam, an 18 W single-mode Verdi laser creates a beam of intense 532 nm laser light

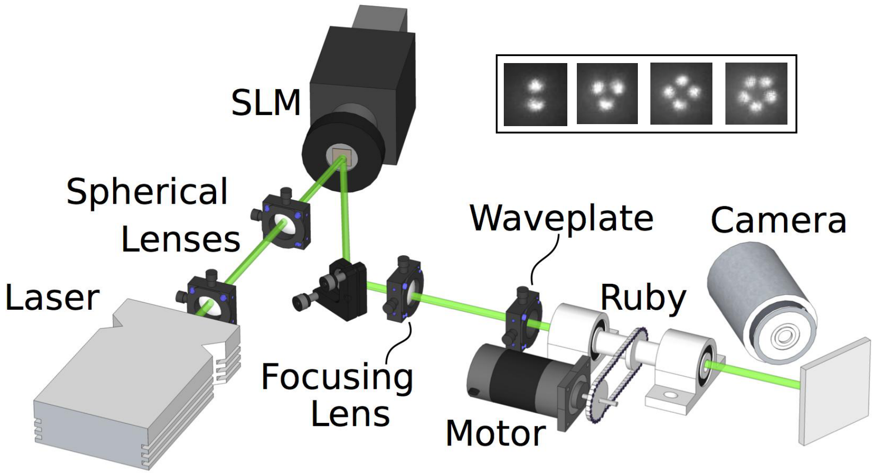

Figure 5.2: 532 nm light passes through two spherical lenses to be expanded before arriving at the spatial light modulator (SLM). The beam is then focused onto the front face of a ruby window, which is spun about its axis by a motor. The light is imaged from the back of the ruby onto a screen and then captured by a camera. Petal patterns are made from superpositions of LG beams with different l values while the ruby window spins at ±19 Hz. Patterns shown in the inset are (from left to right) 2 petals from $l = \pm 1$; 3 petals from $l = (+1, -2)$; 4 petals from $l = \pm 2$; and 5 petals from $l = (+2, -3)$.

with an initial diameter of approximately 4 mm. The beam is expanded by two spherical lenses with focal lengths of 20 mm and 80 mm, and the expanded beam is incident upon a Hamamatsu spatial light modulator (SLM). The SLM creates the desired pattern with careful control of both the phase and intensity of the light. In this experiment, the SLM creates petal patterns from the superposition of two opposite handed LG modes, as shown in Fig. 5.1 for 2 and 4 petals, where the number of petals (N) is given by the difference in the azimuthal mode index of the two beams. In order to interact evenly with the birefringent ruby crystal as it spins, the beam passes through a $\lambda/4$ waveplate to create circularly polarised light. The beam is focused down onto the front face of a standard laser ruby crystal (diameter = 10 mm, length = 90 mm) by a 300 mm focal length spherical lens, as shown in Fig. 5.2.

A stepper motor couples to the ruby mount, and the motor spins the ruby window about its axis up to ≈ ±20 rotations per second. A handheld tachometer found the accuracy of the motor rotation to be 0.05% on the rotation rate in Hz, which causes negligible errors in our experiments, and I use a standard desktop computer to control both the direction and rate of rotation of the ruby. The light is imaged from the back face of the ruby onto a screen, and a camera collects the image to display on a monitor. I use standard National Instruments IMAQ pattern matching software to compare the image under clockwise rotation to that obtained from anticlockwise. This experiment contained the same experimental difficulties described in Chapter 2, including intensity requirement, alignment sensitivity, and image deflection from non-parallel ends of the ruby.

In addition to the IMAQ pattern matching software, I wrote two kinds of pattern matching software. The first would take the template image and rotate it in small increments until it matched the live image pattern, reporting this angle as the rotation angle. This method was relatively accurate, although slightly less accurate than the IMAQ pattern matching software when tested as described below. In the second method, I unwrapped an annular section of the image and measured the intensity of each column of the unwrapped image, as shown in Fig. 2.9. The intensities were plotted and the intensity patterns were matched to sinusoidal signals. As the image rotates, the phase of the sinusoidal figure shifts. By comparing the phase shift, we can deduce the rotation of the image. The experimental results from the second pattern matching software I wrote showed less accurate rotation measurements than other pattern matching softwares tested, and it was therefore discarded.

The built-in National Instruments LabVIEW IMAQ pattern matching software was correct to within 0.1°. I determined this by feeding the programme a static image and then rotating the image manually by a certain angle. By telling the programme to rotate a static image by a certain angle and comparing that angle to the measured rotation given by the pattern matching software, I could determine the accuracy of the software. For example, when rotating the image by 0, 1, 5, 10, 15, 20, and -20 degrees, the LabVIEW pattern matching software measured rotations of 0, 1.03, 5.07, 10.04, 15.09, 20.05, and -20.03 degrees respectively.

The angles of rotation determined by the pattern matching software were slightly affected by the difference in rotation between brighter and dimmer regions of the

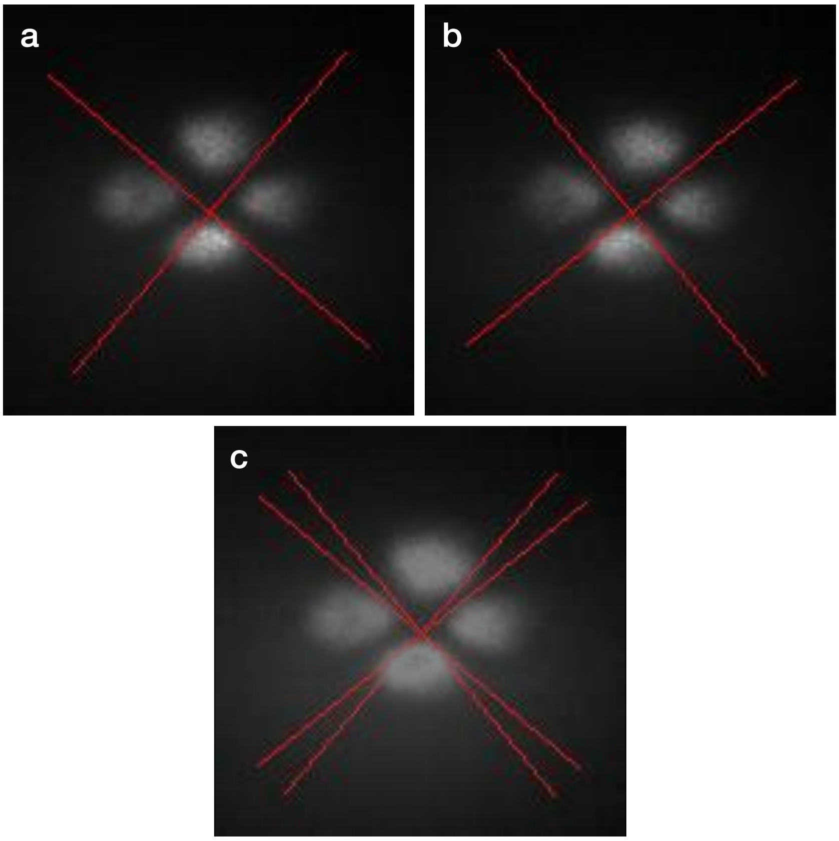

Figure 5.3: Image of pattern with 4 petals, shown rotating (*a*) anticlockwise and (*b*) clockwise. A red X is overlaid on the images to show the orientation of the pattern. (*c*) Images in *a* and *b* are added to show the angle between the orientation of the images while rotating in different directions.

image. The slow-light effect is dependent upon intensity [62], resulting in an S-shaped bending of an elliptical image created from a stretched gaussian beam as shown in Chapter 2. As the bright regions of the OAM pattern experience more slowing, and therefore more rotation, than the dimmer regions, the rotations measured by the pattern matching software may be slight overestimates, if the reported angle is assumed to be the rotation of the complete pattern, not the rotation of the centre of mass of the intensity.

5.6 Results and Discussion

All these reported results relate to intensity patterns formed from the superposition of two beams with OAM indices ℓ_1 and ℓ_2, resulting in rotationally symmetric

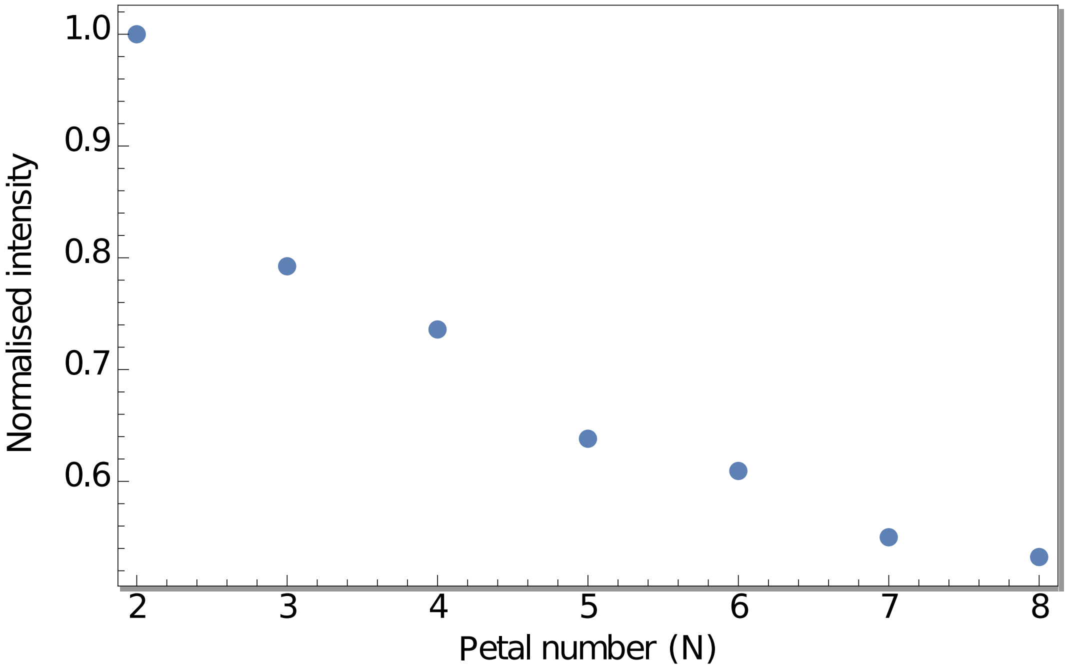

FIGURE 5.4: Normalised peak intensity of petal patterns with $N = 2$ through 8 petals with constant input laser power.

patterns with $N = |\ell_1| + |\ell_2|$ petals, i.e. with N-fold symmetry. In all cases, I report the rotation angle, which is measured by the difference in orientation obtained between clockwise and anticlockwise rotation of the ruby through which the pattern is transmitted.

Because the petal patterns vary in size depending on the value of N, the peak intensity in each pattern varies for constant input laser power, as shown in Fig. 5.4. Intensity has a large effect on the slow-light interaction. Figure 5.5 illustrates the general effect of increasing power on images with different petal numbers. (Data shown in Fig. 5.5 was taken to highlight the need for constant intensities between petal patterns, and as it was collected in one data run, error bars cannot be added.) To overcome the problem of differing peak intensities for patterns with different numbers of petals, the laser power was adjusted between patterns such that the peak intensity was held constant. The laser power was set to 3.90 W, 4.30 W, 4.68 W, and 5.16 W, for N =2, 3, 4, and 5, respectively, to result in constant peak intensities in all four patterns in Fig. 5.6, where the rotation angle of the petal patterns is plotted as a function of rotation speed for various numbers of petals. The experiment was conducted multiple times, resulting in independent

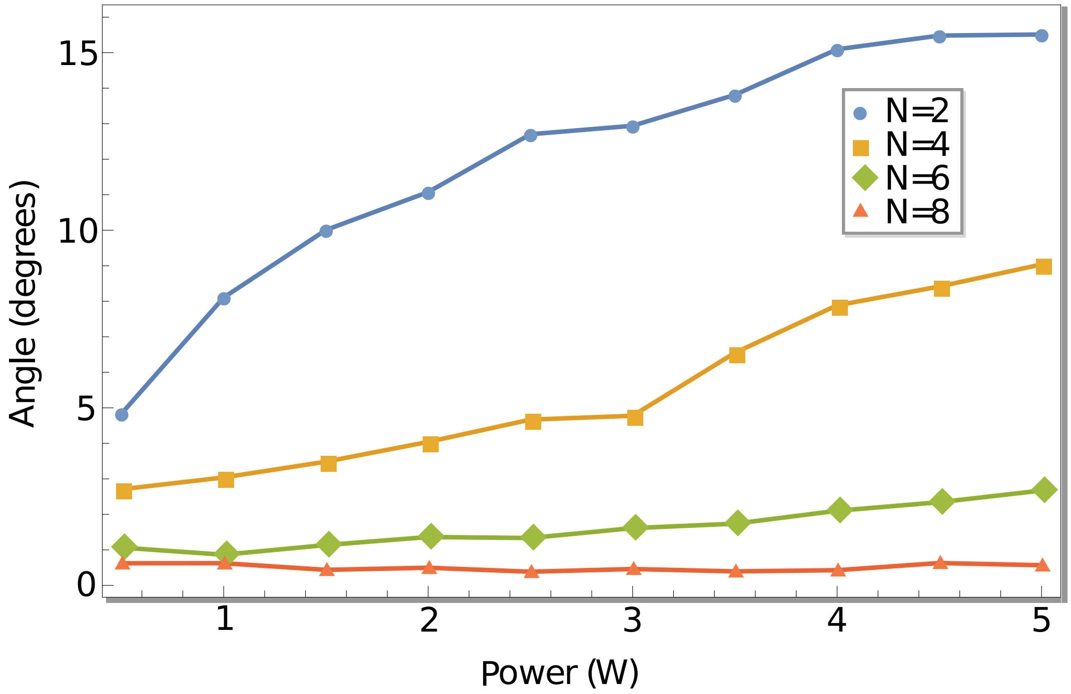

Figure 5.5: Rotation angle of petal patterns with $N = 2$, 4, 6, and 8 petals as a function of input laser power while rotated at a constant speed.

measurements that were averaged to achieve the most accurate rotation angles. Errors were calculated by measuring the standard deviation between independent data collections at each point and are shown in Figs. 5.6 and 5.8. Errors are not constant, due to the minor image distortion upon rotation that affected some points more than others.

Firstly, I note that in all cases an increase in N reduces the rotation angle of the pattern. Secondly, I note that the linear increase in the rotation angle with speed holds only up to a certain point, after which the angle tends to saturate, and that this saturation occurs at lower speeds for higher N. Neither of these trends is in simple agreement with the Jones expression that, by contrast, predicts that the rotation angle increases linearly with rotational speed and should be unchanged by the symmetry of the pattern. For any given setup, Jones assumes a constant n_g, but in this slow-light system, n_g is a complicated function of the experimental parameters, as described below.

These two deviations from the behaviour described by the Jones's equation (Eq. (1.10)) can be understood in terms of the nature of the slow-light effect. In the original

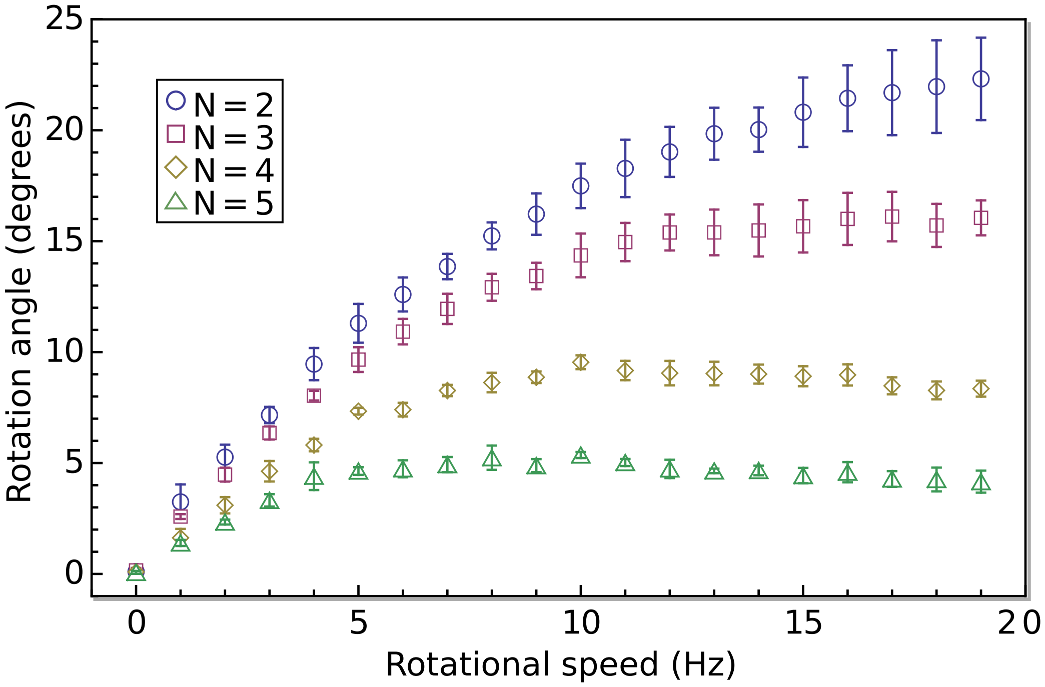

FIGURE 5.6: Rotation angle of petal patterns with $N = 2$ through 5 petals as a function of rotational speed while held at constant peak intensity. Patterns with different N saturate at different rotational speeds. Error bars represent the standard deviations of independent data runs.

slow-light experiments in ruby, the slowing of the light was manifest by the time delay in the sinusoidally modulated intensity of a beam of light [58]. The slowing was largest when the modulation frequency was small compared to the $\approx$ 4 millisecond upper-state lifetime of the ruby. I note that, in this rotational case, the ruby is also subject to a modulated intensity at a frequency equal to the rotation rate of the medium multiplied by the number of petals in the beam. That is to say, each time a given part of the ruby crystal rotates into the position of a bright petal, it will be subject to a high optical intensity, which results in an intensity modulation that depends upon both the rotational frequency of the ruby and N, the petal number. For example, when the ruby is rotating at 10 Hz and is illuminated by a $N = 2$ pattern, each atom in the ruby will experience a fluctuation in the optical intensity at 20 Hz. I note that the change in the rotation speed at which the saturation occurs is consistent with this N-dependent modulation frequency, namely that the saturation frequency scales with the reciprocal of the number of petals. Whereas $N = 2$ saturates at $\approx$ 20 Hz, $N = 3$ saturates at $\approx$ 13

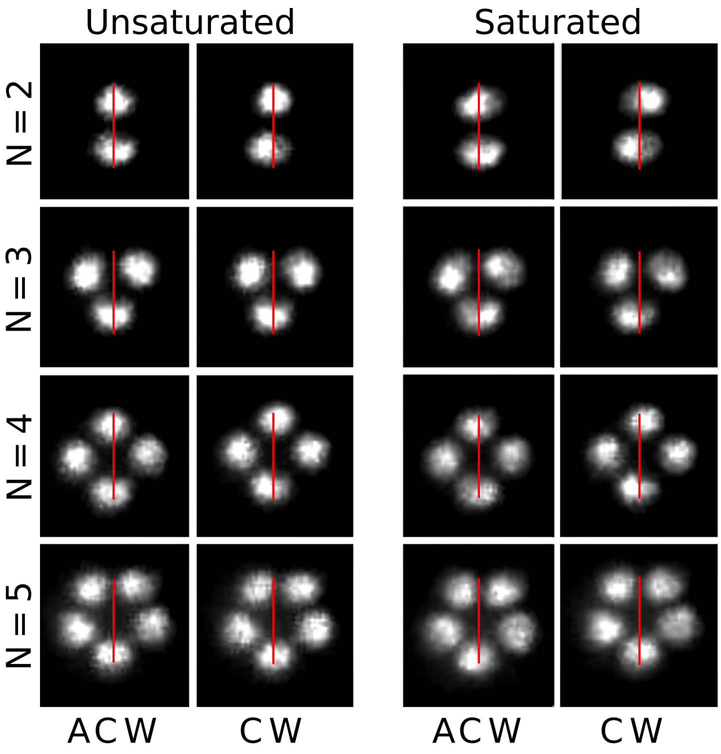

Figure 5.7: Images of N =2 through 5 for unsaturated (low rotational frequency, first column) and saturated (high rotational frequency, second column) modulation frequencies with the ruby spinning anticlockwise (ACW) and clockwise (CW).

Hz, $N = 4$ at ≈ 10 Hz, and $N = 5$ at ≈ 8 Hz, resulting in a modulation frequency of approximately 40 Hz for all four cases. Representative images are shown in Fig. 5.7 for $N = 2$ through 5 in both the unsaturated (low frequency) and saturated (high frequency) modulation while the ruby is rotating anticlockwise and clockwise.

As seen in Fig. 5.6, at a given rotational speed and peak intensity, patterns with more petals rotate through a smaller angle than patterns with fewer petals. Instead of plotting the angle through which the pattern is rotated, I instead plot the fraction of a petal rotated as a function of speed. Figure 5.8 shows that all of the patterns rotate less than 0.15 petals, where a rotation of 1.0 petal would rotate the pattern by 360deg./N, resulting in an identical pattern. The reason for investigating the fraction of a petal rotated by any given pattern is that if

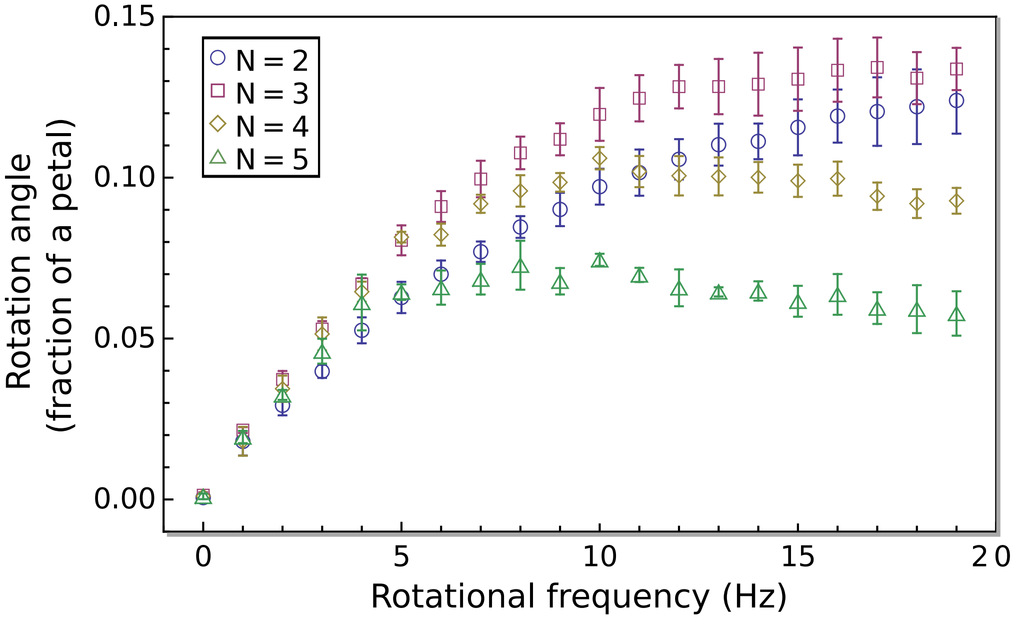

FIGURE 5.8: Fraction of a petal rotated by $N = 2$ through 5 petals as a function of rotational speed while held at constant peak intensity. Error bars represent the standard deviations of independent data runs.

patterns are simply rotated by a fixed amount, rather than a fraction of a petal, increasing the numbers of petals enough could overcome the reduction in overall rotation angle, allowing for a pattern with many petals to rotate by a complete 1.0 petal. However, my observations appear to imply that increasing the number of petals will not result in a pattern where a petal can be completely rotated into an area of darkness, which would correspond to a rotation of 0.5 petals. Irrespective of the number of petals or rotation speed, the maximum rotation observed in this experimental configuration is less than the 0.5 petal distance characterised by the spatial structure of the beam. Further work should be conducted to determine whether the restriction on the image rotation is due to an experimental parameter, such as the differing image size between petal patterns, which affects the peak intensity of the pattern, or the nature of slow light in ruby.

5.7 Summary and Conclusions

Theories on slow light in ruby are highly contentious [63, 64], falling primarily into two models. The first model describes the phenomenon entirely through simple pulse reshaping and states that slow light is not required to explain the observed effects [65]. The second model describes slow light in ruby as group index-dependent process, where the degree of slowing is a complicated function of the functional form of the light field [82]. The observations herein may be used to inform the debate on slow light in ruby, as any successful model would need to account for the results of this experiment. The key features of this work are that the slow-light effect saturates with the modulation frequency, which is a product of the spatial feature of the beam profile and the rotation frequency, and that no petal pattern can rotate through a full spatial feature of the beam profile. These results also show some evidence that all petal patterns rotate through the same angle at low frequencies (see the low frequency region of Fig. 5.8).

I have shown the rotation of OAM-carrying beams that is analogous to the polarisation rotation related to SAM. The rotation of OAM-carrying beams appears to be limited to a fraction of a petal, implying an inability to rotate a petal pattern by the angular extent of a petal, regardless of the number of petals in the beam. These results highlight the complicated nature of slow light in the mechanical Faraday effect for OAM-carrying beams and provide experimental evidence against which theories of slow light in ruby can be measured.

Chapter 6

Conclusions and Future Work

Slow light has become the focus of much research due to its potential for dramatically changing the effective velocity of light. While being able to control the velocity of light has many potential applications, some of which have already begun to be realised, the mechanism by which light is slowed in ruby is not well understood. I have conducted careful experimental work to inform our understanding of slow light in ruby and provide direction for further uses of this phenomenon.

When passing an image through a spinning, self-pumped ruby, I observed rotations of the bright and dark regions by the same amount. Positions where the input beam was dark became bright in the output cross section due to azimuthally dragged light. The observed movement of the bright and dark regions could not be explained by optical bleaching alone, since such effects cannot lead to an increase in energy. Such effects could, however, be explained by narrow band absorption and the associated change in group refractive index. The dark region shows signs of movement when created by either the shadow of a wire or a π-radian phase discontinuity, but use of the phase discontinuity leads to more conclusive and definitive results. The slow-light effect is stronger when the ruby is illuminated by an elliptical beam rather than a round one, as the elliptical beam causes a stronger modulation of the ruby.

It is important to eliminate fluorescence when studying slow light in ruby. The illumination from fluorescence can mask the actual slow-light effect. With fluorescence present, any observed dragging could not confidently be ascribed to the effects of slow light in ruby, as I confirmed through a simple model. With fluorescence removed and other experimental parameters precisely controlled, the measured dragging can be studied to learn about the effects of slow light in ruby.

In the experimental work in Chapter 2, the level of laser power used in the experiment does not fully saturate the ruby, due to power limitations of the laser. This leaves open the question of whether a more powerful laser that could saturate the ruby would increase the slow-light effect and therefore the rotation angle. Future work should be done in this area to fully understand the exact relationship between laser power and the slow-light effect in ruby.

All of the work conducted with a rotating ruby relied upon spatial modulation of the ruby. However, the slow-light effect can also be created in a temporal domain by illuminating a stationary ruby with a temporally modulated signal. Investigating slow light in ruby in the temporal domain was useful when examining a very sharp drop in intensity of the modulating beam. I carefully controlled the experimental parameters to show the existence of a pronounced tail on the trailing edge of the transmitted signal due to the light pulse being slowed as it propagates through the ruby.

Light being delayed beyond the input pulse cutoff is not compatible with a simple model of pulse delay in a time-varying (saturable) absorber. Instead, my experimental evidence supports a more complicated model of slow light in ruby that results in a delay of the transmitted optical energy. The results from this temporally-modulated experiment support the results of my spatially-modulated experiment described above. A distortion of the pulse shape of the temporally-modulated signal shows that individual Fourier components of the square-wave modulated signal are delayed by different amounts when transmitted through ruby, a result that supports early investigations of sine-wave modulated signal delays in ruby.

After establishing that both bright and dark regions of an image can experience slow-light effects and that these effects arise from a mechanism more complicated than simple saturable absorber, I went on to investigate how a more complicated beam would experience slow-light effects. I briefly studied the rotation of complicated patterns created from simple metal masks. However, to more precisely understand how the shape of the pattern affects rotation, I then examined the slow light-enhanced rotation of beams carrying orbital angular momentum.

By studying the rotation of orbital angular momentum-carrying beams, I observed that the slow-light effect saturates with the modulation frequency, which is a product of the spatial feature of the beam profile and the rotation frequency. The rotation of orbital angular momentum-carrying beams appears to be limited to a fraction of a petal, implying an inability to rotate a petal pattern by the angular extent of a full petal, regardless of the number of petals in the beam. The degree of slowing is therefore seen to be a complicated function of the form of the light field, an effect not predicted by previously suggested theories.

The experimental observations made in this thesis may be used to inform the debate on slow light in ruby. These results highlight the complex nature of slow light and provide experimental evidence against which theories of slow light in ruby can be measured. The results presented in this thesis further the understanding of the mechanisms that cause slow light and pave the way for applications dependent on the preservation of complex patterns in slow-light media.

There is still much exciting work to be done in this area, which would both help us gain a deeper understanding of slow light in ruby and provide possible new applications. One possible application that should be investigated is the use of ruby as a slow-light medium to increase the sensitivity of an interferometer. A brief section containing my initial work on a slow light-enhanced interferometer is included in Appendix B. The largest challenge with the slow light-enhanced interferometer experiment is the experimental sensitivity to alignment. If the alignment issues were solved, this could potentially become an application with many uses.

Another interesting experiment to conduct following my work with ruby would be to investigate rotary image drag in alexandrite. Ruby acts as a slow-light medium when modulated at low frequencies, but at high frequencies, the slow-light effect disappears. Alexandrite also acts as a slow-light medium when modulated at low frequencies. However, when modulated at high frequencies, alexandrite becomes a fast-light medium. Recreating the rotary photon drag experiment with alexandrite would then, theoretically, allow for an observation of the image rotation in the opposite direction from that in which the bar would be rotating.

Finally, as my thesis focused on learning about the variables that affect slow light in ruby and attempting to differentiate between proposed mechanisms, there is now much information available for a future project to create a complete theory for how ruby slows the effective velocity of the light by such a large amount.

Appendix A

Slow Light Interferometer

A suggested experimental setup for increasing the sensitivity of a Sagnac interferometer is as follows. As shown in Fig. A.1, 3 W of linearly polarised 532 nm light passes through a nonpolarising beamsplitter. The light travels in both directions around a Sagnac interferometer. In one arm of the interferometer, a spherical 50 mm lens focuses the light into a 4-mm-thick ruby window from each direction. The lenses and ruby are mounted on a moveable stage. When pumped with 532 nm light, the ruby acts as a slow-light medium, and thus as the light passes through the ruby window, its group velocity is effectively slowed to hundreds of meters per second. The ruby and lenses are propelled along the stage, moving longitudinally along the path of the propagating light. The interference pattern from the recombined clockwise and counterclockwise light is collected on a CCD array and recorded for subsequent analysis.

The larger the velocity of the stage, the longer the light traveling in the same direction as the ruby spends in the ruby. Conversely, the light traveling in the opposite direction of the stage spends a shorter time in the ruby the faster the stage is travelling. Thus, a large velocity will maximise the difference in the two effective path lengths, resulting in a large change in the relative phases. A shift in the relative phases results in a change in the interference fringes, which can be observed and measured on a CCD.

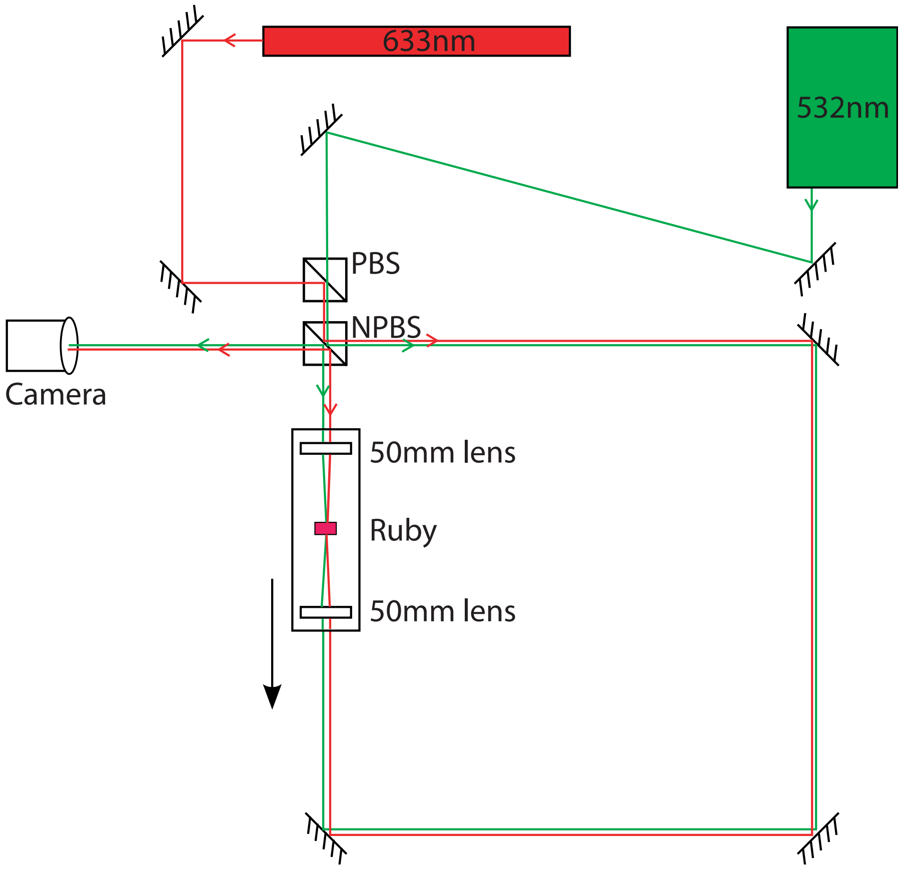

Figure A.1: A 3 W beam of 532 nm light passes in both directions around a Sagnac interferometer. The light is focused on a ruby window, which moves longitudinally along one arm of the interferometer. Light passing clockwise through the interferometer spends less time in the ruby than light passing counterclockwise through the interferometer. The interference pattern of the recombined light is imaged onto a camera. Red (633 nm) laser light follows the same path as the 532 nm light and is viewed on the same camera.

A large velocity would create a large difference between the two path lengths and thus between the phases of the two beams. However, if that velocity were constant, there would be a constant difference between the phases. A constant phase difference would just look like a stationary fringe pattern. Therefore, I investigate the period when the stage is accelerating. The faster the velocity changes, the more quickly the fringes move. By recording high speed video, the movement of fringes could be tracked even when they are shifting quickly.

As deflection and misalignment can easily obscure the desired results in this experiment, a large acceleration is critical. The faster the acceleration, the faster the fringes would move, allowing for the a large fringe movement to be captured over a shorter amount of time. Greater acceleration would therefore reduce the amount of displacement experienced by the fringes due to misalignment.

Because the slow-light effect only occurs when green light interacts with the ruby, a red (633 nm) laser beam could be coupled into the interferometer as well. As the red light would not be slowed by the ruby, the red fringes could be used as a control against which the motion of the green fringes can be compared. As the green and red fringes would both be affected in the same way by the motion of the ruby, that motion could be corrected for in the green fringes, leaving only the change in fringes due to the slow-light effect.

I explored a few initial methods of moving the ruby along the arm of the interferometer. The first method I investigated involved connecting the stage containing the lenses and ruby to rails and a strong spring. The stage was pulled back and then released, and the fringes were recorded as the stage was pulled along the rails by the spring. This method provided a large acceleration, its greatest benefit.

However, this method also had a major drawback. The alignment of the interferometer is crucial. The interferometer has to be aligned such that the laser in both directions travels along exactly the same path as the ruby on the rails. As the ruby moves along the rails, even the slightest misalignment will disrupt the fringes to a large enough degree to render the desired effect unmeasurable. The use of a spring resulted in a large acceleration, however in addition to the difficulty in alignment, the spring did not pull the stage evenly, causing a misalignment of the stage as it moved along the rails. The alignment problems, both from aligning the interferometer and from the uneven pulling of the stage, were great enough that this method was discarded.

In an attempt to achieve a large acceleration, the stage with the ruby and lenses was then mounted vertically. The entire interferometer was placed on end, so that the ruby would fall along one arm of the interferometer simply under the force of gravity.

The main flaw with the experimental setup where the ruby falls with an acceleration due to gravity is that the ruby travels a long distance, where even a slight misalignment would result in a large deflection of the fringes. One potential solution for reducing the misalignment of the system involves the use of a piezoelectric material. A piezoelectric material is a specific type of electrostrictive material, a material that will compress with the application of electrical current. Piezoelectric materials will also, inversely, generate an electrical current when subjected to an applied force. By applying an alternating current, the piezoelectric material could be driven in a resonant way, resulting in a very fast compression/relaxation cycle. In order to reach the acceleration needed, the piezoelectric material would need to be driven at a very high frequency, as the distance over which it would be moving is very small.

To determine how fast the ruby would need to travel to measure a fringe movement, I consider that when the ruby is accelerating, the fringes will shift as the path lengths change. Ideally, the fringes would move by one wavelength, that is one maxima will shift to the position of the next maxima. To calculate the required speed, I use the following equation:

$$\Delta\phi = \frac{v_r L n_g n_\phi}{c\lambda}. \tag{A.1}$$

In Eq. A.1, $\Delta\phi$ is the change in phase due to the path length difference, v_r and L are the velocity and length of the ruby, and n_g and n_ϕ are the group and phase indices of the ruby. Rearranging Eq. A.1 for the velocity of the ruby gives

$$v_r = \frac{\Delta\phi c\lambda}{L n_g n_\phi}, \tag{A.2}$$

and by substituting in Eq. 1.5, the result is

$$v_r = \frac{\Delta\phi v_g \lambda}{L n_\phi}. \tag{A.3}$$

A fringe movement of one wavelength would occur when $\Delta\phi = 2\pi$. Under the correct experimental parameters, ruby has an expected $n_g \approx 10^5$ resulting in $v_g \approx 3000$. Calculations using $\lambda = 532$nm, $L = 4 \times 10^{-3}$m and $n_\phi = 1.5$ show that the ruby would have to reach a velocity of approximately 1.67 m/s to see a fringe movement of one wavelength. For comparison, with a piece of glass in one arm of the interferometer instead of ruby, in order to see a fringe movement of one wavelength, the glass would have to travel at approximately 1.1×10^5 m/s. At 1.67 m/s, a glass bar would only cause a displacement of approximately 0.015% of a fringe.

Initial experimental work was conducted, and a video of the recombined light from the interferometer was recorded, with representative images from the video taken with the dropping method shown in Fig. A.2.

In initial experimental results, both the red and green fringes experience some wobbling due to the motion of the stage. However, as can be seen in Fig. A.2, there is a marked increase in movement of the green fringes with comparison to the red, with the green fringes shifting by approximately 0.7 wavelengths, which agrees with my calculation that the stage in this experiment reached a maximum velocity of just above 1 m/s. It has been shown experimentally that the conditions required for slowing light are quite specific, and alignment of the lenses focusing into the ruby is very sensitive [58]. If misaligned, the light does not experience a strong slow-light effect. By maximising the experimental conditions required for a stronger slow-light effect, i.e. much slower propagation of the 532 nm laser light through the ruby, the sensitivity of the spectrometer could be increased even more. In other experiments (see Chapter 2) I have reached effective group velocities of approximately 8.3 m/s. If one was able to achieve that large of a slow-light effect with this setup, the sensitivity could be increased by about 99.7 percent. With current velocity limits, this would still be an increase from 0.7 fringe wavelength movements to over 200 fringe wavelength movements. In the future, a high-speed camera could be used to track the position of the ruby and calculate its exact speed.

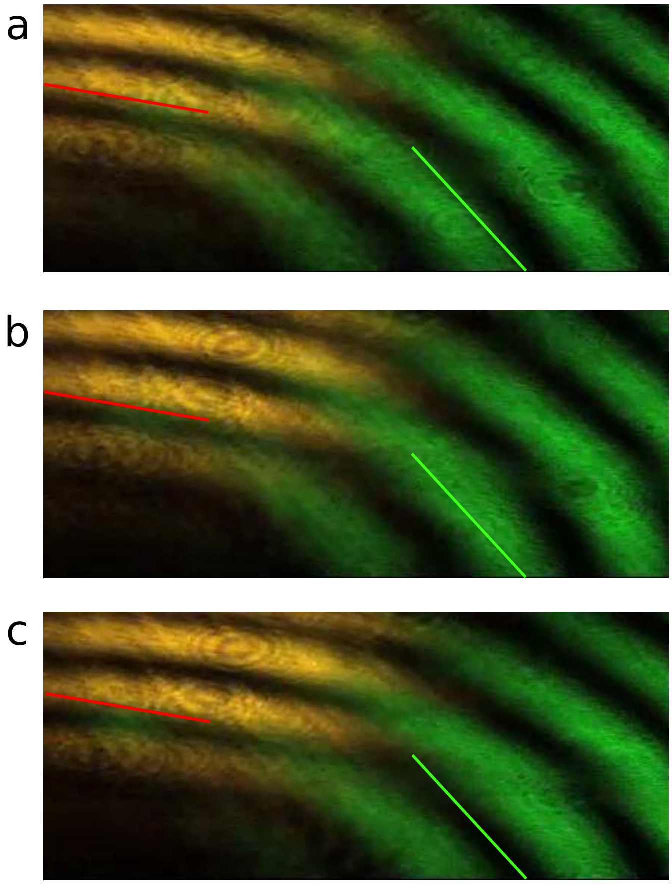

Figure A.2: Fringe positions (*a*) while the ruby was stationary, (*b*) while the ruby was accelerating, and (*c*) while the ruby was moving at its highest velocity. Green fringes are seen to shift by approximately 0.7 fringes, while red fringes do not. Red and green reference lines are added in all three images at the same position and orientation to show the relative movement of the fringes.

The main drawbacks to this experiment has been the sensitivity of the alignment for the slow-light effect to occur and the potential for misalignment while the ruby is moving. These results are preliminary, and further work on this topic would have to be conducted to eliminate all sources of possible error before any decisive claims could be made about the exact effect of the slow-light material on the interferometric sensitivity.

Appendix B

Additional Rotated Patterns

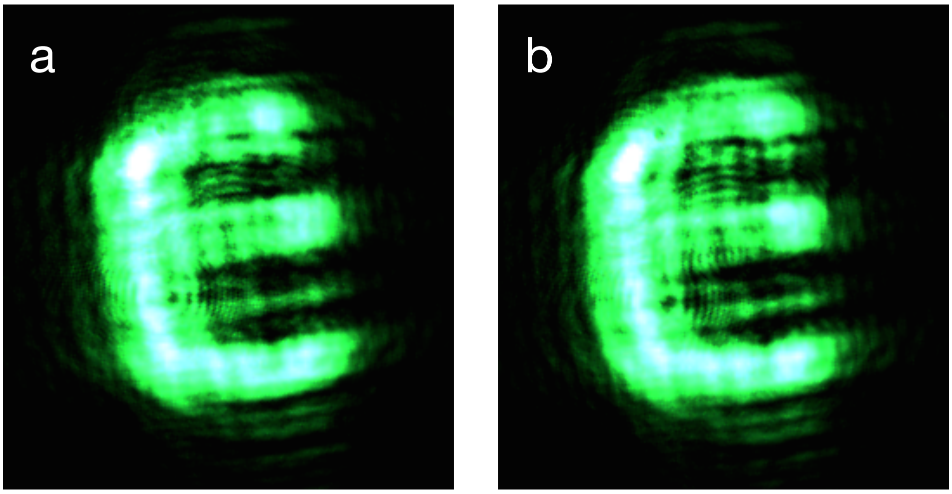

Figure B.1: Rotation of pattern of an E made by a metal mask. Images shown were taken when the ruby was rotating (a) anticlockwise and (b) clockwise.

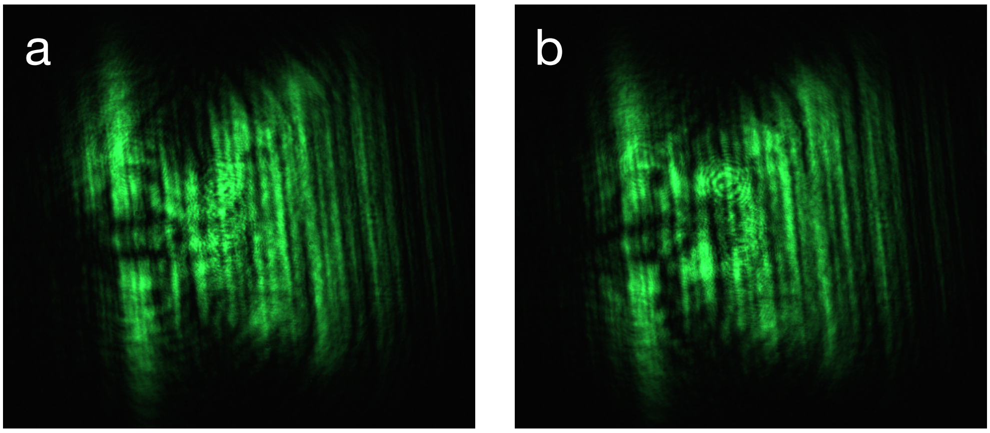

Figure B.2: Rotation of pattern made by a metal mask. Images shown were taken when the ruby was rotating (a) anticlockwise and (b) clockwise.

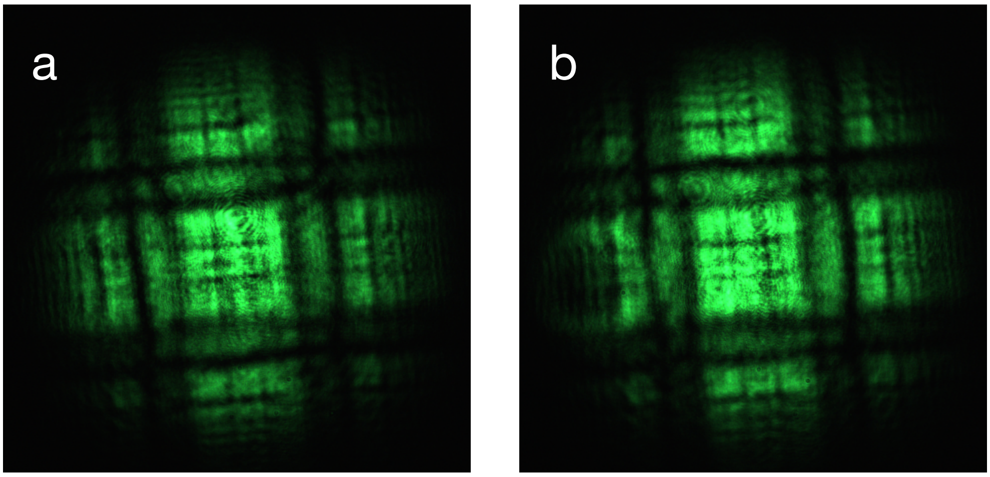

Figure B.3: Rotation of a square pattern made by a metal mask. Images shown were taken when the ruby was rotating (a) anticlockwise and (b) clockwise.

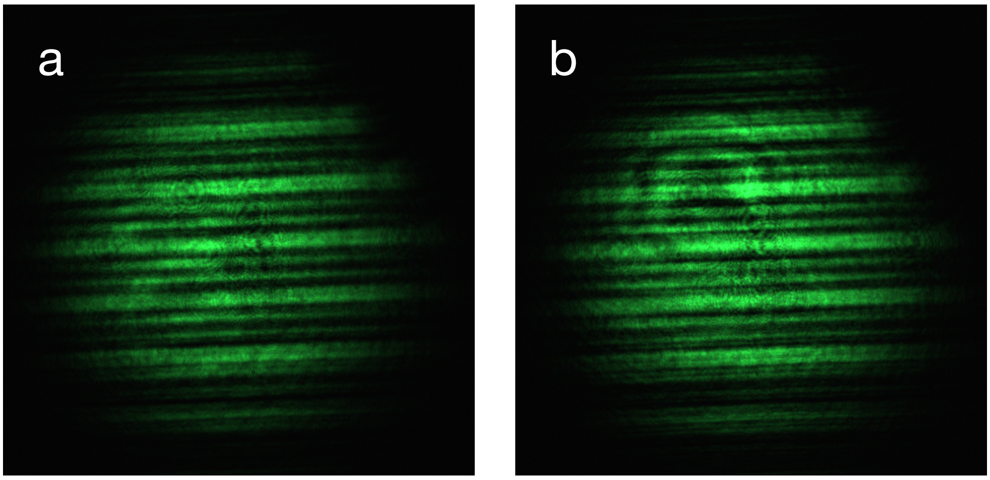

Figure B.4: Rotation of a pattern of horizontal lines made by a metal mask. Images shown were taken when the ruby was rotating (a) anticlockwise and (b) clockwise.

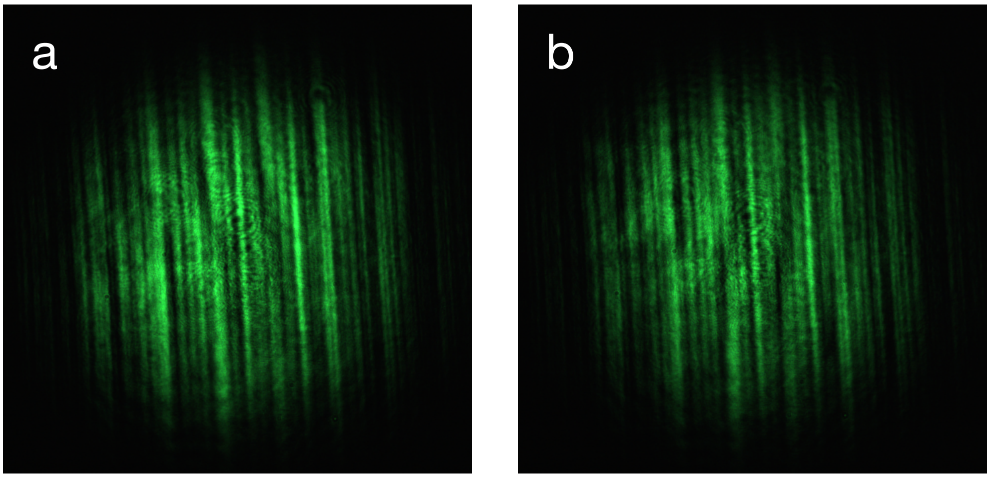

Figure B.5: Rotation of a pattern of vertical lines made by a metal mask. Images shown were taken when the ruby was rotating (a) anticlockwise and (b) clockwise.

Bibliography

[1] E. B. Aleksandrov and V. S. Zapasskii. Chasing 'slow light'. *Phys Uspekhi*, 49:1067–1075, 2006.

[2] Emil Wolf, editor. *Progress in Optics*, volume 43. Elsevier, 2002.

[3] J.A. Carruthers and T Bieber. Pulse velocity in a self-locked he-ne laser. *Journal of Applied Physics*, 40:426–428, 1969.

[4] A Frova, M.A. Duguay, C.G.B. Garrett, and S.L. McCall. Pulse delay effects in the he-ne laser mode-locked by a ne absorption cell. *Journal of Applied Physics*, 40:3969–3972, 1969.

[5] F.R. Faxvog, C.N.Y. Chow, T Bieber, and J.A. Carruthers. Measured pulse velocity greater than c in a neon absorption cell. *Applied Physics Letters*, 17(192-193), 1970.

[6] Lee Casperson and Amnon Yariv. Pulse propagation in a high-gain medium. *Physical Review Letters*, 26:293–295, 1971.

[7] C.G.B. Garrett and D.E. McCumber. Propagation of a gaussian light pulse thorugh an anomalous dispersion medium. *Physical Review A*, 1:305–313, 1970.

[8] S Chu and S Wong. Linear pulse propagation in an absorbing medium. *Physical Review Letters*, 48(11):738–741, 1982.

[9] A. Katz and R.R. Alfano. Pulse propagation in an absorbing medium. *Physical Review Letters*, 49:1292, 1982.

[10] Bernard Segard and Bruno Macke. Observation of negative velocity pulse propagation. *Physics Letters*, 109(5):213–216, 1985.

[11] Min Xiao, Yong qing Li, Shao zheng Jin, and Julio Gea-Banacloche. Measurement of dispersive properties of electromagnetically induced transparency in rubidium atoms. *Physical Review Letters*, 74:666–669, 1995.

[12] A Kasapi, Maneesh Jain, G.Y. Yin, and S.E. Harris. Electromagnetically induced transparency: propagation dynamics. *Physical Review Letters*, 74 (13):2447–2450, 1995.

[13] L. V. Hau, S. E. Harris, Z Dutton, and C. H. Behroozi. Light speed reduction to 17 metres per second in an ultracold atomic gas. *Nature*, 397:594–598, 1999.

[14] Michael M. Kash, Vladimir A. Sautenkov, Alexander S. Zibrov, L. Hollberg, George R. Welch, Mikhail D. Lukin, Yuri Rostovtsev, Edward S. Fry, and Marlan O. Scully. Ultraslow group velocity and enhanced nonlinear optical effects in a coherently driven hot atomic gas. *Physical Review Letters*, 82 (26):5229–5232, 1999.

[15] Chien Liu, Zachary Dutton, Cyrus H Behroozi, and Lene Vestergaard Hau. Observation of coherent optical information storage in an atomic medium using halted light pulses. *Nature*, 409:490–493, 2001.

[16] Joseph E Vornehm and Robert W. Boyd. *Tutorials in Complex Photonic Media: Slow and Fast Light*, chapter 19, pages 647–685. SPIE Digital Library, 2009.

[17] Leon Brillouin. Uber die fortpflanzung des lichtes in dispergierenden medien. *Ann. Phys.*, 349:203–240, 1914.

[18] A Icsevgi and W E Lamb. Propagation of light pulses in a laser amplifier. *Physical Review*, 185(2):517–545, 1969.

[19] N G Basov, R V Ambartsumyan, V S Zuev, P G Kryukov, and V S Letokhov. Nonlinear amplification of light pulses. *Soviet Physics Jetp*, 23(1):16–22, 1966.

[20] Eric L Bolda, John C Garrison, and Raymond Y Chiao. Optical pulse propagation at negative group velocities due to a nearby gain line. *Physical Review A*, 49(4):2938–2947, 1994.

[21] A.M. Steinberg and R.Y. Chiao. Dispersionless, highly superluminal propagation in a medium with a gain doublet. *Physical Review A*, 49:2071–2075, 1994.

[22] L.J. Wang, A. Kuzmich, and A. Dogariu. Gain-assisted superluminal light propagation. *Nature*, 406:277–279, 2000.

[23] Raymond Y Chiao, Alexander E. Kozhekin, and Gershon Kurizki. Tachyonlike excitations in inverted two-level media. *Physical Review Letters*, 77: 1254–1256, 1996.

[24] R Bonifacio and L A Lugiato. Cooperative radiation processes in two-level systems: superfluorescence. *Physical Review A*, 11:1507, 1975.

[25] George C. Sherman and Kurt Edmund Oughstun. Description of pulse dynamics in lorentz media in terms of the energy velocity and attenuation of time-harmonic waves. *Physical Review Letters*, 47(20):1451–1454, 1981.

[26] G Diener. Superluminal group velocities and information transfer. *Physics Letters A*, 223(5):327–331, 12 1996.

[27] A. M. Akulshin, S. Barreiro, and A. Lezama. Steep anomalous dispersion in coherently prepared rb vapor. *Phys. Rev. Lett.*, 83:4277–4280, Nov 1999.

[28] Bilha Segev, Peter W. Milonni, James F. Babb, and Raymond Y Chiao. Quantum noise and superluminal propagation. *Physical Review A*, 64: 022114–022128, 2000.

[29] A. Kuzmich, A. Dogariu, L. J. Wang, P. W. Milonni, and R. Y. Chiao. Signal velocity, causality, and quantum noise in superluminal light pulse propagation. *Physical Review Letters*, 86(18):3925–3929, 04 2001.

[30] C Monat, M de Sterke, and B J Eggleton. Slow light enhanced nonlinear optics in periodic structures. *Journal of Optics*, 12:104003 1–17, 2010.

[31] G. M. Gehring, A Schweinsberg, C Barsi, N Kotinski, and R. W. Boyd. Observation of backward pulse propagation through a medium with a negative group velocity. *Science*, 312:895–897, 2006.

[32] A. Schweinsberg, N. N. Lepeshkin, M. S. Bigelow, R. W. Boyd, and S. Jarabo. Observation of superluminal and slow light propagation in erbium-doped optical fiber. *Europhysics Letters*, 73:218–224, 2006.

[33] Yurii A Vlasov, Martin O'Boyle, Hendrik F. Hamann, and Sharee J. McNab. Active control of slow light on a chip with photonic crystal waveguides. *Nature*, 438:65–69, 2005.

[34] Juntao Li, Liam O'Faolain, and Thomas F. Krauss. Four-wave mixing in slow light photonic crystal waveguides with very high group index. *Optics Express*, 20(16):17474–17479, 2012.

[35] Chunle Xiong, Christelle Monat, Matthew J. Collins, Laurent Tranchant, David Petiteau, Alex S. Clark, Christian Grillet, Graham D. Marshall, Michael J. Steel, Juntao Li, Liam O'Faolain, Thomas F. Krauss, and Benjamin J. Eggleton. Characteristics of correlated photon pairs generated in ultracompact silicon slow-light photonic crystal waveguides. *Ieee Journal of Selected Topics in Quantum Electronics*, 18(6):1676–1683, 2012.

[36] M CastellanosMunoz, A Y Petrov, L O'Faolain, J Li, T. F. Krauss, and M Eich. Optically induced indirect photonic transitions in a slow light photonic crystal waveguide. *Physical Review Letters*, 112(5), 2014.

[37] Joyce K S Poon, Philip Chak, John M Choi, and Amnon Yariv. Slowing light with fabry-perot resonator arrays. *Journal of the Optical Society of America B*, 24(11):2763–2769, 2007.

[38] David D Smith, Hongrok Chang, Kirk A Fuller, A T Rosenberger, and Robert W Boyd. Coupled-resonator-induced transparency. *Physical Review A*, 69:063804, 1–6, 2004.

[39] Jianqiang Gu, Ranjan Singh, Xiaojun Liu, Xueqian Zhang, Yingfang Ma, Shuang Zhang, Stefan A. Maier, Zhen Tian, Abul K. Azad, Hou-Tong Chen, Antoinette J. Taylor, Jiaguang Han, and Weili Zhang. Active control of electromagnetically induced transparency analogue in terahertz metamaterials. *Nat Commun*, 3:1151, 10 2012.

[40] C.J. Chang-Hasnain, P.-C. Ku, J. Kim, and S.-L. Chuang. Variable optical buffer using slow light in semiconductor nanostructures. *Proceedings of the IEEE*, 91:1884–1897, 2003.

[41] R. W. Boyd. Material slow light and structural slow light: similarities and differences for nonlinear optics [invited]. *Journal of the Optical Society of America B*, 28:A38–A44, 2011.

[42] D.C. Hutchings, M. Sheik-Bahae, D.J. Hagan, and E.W. Van Stryland. Kramers-kronig relations in nonlinear optics. *Optical And Quantum Electronics*, 24:1–30, 1992.

[43] S.E. Harris, J.E. Field, and A. Kasapi. Dispersive properties of electromagnetically induced transparency. *Physical Review A*, 46(1):29–32, 1992.

[44] Luc Thévenaz. Slow and fast light in optical fibres. *Nature Photonics*, 2: 474–481, 2008.

[45] Yoshitomo Okawachi, Matthew S. Bigelow, Jay E Sharping, Zhaoming Zhu, Aaron Shweinsberg, Daniel J. Gauthier, Robert W. Boyd, and Alexander L. Gaeta. Tunable all-optical delays via brillouin slow light in an optical fiber. *Physical Review Letters*, 94:153902, 2005.

[46] Kwang Yong Song, Miguel Gonzalez Herraez, and Luc Thévenaz. Gain-assisted pulse advancement using single and double brillouin gain peaks in optical fibers. *Optics Express*, 13(24):9758–9765, 2005.

[47] Kwang Yong Song, Miguel Gonzalez Herraez, and Luc Thévenaz. Long optically controlled delays in optical fibers. *Optics Letters*, 30(14):1782–1784, 2005.

[48] Jay E Sharping, Yoshitomo Okawachi, and Alexander L. Gaeta. Wide bandwidth slow light using a raman fiber amplifier. *Optics Express*, 13(16):6092–6098, 2005.

[49] Mike van der Poel, Jesper Mork, and Jorn M. Hvam. Controllable delay of ultrashort pulses in a quantum dot optical amplifier. *Optics Express*, 13(20): 8032–8037, 2005.

[50] E. Baldit, K. Bencheikh, P. Monnier, J.A. Levenson, and V. Rouget. Ultraslow light propagation in an inhomogeneously broadened rare-earth ion-doped crystal. *Physical Review Letters*, 95:143601, 2005.

[51] Pengfei Wu and D.V.G.L.N. Rao. Controllable snail-paced light in biological bacteriorhodopsin thin film. *Physical Review Letters*, 95:253601, 2005.

[52] Chandra S. Yelleswarapu, Reji Phillip, Francisco J. Aranda, Brian R. Kimball, and D.V.G.L.N. Rao. Slow light in bacteriorhodopsin solution using coherent population oscillations. *Optics Letters*, 32(13):1788–1790, 2007.

[53] Chandra S. Yelleswarapu, Samir Laoui, Reji Phillip, and D.V.G.L.N. Rao. Coherent population oscillations and superluminal light in a protein complex. *Optics Express*, 16(6):3844–3852, 2008.

[54] Sonia Melle, Oscar G. Calderon, F. Carreno, Eduardo Cabrera, M.A. Anton, and S. Jarabo. Effect of ion concentration on slow light propagation in highly doped erbium fibers. *Optics Communications*, 279:53–63, 2007.

[55] P. Palinginis, S. Crankshaw, F. Sedgwick, E.-T. Kim, M. Moewe, C.J. Chang-Hasnain, H. Wang, and S.-L. Chuang. Ultraslow light (¡200 m/s)

propagation in a semiconductor nanostructure. *Applied Physics Letters*, 87: 171102, 2005.

[56] Pei-Cheng Ku, Forrest Sedgwick, Connie J. Chang-Hasnain, Phedon Palinginis, Tao Li, Hailin Wang, Shu-Wei Chang, and Shun-Lien Chuang. Slow light in semiconductor quantum wells. *Optics Letters*, 29(19):2291–2293, 2004.

[57] Lloyd W. Hillman, Robert W. Boyd, Jerzy Krasinski, and C.R. Stroud. Observation of a spectral hole due to population oscillations in a homogeneously broadened optical absorption line. *Optics Communications*, 45(6):416–419, 1983.

[58] Matthew S. Bigelow, Nick N. Lepeshkin, and Robert W. Boyd. Observation of ultraslow light propagation in a ruby crystal at room temperature. *Physical Review Letters*, 90(11):113903 1–4, 2003.

[59] Matthew S. Bigelow, Nick N. Lepeshkin, and Robert W. Boyd. Superluminal and slow light propagation in a room temperature solid. *Science*, 301:200–202, 2003.

[60] H. Su and S. L. Chuang. Room temperature slow and fast light in quantum-dot semiconductor optical amplifiers. *Applied Physics Letters*, 88:061102, 2006.

[61] Sonja Franke-Arnold, Graham Gibson, Robert W. Boyd, and Miles J. Padgett. Rotary photon drag enhanced by a slow-light medium. *Science*, 333: 65–67, 2011.

[62] E Wisniewski-Barker, G Gibson, S Franke-Arnold, Z Shi, R W Boyd, and M J Padgett. Evidence of slow-light effects from rotary drag of structured beams. *New Journal of Physics*, 15(8):083020, 2013.

[63] G G Kozlov, S V Poltavtsev, I I Ryzhov, and V S Zapasskii. Comment on ‘evidence of slow-light effects from rotary drag of structured beams’. *New Journal of Physics*, 16:038001, 2014.

[64] E Wisniewski-Barker, G Gibson, S Franke-Arnold, Z Shi, R W Boyd, and M J Padgett. Reply to comment on 'evidence of slow-light effects from rotary drag of structured beams'. *New Journal of Physics*, 16:038002, 2014.

[65] V S Zapasskii and G G Kozlov. A saturable absorber, coherent population oscillations, and slow light. *Optics and Spectroscopy*, 100(3):419–424, 2006.

[66] V. S. Zapasskii and G. G. Koslov. Slow light and slow current. *Optics and Spectroscopy*, 104:95–98, 2008.

[67] R M Camacho, M V Pack, J C Howell, A Schweinsberg, and R W Boyd. Wide-bandwidth, tunable, multiple-pulse-width optical delays using slow light in cesium vapor. *Physical Review Letters*, 98(15), 2007.

[68] Daryl M Beggs, Isabella H Rey, Tobias Kampfrath, Nir Rotenberg, L Kuipers, and Thomas F. Krauss. Ultrafast tunable optical delay line based on indirect photonic transitions. *Physical Review Letters*, 108:213901 1–4, 2012.

[69] Jay E Sharping, Yoshitomo Okawachi, James van Howe, Chris Xu, Yan Wang, Alan E. Willner, and Alexander L. Gaeta. All-optical, wavelength and bandwidth preserving, pulse delay based on parametric wavelength conversion and dispersion. *Optics Express*, 13(20):7872–7877, 2005.

[70] Yoshitomo Okawachi, Jay E Sharping, Chris Xu, and Alexander L. Gaeta. Large tunable optical delays via self-phase modulation and dispersion. *Optics Express*, 14(25):12022–12027, 2006.

[71] Zhimin Shi, Robert W. Boyd, Daniel J. Gauthier, and C.C. Dudley. Enhancing the spectral sensitivity of interferometers using slow-light media. *Optics Letters*, 32(8):915–917, 2007.

[72] F. Zimmer and M. Fleischhauer. Sagnac interferometry based on ultraslow polaritons in cold atomic vapors. *Physical Review Letters*, 92:253201, 2004.

[73] Selim M. Shahriar, Gour Pati, Venkatesh Gopal, Renu Tripathi, George Cardoso, Prabhakar Pradhan, Mary Messal, and Ranjit Nair. Precision

rotation sensing and interferometry using slow light. In *Quantum Electronics and Laser Science Conference*, number JWB97. Optical Society of America, 2005.

[74] Zhimin Shi, Robert W. Boyd, Ryan M. Camacho, Praveen K. Vudyasetu, and John C. Howell. Slow-light fourier transform interferometer. *Physical Review Letters*, 99:240801 1–4, 2007.

[75] M. S. Shahriar, G. S. Pati, R. Tripathi, V. Gopal, M. Messall, and K. Salit. Ultrahigh enhancement in absolute and relative rotation sensing using fast and slow light. *Physical Review A*, 75:053807, 2007.

[76] M. Salit, G. S. Pati, K. Salit, and M. S. Shahriar. Fast-light for astrophysics: super-sensitive gyroscopes and gravitational wave detectors. *Journal of Modern Optics*, 54(16-17):2425–2440, 2007.

[77] M H Fizeau. On the effect of the motion of a body upon the velocity with which it is traversed by light. *Philosophical Magazine and Journal of Science*, 19(127):245–260, 1860.

[78] R.V. Jones. 'fresnel aether drag' in a transversely moving medium. *Proceedings of the Royal Society of London A*, 328:337–352, 1972.

[79] R. V. Jones. Rotary 'aether drag'. *Proceedings of the Royal Society of London A*, 349:423–439, 1976.

[80] Enrico Fermi. Sul trascinamento del piano di polarizzazione da parte di un mezzo rotante. *Rend Lincei*, 32:115–118, 1923.

[81] M. A. Player. On the draggin of the plane of polarization of light propagating in a rotating medium. *Proceedings of the Royal Society of London A*, 349: 441–445, 1976.

[82] Robert W. Boyd. Slow and fast light: fundamentals and applications. *Journal of Modern Optics*, 56(18-19):1908–1915, 2009.

[83] T F Krauss. Why do we need slow light? *Nature Photonics*, 2:448–450, 2008.

[84] R. S. Tucker, P. C. Ku, and C. J Chang-Hasnain. Slow-light optical buffers: Capabilities and fundamental limitations. *Journal of Lightwave Technology*, 23:4046–4066, 2005.

[85] J. Keaveney, I.G. Hughes, A. Sargsyan, D. Sarkisyan, and C.S. Adams. Maximal refraction and superluminal propagation in a gaseous nanolayer. *Physical Review Letters*, 109:233001, 2012.

[86] O Kocharavskaya, Y Rostovtsev, and Scully M. Stopping light via hot atoms. *Physical Review Letters*, 86:628–631, 2001.

[87] Giovanni Piredda and Robert W. Boyd. Slow light by means of coherent population oscillations: laser linewidth effects. *Journal of European Optical Society*, 2:07004, 1–4, 2007.

[88] G Nienhuis, J P Woerdman, and Ian Kuscer. Magnetic and mechanical faraday effects. *Physical Review A*, 46(11):7079–7092, 1992.

[89] Jorg B Götte, Stephen M Barnett, and Miles Padgett. On the dragging of light by a rotating medium. *Proceedings of The Royal Society of London A*, 463(2085):2185–2194, 2007.

[90] Miles Padgett, Graeme Whyte, John Girkin, Amanda Wright, Les Allen, Patrik Öhberg, and Stephen M. Barnett. Polarization and image rotation induced by a rotating dielectric rod: an optical angular momentum interpretation. *Optics Letters*, 31(14):2205–2207, 2006.

[91] Emma Wisniewski-Barker, Graham M. Gibson, Sonja Franke-Arnold, Zhimin Shi, Paul Narum, Robert W. Boyd, and Miles J. Padgett. Experimental investigation of the transient dynamics of slow light in ruby. *New Journal of Physics*, 16:123054, 1–6, 2014.

[92] D F Nelson and M D Sturge. Relation between absorption and emission in the region of the r lines of ruby. *Physical Review*, 137(4A):1117–1130, 1965.

[93] J Poynting. The wave motion of a revolving shaft, and a suggestion as to the angular momentum in a beam of circularly polarised light. *Proceedings of the Royal Society of London A Series A*, 82:560–567, 1909.

[94] R Beth. Mechanical detection and measurement of the angular momentum of light. *Physical Review*, 50:115–125, 1936.

[95] L. Allen, M.W. Beijersbergen, R.J.C. Spreeuw, and J.P. Woerdman. Orbital angular momentum of light and the transformation of laguerre-gaussian laser modes. *Physical Review A*, 45:8185–8189, 1992.

[96] C G Darwin. Notes on the theory of radiation. *Proceedings of the Royal Society of London Series A*, 136:36–52, 1932.

[97] Les Allen, Miles J Padgett, and M Babiker. The orbital angular momentum of light. *Progress in Optics*, 39:291–372, 1999.

[98] Miles J Padgett and Les Allen. The poynting vector in laguerre-gaussian laser modes. *Optics Communications*, 121:36–40, 1995.

[99] J. Leach, S. Keen, M. J. Padgett, C. Saunter, and G. D. Love. Direct measurement of the skew angle of the poynting vector in a helically phased beam. *Optics Express*, 14(25):11919–11924, 2006.

[100] M Berry, J Nye, and F Wright. The elliptic umbilic diffraction catastrophe. *Philosophical Transactions of the Royal Society of London*, 291:453–484, 1979.

[101] P. Coullet, L. Gil, and F. Rocca. Optical vortices. *Optics Communications*, 73(5):403–408, 1989.

[102] J F Nye and M Berry. Dislocations in wave trains. *Proceedings of the Royal Society of London A*, 336:165–190, 1974.

[103] J M Vaughan and D Willetts. Interference properties of a light-beam having a helical wave surface. *Optics Communications*, 30:263–267, 1979.

[104] L A Lugiato, F Prati, L M Narducci, and G.-L Oppo. Spontaneous breaking of the cylindrical symmetry in lasers. *Optics Communications*, 69:387–392, 1989.

[105] C Tamm. Frequency locking of 2 transverse optical modes of a laser. *Physical Review A*, 38:5960–5963, 1988.

[106] C Tamm and C Weiss. Bistability and optical switching of spatial patterns in a laser. *Journal of the Optical Society of America B*, 7:1034–1038, 1990.

[107] G D'Alessandro and G.-L Oppo. Gauss-laguerre modes: a sensible basis for laser dynamics. *Optics Communications*, 88:130–136, 1992.

[108] D McGloin and K Dholakia. Bessel beam: diffraction in a new light. *Contemporary Physics*, 46:15–28, 2005.

[109] J C Gutierrez-Vega, M D Iturbe-Castillo, and S Chavez-Cerda. Alternative formulation for invariant optical fields: Mathieu beams. *Optics Letters*, 25: 1493–1495, 2000.

[110] M A Bandres and J C Guttierrez-Vega. Ince-gaussian beams. *Optics Letters*, 29:144–146, 2004.

[111] M Harris, C A Hill, and J M Vaughan. Optical helices and spiral interference fringes. *Optics Communications*, 106:161–166, 1994.

[112] M. Padgett, J. Arlt, N. Simpson, and L. Allen. An experiment to observe the intensity and phase structure of laguerre-gaussian laser modes. *American Journal of Physics*, 64(1):77–82, 1996. 12.

[113] M S Soskin, V N Gorshkov, M V Vasnetsov, J T Malos, and N R Heckenberg. Topological charge and angular momentum of light beams carrying optical vortices. *Physical Review A*, 56:4064–4075, 1997.

[114] Johannes Courtial, Kishan Dholakia, D A Robertson, Les Allen, and Miles J Padgett. Measurement of the rotational frequency shift imparted to a rotating light beam possessing orbital angular momentum. *Physical Review Letters*, 80(15):3217–3219, 1998.

[115] I V Basistiy, V V Slyusar, Marat S Soskin, M V Vasnetsov, and A Ya Bekshaev. Manifestation of the rotational Doppler effect by use of an off-axis optical vortex beam. *Optics Letters*, 28(14):1185–1187, 2003.

[116] L. Allen, M. Babiker, and W. L. Power. Azimuthal doppler-shift in light-beams with orbital angular-momentum. *Optics Communications*, 112(3-4): 141–144, 1994.

[117] S. Franke-Arnold, S. M. Barnett, E. Yao, J. Leach, J. Courtial, and M. Padgett. Uncertainty principle for angular position and angular momentum. *New Journal of Physics*, 6:1–8, 2004.

[118] M.W. Beijersbergen, L. Allen, H.E.L.O. van der Veen, and J.P. Woerdman. Astigmatic laser mode converters and transfer of orbital angular momentum. *Optics Communications*, 96:123–132, 1993.

[119] N González, G Molina-Terriza, and Juan P Torres. How a dove prism transforms the orbital angular momentum of a light beam. *Optics Express*, 14 (20):9093–9102, 2006.

[120] Marco W Beijersbergen, R P C Coerwinkel, Martin Kristensen, and J P Woerdman. Helical-wavefront laser beams produced with a spiral phaseplate. *Optics Communications*, 112(5-6):321–327, 1994.

[121] G. A. Turnbull, D. A. Robertson, G. M. Smith, L. Allen, and M. J. Padgett. Generation of free-space laguerre-gaussian modes at millimetre-wave frequencies by use of a spiral phaseplate. *Optics Communications*, 127(4-6): 183–188, 1996. 15.

[122] S S R Oemrawsingh, J A W van Houwelingen, E R Eliel, J P Woerdman, E J K Verstegen, J G Kloosterboer, and G W t Hooft. Production and characterization of spiral phase plates for optical wavelengths. *Applied Optics*, 43(3):688–694, 2004.

[123] K Sueda, G Miyaji, N Miyanaga, and M Nakatsuka. Laguerre-gaussian beam generated with a multilevel spiral phase plate for high intensity laser pulses. *Optics Express*, 12(15):3548–3553, 2004.

[124] V Yu Bazhenov, M V Vasnetsov, and Marat S Soskin. Laser-beams with screw dislocations in their wave-fronts. *JETP Letters*, 52(8):429–431, 1990.

[125] V Yu Bazhenov, Marat S Soskin, and M V Vasnetsov. Screw Dislocations in Light Wavefronts. *Journal Of Modern Optics*, 39(5):985–990, May 1992.

[126] Norman R Heckenberg, Richard McDuff, C P Smith, Halina Rubinsztein-Dunlop, and M J Wegener. Laser beams with phase singularities. *Optical And Quantum Electronics*, 24(9):S951–S962, 1992.

[127] J. E. Curtis, B. A. Koss, and D. G. Grier. Dynamic holographic optical tweezers. *Optics Communications*, 207(1-6):169–175, 2002. 22.

[128] Jonathan Leach, Mark R Dennis, Johannes Courtial, and Miles J Padgett. Vortex knots in light. *New Journal Of Physics*, 7:55–55, January 2005.

[129] J Masajada and B Dubik. Optical vortex generation by three plane wave interference. *Optics Communications*, 198(1-3):21–27, 2001.

[130] K. O'Holleran, M. J. Padgett, and M. R. Dennis. Topology of optical vortex lines formed by the interference of three, four, and five plane waves. *Optics Express*, 14(7):3039–3044, 2006. 18.

[131] J F Nye. *Natural Focusing and Fine Structure of Light: Caustics and Wave Dislocations*. Taylor & Francis, 1999.

[132] G. A. Swartzlander and R. I. Hernandez-Aranda. Optical rankine vortex and anomalous circulation of light. *Physical Review Letters*, 99(16):4, 2007.

[133] B. Thidé, H. Then, J. Sjoholm, K. Palmer, J. Bergman, T. D. Carozzi, Y. N. Istomin, N. H. Ibragimov, and R. Khamitova. Utilization of photon orbital angular momentum in the low-frequency radio domain. *Physical Review Letters*, 99(8):1–4, 2007.

[134] S. Sasaki and I. McNulty. Proposal for generating brilliant x-ray beams carrying orbital angular momentum. *Physical Review Letters*, 100(12):4, 2008.

[135] Masaya Uchida and Akira Tonomura. Generation of electron beams carrying orbital angular momentum. *Nature*, 464(7289):737–739, April 2010.

[136] J Verbeeck, H Tian, and P Schattschneider. Production and application of electron vortex beams. *Nature*, 467(7313):301–304, September 2010.

[137] H He, M E J Friese, Norman R Heckenberg, and Halina Rubinsztein-Dunlop. Direct observation of transfer of angular momentum to absorptive particles from a laser beam with a phase singularity. *Physical Review Letters*, 75(5): 826–829, 1995.

[138] Miles Padgett and Les Allen. Optical tweezers and spanners. *Physics World*, 10(9):35–38, 1997.

[139] N B Simpson, Kishan Dholakia, Les Allen, and Miles J Padgett. Mechanical equivalence of spin and orbital angular momentum of light: an optical spanner. *Optics Letters*, 22(1):52–54, 1997.

[140] A T O'Neil, I MacVicar, Les Allen, and Miles J Padgett. Intrinsic and extrinsic nature of the orbital angular momentum of a light beam. *Physical Review Letters*, 88(5):053601, 2002.

[141] Johannes Courtial and Miles J Padgett. Limit to the orbital angular momentum per unit energy in a light beam that can be focussed onto a small particle. *Optics Communications*, 173:269–274, 2000.

[142] M E J Friese, Halina Rubinsztein-Dunlop, J Gold, P Hagberg, and D Hanstorp. Optically driven micromachine elements. *Applied Physics Letters*, 78(4):547–549, 2001.

[143] K Ladavac and David Grier. Microoptomechanical pumps assembled and driven by holographic optical vortex arrays. *Optics Express*, 12(6):1144–1149, 2004.

[144] Jonathan Leach, Hasan Mushfique, Roberto Di Leonardo, Miles Padgett, and Jonathan M Cooper. An optically driven pump for microfluidics. *Lab On A Chip*, 6(6):735–739, 2006.

[145] M Babiker, William L Power, and Les Allen. Light-induced torque on moving atoms. *Physical Review Letters*, 73(9):1239–1242, 1994.

[146] T Kuga, Y Torii, N Shiokawa, T Hirano, Y Shimizu, and H Sasada. Novel optical trap of atoms with a doughnut beam. *Physical Review Letters*, 78 (25):4713–4716, 1997.

[147] T. A. Nieminen, A. B. Stilgoe, N. R. Heckenberg, and H. Rubinsztein-Dunlop. Angular momentum of a strongly focused gaussian beam. *Journal of Optics A*, 10(11):6, 2008.

[148] Y Zhao, J S Edgar, G D M Jeffries, D McGloin, and D T Chiu. Spin-to-orbital angular momentum of a strongly focused optical beam. *Physical Review Letters*, 99:073901, 2007.

[149] M. E. J. Friese, T. A. Nieminen, N. R. Heckenberg, and H. Rubinsztein-Dunlop. Optical alignment and spinning of laser-trapped microscopic particles. *Nature*, 394(6691):348–350, 1998. 6.

[150] D. L. Andrews, L. C. D. Romero, and M. Babiker. On optical vortex interactions with chiral matter. *Optics Communications*, 237(1-3):133–139, 2004. 20.

[151] F Araoka, T Verbiest, K Clays, and A Persoons. Interactions of twisted light with chiral molecules: An experimental investigation. *Physical Review A*, 71 (5):055401, 2005.

[152] L. Allen and Miles Padgett. Equivalent geometric transformations for spin and orbital angular momentum of light. *Journal of Modern Optics*, 54(4): 487–491, 2007.

[153] M J Padgett and L Allen. The angular momentum of light: optical spanners and the rotational frequency shift. *Optical And Quantum Electronics*, 31(1): 1–12, 1999.

[154] E J Galvez, P R Crawford, H I Sztul, M J Pysher, P J Haglin, and R E Williams. Geometric phase associated with mode transformations of optical beams bearing orbital angular momentum. *Physical Review Letters*, 90(20): 203901, 2003.

[155] B A Garetz and Arnold S. Variable frequency-shifting of circularly polarized laser-radiation via a rotating half-wave retardation plate. *Optics Communications*, 31(1):1–3, 1979.

[156] B A Garetz. Angular dopper effect. *Journal Of The Optical Society Of America*, 71(5):609–611, 1981.

[157] G. Nienhuis. Doppler effect induced by rotating lenses. *Optics Communications*, 132:8–14, 1996.

[158] J Courtial, D A Robertson, and K Dholakia. Rotational frequency shift of a light beam. *Physical Review Letters*, 81(22):4828–4830, 1998.

[159] DP Ghai, P Senthilkumaran, and R S Sirohi. Single-slit diffraction of an optical beam with phase singularity. *Optics and Lasers in Engineering*, 47 (1):123–126, 2009.

[160] H I Sztul and R R Alfano. Double-slit interference with Laguerre-Gaussian beams. *Optics Letters*, 31(7):999–1001, 2006.

[161] J M Hickmann, E J S Fonseca, W C Soares, and S Chávez-Cerda. Unveiling a Truncated Optical Lattice Associated with a Triangular Aperture Using Light's Orbital Angular Momentum. *Physical Review Letters*, 105:053904, July 2010.

[162] G Molina-Terriza, Juan P Torres, and Lluis Torner. Management of the angular momentum of light: Preparation of photons in multidimensional

vector states of angular momentum. *Physical Review Letters*, 88(1):013601, 2002.

[163] A. Mair, A. Vaziri, G. Weihs, and A. Zeilinger. Entanglement of the orbital angular momentum states of photons. *Nature*, 412(6844):313–316, 2001. 26.

[164] Eric Yao, Sonja Franke-Arnold, Johannes Courtial, Miles J Padgett, and Stephen M Barnett. Observation of quantum entanglement using spatial light modulators. *Optics Express*, 14(26):13089–13094, 2006.

[165] S N Khonina, V V Kotlyar, R V Skidanov, V A Soifer, P Laakkonen, and J Turunen. Gauss-Laguerre modes with different indices in prescribed diffraction orders of a diffractive phase element. *Optics Communications*, 175(4-6):301–308, 2000.

[166] Olof Bryngdahl. Geometrical transformations in optics. *Journal of the Optical Society of America*, 64:1092–1099, 1974.

[167] W J Hossack, A M Darling, and A Dahdouh. Coordinate transformations with multiple computer-generated optical-elements. *Journal Of Modern Optics*, 34(9):1235–1250, 1987.

[168] Gregorius C G Berkhout, Martin P J Lavery, Johannes Courtial, Marco W Beijersbergen, and Miles J Padgett. Efficient sorting of orbital angular momentum states of light. *Physical Review Letters*, 105(15):153601, 2010.

[169] M. P. J. Lavery, D. J. Robertson, G. C. G. Berkhout, G. D. Love, M. J. Padgett, and J. Courtial. Refractive elements for the measurement of the orbital angular momentum of a single photon. *Optics Express*, 20(3):2110–2115, 2012.

[170] L. Torner, J. P. Torres, and S. Carrasco. Digital spiral imaging. *Optics Express*, 13(3):873–881, 2005.

[171] Barry Jack, Jonathan Leach, J Romero, Sonja Franke-Arnold, Monika Ritsch-Marte, Stephen M Barnett, and Miles J Padgett. Holographic ghost

imaging and the violation of a bell inequality. *Physical Review Letters*, 103 (8):083602, 2009.

[172] S. Furhapter, A. Jesacher, S. Bernet, and M. Ritsch-Marte. Spiral interferometry. *Optics Letters*, 30(15):1953–1955, 2005.

[173] Grover A Swartzlander. Peering into darkness with a vortex spatial filter. *Optics Letters*, 26(8):497–499, 2001.

[174] J. H. Lee, G. Foo, E. G. Johnson, and G. A. Swartzlander. Experimental verification of an optical vortex coronagraph. *Physical Review Letters*, 97 (5):4, 2006.

[175] Grover Swartzlander, Erin Ford, Rukiah Abdul-Malik, Laird Close, Mary Peters, David Palacios, and Daniel Wilson. Astronomical demonstration of an optical vortex coronagraph. *Optics Express*, 16(14):10200–10207, 2008.

[176] Allison M. Yao and Miles J. Padgett. Orbital angular momentum: origins, behavior and applications. *Advances in Optics and Photonics*, 3:161–204, 2011.

[177] Grover A Swartzlander and C T Law. Optical vortex solitons observed in Kerr nonlinear media. *Physical Review Letters*, 69(17):2503–2506, 1992.

[178] J Arlt, K Dholakia, L Allen, and M J Padgett. Efficiency of second-harmonic generation with bessel beams. *Physical Review A*, 60(3):2438–2441, 1999.

[179] R Ghosh and L Mandel. Observation of nonclassical effects in the interference of two photons. *Physical Review Letters*, 59(17):1903–1905, 1987.

[180] A. Vaziri, G. Weihs, and A. Zeilinger. Experimental two-photon, three-dimensional entanglement for quantum communication. *Physical Review Letters*, 89(24):4, 2002.

[181] G Molina-Terriza, A Vaziri, J Řeháček, Z Hradil, and A Zeilinger. Triggered qutrits for quantum communication protocols. *Physical Review Letters*, 92 (16):167903, 2004.

[182] Julio T Barreiro, Nathan K Langford, Nicholas A Peters, and Paul G Kwiat. Generation of hyperentangled photon pairs. *Physical Review Letters*, 95(26): 260501, December 2005.

[183] E Nagali, F. Sciarrino, F D Martini, L Marrucci, B Piccirillo, E Karimi, and E Santamato. Quantum information transfer from spin to orbital angular momentum of photons. *Physical Review Letters*, 103:013601, 2009.

[184] Ebrahim Karimi, Jonathan Leach, Sergei Slussarenko, Bruno Piccirillo, Lorenzo Marrucci, Lixiang Chen, Weilong She, Sonja Franke-Arnold, Miles Padgett, and Enrico Santamato. Spin-orbit hybrid entanglement of photons and quantum contextuality. *Physical Review A*, 82(2):02215, August 2010.

[185] Jonathan Leach, Barry Jack, J Romero, Monika Ritsch-Marte, Robert W Boyd, A K Jha, Stephen M Barnett, Sonja Franke-Arnold, and Miles J Padgett. Violation of a bell inequality in two-dimensional orbital angular momentum state-spaces. *Optics Express*, 17(10):8287–8293, 2009.

[186] Jonathan Leach, Barry Jack, J Romero, A K Jha, A M Yao, Sonja Franke-Arnold, D G Ireland, Robert W Boyd, Stephen M Barnett, and Miles J Padgett. Quantum correlations in optical angle-orbital angular momentum variables. *Science*, 329(5992):662–665, August 2010.

[187] J Romero, Jonathan Leach, Barry Jack, Mark R Dennis, Sonja Franke-Arnold, Stephen M Barnett, and Miles Padgett. Entangled optical vortex links. *Physical Review Letters*, 106(10):100407, March 2011.

[188] V. D. Salakhutdinov, E R Eliel, and W Löffler. Full-field quantum correlations of spatially entangled photons. *Physical Review Letters*, 108:173604, April 2012.

[189] S. Barreiro and J. W. R. Tabosa. Generation of light carrying orbital angular momentum via induced coherence grating in cold atoms. *Physical Review Letters*, 90(13):4, 2003.

[190] Michael Faraday. Experimental researches in electricity - nineteenth series. *Philosophical Transactions of the Royal Society of London*, 136:1–20, 1846.

[191] J G Dawber. The Faraday Effect, Magnetic Rotatory Dispersion and Magnetic Circular Dichroism. *Analyst*, 89:755–762, 1964.

[192] Sonja Franke-Arnold, Jonathan Leach, Miles J Padgett, V E Lembessis, D Ellinas, A J Wright, J M Girkin, Patrik Ohberg, and A S Arnold. Optical ferris wheel for ultracold atoms. *Optics Express*, 15(14):8619–8625, 2007.