Three-dimensional Light Sculptures and their Interaction with Atomic Media

An experimentalist's guide

Adam Selyem

Submitted in fulfilment of the requirements for the Degree of Doctor of Philosophy
to the
School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

September 1, 2019

Author's declaration

The work described in this Thesis was carried at the University of Glasgow under the supervision of Dr. Sonja Franke-Arnold, School of Physics and Astronomy, in the period October 2015 to March 2019. The author hereby declares that the work described in this Thesis is his own, except where specific references are made. It has not been submitted in part or in whole to any other university for a degree.

Author's signature:

_____

Adam Selyem
September 1, 2019

Contents

List of Figures

Abstract

In recent years great progress was made in the spatial control of light with dynamic phase and amplitude modulators such as spatial light modulators and digital micromirror devices. In this work we describe the theory and practice of light shaping with such devices, detailing the spatial control of amplitude, phase and polarisation of coherent laser beams. We use our expertise in generating and measuring light fields with spatially dependent polarisation structures to characterise the correlations between spatial modes and polarisation in such light fields. We do this by adapting concurrence, a quantum measure of entanglement, to these classical correlations.

One of the most promising application of coherent laser light is in the control of atomic media via atom-light interactions. In this work we describe the construction of simple external cavity diode lasers designed for the generation of resonant light for atomic physics applications. We exploit these lasers and spatial light modulators to create and measure three-dimensional atomic population structures in a warm rubidium vapour.

We also implement a magneto-optic and a dynamic dark spontaneous-force optical trap for rubidium. These traps produce dense ($\sim 10^{11}$ cm^{-3}) and cold (~ 100 μK) clouds of rubidium atoms. We develop the theory of spatially dependent electromagnetically induced transparency in such traps using rate equations. We find that the absorption of linearly polarised light depends on the relative direction of a magnetic field and the light polarisation. We use the cold atom clouds to measure the direction of magnetic fields by using this dependence and laser beams with structured polarisation.

Acknowledgements

If six years ago you asked me what I would be doing now I would have said something along the lines of "I don't know, but definitely *not* optics!" My opinion probably wouldn't have changed without Sonja Franke-Arnold, who effortlessly convinced me that actually, optics and atoms together are the coolest thing in the world. One couldn't ask for a better supervisor, and I'll always be grateful to her for seeing something in me that I'm not sure I see even now, and giving me the opportunity that gave me the best three years of my life, at least so far.

I've learned much in the last five years, and in that time I've learned the most from Neal Radwell, who taught me everything I know about experimental science. Not only that, but he's also been a great friend. I'm not sure he knows how much comfort the many-many hours of playing with plastic spaceships and talking about games has brought, but I'm sincerely very grateful for it all.

It wasn't just two people who made every weekday in the dungeon of the Kelvin building so rewarding. I'm not sure I realise the full extent of just how lucky I was to be part of the Optics group. Every single member of the group deserves thanks for being extremely welcoming, open and friendly, providing the perfect work environment. Thanks for making going to work feel like going home! In particular, thanks to Tom Clark for being a very fun and helpful lab mate, and enthusiastic collaborator. Special thanks go to the little Optics board gaming group, particularly Daan Stellinga, Ryan Hawley, Euan Cowie and Francesco Castellucci; I've had a ton of fun, I hope you did, too.

Very special thanks to Georgi "Morwoen" Tsaklev, Reni Mihaylova and Kirsty Hair for providing an amazing home away from home that was always great to return to after work. You've all been great friends, and your support means a lot. Cheers guys!

It seems like living with Bulgarians is my thing, and never has that worked out better than with the lovely Dessy Arabadzhiyska. No-one has made me feel more appreciated. Thank you so much for being there for me! And thanks for all the gerbils.

Végül, de semmiképpen sem utolsósorban, köszönet jár szüleimnek és nagyszüleimnek, akik nélkül sose juttotam volna el idáig. Nagyon hálás vagyok minden támogatásotokért és azért, hogy végig hittetek bennem. Köszönök mindent!

To the reader

When writing this thesis it has been my goal to collect all knowledge and intuition I gained over the last five years of doing lots of experiments but not much maths. I hope that the information contained within will be useful for introducing students to the practicalities of structured light and atomic physics. Experimental skills take a while to build up in any case, but with good intuition the process is quicker. I hope I manage to convey why I think about things the way I do, and through this I might help the reader gain/deepen understanding about the topics in this thesis. Because of my intention to make this thesis somewhat of a tutorial for new students I am sure I go into apparently too much detail about some well-known optics, for example. I remember I was confused about very basic concepts, so I think erring on the side of too much information is perhaps better. But let me stress that I believe that trying things in the lab is the best and fastest way of learning and building up intuition.

Before we begin, let me say a few things about notation. There are lots of figures throughout this thesis that show optical phase in hue. The colour map in these figures is *always* the one shown figure 1 below.

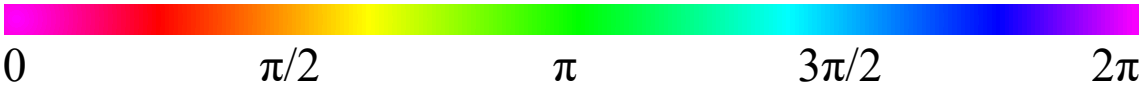

Figure 1: Phase to hue map.

There are also lots of figures with plots of polarisation across beam profiles. The polarisation in these plots is represented by a coloured polarisation ellipse drawn according to the definitions in figure 1.1. The colour scheme is shown in figure 2 below.

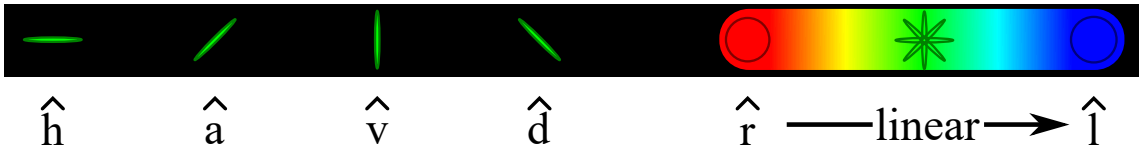

Figure 2: Polarisation plot legend.

Lastly, I would like to make a comment about footnotes. I like to put information that I feel supplements the main text but does not necessarily fit into the sentence structure into footnotes. I also put tangential comments and observations there, and also things that I find fun, funny or interesting that have nothing to do with the main text[1].

[1] Cave, hic sunt iocus.

Introduction

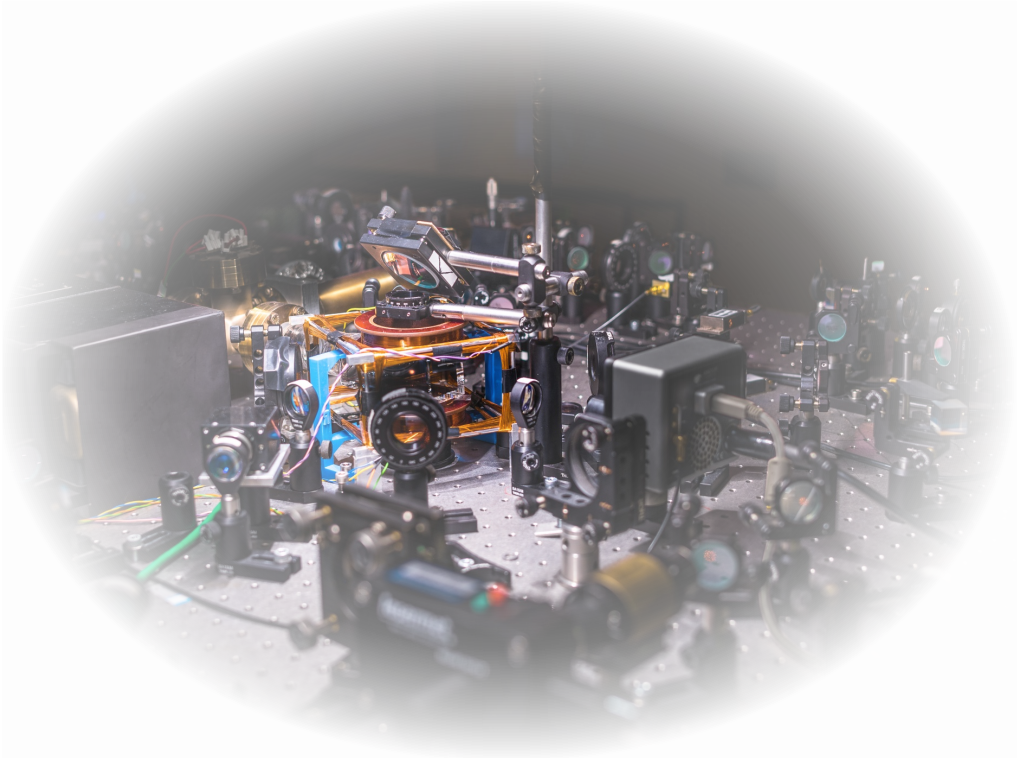

Introduction

Light is everywhere. Practically all interactions our minds have with the outside world through our bodies are mediated electromagnetically via photons, whether they are virtual or real. Every detector technology we have works mainly by electromagnetic interaction. It is not surprising that humanity's most advanced physical understanding is in the field of electromagnetism [1]. But our understanding of light is not just theoretical either. Recently we have gained exquisite control over properties of light. With tunable lasers [2] we can produce coherent electromagnetic radiation in a wide range of frequencies, a technology which has truly wide-ranging applications from communication [3] to gas sensing [4]. With pulsed lasers we can measure distances to objects in the environment extremely precisely by measuring the time-of-flight of pulses scattered from them [5, 6]. Because of the relatively narrow frequency band emitted even by pulsed lasers such measurements are pretty much independent from background illumination as well. Through frequency and pulse length (which are linked to each other through an uncertainty relation, connecting the pulse length to the inverse bandwidth [7]) we have control over light in the time domain, although for the work in this thesis this will not be very important.

Of much more relevance is the spatial structure of light. We are used to some rudimentary form of structured light in everyday life, since our eyes detect spatial information from the environment because the colour and intensity of light arriving to our retinas carries spatial information. But there is more to light than colour and intensity. Polarisation is undetectable to us, but apparently bees, for example, see polarisation and use it to navigate [8]. We exploit polarisation in our liquid crystal display technology and communication [9], just to name a few. The phase of electromagnetic radiation is another property we cannot see, and it only becomes important for coherent radiation. There it leads to interference effects and affects propagation. Both polarisation and phase are properties of light that change with spatial position, leading to interesting effects on propagation.

Considering beams of coherent, monochromatic laser light, which we will do throughout this thesis, intensity, phase and polarisation can all vary in a plane transverse to the propagation direction, more or less independent from each other. Of course, constraints need to be placed on them to be physically realisable. They also evolve along the propagation direction in a well defined manner, leading to structures in three dimensions.

In Part I of this thesis we will think about these properties of light, and look at ways of controlling their structure in experiments. In chapter 1 we introduce the basics of the theory of electromagnetic waves, including amplitude (which relates to intensity), phase and polarisation. We look at what structures they take in conventional laser beams, and then we consider ways of manipulating them with common, simple optical elements, such as lenses, gratings and polarisation optics.

In chapter 2 we introduce two devices that have become rather popular in the optics community in

recent years because of their versatility. Spatial light modulators (SLMs) and digital micromirror devices (DMDs) are digital devices originally developed for overhead projectors. They offer high resolution control over the phase of coherent light, which we show to be sufficient for control of amplitude and polarisation as well, with some cleverness. We go into quite a bit of detail about experimental techniques involved in their calibration and operation. We also discuss the design of digital holograms that these devices can display.

There is, of course, purpose behind gathering all this knowledge. In chapter 3 we use our expertise in the creation and characterisation of light fields to investigate correlations between space and polarisation in beams with structured polarisation. With the help of our collaborators we show that concurrence, a measure of entanglement in quantum mechanics, can be adapted to describe such classical correlations. We also relate concurrence to easy-to-measure Stokes parameters. We hope that this technique can enable quantum-like technologies that rely on correlations, but not non-locality.

But light does not exists only in a vacuum. Matter is also everywhere, and without interaction between light and matter we would not know anything about the universe. In everyday life we usually experience the effects of light interacting with bulk solids[2], but the interactions between single atoms (or ensembles atoms in a vapour) can become excitingly different. The electronic structure of atoms of the same species is identical, which is useful in two ways. If we know how one reacts to external influences, we know how all of them will react. They can also act collectively, although this we will not consider in detail here.

Our models for the electronic structure of atomic species with a few electrons in their outermost shell are excellent. This understanding is exploited in, for example, the definition of time via atomic time standards [10], precision measurements of gravity [11] and magnetic fields [12], and quantum computing [13]. In this thesis we are more interested in the interaction of extended atomic vapours with spatially structured light.

First we need to understand how atoms react to light. Chapter 4 introduces the electronic structure of rubidium, our atom species of choice. We then describe the most basic atom-light interactions, absorption and emission. These processes form the basis of all of the atomic physics discussed in later chapters. As we will see, the frequency of light is of great importance in such interactions, because of the need to match the energy of photons in the light beam to the separation of energy levels in the atom.

For this reason in chapter 5 we detail the construction and operation of sources of frequency-tunable laser light in the form of external cavity diode lasers (ECDLs). Constructed from a laser diode and a grating, ECDLs are cheap and simple to construct. We describe two types of ECDLs that were used in our experiments. We also consider how to stabilise their frequency to that of atomic transitions via Doppler-free spectroscopy. We also discuss the use of acousto-optic modulators (AOMs) for frequency control and switching.

Chapter 6 brings together our beam shaping expertise from Part I and our understanding of atomic physics (and a bit of rudimentary thermodynamics) for an experiment where we controlled the occupation of atomic states in 3D in a warm rubidium vapour. We describe the population structures theoretically using rate equations. In the experiment we used a concept similar to electron shelving to measure the populations we inscribed in the vapour, and we get 3D information

[2]Or liquids (I guess glass also counts here even though it looks more solid). And a bit of plasma, if we look up at a clear sky. Not very frequently in Scotland, then.

by tomographic reconstruction from fluorescence images. We hope that this work will be useful for designing 3D quantum memories.

Most high-precision atomic physics experiments benefit from long coherence times that arise at low temperatures and densities. We are more interested in using the slow movement of cold atoms during their interaction with structured light. For this reason in chapter 7 we introduce methods for cooling and trapping a cloud of atoms in a magneto-optical trap (MOT) and a spontaneous force optical trap (SpOT). We consider both the theory and practice of operating such traps, with particular attention to controlling magnetic fields.

These magnetic fields play a crucial role in the experiments described in chapter 8. These experiments involve controlling the transparency of a cloud of atoms to resonant laser light via electromagnetically induced transparency (EIT). We present two theoretical frameworks we used to explore the absorption of a beam with structured polarisation in the presence of a magnetic field in a 4-level tripod atomic structure. In the specific example of rubidium 87 such a level structure is realised with three Zeeman sublevels of a hyperfine ground state, along with one excited state. First we consider the atoms quantised along the propagation direction of the beam. In this frame a transverse component of the magnetic field couples the three ground states, making the absorption dependent on the spatially varying phase difference between the two circular polarisation components of the laser beam. We find the absorption by evaluating Fermi's golden rule.

In the second model we consider the atoms quantised along the magnetic field direction. In this context we describe the absorption in terms of optical pumping between the three ground states, and derive an expression for the absorption from rate equations. We show that the absorption patterns depend on the magnetic field direction, allowing us to construct an atomic compass. We expect that this experiment will be extended to utilise EIT for storage of spatially structured quantum information.

In the final chapter we briefly introduce two small projects we took part in that are tangentially related to the core of our work. In one experiment we used a single trapped ytterbium ion in an attempt to measure the longitudinal polarisation component that arises when a beam with structured polarisation is strongly focused. Out of interest we introduce the theory of ion trapping and cooling, which operates on similar principles to the atom traps described in chapter 7, but afford even better control. Lastly we describe a short investigation into the diffraction of structured polarisation from double slits in different geometries.

Articles published (or submitted) by the author during the work described in this thesis

N. Radwell, R. F. Offer, A. Selyem, and S. Franke-Arnold, Optimisation of arbitrary light beam generation with spatial light modulators, *Journal of Optics*, vol. 19, no. 9, p. 095605, 2017.

A. Selyem, S. Fayard, T. W. Clark, A. S. Arnold, N. Radwell, and S. Franke- Arnold, Holographically controlled three-dimensional atomic population patterns, *Optics Express*, vol. 26, no. 14, pp. 18 51318 522, 2018.

N. Radwell, A. Selyem, L. Mertens, M. P. Edgar, and M. J. Padgett, Hybrid 3d ranging and velocity tracking system combining multi-view cameras and simple lidar, *Scientific Reports*, vol. 9, no. 1, p. 5241, 2019.

A. Selyem, C. Rosales-Guzmán, S. Croke, A. Forbes, and S. Franke-Arnold, Basis independent tomography of complex vectorial light fields by stokes projections, *arXiv preprint* arXiv:1902.07988, 2019.

Part I

Light

Chapter 1

Amplitude, phase and polarisation

1.1 Introduction

In its classical description, light is considered a travelling transverse wave of the electromagnetic field. This has famously been derived from first principles by Maxwell. Maxwell's equations, in the absence of any charges, can be written as

$$\begin{aligned} \nabla \cdot \mathbf{E} &= 0 & \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t}, \\ \nabla \cdot \mathbf{B} &= 0 & \nabla \times \mathbf{B} &= \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}, \end{aligned} \tag{1.1}$$

where $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic vector fields, respectively, and μ_0 and ϵ_0 are the permeability and permittivity of free space, respectively. From these equations, using simple vector calculus, one can obtain equation (1.2),

$$\frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} - \nabla^2 \mathbf{E} = 0, \tag{1.2}$$

where we have defined $c = \mu_0 \epsilon_0 \approx 2.998 \times 10^8$ ms^{-1}, the speed of light in vacuum. We recognise this equation as a wave equation for the electric field. An equation identical in form to (1.2) exists for the magnetic field as well, but for all purposes in this thesis it is sufficient to consider the electric field component. Some solutions of this differential equation take the form of a wave travelling in the $\hat{r}$ direction,

$$\mathbf{E}(\mathbf{r}, t) = E(\mathbf{r}) \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega t + \phi_0(\mathbf{r}))]\hat{e}(\mathbf{r}), \tag{1.3}$$

where $\mathbf{r} = [x, y, z]$ and t are spatial and temporal coordinates. $\mathbf{k}$ is the wave vector, whose magnitude is $|\mathbf{k}| = \omega/c = 2\pi/\lambda$, and its direction is perpendicular to surfaces of constant phase $\phi = \mathbf{k} \cdot \mathbf{r} - \omega t + \phi_0(\mathbf{r})$. Phase should not be confused with the phase offset $\phi_0(\mathbf{r})$, the phase at $x = y = z = t = 0$. $\hat{e}$ is the unit vector defining the direction in which the electric field oscillates, known as the polarisation direction. These are all important parameters, but in the following we will pay particular attention to the control of the spatially dependent terms, the field amplitude $E(\mathbf{r})$, the phase offset $\phi(\mathbf{r})$ and the polarisation $\hat{e}(\mathbf{r})$. This is because we notice that fixing the values of $E_{x,y,z}$ and $\phi_{x,y,z}$ at an arbitrary position 0 along the propagation axis z fully determines their values at any position in space.

We also note that linear superpositions of solutions of the form of equation (1.3) are also solutions

to equation (1.2). There are some interesting consequences to this fact. Let us consider the sum of a set of a very simple forms of equation (1.3) in 1 dimension, at an arbitrary time $t = 0$,

$$\mathbf{E}(z) = \Sigma_i E_i \sin(2\pi z/\lambda + \phi_i)\hat{e_i}, \tag{1.4}$$

where we have assumed for simplicity that the wavelengths are all the same (which will be the case throughout this thesis).

1.2 Polarisation

In the very simple case where $\hat{e_i}$ are the same in equation (1.4), the waves simply interfere (complex amplitudes are added). If, however, $\hat{e_i}$ have orthogonal components, the total electric field oscillation will have a different direction, perhaps even changing as the waves propagate. Since in the paraxial limit the electric field oscillates in the plane transverse to the propagation direction, it can be decomposed into two orthogonal components, usually in the x and y directions, with amplitudes $E_{x,y}$ and phases $\phi_{x,y}$,

$$\mathbf{E} = \begin{bmatrix} E_x e^{\phi_x} \\ E_y e^{\phi_y} \end{bmatrix}. \tag{1.5}$$

We define six polarisation basis vectors using this principle:

$$\begin{aligned}
\hat{h} &= \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \hat{v} &= \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \\
\hat{a} &= \frac{1}{\sqrt{2}}(\hat{h} - \hat{v}) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}, & \hat{d} &= \frac{1}{\sqrt{2}}(\hat{h} + \hat{v}) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \\
\hat{l} &= \frac{1}{\sqrt{2}}(\hat{h} - i\hat{v}) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}, & \hat{r} &= \frac{1}{\sqrt{2}}(\hat{h} + i\hat{v}) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix},
\end{aligned} \tag{1.6}$$

where the inclusion of the imaginary number i corresponds to a phase shift of $\pi/2$, and we assert that the polarisation direction is transverse to the propagation direction, so we omit the z-components of the vectors. These basis vectors are illustrated in figure 1.1. In Jones calculus matrices describe the action of optical elements on the polarisation of light [14]. These matrices multiply the electric field vectors written in the form of equation (1.5) to produce the output electric field vector. The most common polarising elements are polarisers and waveplates, described in the next few paragraphs.

When thinking about light interacting with polarisation optics it is useful to decompose the polarisation of the incoming light field into parallel and perpendicular components with respect to the axis (axes) of the optic. Polarisers transmit the polarisation component of an incoming light field that is parallel to their axis, and absorb (or reflect) the perpendicular polarisation. In experiments polarisers are characterised by an extinction ratio, which describe how much of the "wrong" polarisation is transmitted - the lower this number, the better, but no real polariser block all of the orthogonal polarisation. They may also absorb some of the polarisation component they are meant to transmit, and could also introduce phase shifts. Global phase shifts common to both polarisation components would not matter much unless a polariser was placed into an interferometer. In Jones calculus, a polariser with its axis oriented horizontally is represented by the matrix $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$.

Waveplates, or retarders, are usually made from a birefringent material with a fast and a slow axis. As the names suggest, light whose polarisation is parallel to the fast axis travels faster through

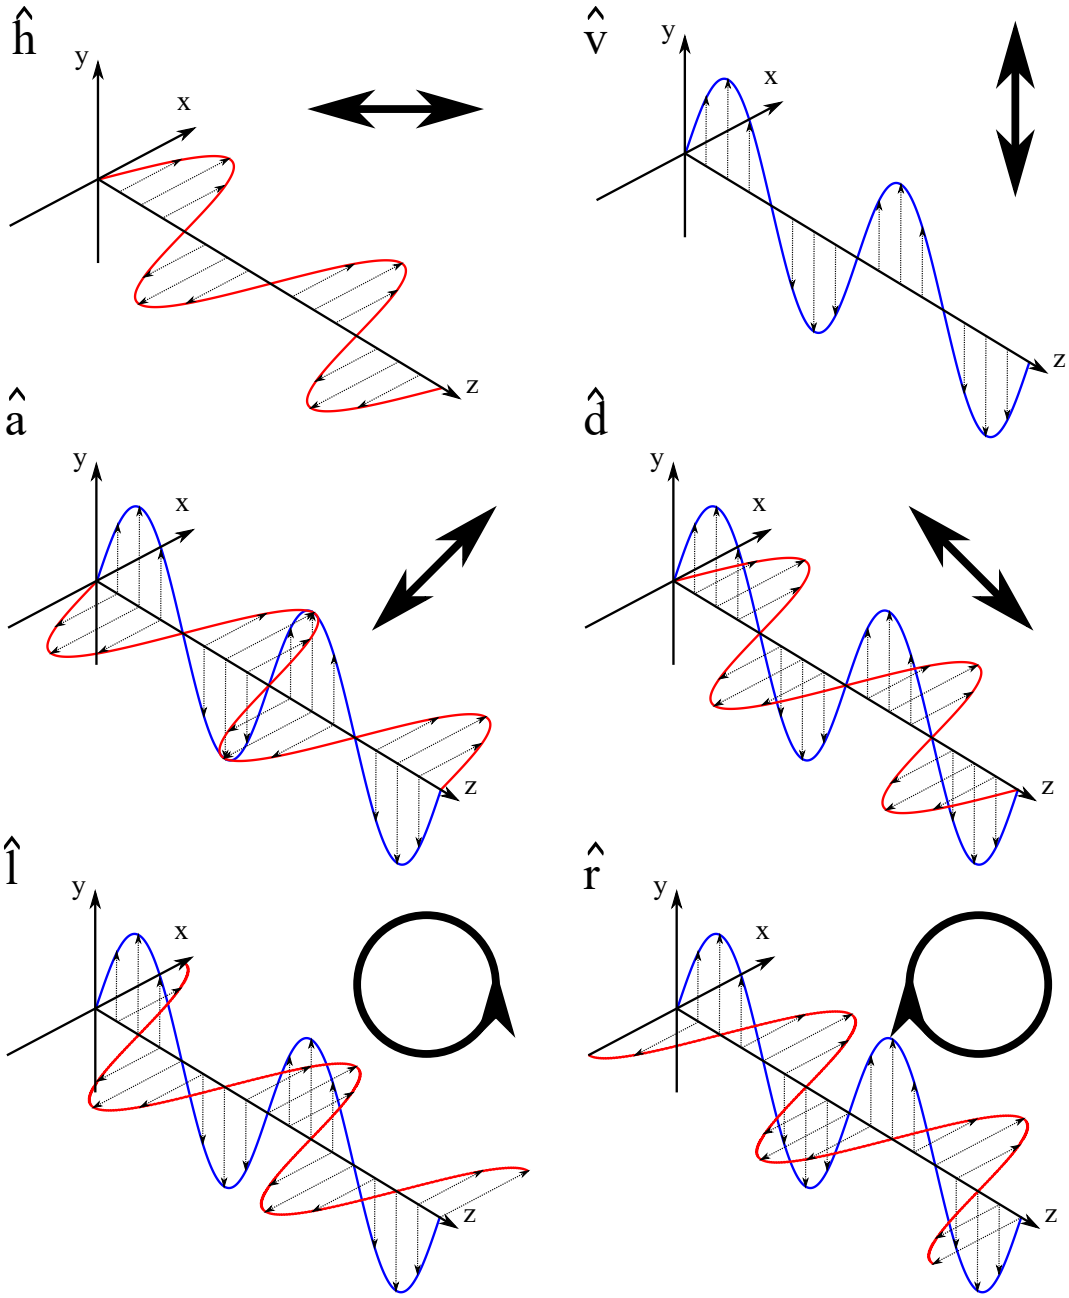

Figure 1.1: Definitions of polarisation directions used throughout this thesis. Polarisation directions are viewed along the propagation direction, but the naming convention follows the what the observer looking against the propagation would see.

the waveplate than the orthogonal polarisation. By appropriately choosing the thickness of the material one can introduce any phase shift to one polarisation component with respect to the other. This phase shift is given by

$$\Delta\phi = \frac{2\pi\Delta n L}{\lambda}, \tag{1.7}$$

where $\Delta n = n_{\text{slow}} - n_{\text{fast}}$ is the difference between the refractive indices of the material along the slow and fast axes, and L is the thickness of the waveplate.

In practice one finds two types of waveplates, quarter- and half-wave plates (abbreviated as $\lambda/4$ and $\lambda/2$ plates, respectively). A $\lambda/4$ plate a introduces $\pi/2$ phase shift. This turns an input linear polarisation at 45° to the plate axes into circular polarisation, input circular polarisations into linear, and input linear polarisations at angles other than $m \times 45°$ ($m \in \mathbb{Z}$) are turned into elliptical polarisation. A $\lambda/4$ plate at an angle α to the horizontal is represented by the matrix

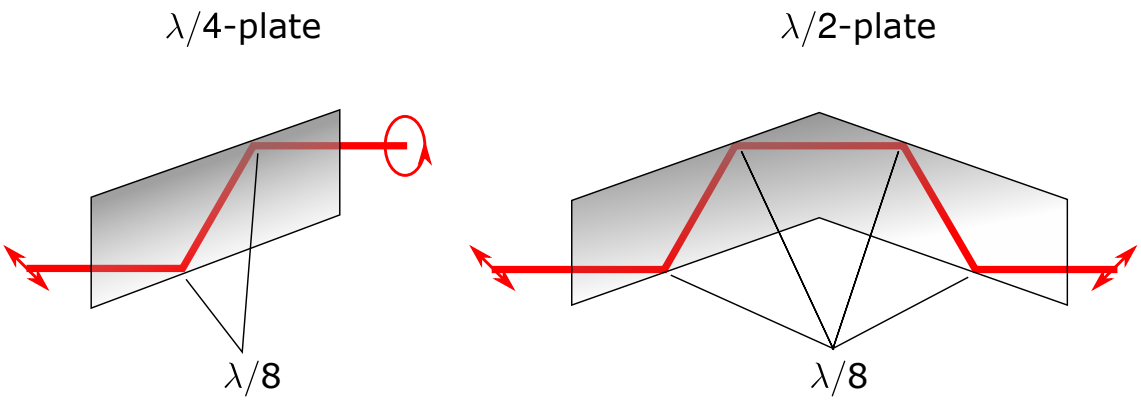

Figure 1.2: Quarter-wave and half-wave Fresnel rhombs. Their action is illustrated for input light at 45°.

$$\exp\tfrac{-i\pi}{4}\begin{bmatrix}\cos^2\alpha + i\sin^2\alpha & (1-i)\sin\alpha\cos\alpha \\ (1-i)\sin\alpha\cos\alpha & \sin^2\alpha + i\cos^2\alpha\end{bmatrix}.$$

A $\lambda/2$ plate introduces a π phase shift between the polarisation components along its axes - this results in a rotation of the input linear polarisation by an angle given by 2α, where α is the angle between the input transverse polarisation and the waveplate fast axis. If the input polarisation is elliptical or circular, the handedness is also reversed. The Jones matrix for a $\lambda/2$ plate is $\exp\frac{-i\pi}{2}\begin{bmatrix}\cos 2\alpha & \sin 2\alpha \\ \sin 2\alpha & \cos 2\alpha\end{bmatrix}$.

In practice, for optical wavelengths the material thickness required for a waveplate to shift exactly $\lambda/4$ (or $\lambda/2$) would be so thin (a few 100s of nanometres) that handling them would be impossible. Commercially available birefringent waveplates usually come instead in two varieties, multi-order and zero-order. A multi-order waveplate is a single slab of birefringent material cut to a thickness such that $\Delta\phi = m2\pi + \lambda/4$ for a given wavelength. A zero-order waveplate is constructed of two multi-order waveplates with their axes aligned 90° with respect to each other. This leads to a total phase shift that is the difference of the shift due to each component waveplate, and can be engineered such that $\Delta\phi = \lambda/4$. These two flavours of waveplates are equivalent for the wavelength for which they are designed, but for a multi-order waveplate the cumulative slightly wrong shifts at a wavelength even minimally different from the design one leads to behaviour that is hard to predict, and almost certainly wrong. In contrast, zero-order waveplates perform pretty well for a relatively larger range of wavelengths (although it is still a bad idea to use them more than a few tens of nanometres off the design wavelength).

There exists an entirely different class of retarders, however, which are effective over a huge range of wavelengths. These are Fresnel rhombs, which exploit the phase shift arising on total internal reflection between the s and p polarisation components (which are perpendicular and parallel to the reflecting surface, respectively). This phase shift is controlled by the rhombs' internal angle, which is designed such that the phase shift is very close to $\lambda/8$ for visible light for most glasses, so two reflections can be used to generate a $\lambda/4$ retardance, or four reflections for $\lambda/2$. Such devices are shown in figure 1.2.

1.3 Measuring polarisation: Stokes parameters

We measure polarisation using Stokes measurements. This involves projecting the light field to be measured onto each of at least 4 of the polarisation basis vectors (e.g. $\hat{h}, \hat{v}, \hat{d}, \hat{r}$), and measuring

intensity in these projection. In this work we usually project onto all basis vectors for conceptual simplicity. From the measured intensities we can calculate the four Stokes parameters,

$$\begin{aligned} S_0 &= I = I_h + I_v = I_d + I_a = I_r + I_l \\ S_1 &= Q = I_h - I_v \\ S_2 &= U = I_d - I_a \\ S_3 &= V = I_r - I_l. \end{aligned} \tag{1.8}$$

From these expressions it is straightforward to see why four measurements are sufficient, since, for example, $S_2 = I_d - (S_0 - I_d)$. So we need to measure in any pair of orthogonal polarisations to get S_0, and then in one each from the other two pairs. Another way of thinking about this is that there are four parameters of the light field that we need to determine (E_x, E_y, ϕ_x and ϕ_y), so we need four measurements that give information about each one. We can then construct a Stokes vector

$$\mathbf{S} = \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix}, \tag{1.9}$$

which is usually normalised such that $S_0 = 1$. It can represent any polarisation state (including partially polarised states, for which $\sqrt{S_1^2 + S_2^2 + S_3^2} < 1$). Pure polarisation states live on the surface of a sphere known as the Poincaré sphere, shown in figure 1.3, while partially polarised states are inside the sphere. Pure polarisation states can be written also in the Jones formalism, although there is no unique way to compute a Jones vector from a Stokes vector, because of the existence of an arbitrary global phase. If we set $\phi_x = 0$, then

$$\begin{aligned} |E_x| &= \sqrt{\frac{S_0 + S_1}{2}}, & \phi_x &= 0 \\ |E_y| &= \sqrt{\frac{S_0 - S_1}{2}}, & \phi_y &= \tan^{-1}\left(\frac{-S_3}{-S_2}\right). \end{aligned} \tag{1.10}$$

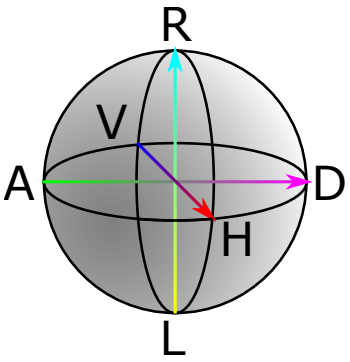

Figure 1.3: The Poincaré sphere, showing the relations between polarisation basis vectors. The axes shown in colour also correspond to Stokes parameters S_{1-3}.

1.4 Amplitude and phase

So far we have talked quite a bit about the relative phase between orthogonal polarisation components, but variations of the phase across a beam profile is also very much worthy of attention, since this is the property that the most ubiquitous conventional optical elements modify, and it has a huge effect on the propagation of light. Amplitude is perhaps the most obvious property of light, being the one that our eyes are sensitive to[1]. It is also in a sense inseparable from phase due to diffraction.

For now let us consider a uniformly polarised light beam, $\hat{e}_{\mathbf{r}} = \hat{e}_{\mathbf{0}}$. Then the simplest solution of equation (1.2) is

$$\mathbf{E}(x, y, z) = E_0 \sin(2\pi z/\lambda + \phi_0)\hat{e}_0, \tag{1.11}$$

[1]Strictly speaking most detectors, including our eyes, are sensitive to the intensity $I = |E_{\mathbf{r}} \exp(i\phi_{\mathbf{r}})|^2$, which is why the phase information is lost.

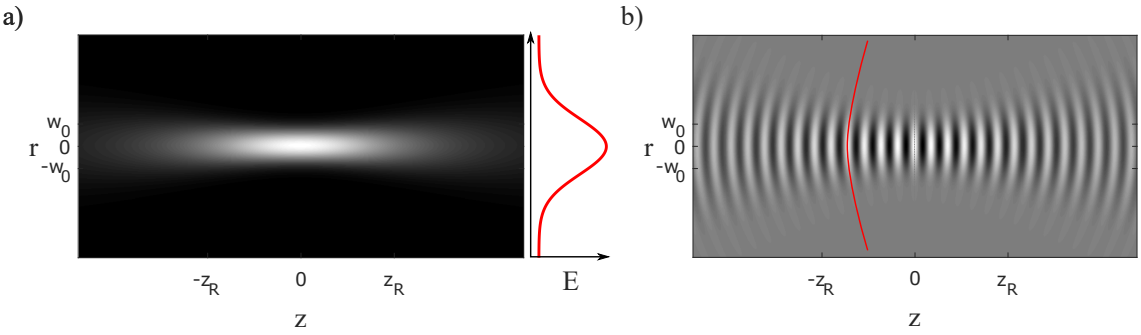

Figure 1.4: Propagation of a Gaussian beam, showing a) intensity and b) electric field amplitude, with $\lambda = 4/3w_0$. An example of a curved phase front is shown in red.

which has one amplitude and phase in any transverse plane, and is known as a plane wave solution. Unfortunately, while easy to think about, it is not physically realisable. This is easy to see from the fact that the energy content in any transverse plane is infinite, because the electric field amplitude is finite extending to infinity. Therefore in realistic solutions the amplitude must go to zero at infinity. The most common solution that behaves like this is the well-known Gaussian beam [15], written here in radial coordinates,

$$\mathbf{E}(r,z) = E_0 \frac{w_0}{w_z} \exp\left(\frac{-r^2}{w_z^2}\right) \exp\left(-i\left(kz + k\frac{r^2}{z\left(1+z_R^2/z^2\right)} - \psi_z\right)\right)\hat{e}_0, \tag{1.12}$$

where w_z is the radius in a transverse plane where the amplitude falls to $1/e$ of its axial value, and its value at $z = 0$ is the waist w_0. $\psi_z = \arctan(z/z_R)$ is known as the Gouy phase arising from a change in phase velocity near the waist and $z_R = \pi w_0^2/\lambda$ is the Rayleigh range. The amplitude cross-section of such beams is a Gaussian that changes only in size as it propagates due to diffraction. It should be noted that expressions of this type describe modes of propagation only in the paraxial limit, when the angle that light rays make with the optical axis are small. Phase fronts, surfaces of constant phase, form parabolic surfaces with the exception of the waist plane, where the phase front is a plane. This is shown in figure 1.4. In general the smaller the waist, the larger the wavefront curvature and hence the divergence of the beam. This is a very general statement for coherent light: a small light structure (whether it is a beam, light through an aperture or anything else) always diverges faster than a larger one, provided that they don't have some weird phase structure[2].

From figure 1.4 we can already see the intimate relationship between beam size and phase front curvature for Gaussian beams. When thinking about changing beam size one optical element in particular comes to mind, the lens. In fact, the operation of a lens is exactly the application of a curvature to the phase front of the incoming beam, which results in the focusing (or defocusing) of the beam. Examples of the phase profile associated with a lens are shown in figure 1.5. A particularly interesting application of the fact that all that matters for a lens to function is the phase curvature is a Fresnel lens. Such a lens is designed to minimise the use of glass in the formation of the lens, and looks very much like figure 1.5c (although the segments of course have many wavelengths of thickness, because of both the refractive index of glass responsible for the phase shifting and practicality).

For subsequent parts of this thesis it is crucial to understand the action of a lens in slightly more detail. In particular, we need to look at the relationship between the focal planes in front and behind a lens. Optical fields can be decomposed into plane waves with unique **k**-vectors,

[2]... and the structure in question is not much smaller than the wavelength of light, but in that case the paraxial approximation breaks down anyway and more careful treatment is required.

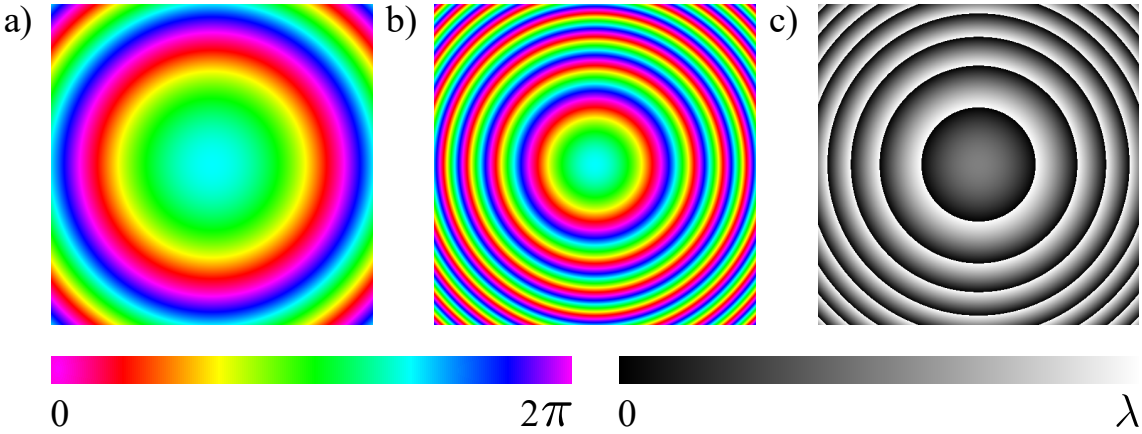

Figure 1.5: Phase profiles of different lenses. The profiles are plotted over 1 mm^2. Focal lengths are a) 200 mm and b,c) 50 mm. The lens profile modulo wavelength is shown in c).

determining their propagation direction. It can be seen from simple ray diagrams that a lens maps each incoming $\mathbf{k}$ to a unique position in its back focal plane. The reverse is also true: it maps each position in its front focal plane to a unique direction in its back focal plane. This already indicates that these planes are Fourier transforms of each other, and this can be shown rigorously. Thus a lens maps position space onto direction space (with a scaling factor proportional to its focal length). It is then no surprise that a large Gaussian beam arriving at a lens focuses to a smaller Gaussian beam - that is exactly what a Fourier transform does! What may be slightly surprising at first sight is that a small Gaussian beam focuses to a large one, but even this is not unexpected when one considers the effect of diffraction mentioned above. These concepts are beautifully intertwined, and I hope that the reader appreciates this even without showing rigorous mathematics; intuition that takes all this into account takes a while to build up, but is a powerful tool indeed.

On a more practical note, aligning lenses with laser beams is quite easy. A lens should always be centered on the laser beam, perpendicular to its propagation direction. The latter can be a bit hard to judge, but usually small deviations, on the scale of a few degrees, do not matter much in practice. To center the lens on the beam, place a screen (piece of paper, detector card, pinhole or even a camera) relatively far from the intended lens position further along the beam. Note the position where the beam hits the screen without the lens. Now insert the lens in such a way that the beam is centered on the same position, and the lens is not tilted appreciably away from being perpendicular to the beam. Determining the centre of the beam after the lens has been inserted might be tricky if the beam expanded a lot by the time it reaches the screen. If this is the case, one can work backwards towards the laser source when placing lenses, starting from the one closest to the monitoring screen, always maintaining the beam position on the screen.

1.5 Imaging

It is also well known that lenses are capable of imaging, i.e. relaying optical fields from one plane to another. Two lenses can relay any optical field from one plane to another if positioned properly, but one lens can only reproduce intensities correctly. As discussed above, a lens Fourier transforms the optical field at its front focal plane to its back focal plane (also called far field or Fourier plane), and placing another lens such that its front focal plane coincides with the Fourier plane of the first allows us to undo that first Fourier transform, up to an unimportant global phase. This configuration is known as a telescope, or a 4f imaging system (because of the 4 focal lengths between the object and image planes, see figure 1.6). It can change the transverse scale

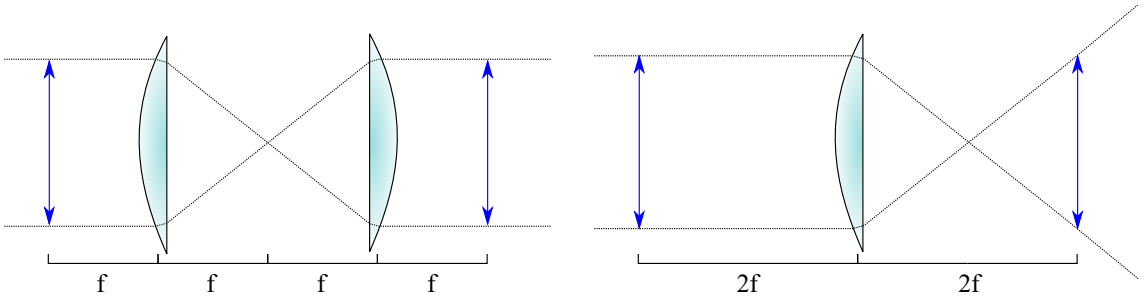

Figure 1.6: Imaging with lenses. Both the telescope (left) and single lens (right) imaging systems are shown in the 1:1 magnification configurations.

of the formed image proportionally to the ratio of the focal lengths of the two lenses, producing a magnification of f_2/f_1, where $f_{1,2}$ are the focal lengths of the first and second lens, respectively.

A single lens of focal length f can also reproduce the intensity distribution that is a distance p in front of the lens at a plane q behind it, with a magnification of q/p, obeying $1/f = 1/p + 1/q$. This is, in principle[3], how single lens reflex (SLR) cameras form images. We emphasise the fact that this configuration only reproduces the intensity correctly, because the phase acquires an extra curvature due to the lens, and so the propagation properties of the light field will be different after the image plane. This is not the case in the telescope configuration, which reproduces phase as well as intensity correctly.

Of course, in reality imaging systems can not reproduce the object plane exactly in the image plane because of the finite size of the lenses, and their imperfections. The former leads to the limited spatial resolution of the imaging system, often expressed as the Rayleigh resolution limit, which states that two point sources are resolvable if their angular separation is

$$\theta \geq 1.22\frac{\lambda}{D}, \tag{1.13}$$

where D is the diameter of the lens. This expression is valid as long as the resolution is limited by the finite size of the aperture.

The other reason lenses in practice often do not image perfectly is that they usually have some sort of aberration. The most common of these is spherical aberration, arising because most lenses are ground to have a spherical surface as opposed to the ideal hyperbolic surface[4]. The effect is that light rays entering the lens far off-axis do not focus to the same position that rays closer to the axis do. A spherical lens thus behaves very close to a hyperbolic one only close to its center because the two shapes are very similar there. Over what range this is true depends on the focal length and diameter of the lens; standard 1 inch diameter long focal length lenses ($f > 40$) tend to be very good approximations of the hyperbolic surface over their entire diameter, but short focal length lenses should be used with caution.

1.6 Mirrors

Like a lens, a simple tilted mirror is also a spatially dependent phase shifter, because different parts of the beam arriving at a tilted mirror encounter the mirror surface at different phases.

[3]In reality SLR lenses are much more complicated designs, even fixed focal length lenses contain many lens elements to reduce aberrations.

[4]The reason for this is simply that it is much easier to do.

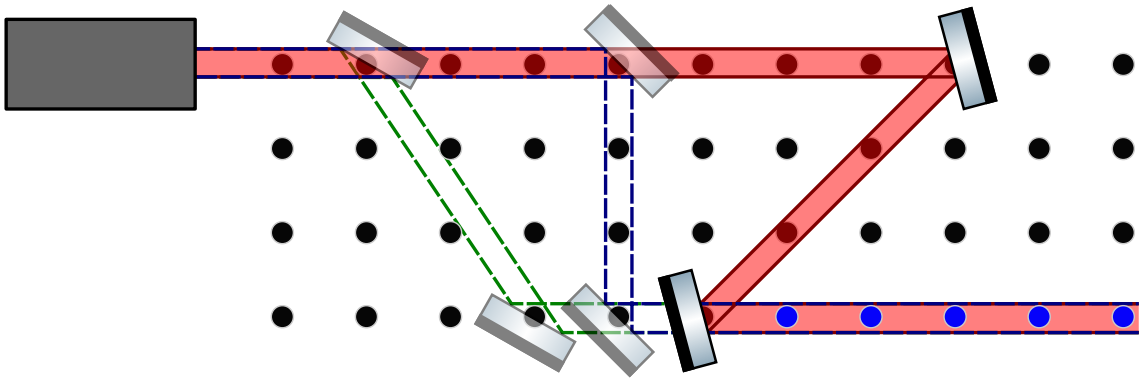

Figure 1.7: An example of using two mirrors to direct a laser beam over a set of holes in an optical table (marked blue). The angle and position of the first mirror selects the position in the $x - y$ plane where the second mirror needs to be placed. The angle for the second mirror is also determined by the first. A few options are shown in different colours.

Mirrors themselves are extremely ubiquitous on optical benches because of logistics reasons. A pair of mirrors allow the precise alignment, both in position and in direction, of laser beams through any point of choice in space[5]. Two mirrors are usually needed because there are two degrees of freedom (position and direction) of the laser beam that they need to modify, which would be coupled if only one mirror was used. This is illustrated in figure 1.7. As the figure shows, there is usually a large number of options for the placement and angles of the mirrors. The choice from these options the experimentalist has to make is informed by other optical elements on the bench and properties of the actual mirrors.

Mirrors, especially dielectric ones, cause phase shifts between s- and p-polarised components (that is, polarisation components perpendicular and parallel to the plane of reflection). Some of this is inevitable; even an ideal mirror phase shifts reflected light by π. This is not an issue as long as the polarisation is purely s or p in the plane of reflection. This usually means horizontal or vertical polarisations if the mirrors deflect beams in the plane of the optical bench surface, the most common scenario. A second mirror can undo the ideal phase shift, but if the mirrors are not ideal then this compensation is not perfect, and the output polarisation can be altered. For this reason mirrors are designed to be used at certain angles (usually either 45° to the beam path, or nearly perpendicular to it) at which such phase shifts are minimised.

Dielectric mirrors are designed to operate within a narrow band of wavelengths and for a small range of angles, usually either close to perpendicular or at 45° to the beam propagation, where they provide exceptional reflectivity. Metallic ones perform better at larger range of angles and have much less stringent requirements on their operating wavelength, but can be harder to handle. Metallic mirror surfaces are susceptible to oxidisation, especially in high humidity, which severely reduces their reflectivity. They are therefore supplied with protective coating, which is very delicate and so their reflective surface should not be touched. Dielectric mirrors are usually not permanently damaged by touch[6].

When designing optical paths it is generally a good idea to try and make sure that beams are travelling parallel to lines of holes in the optical table, and keep its height constant. There is no physical reason for this, it just makes it easier to follow where beams are going in a complicated setup.

[5]Provided that they physically fit where they need to go to allow the required alignment, in the sense that they do not overlap other optical elements.

[6]Although they do need cleaning if such a thing happens.

1.7 Gratings

It is well-known that diffraction gratings disperse polychromatic light into a spectrum that appears in discrete diffraction orders. Throughout most of this thesis (apart from a brief discussion of external cavity diode lasers in chapter 5) we ignore this dispersion because we work with monochromatic CW laser light. We note that when dealing with polychromatic light, for example ultrashort pulses, dispersive elements introduce further complications such as pulse front tilt [16]. Here instead we exploit the fact that gratings can change the propagation direction of light.

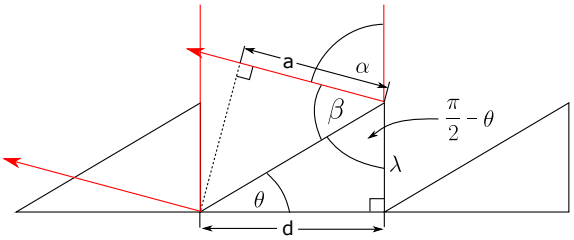

Figure 1.8: Rays (red) reflected from a reflective diffraction grating (black) of depth λ, period d and blazing angle θ.

With a little effort it can be shown that the first diffracted order of blazed grating with depth λ and blazing angle θ appears at the same angle as if it was reflected off a mirror at the same angle. The situation is shown in figure 1.8. Two rays (shown in red) interfere constructively in the far field if their path difference is a multiple of the wavelength. We have chosen two rays that arrive on the grating surface with λ path difference, so that we only need to consider the reflected path difference, denoted a. Thus we require that, for the first diffracted order, $a = \lambda$. Through the appearance of similar triangles this implies that $\beta = \pi/2 - \theta$. Then the first diffracted order is deflected by an angle $\alpha = \pi - (\beta + \pi/2 - \theta) = \pi - 2(\pi/2 - \theta) = -2\theta$, which is exactly the expression for a mirror tilted at an angle θ.

1.8 Propagation

The reason we pay so much attention to the phase and amplitude of a beam in a transverse plane is that the field at any value of z is fundamentally linked to any plane at a different z. One can obtain the field in the transverse plane at z from a known one at some z_0 by propagation. We do this by exploiting the fact that the Fourier plane of a beam is equivalent to its plane wave decomposition, that is, every point in the Fourier plane contains the amplitude and phase of one of the plane wave components of the beam [17]. Thus we can take the 2D Fourier transform of a transverse plane of a beam we wish to propagate, advance the phase of the plane wave components appropriately, and take the inverse Fourier transform of the modified far field to get the propagated beam. This method is known as angular spectrum wave propagation [18]. What is left for us to do is to describe how to calculate the extra phase for the different plane wave components.

Let us reiterate that every point in the Fourier plane corresponds to one of the plane wave components of the optical field at z_0. The plane wave components are all characterised by one unique $\mathbf{k}$-vector with a unique set of k_x, k_y and k_z components. Since we wish to propagate the beam in the z direction we need to advance the phase by $\Delta z k_z$, where Δz is the propagation distance $z - z_0$. Thus we need to determine the z component of the wavenumber, k_z. We can easily determine the k_x and k_y components at each position (x, y) in the far field as

$$k_x(x,y) = \frac{2\pi}{x - x_0}, \quad k_y(x,y) = \frac{2\pi}{y - y_0}, \tag{1.14}$$

which is also a statement of the fact that a larger transverse distance from the beam centre in the

far field corresponds to higher spatial frequencies in the near field. Then, knowing that $|\mathbf{k}| = 2\pi/\lambda$,

$$k_z(x,y) = \sqrt{\left(\frac{2\pi}{\lambda}\right)^2 - \left(k_x^2(x,y) + k_y^2(x,y)\right)}. \tag{1.15}$$

We now have all components necessary to perform the entire propagation:

$$E(x,y,z) = \mathscr{F}^{-1}\left(\mathscr{F}\left(E(x,y,z_0)\right) e^{i\Delta z k_z(x,y)}\right). \tag{1.16}$$

This expression works for both positive and negative Δz (and the trivial case $\Delta z = 0$), meaning that we can use it to determine the field in a preceding plane as well.

Of course propagation this way assumes that there is no effect that changes the phase or amplitude between the two planes at z_0 and z. We can, however, use it to propagate a field from one optical component to the next, and in the planes of the optical components we can apply their action to the field there. This allows us to simulate entire optical benches with many optical elements fairly efficiently, especially in numerical simulations using fast Fourier transforms (FFTs). One can make a further small speed improvement by restricting ourselves to the paraxial limit. This allows us to take the Taylor expansion of equation (1.15) about $(k_x, k_y) = (0,0)$ to get

$$k_z \approx \frac{2\pi}{\lambda} - \frac{k_x^2 + k_y^2}{4\pi/\lambda}, \tag{1.17}$$

which eliminates the need to take square roots, which is relatively computationally expensive. Throughout this work used this expression to perform beam propagation simulations.

1.9 Higher order modes

So far we have considered very simple fundamental solutions to the wave equation (1.2) in the paraxial limit. Higher order modes are something of a speciality of the Optics group at the University of Glasgow, and central to most of the techniques described later on, either as a motivation or as a tool. The two most commonly used sets of modes are the Hermite-Gaussian and Laguerre-Gaussian modes [19], and no thesis from the Optics group can be complete without their terrifying[7]-looking expressions:

$$\begin{aligned}\mathbf{HG}_n^m(x,y,z) =& E_0 \frac{w_0}{w(z)} H_m\left(\frac{\sqrt{2}x}{w(z)}\right) H_n\left(\frac{\sqrt{2}y}{w(z)}\right) \exp\left(-\frac{x^2+y^2}{w^2(z)}\right) \times \\ & \exp\left(-i\frac{k(x^2+y^2)}{2R(z)}\right) \exp\left(-ikz\right) \exp\left(i\psi(z)\right) \hat{e}, \qquad n,m \in \mathbf{N},\end{aligned} \tag{1.18}$$

$$\begin{aligned}\mathbf{LG}_p^\ell(r,\phi,z) =& E_0 \frac{w_0}{w(z)} \left(\frac{r\sqrt{2}}{w(z)}\right)^{|\ell|} \exp\left(-\frac{r^2}{w^2(z)}\right) L_p^{|\ell|}\left(\frac{2r^2}{w^2(z)}\right) \exp\left(-i\ell\phi\right) \times \\ & \exp\left(-i\frac{kr^2}{2R(z)}\right) \exp\left(-ikz\right) \exp\left(i\psi(z)\right) \hat{e}, \qquad \ell \in \mathbf{Z}, p \in \mathbf{N},\end{aligned}$$

where all the symbols carry their previously used meanings, and the new ones are R_z the wavefront curvature at z, the Hermite polynomials of order j, $H_j(x)$, generalised Laguerre polynomials $L_p^{|\ell|}(x)$, and the Gouy phase picks up a new definition, $\psi(z) = (N+1)\tan^-1(z/z_R)$. The value of the combined mode number N for the two classes of beams is different, for the Hermite-Gaussian (HG) case $N = m + n$ while for the Laguerre-Gaussian (LG) case it is $N = |\ell| + 2p$. Note that the HG

[7] At least to me. I mean look at them, they don't even fit on a single line!

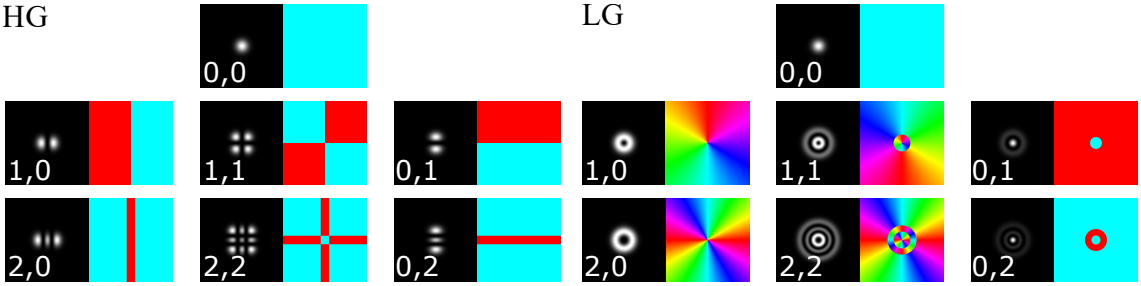

Figure 1.9: Zoo of higher order Gaussian modes, with intensity on the left and phase on the right. The indices in the corners are m, n for HG modes and ℓ, p for LG modes.

modes are described in Cartesian coordinates and carry rectangular symmetry, while the LG modes are in cylindrical coordinates and so have cylindrical symmetry. This can be seen in the case of a few example modes shown in figure 1.9.

Since these solutions are eigenmodes of propagation, they only change in transverse size on propagation (and they of course pick up a global phase, which is usually unimportant, except in interferometric setups). It is also worth noting that modes with different indices within either set are orthogonal to each other. Both mode sets are complete, which means that any light field can be decomposed into a (potentially infinite) subset of them, in the paraxial approximation.

An interesting relation between phase and intensity can be seen in all HG and LG modes with non-zero mode number. In points within a light field where the phase is discontinuous (e.g. phase jumps in HG modes and p-modes, LG modes with non-zero p, and optical vortices at the centre of ℓ-modes) the intensity is exactly zero.

A particularly important property of the Laguerre-Gaussian modes with non-zero ℓ is that they carry optical orbital angular momentum (OAM) [20, 21]. This is due to the helical phase fronts introduced by the azimuthally varying phase term, $e^{-i\ell\phi}$. The reason this is interesting is that OAM is a quantity that is conserved in most interactions with matter, with a few exceptions, in which case it is transferred to spin angular momentum (SAM). We will talk about this spin-orbit exchange in more detail later on. Spin angular momentum is related to the polarisation of light, specifically it is associated with circular polarisation. It is interesting to note that SAM can take two values, $\mp\hbar$ corresponding to left- and right-handed polarisation (and their linear combinations), while OAM can take any integer value (and their linear combinations). In other words, SAM (and so polarisation) spans a 2-dimensional Hilbert space, while OAM spans an infinite-dimensional one, a fact that could be useful for high-bandwidth quantum information transfer and storage.

1.10 In summary

We have briefly introduced light as an oscillation of the electromagnetic field and described it in terms of polarisation, amplitude and phase. We looked at some examples of manipulating polarisation (polarisers and waveplates) and the changes of propagation due to structured phase (focusing by a lens and deflection by a mirror/grating). We saw examples of higher order modes of propagation with structured phase and intensity. In the following, we will look at some methods of controlling phase, intensity and polarisation, mostly using dynamic optical elements but also introducing a few static ones.

Chapter 2

Shaping light

2.1 Introduction

Let us remind ourselves of one of the key points of the previous sections: the entire three-dimensional structure of phase, intensity and polarisation of a laser beam is defined by one single cross-section of the beam. This is because light behaves according to Maxwell's equations, so if we know the electric field in one plane, say, $z = 0$, we can easily find out what the field looks like some distance z away by propagation. As long as we restrict ourselves to the paraxial limit, this is particularly easy. All this means is that if we could arbitrarily control the phase and intensity of a beam in one plane, we could control its 3-dimensional structure as its transverse structure changes on propagation. This is exactly what spatial light modulators (SLMs) and digital micromirror devices (DMDs) allow us to do. We published some of the considerations about the limitations on beams that can be practically produced by such devices in [22].

2.2 Phase-only Spatial Light Modulators

The section title spoils the key feature of the SLMs we consider here, namely that they offer direct control over the phase of light. As we will see, this is enough for simultaneous control of phase, amplitude and polarisation. The theoretical discussions in this chapter are described in [23] and [22], and references therein.

2.2.1 Phase-only Holograms

It is obvious that a phase-only SLM provides control of the phase of a laser beam reflected from it, but the ability to control amplitude is more subtle. This is necessary if we wish to be able to create an arbitrary desired complex field $\mathbf{E}_{\text{des}}(x, y) = A_{\text{des}}(x, y) \exp(i\phi_{\text{des}}(x, y))\hat{e}$ in a transverse plane, where we restricted ourselves to shaping a single polarisation component $\hat{e}$ for now, which we will drop from the following discussion for simplicity.

After the interaction with the SLM we wish to spatially separate the newly structured beam from the remnants of the unstructured light, so that we can use the desired beam without any contamination. This is most easily achieved by deflecting the structured light, changing its wavevector,

$$E_{\text{in}}(x, y) \exp(i\mathbf{k'}_{\text{in}} \cdot \mathbf{r}) \exp(i\mathcal{H}(x, y)) = E_{\text{des}}(x, y) \exp(i\mathbf{k}_{\text{des}} \cdot \mathbf{r}), \tag{2.1}$$

where E_{in} is the input complex field travelling in the $\mathbf{k'}_{\text{in}}$ direction, the desired field propagates in the direction given by $\mathbf{k}_{\text{des}}$ and $\mathcal{H}$ is the phase-only hologram displayed on the SLM. Here we describe the process considering the SLM as a transmissive element, although the ones used in our experiments are the reflective type - this makes no conceptual difference. Figure 2.1 shows the relevant directions.

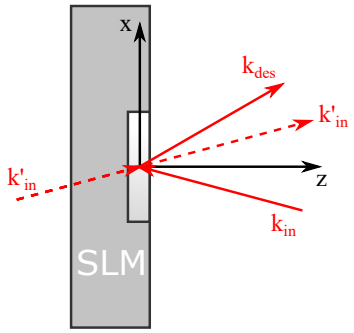

Figure 2.1: Propagation directions of diffracted and input beams on a reflective-type SLM.

We can easily express the action of the phase-only hologram as

$$\exp(i\mathcal{H}(x,y)) = \frac{E_{\text{des}}(x,y)}{E_{\text{in}}(x,y)} \exp(i(\mathbf{k}_{\text{des}} - \mathbf{k'}_{\text{in}}) \cdot \mathbf{r}) \tag{2.2}$$
$$= A_{\text{rel}}(x,y) \exp(i\phi_{\text{rel}}(x,y)),$$

where $A_{\text{rel}} = A_{\text{des}}/A_{\text{in}}$ is the relative amplitude and $\phi_{\text{rel}} = \text{mod}(\phi_{\text{des}} - \phi_{\text{in}} + (\mathbf{k'}_{\text{des}} - \mathbf{k'}_{\text{in}}) \cdot \mathbf{r}, 2\pi)$ is the relative phase, including the direction change. As discussed in section 1.6, a grating with 2π phase depth behaves like a tilted mirror, deflecting the beam impinging on it. We can therefore achieve the rotation of the k-vectors by including the proper grating phase, given by

$$\phi_g(x,y) = \frac{2\pi x}{d}, \quad \text{thus} \tag{2.3}$$
$$\phi_{\text{rel}} = \text{mod}(\phi_{\text{des}} - \phi_{\text{in}} + \phi_g, 2\pi),$$

where $d = 2\pi/(\mathbf{k}_{\text{des}} - \mathbf{k'}_{\text{in}})$ is the grating period. It is clear from equation (2.2) that we cannot simply use the relative phase ϕ_{rel} as the phase-only hologram $\mathcal{H}$, because while it would give the correct phase to the beam in the image plane of the SLM, it would not reproduce the correct intensity pattern. Instead, the intensity distribution of the input beam would reappear. Additional amplitude shaping is thus required, but since the phase-only SLMs only give access to phase control this has to be achieved through the modulation of the relative phase.

Reducing the depth of a blazed diffraction grating while keeping the period constant does not change the angle of diffraction, but it does change the diffraction efficiency. Shallower gratings diffract less light into the desired first order and leave more light in the zeroth order. We can spatially control the grating depth, and thus amplitude in the first grating order, by multiplying the relative phase by some function of the relative amplitude,

$$\mathcal{H}(x,y) = f\left(A_{\text{rel}}(x,y)\right) \phi_{\text{rel}}(x,y). \tag{2.4}$$

One can then place a spatial filter in the Fourier plane of the SLM that blocks the zeroth grating order, and would then be left with the desired beam. There are multiple choices one could make for the function f. The simplest method is to simply scale the relative phase by the relative amplitude, $f(A_{\text{rel}}) = A_{\text{rel}}$. This expression assumes that diffraction efficiency depends linearly on grating depth, which is incorrect. Yet, surprisingly enough, this method works rather well in practice.

A more rigorous treatment was proposed by Davis *et al.* [24]. They calculated that the first-order

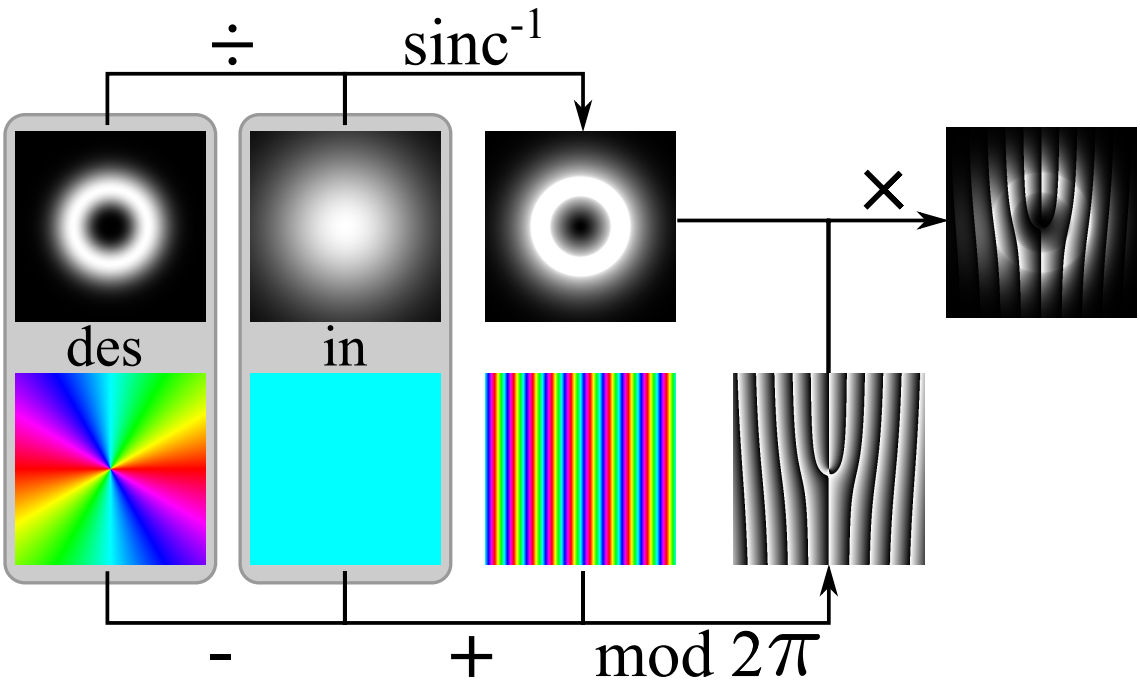

Figure 2.2: Calculation of a hologram designed to generate an $\ell = 2$ LG mode from a simple Gaussian input mode. The indicated operations are performed from left to right. The scales of grayscale images are 0-$I_{\max}$ for intensity images and 0-2π for phase maps.

diffracted light amplitude from a phase grating is given by

$$A(x,y) = \exp\left(i(1 - f(A(x,y)))\pi\right) \times \operatorname{sinc}(\pi(1 - f(A(x,y)))), \tag{2.5}$$

where $\operatorname{sinc}(x) = \sin(x)/x$, with $\operatorname{sinc}(0) = 1$ at $x = 0$. We can find the function f, omitting the exponential term, as

$$f(A_{\text{rel}}(x,y)) = 1 - \frac{1}{\pi}\operatorname{sinc}^{-1}(A_{\text{rel}}(x,y)). \tag{2.6}$$

The sinc function is not directly invertible, but the inverse function sinc^{-1} can be defined on the domain $[0, \pi]$. In the rest of this work this will be our preferred method for amplitude shaping with phase-only SLMs, essentially applying a spatially varying amplitude mask to the phase grating. There are other methods, which have their own strengths in different scenarios [25], but for our purposes this simple solution works extremely well. Our hologram calculation process is illustrated in figure 2.2.

2.2.2 SLM Phase Response

The pixelated device we used during this work was a liquid crystal on silicon (LCOS) type. As the name suggests, each pixel contains long, thin liquid crystals whose orientation can be controlled by applying an electric field through them. For this purpose each pixel has electrodes at the front and back surfaces which minimally affect light transmission. When maximum field is applied, the long axis of the crystals is perpendicular to the polarisation direction, such that the optical path through them is minimal (see figure 2.3a). As the electric field is decreased the crystals rotate to be parallel to the polarisation direction, maximising the optical path when minimum field is applied (figure 2.3b). We would also like to note that the phase across a single pixel is not uniform, but shows something approaching a sinusoidal variation.

The SLM used here is a reflective type, which means that the back surface of pixels is a mirror. This means that light travels through the liquid crystals twice. If the optical retardation between the

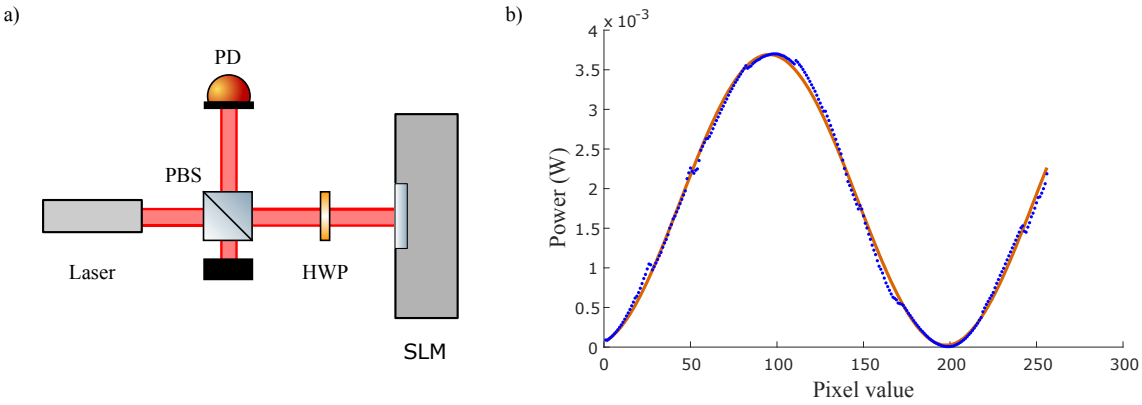

Figure 2.4: Schematic diagram (a) of the setup used to measure the SLM lookup table and measured power as a function of pixel value displayed on the SLM (b). Measured data is shown in blue, and the red curve is a sine fit to the data. The period of the fit is 204.3. PBS: polarising beam splitter, PD: photodiode, HWP: half-wave plate.

minimal and maximal path through the pixels is at least half of the wavelength of the input light then this system gives arbitrary control of the phase of the input light at each pixel independently. In practice the crystal orientation in the devices used in this work can be set to one of $2^8 = 256$ settings, but this range may not be fully utilised for a given wavelength.

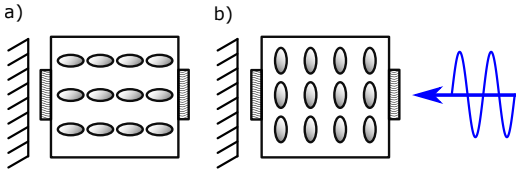

Figure 2.3: Schematic diagram of liquid crystals in a pixel of a LCOS SLM, with a) maximum electric field applied and b) no electric field applied. Light input direction and polarisation is shown by the arrow, and mirrors on the back surface of the pixel are shown as hatched areas.

Let us consider such an SLM that applies a phase difference of exactly 2π between the 0 and 255 setting of a pixel to light of wavelength 800nm. This means that the extra optical path introduced at a pixel value of 255 in a double pass through the device is 800nm. As a result, using 780nm light, between the settings of 0 and 255 this device would apply ${}^{800}/_{780} \times 2\pi > 2\pi$. This fact means that the control software needs to take into account the wavelength of light used. In practice this was done by generating a lookup table to convert from phase to 8-bit value to be sent to the SLM. The procedure for doing so is outlined below. To understand how this works it is important to note that the device is polarisation sensitive, it interacts only with horizontally polarised light and leaves vertically polarised light untouched.

A schematic diagram of the simple experimental setup used to measure the phase response lookup table is shown in figure 2.4a. The laser light is transmitted through a polarising beam splitter (PBS), which transmits horizontally ($|H\rangle$) polarised light and reflects vertically ($|V\rangle$) polarised light. The reflected light is blocked. The transmitted, now $|H\rangle$ polarised light passes through a half-wave plate (HWP) with its fast axis at 22.5° with the horizontal such that it turns the polarisation to $|A\rangle$. Since the SLM only acts on the $|H\rangle$ component, it can introduce a phase shift between the $|H\rangle$ and $|V\rangle$ components of this $|A\rangle$-polarised beam. In this setup we set the SLM to apply the same phase shift over the entire surface. When this phase shift is 0, the SLM just acts as a mirror which in this case turns the light from $|A\rangle$ to $|D\rangle$ because it only reverses the propagation direction. The new $|D\rangle$ polarisation passes once more through the HWP which rotates the polarisation back into $|H\rangle$. All light with this polarisation is transmitted through the PBS, not reaching the photodiode, corresponding to minimal signal. The signal from the photodiode is recorded by the computer controlling the SLM.

When the SLM applies a phase shift of π between the $|H\rangle$ and $|V\rangle$ components (and after reflection), the input $|A\rangle$ light turns back into $|A\rangle$, which turns into $|V\rangle$ after the HWP and is reflected by the PBS. This corresponds to maximal signal from the photodiode. Between these extremes the amount of light reaching the photodiode varies sinusoidally. This is shown in figure 2.4b.

During the experiment the value displayed on the SLM is ramped from 0 to 255, while recording the photodiode signal corresponding to each value. By fitting the resulting curve with a sine we can determine the actual phase shift introduced by a given 8-bit value sent to the SLM, and such can determine the lookup table for the laser wavelength used. In fact for the SLMs we used it was sufficient to find the period of the sine because the phase response of these SLMs is (very nearly) linear with pixel value. This is not, in general, true for all SLMs. The period of the fitted sine corresponds to the pixel value at 2π phase shift. The data shown in figure 2.4 taken using a Hamamatsu LCOS X13138-02. The nearly linear response of the SLM can be seen from the close match between the data and the fitted sine. The period of the fit is 204.3, so this corresponds to a phase shift of 2π (and ~ 102 would correspond to π, etc.)

Notice that the minimum (~ 0) phase shift does not occur for a pixel value of 0 (it's close to 200). This is not an issue, because for most applications we are only interested in *relative* phase shifts introduced by the SLM between parts of a beam. The absolute phase shift only plays a role in a few measurements, such as the one presented above, or interferometers in which the SLM is only interacting with one arm.

The measurement described here assumes that the phase shift applied by the SLM for a given pixel value is the same for all pixels (or at least the ones that the laser beam interacts with). This is a decent assumption, especially for small beams, but it is not exactly true. SLMs are prone to aberrations mostly due to their surface not being optically flat, although the phase response of all pixels are not identical either. We assume the latter to be a small enough effect to be neglected. In any experiment we performed so far we have never seen any strange behaviour that would need to be attributed to this, mainly because pixels with reduced phase throw are located near the edges of the SLM, which we do not often utilise. However, the former is significant and needs to be taken into account.

2.2.3 Aberrations

Phase aberrations are commonly described in terms of the Zernike polynomials Z_n^m [26]. The most frequently encountered examples are shown in figure 2.5. Z_0^0 is known as piston, and describes the static phase offset we encountered two paragraphs ago. The second row in figure 2.5, $Z_1^{\pm 1}$ are tip-tilt, which we recognise as examples of unwanted deflection. So far all of these can almost always be ignored because tip-tilt is usually automatically corrected in further alignment by the experimenter constructing an optical bench.

Of more practical importance are the ones in the third row. These are known as defocus (Z_2^0) and astigmatism ($Z_2^{\pm 2}$). We recognise the phase structure of defocus as the same as that of a lens, and it is due to a bulge in the SLM surface. Astigmatism behaves like a strange lens which has two focal lengths for two orthogonal directions; a more familiar example of such an optical element would be a cylindrical lens. These result from bends in the SLM surface; the archetypal phase patterns look like saddle-shaped bends.

Because of the need for a far-field filter in SLM-based light shaping setups it is technically possible

in most experiments to correct defocus by moving the two telescope lenses closer to or further away from the SLM. In practice this is very hard to do right. The most common SLM aberrations are (mostly) combinations of astigmatism and defocus, so shifting the lenses cannot solve all problems anyway. We could conceivably insert cylindrical and spherical lenses with the correct focal lengths to correct the aberrations, but we already have access to a device that can arbitrarily control phases in the SLM plane - the SLM itself! In the following we describe two methods used to determine the aberrations of the SLM surface, and we discuss how to correct them.

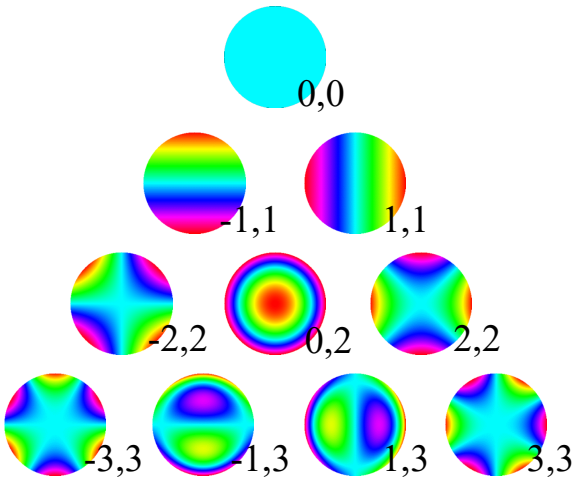

Figure 2.5: The first few Zernike polynomials, corresponding to the most common aberrations. Indices are m,n.

Our goal is to build up an image of the phase profile of the SLM. To do this we use concepts of computational imaging. We illuminate the entire SLM surface and display a simple blazed grating with 2π phase depth in a small section of the SLM, a superpixel. In the far field of the SLM we then see two spots[1]. One very bright spot corresponds to the zeroth order of the grating and the unaffected reflection of the parts of the SLM that do not contain the grating superpixel, and a much dimmer one contains the first order diffraction from the superpixel. We can then scan this superpixel across the SLM surface, and by monitoring the power in the first order diffracted spot we can measure the relative optical power arriving in any given superixel area. We normalise the first order diffracted powers to the brightest superpixel, and through interpolation we can build up a smooth distribution of the peak-normalised optical power arriving on the SLM. This gives us a measurement of the input beam intensity, but does not tell us anything about the phase yet.

Since we are interested in the relative phase between different parts of the SLM, we can simultaneously display two superpixels in different positions on the SLM, and use interference to deduce their relative phase. We usually place one superpixel in the centre of the SLM because this area receives the most optical power. We then do not change this superpixel for the duration of the experiment. A second superpixel is scanned across the SLM. Recall that in the Fourier plane spatial position turns into angle and angle information turns into spatial position. This means that the first diffracted orders of the two superpixels arrive at the same position in the far field (since their grating periods are the same), but they will have an angle between them (because they come from different positions), see figure 2.6. This results in the formation of linear interference fringes. We then apply a series of global phase shifts to the second superpixel. This moves the interference fringes. We measure the intensity in a small point within the interference pattern, and observe a sinusoidal variation as the interference fringe moves across the measurement point. We then fit this sinusoid, and extract the phase from the fit. This is in fact the relative phase between the center area of the SLM and the second superpixel position. Of course this is an average phase within the superpixel, but it is reasonable to assume that the phase is close to flat across small superpixels. In practice we use 20-by-20 pixel superpixels, in order to maximise resolution and fringe visibility, and this method works well. We then fit the measured phase with the Zernike polynomials to obtain the phase map of the light on the SLM.

[1]... in principle. In practice there will be more spots due to higher orders of diffraction, but power in those quickly becomes negligible.

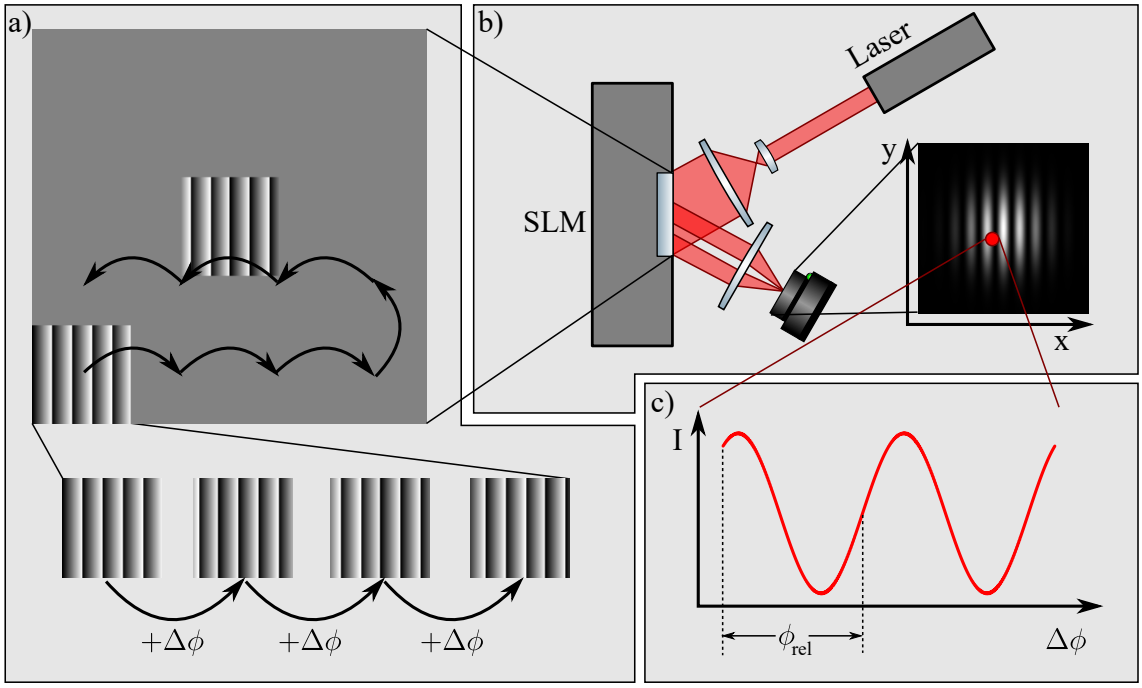

Figure 2.6: SLM aberration measurement process. a) Superpixels containing a phase grating are scanned and phase-shifted. b) The diffracted orders interfere in the far field. c) As the scanned superpixel is phase-shifted, the interference fringes move and the intensity in a camera pixel varies sinusoidally. The relative phase ϕ_{rel} between of the stationary central superpixel and the scanned one can be extracted from the fit.

There is a slightly different way of using superpixels to determine the SLM phase map [27], for which we need to assume that the aberrations are relatively well-behaved, i.e. reasonably planar over superpixels. This can usually be assured by using superpixels of sufficiently small size. Consider a superpixel containing a grating with some period. We could say that in the far field the light diffracted from this grating is displaced from the zeroth order propagation direction by an amount related to the grating period. However, this statement assumes that the SLM phase over this superpixel is flat. If, however, it contains some extra tip or tilt then the diffracted spot will be displaced further accordingly. We can then tile the entire SLM surface with superpixels with different known gratings and get information of excess tip-tilt over each superpixel. We do this by knowing how much the diffracted light from a superpixel should be displaced from the zeroth order due to the computer-generated grating and how much it actually is displaced in the experiment, and retrieving the extra change in k-vector. This can, in principle, be calculated knowing the grating period and lens focal length, but in reality there is a way to take a calibration image, which is more reliable in practice because it does not assume perfect alignment.

For this we calculate a multiplexed grating that contains all gratings that we used to tile the SLM surface. This is simply done like this:

$$\mathcal{G}(x,y) = \arg\left(\sum_{j}^{n} \exp\left(i\phi_{g,j}(x,y)\right)\right), \tag{2.7}$$

where we used j superpixels containing the grating phases $\phi_{g,j}$. Such a multiplexed grating creates an array of spots in the far field, each corresponding to one grating in the set. The important difference between this set of diffracted spots and the one due to the tiled gratings is that now all gratings sample phase aberrations of the entire SLM, not just a subset. This means that they will be displaced by the same amount due to global tip-tilt, and their shapes will be aberrated due to higher order aberrations. Crucially, they will not be differentially displaced from the the positions

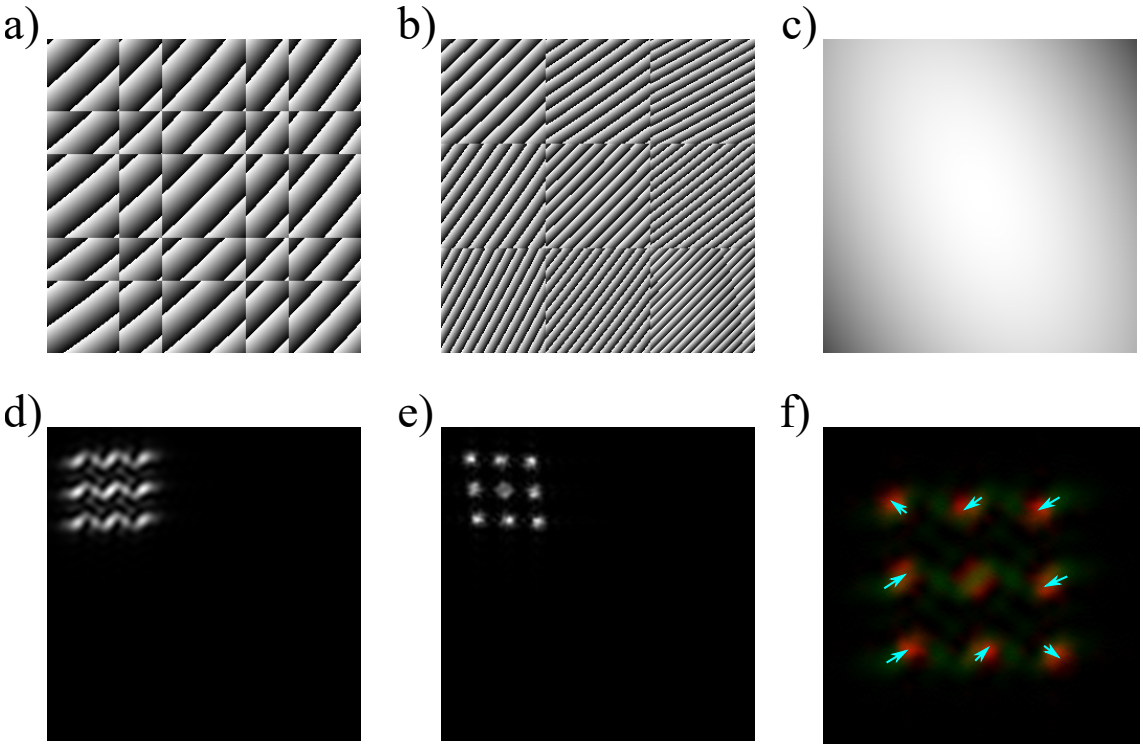

Figure 2.7: Operating principle of the Shack-Hartmann aberration measurement scheme. The SLM displays a) multiplexed and b) tiled holograms in sequence. These are shown with the aberration phase included. Correspondingly, in the far field d) aberrated and e) displaced spots appear. Here we simulated an aberration consisting of two cylindrical lens phases with different focal lengths, seen in c). This is a fairly common type of aberration. The displacement of the spots is shown by cyan arrows in f), where the far field of the multiplexed gratings is shown in green and that of the tiled gratings in red.

dictated by the gratings. This is shown in figure 2.7.

We can thus compare the far field images generated by the tiled and multiplexed gratings, find the centres of masses of the diffracted spots in each, and calculate displacement vectors due to the excess tip-tilt over each superpixel. This gives gradients of the SLM phase over each superpixel, which can be translated into a rough estimate of the SLM aberration map, which can then be fitted with Zernike polynomials to get a smooth phase map. This method, which we call the Shack-Hartmann method[2] has the advantage over the scanning one in acquisition speeds, since only two images are needed.

The disadvantages are that because of the need to estimate displacement vectors we need to calculate centres of masses of spots that do not have uniform shapes: recall that because the multiplexed grating samples the entire SLM its far field spots have differents shapes than in the tiled case. This issue is further complicated by noise in the images, which is notoriously bad for calculating centre of mass. There can also be an issue with spot registration; if two spots are close together in the far field of the tiled grating it can be difficult to know which one corresponds to which spot from the multiplexed grating. Also, since we never measured actual phase, but phase gradients, building up an estimate over the SLM surface is a bit more tricky, and needs an assumption of relative smoothness (which, to be fair, is usually a perfectly reasonable assumption). Ultimately the increase in acquisition speed could be worth exploring solutions to these issues, especially if one needs to perform aberration measurements often.

[2]Because of similarities to the wavefront-measurement sensors used in adaptive optics in astronomy and SLR autofocus sensors.

Note that the measurements outlined above does not only measure the SLM aberrations but also any aberrations in the input beam as well. It is very common to see a relatively large defocus aberration that does not come from the SLM, but some slightly incorrectly aligned lenses before the SLM. This is good because it allows us to correct any imperfection in the pre-SLM system so can be used in a complicated setup to save alignment effort, but it is also bad because it needs to be done every time the setup changes. I would personally recommend carefully aligning all optical elements and measuring SLM aberrations once, and correcting for those same aberrations even when the setup changes; I feel that it is better to have well-aligned systems than to try to measure aberrations every time something changes[3].

The correction itself is very simple. We just need to add the aberration phase to the input phase in equation (2.2):

$$\exp(i\mathcal{H}(x,y)) = \frac{E_{\text{des}}(x,y)}{A_{\text{in}}(x,y)\exp(i\phi_{\text{abr}}(x,y))}\exp(i(\mathbf{k}_{\text{des}} - \mathbf{k'}_{\text{in}})\cdot\mathbf{r}), \tag{2.8}$$

where ϕ_{abr} is the measured aberration phase, *including* the input beam phase. The logic of this is clear, since in the last paragraph we discussed how the method can be used to measure the phase profile of the entire beam line before and including the SLM; so the correction is simply to explicitly include our knowledge of the aberrations of the light field at the SLM. We could also equivalently subtract ϕ_{abr} from the desired phase. Of course the amplitude function from equation 2.4 should be factored in as well.

2.2.4 Spatial filtering

As mentioned previously, we use gratings to separate the phase- and amplitude-shaped desired beam from the unwanted remnants of the input beam in the far field of the SLM. This is achieved by placing a lens its focal length away from the SLM. However, this is obviously not sufficient to obtain the desired beam, as beyond the far field all grating orders diverge because of the phase curvature introduced by the lens. We need a second lens to complete a telescope and collimate the beams again, but this comes with the issue that the telescope forms an image of the SLM plane, recombining the zeroth order with the first, which is not what we want. We can, however, place a far field filter in the Fourier plane of the telescope. Such a filter ideally blocks all light from all orders apart from the first. In practice we use variable aperture irises to transmit the first order. The effects of changing the far field filter size, discussed below, can be seen in figure 2.8.

An issue with this idea is that beams never quite fall off to zero power, so it is impossible to filter orders perfectly. Some light from the zeroth order always remains, and some light from the first order will be cut off. The smaller the filter aperture around the first order, the smaller the influence of the remaining light from the zeroth order will be. However, recall that in the far field the Fourier transform of the spatial structure of the desired beam appears. Thus a small far field filter acts as a low-pass filter, inevitably truncating high spatial frequency components. The smaller the aperture the stronger this effect. It turns out that for most beams of interest in this work this is a good thing, since we rarely want beams with hard edges or high frequency details. In fact far field filters in the form of few tens of μm diameter pinholes are often used to clean up the messy, non-Gaussian beams that many lasers output. An example of this is shown in figure 2.9.

The larger the separation between the grating orders in the far field the easier it is to filter out

[3]Which happens very frequently.

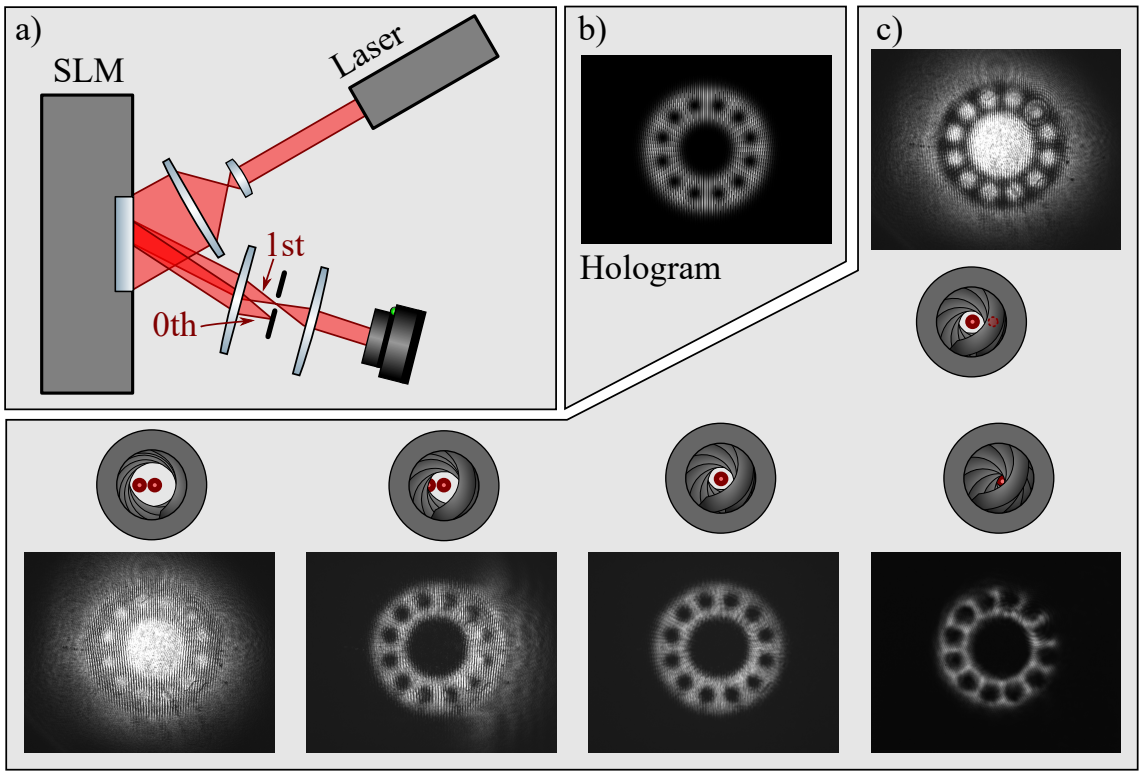

Figure 2.8: Effects of different far field filter sizes. a) Experimental setup used to record the images. The camera is close to, but not quite in the image plane of the SLM. b) The hologram here is designed to generate a superposition of LG_0^8 and LG_0^{20} modes. c) Bottom row, left to right: aperture fully open; aperture cutting (but not blocking entirely) the 0th order; aperture transmitting the 1st order only; aperture too small, high spatial frequencies truncated. Top right: 0th order transmitted. Notice the vertical interference fringes where the 0th and 1st orders overlap; this is an image of the grating.

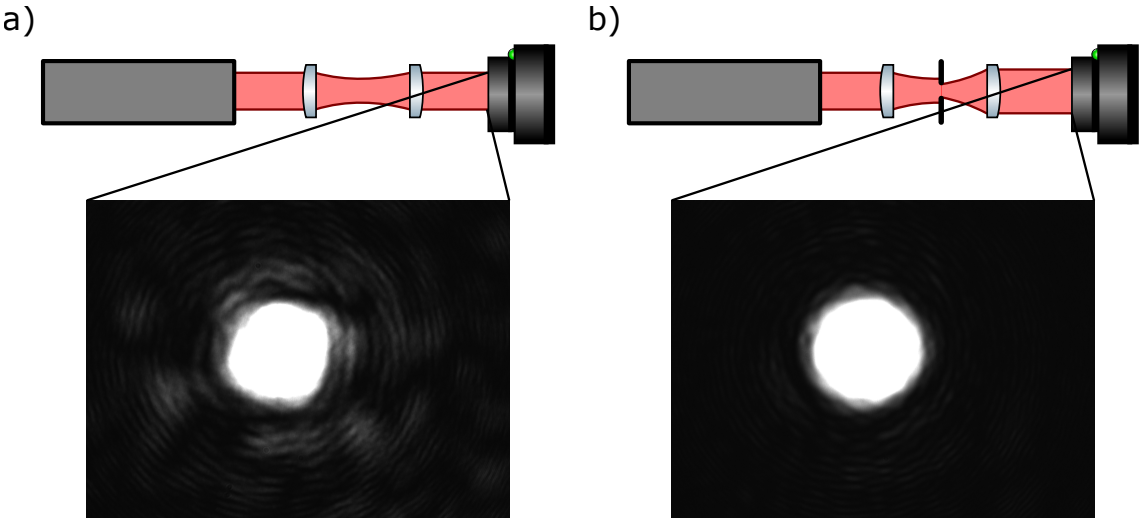

Figure 2.9: Filtering a He-Ne laser output mode. The images are deliberately overexposed to highlight a) the unwanted structure in the unfiltered beam and b) their absence after filtering. Notice that the filtered beam has a slightly larger diameter because the filter aperture is smaller than the input beam waist and the Fourier relationship between the far and near fields.

the unwanted orders. Since angles in the near field turn into positions in the far field, we desire gratings that introduce large angles between the k-vectors of the outgoing orders, so large θs in figure 1.8. The wavelength in any experiment is fixed by the laser, so the only way to change θ is by varying the grating period. There is an inverse relationship between far field separation and grating period: short periods lead to large separation and vice versa.

So it seems logical that we desire as small a grating period as we can get. The limit is set by the SLM pixel size. The smallest period we can have is two pixels per 2π phase cycles. This, however, is a very poor approximation of the blazed grating phase; in fact it looks like a binary grating. We investigated this, along with a number of other limitations of SLMs, in [22]. We find that the gratings appear too discretised below approximately 5 pixels per period and the grating efficiency is severely reduced. Peak efficiency is reached around 5 pixels per period, so this is the most practical choice of grating period. More pixels do not gain much in efficiency as the grating phase is already well represented, but the risk of contamination by the 0th order increases. This gives a distinct advantage to high resolution SLMs with small pixels, because the 5 pixel period is represented in a smaller physical size.

The last thing to consider is the fact that since the way the method we described for intensity shaping works by 'throwing away' unwanted light from the input beam, the total power in the generated beam suffers. This gets worse for desired beams that have a small overlap with the input mode. The worst case scenario would be if the input beam had zero intensity where the desired beam has its maximum. In a sense, the desired beam needs to 'fit under' the input beam; this is where the analogy to sculpting in the title of this thesis comes from. Figure 2.10 illustrates the idea. In extreme scenarios, where the overlap between the input and desired beam profiles is small and one would get essentially no power in the desired beam, it may be desirable to not use the measured input beam for the intensity correction. Instead, one could use a thresholded version of the input profile, shown in the second row of figure 2.10. Of course the intensity profile of the output beam will be wrong by doing this, but if the phase profile of the desired beam is more important for the experiment, this might be a reasonable tradeoff.

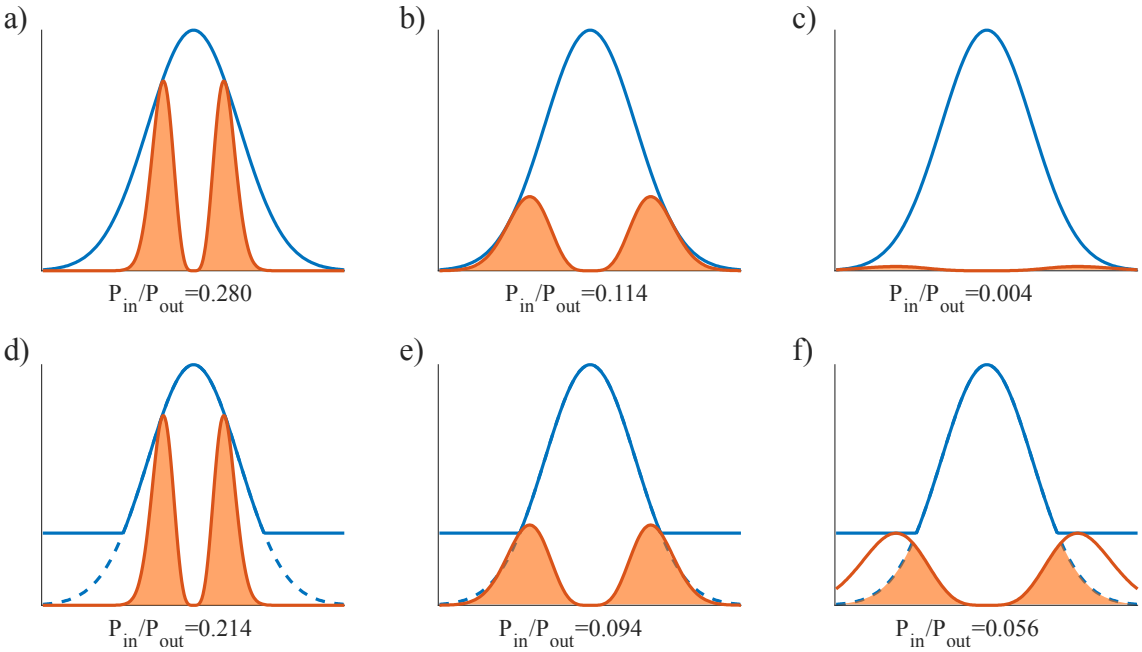

Figure 2.10: Beam sculpting using a-c) the measured input beam and d-f) thresholded input beam. The input beam is always the same Gaussian, and the desired beam is a LG_0^2 mode with a waist that is a,d) 1/3, b,e) 2/3 and c,f) 3/3 of the input waist. The ratios of the input and output powers in each scenario are shown under the graphs.

2.3 Digital Micromirror Devices

Digital micromirror devices feature a 2D array of μm-sized mirrors that can be independently electronically switched to point in two different directions. These devices act on the amplitude of the incident light, in contrast with the phase modulation of SLMs. The square mirrors in a DMD are arranged in a diamond pattern, with a hinge along one of their diagonals. They can flip along this diagonal to point in one of two directions, $\pm 12°$ to the DMD normal. These very same devices are used in current overhead projectors, in conjunction with RGB LEDs, to produce full colour images at video rates. This is achieved by pulse width modulation of the red, green and blue channels in sequence at a very high rate (even the cheapest DMDs run at 5 kHz, with top-of-the-line ones exceeding 20 kHz) with spatial resolution. To the human eye, with a relatively long integration time, this appears as a single, coloured image.

Note that because of the diagonal arrangement of DMD pixels the rows of pixels are twice as dense as the columns are (see figure 2.11), which means that images appear stretched on a DMD. This needs to be taken into account when designing holograms.

2.3.1 Binary holograms

The DMD acts as a spatially dependent binary intensity modulator; each pixel can be 'on' or 'off', sending light either in the desired direction, towards the rest of the experiment, or away from it. It is not immediately obvious that such a device can provide the same flexibility that a phase modulator SLM can. It turns out DMDs are, in fact, capable of this.

Light in a given order diffracted from a uniform diffraction grating has a flat wavefront. However, modulating the grating slit positions leads to a curvature of the wavefront. In other words, spatially dependent lateral shifts of the grating slits introduces spatially dependent phase shifts. This is because if a grating is shifted in the direction parallel to its surface by 1 grating period the output

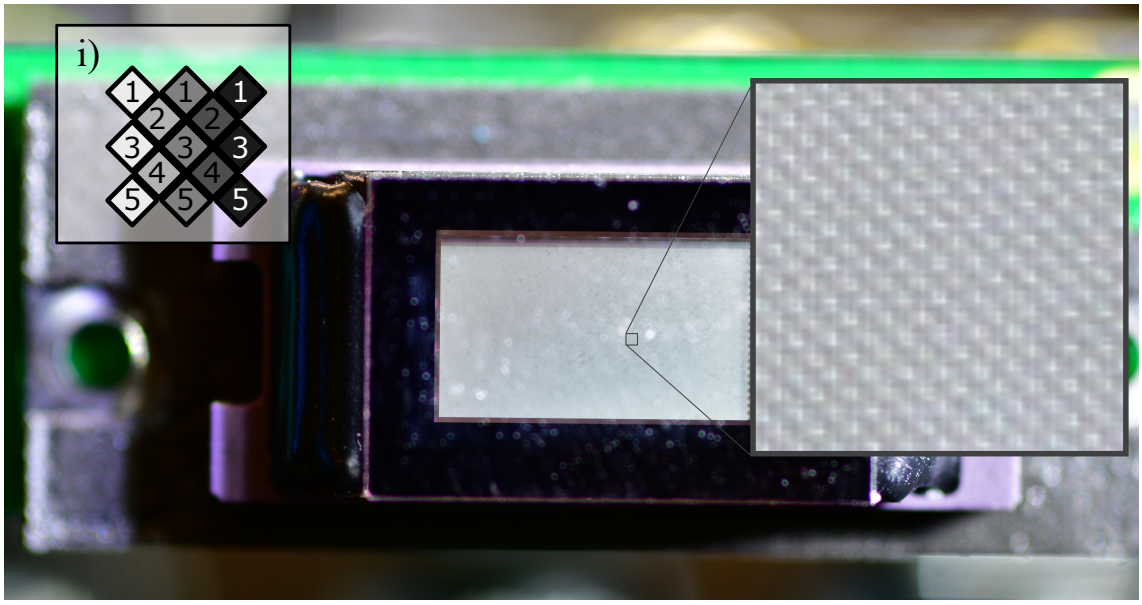

Figure 2.11: Photograph of a (rather dusty) DMD surface, where the micromirrors are resolved. The bright spots in the inset are mirrors pointing to the right, and dark regions are mirrors pointing to the left. These mirrors have approximately 7.6μm side length. The inset (i) shows the pixel arrangement, where numbers show the row index and shading shows columns.

light acquires a 2π phase shift [28]. Shifting a slit with respect to the rest of the grating applies this effect in a spatially dependent manner. Spatial control of the intensity is possible by spatially varying the slit widths[4], simply because this modulates the amount of transmitted light. On the DMD surface slits translate to 'on' pixels, surrounded by 'off' pixels. The binary hologram that generates a desired optical field $A_{\mathrm{des}}(x,y)\exp(i\phi_{\mathrm{des}}(x,y))$ in the image plane of the DMD is given by [29]

$$\mathcal{H}(x,y) = \frac{1}{2} + \frac{1}{2}\mathrm{sgn}\left(\cos(\phi_{\mathrm{rel}}(x,y)) + \cos(\sin^{-1}(A_{\mathrm{rel}}(x,y)))\right), \tag{2.9}$$

where sgn stands for the sign function. Note that the binary nature of the hologram leads to a reduced efficiency as it does not diffract preferentially to either the +1st or -1st order.

Furthermore, DMDs present an additional challenge in that the pixels, being physically separate mirror surfaces, also act as an additional grating, with a period corresponding to the pixel size. The orders of this overall grating each contain the diffraction pattern generated by the computer-generated grating, and for practical reasons only the brightest order of the pixelation grating is used in experiments. An example is shown in figure 2.12. This leads to further loss of efficiency, although it does not degrade beam quality. Practically, it is difficult to achieve more than 10% efficiency with a DMD.

DMDs also suffer from the same types of aberrations that SLMs do, and because of the use of gratings for amplitude and phase shaping the generated beams also require spatial filtering. The techniques described above for SLMs work just as well here. For obvious reasons there is no need to calibrate the phase response of the DMD for different wavelengths.

[4]The pulse width modulation (PWM) trick for intensity control that works very well for the human eye is not capable of generating coherent beams with amplitude shaping. A proper binary hologram works, in principle, for a single photon, but PWM does not.

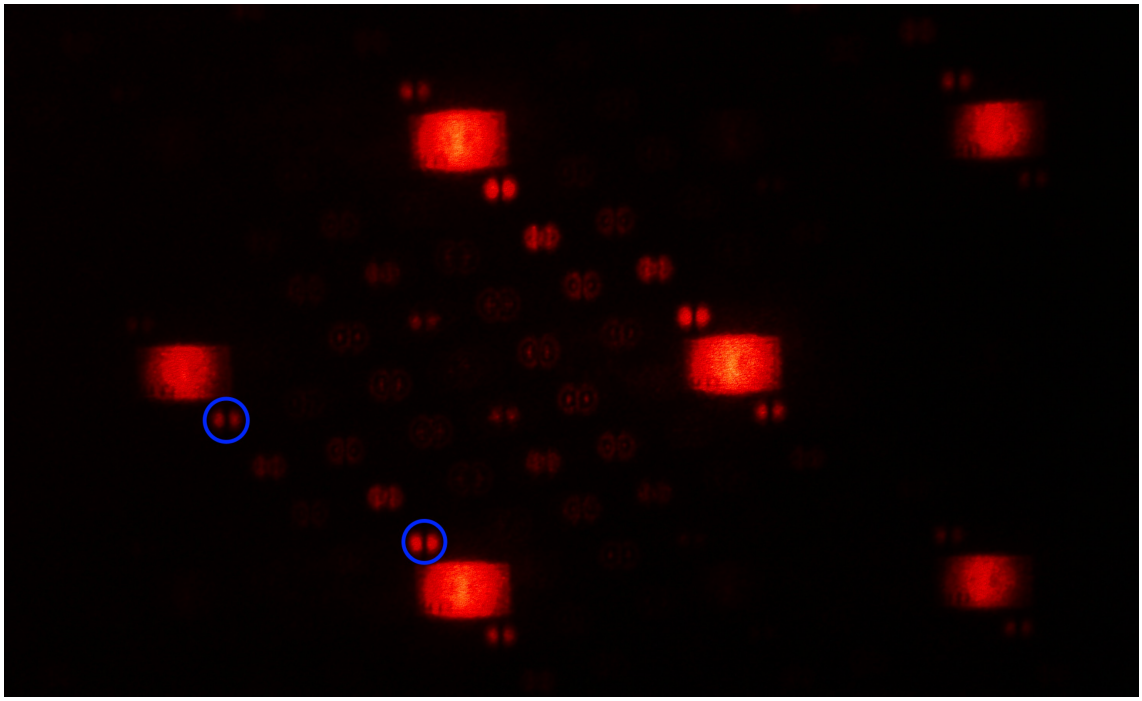

Figure 2.12: A subset of diffraction orders from a DMD using a He-Ne laser. The large, roughly square-shaped light beams show the pixelation orders of the DMD itself. Surrounding each of these zeroth orders of the computer-generated hologram one can see the first and higher diffraction orders of the hologram. Some 1st and -1st orders (intended to be HG_0^1 modes) are circled in blue.

2.4 Spatially dependent polarisation

So far all of our discussion was restricted to generating light fields with spatially varying phase and amplitude, but we did not consider the spatial control of polarisation. As we discussed earlier, an SLM is polarisation sensitive; it only acts on a single polarisation, usually $\hat{h}$, so it seems that spatially dependent polarisation control with such a device is doomed to fail. However, just as a DMD is capable of phase control even though it does not appear to be at first glance, polarisation control is also possible with a bit of cleverness.

2.4.1 Using an SLM

Recall the definitions of the polarisation basis vectors in equation (1.6). It can be seen that any basis vector can be expressed in terms of a pair of others. For example, $\hat{a} = 1/\sqrt{2}(\hat{h} - \hat{v})$ and $\hat{r} = 1/\sqrt{2}(\hat{h} + i\hat{v})$. This is because pure polarisation states span a two-dimensional vector space.

The point really is that if we could, at every point in a plane z, generate arbitrarily weighted coherent superpositions of two orthogonal polarisation components with an arbitrary phase between them, we could generate any polarisation at any point in space. For this we need two orthogonally polarised amplitude and phase shaped laser beams (that are coherent with each other) and overlap them.

We can use an SLM to generate two (or, in principle, any number of) structured beams propagating in different directions. This is actually very easy. We need to multiplex gratings with different periods, each containing the required phase and amplitude information for the different desired

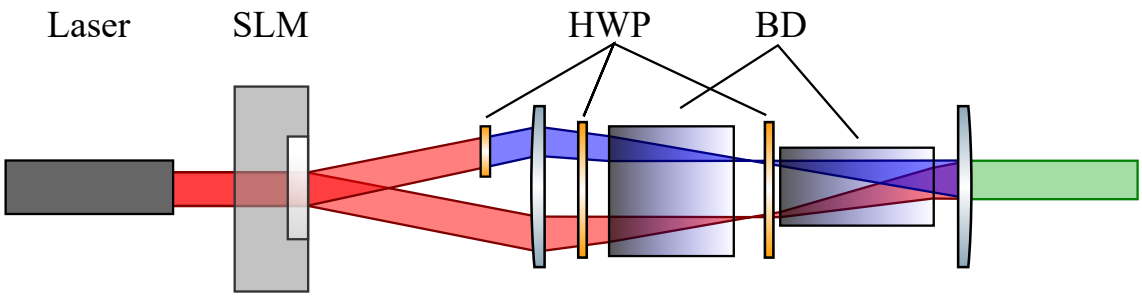

Figure 2.13: Practical polarisation shaping with an SLM. The grating 0th order is not shown for clarity. The SLM is drawn in transmission, although in reality it is reflective. HWP: half-wave plate, BD: calcite beam displacer.

beams, into a single hologram. For n beams this hologram is defined as

$$\mathcal{H} = \arg\left(\sum_{j}^{n} A_{rel,j} \times \exp\left(i \bmod\left(\phi_{des,j} - \phi_{in} + \phi_{g,j}, 2\pi\right)\right)\right), \tag{2.10}$$

where the spatial dependence of the terms is implicitly considered.

An SLM acts only on horizontal polarisation. In order to control polarisation with an SLM, we need to decompose the desired vector field into horizontally and vertically polarised complex scalar fields, and generate a hologram using equation (2.10) above to direct them in two different directions. We then Fourier transform the optical field after the hologram using a lens, just like in the standard setup for an SLM. In the far field the two first orders of the two computer-generated gratings are spatially separated because $\phi_{g,j}$ are different. We place a half-wave plate into the beam corresponding to the beam component that will be vertically polarised to turn it from horizontal to vertical polarisation. We then overlap the two, now orthogonally polarised, beam components using a beam displacer. The resulting vector beam[5] is then collimated by a second lens completing the telescope, and the desired polarisation structure appears in the image plane of the SLM.

There are a few practical considerations that somewhat alter the description in the preceding paragraph. First, one needs to carefully choose the two gratings to separate the two first orders. They should be far enough from the zeroth order to be efficiently filterable. Grating efficiencies vary in reality with grating period, so the two gratings should have the same period but pointing in different directions symmetric around the zeroth order to minimise efficiency differences between the two beam components.

Second, ideally one would like to perform all manipulations of the two beam components in the far field, where they are maximally separated and where their propagation directions are the same. This is not practically achievable because of the finite sizes of the optical elements needed. Calcite beam displacers are particularly long, few centimetres in length. The practical setup is shown in figure 2.13.

Notice that we use two calcite beam displacers in order to make alignment easier. These optical elements displace light with polarisation parallel to one of their axes and leave orthogonally polarised light unaffected. The half-wave plates before the beam displacers are used to rotate the polarisation of one of the beams to match the axes of the displacer (and turn the other beam to be untouched). Two displacers are used for flexibility. The alignment of the displacers is quite

[5]Throughout this thesis we refer to polarisation-shaped beams as vector beams. Technically all beams are vectorial, of course, since even a uniformly polarised beam has a polarisation vector associated with it.

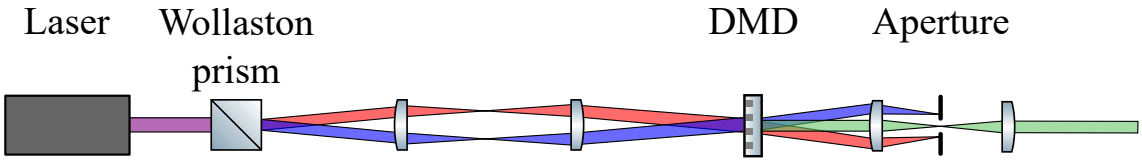

Figure 2.14: Practical polarisation shaping with a DMD. The zeroth orders of the grating are not shown for clarity. The DMD is drawn in transmission, although in reality it is reflective.

time-consuming and counterintuitive in the first instance, although it does get easier with practice.

The two polarisation components also travel slightly different distances. We need to make sure that the optical path difference is shorter than the coherence length of the laser, otherwise the resulting beam becomes not a vector beam, but an incoherent superposition of horizontally and vertically polarised beams. A much simpler setup can be utilised using a DMD.

2.4.2 Using a DMD

In our DMD-based polarisation shaping device we exploit the fact that the micromirrors are not polarisation sensitive. Because of this, they can interact with two orthogonal polarisation components, shape them appropriately and overlap them, all in a single plane. Our way to do this is illustrated in figure 2.14.

First, we split a diagonally polarised laser beam into horizontal and vertical components using a Wollaston prism. Such a prism is a lot like a polarising beam splitter, but the output propagation angle between the two orthogonally polarised beams is not 90°, but much smaller[6]. The plane where the beams are split is imaged onto the DMD using a telescope. The magnification of the telescope can be used to control the size of the beams on the DMD, but also the angle between them. In practice, along the propagation direction there is no single transverse plane where the beam-splitting happens, since the interface within the prism is diagonal. This is not really a problem, however, because all we really need on the DMD is a large spatial overlap between the two orthogonally polarised beams; they do not need to be completely overlapped.

We once again need to use a hologram corresponding to two multiplexed gratings, just like in the SLM case. However, in this case there is an extra consideration to the choice of the grating periods and directions which was not present in the SLM case. Because we do not need different optical elements in the two beams that are diffracted from the DMD (since the polarisations are already orthogonal) we can overlap the two appropriate first diffracted orders immediately after the DMD. We do this by choosing grating periods and directions in such a way that the input angle between the two polarisation components is exactly compensated in the first orders of the two gratings. There is no need for any further optical elements, apart from far field filtering, the desired vector beam appears immediately behind the DMD.

Since the aligment of overlapping of the two polarisation components is done entirely by the design of the multiplexed grating, in other words, in software, the efforts required of the experimenter are minimal. Setting up the two telescopes shown in figure 2.14, as well as positioning the far field filter, are the only experimental challenges, and none of these are very hard or time-consuming with a little bit of experience. One can determine an appropriate pair of grating periods and directions by placing a camera in the far field of the DMD and iteratively modifying the grating parameters

[6]It depends on the material. 1° for quartz, 1.334° for MgF_2 and 20° for calcite.

Table 2.1: Comparison of some experimentally relevant properties of SLMs and DMDs.

	SLM	DMD
Shaping accuracy	very good	very good
Output power	high (80%)	low (10%)
Refresh rate	video (60 Hz)	fast (several kHz)
Surface flatness	astigmatic	a bit less astigmatic
Pixel flatness	sinusoid	very flat
Ease of polarisation control	difficult	easy

until an appropriate overlap of the first orders is achieved at an acceptable separation from the zeroth orders[7]. An example of the DMD far field is shown in figure 2.15.

Some experimentally relevant properties of SLMs and DMDs are compared in table 2.1.

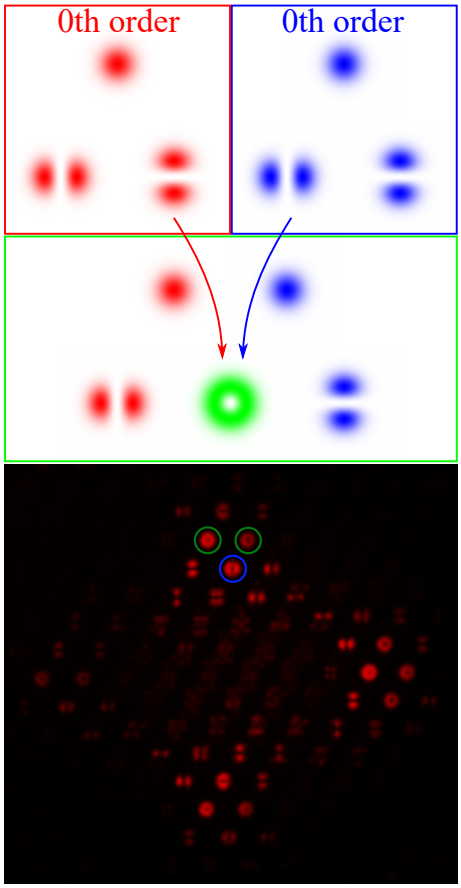

Figure 2.15: DMD far field when generating a radially polarised beam from horizontally (red) and vertically (blue) polarised input beams. In the experimental image (bottom) desired beam is circled in blue. The zeroth orders are circled in green.

2.4.3 Static devices: Q-plates and Fresnel cones

Hopefully based on the discussion in this chapter so far it is obvious that there are many advantages to using dynamic devices for beam shaping. Most importantly, the action of the dynamic optical elements can be changed without requiring realignment of entire optical benches. Aberration compensation can be done easily *in situ*. Experiments can be automated with computers (one example for this is the measurement of aberrations and phase response of an SLM, but we will see more later). They can also simplify alignment, as seen in the case of the DMD polarisation shaping. However, they are limited in resolution, so creating fields with high spatial frequencies in phase or amplitude are problematic or impossible. Efficiency also suffers because of the pixelation of the grating phases. Because of these reasons in some cases it might be advantageous to use static devices. In the following we describe two examples of such devices of particular interest: Q-plates and Fresnel cones.

A Q-plate is a liquid crystal device, like an SLM. It is characterised by a number q and it affects the two circular polarisations differently: it imprints an orbital angular momentum (LG) phase of $+\ell = 2q$ and $-\ell = -2q$ onto $\hat{r}$ and $\hat{l}$, respectively [30]. Thus they can be used to generate circularly polarised beams with OAM, although they are not capable of amplitude modulation, so they cannot make proper LG beams. If, however, we send a linearly polarised beam through a Q-plate, then the output is linearly polarised, with the polarisation direction rotating with azimuthal angle. Examples of this are shown in figure 2.16. A similar effect can be achieved with segmented half-wave plates, where the segments have differently oriented axes, but a Q-plate varies nearly continuously.

[7] of which there are two now, because of the two input beams.

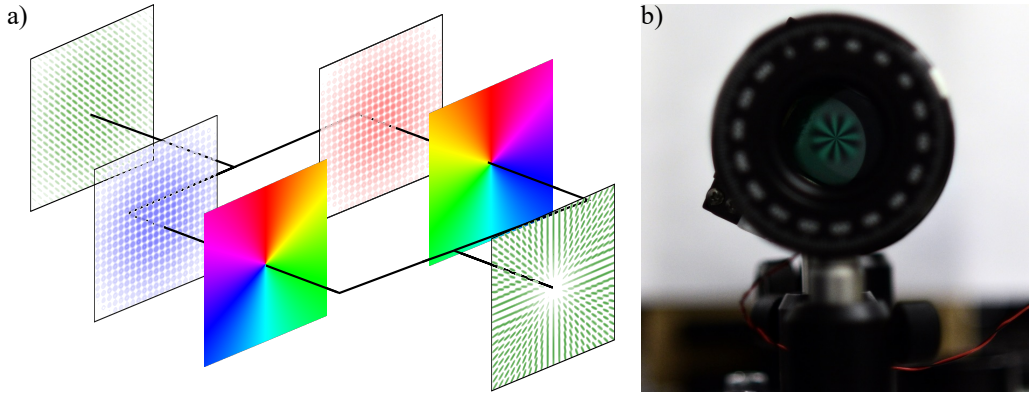

Figure 2.16: a) The action of a $q = 1/2$ Q-plate on an $|D\rangle$ polarised input beam. The linear input beam (indicated by green polarisation lines) is decomposed into left- and right-handed polarisation components (shown in blue and red circles, respectively). These components acquire $\pm 2q$ OAM (shown by the helical phase maps). In this example the coherent superposition of these components gives a radially polarised beam. b) Photograph of a $q = 2$ plate (green hue) between two orthogonal polarisers, illustrating the polarisation structure.

Since Q-plates work by applying spatially dependent phase shifts, their operation is wavelength-dependent. For this reason they come with electrodes, just like an SLM pixel, where a voltage can be applied. In the case of the Q-plate the standard operation requires supplying an AC voltage with a frequency of around 1 kHz. The amplitude of this signal determines the operating wavelength. The Q-plates we used (see chapter 8) are luckily designed exactly for the wavelength used in the experiments so do not require this AC voltage to work well.

We find it interesting to mention that it is possible to obtain both OAM and structured polarisation from an extremely simple optical element, a glass cone with cone angle of 90°. Although we did not use a cone in our experiments we include them here for completeness and interest.

Light entering a 90° cone from the flat side undergoes total internal reflection twice and leaves propagating in the inverse direction. In other words the cone is a retroreflector. If a left-hand circularly polarised beam enters such a cone, after the two total internal reflections it leaves left circularly polarised, but travelling in the opposite direction. Thus, in the original frame of reference, it loses 2 units of angular momentum. Since the cone is cylindrically symmetric this angular momentum cannot be imparted by a torque to the cone. Instead it is transferred to orbital angular momentum, such that the beam leaving the cone carries 2 units of OAM, that is, has a phase term of the form $\exp(i2\phi)$, where ϕ is the azimuthal coordinate.

It has been known since the early 19th century that on total internal reflection the polarisation components perpendicular and parallel to the plane of reflection undergo different phase shifts. The differential phase shift between these s and p components are the reason Fresnel rhombs (described in section 1.2) work. It so happens that in most glasses the relative phase shift between s and p is very close to $\pi/4$. Because of the cylindrical geometry of the glass cone, the s and p decomposition of an incoming light beam is azimuthally changing [25]. This leads to uniform input polarisations turning into azimuthally varying polarisations as they exit the cone. Of most interest are beams produced from circularly polarised input beams. Circular polarisation is always equally decomposed into s and p components, but the axis of the decomposition is azimuthally varying. After the two $\pi/4$ phase shifts between the components the output light is always linearly polarised, but the direction of the polarisation is spatially dependent. Such beams exhibit a swirly polarisation pattern, but this can be easily turned into radially or azimuthally polarised beams using a half-wave plate.

2.5 In summary

In this chapter we looked in great detail at the practical details of using SLMs and DMDs to structure light. We talked about correctly designing phase-only holograms that allow us to generate beams with arbitrary spatially dependent phase and amplitude. We also saw how to measure and correct aberrations of optical systems including an SLM. We described how to design binary holograms for use with a DMD to achieve the same capabilites that an SLM has (albeit at lower powers, but much higher refresh rates). We showed two setups for spatially controlling the polarisation of laser beams using dynamic elements, and introduced two static elements that can be used to generate useful polarisation structures at the expense of amplitude control. In the next chapter we will apply some of this knowledge to investigate correlations that mimic some properties of quantum entanglement.

Chapter 3

Measuring concurrence in structured polarisation

3.1 Introduction

In order to illustrate the power of structured light we investigate correlations within vector beams as a classical analogue to quantum entanglement. At the time of writing the work in this chapter has been submitted for publication[1], although not reviewed. Here we will outline a method to calculate concurrence, a measure of entanglement in quantum mechanics, for vector beams. In this situation concurrence will be used to characterise correlations between spatial modes and polarisation. We then describe an experiment where we generate various beams with different degrees of correlations and measure their concurrence, showing that concurrence is indeed an accurate and practical measure for these correlations. I have built the experiment setup in Glasgow, used it take data and I have analysed the data.

3.2 Correlations and Concurrence

The inherent correlations between the spatial and polarisation degrees of freedom in a vector beam shares mathematical properties with quantum entanglement, with the notable exception of non-locality [32, 33, 34, 35, 36]. This means that the kinds of correlations we are about to discuss are inherently local, because the correlations exist between degrees of freedom of the same photon, not degrees of freedom of a set of photons. In any case, in the experiment we do not work with individual photons.

In particular, we can determine spatial location based on polarisation measurements using these correlations [37], which have been used as a resource for wide ranging applications [38] including laser material processing [39], optical manipulations [40, 41], high resolution microscopy [42], and classical and quantum communications and information [43, 44, 45, 46, 47].

Vector beams can be quantified and analyzed using a variety of approaches, including geometric phase measurement [48], shear interferometry [49], Bell violations [50, 34, 51], and quantum state tomographies via projection measurements [52]. Projection measurements only work for a subset

[1] For the time being see [31].

of vector beams, because one needs to project into a subset of the infinite set of spatial modes. This requires *a priori* knowledge of the possible spatial modes involved. However, the correlations exist no matter what the measurement basis is, which suggests that there ought to be a method for quantifying them without the need to project onto spatial modes.

In the following we show that concurrence can be expressed in terms of Stokes parameters, allowing us to characterise vector beams without knowledge of its spatial mode degree of freedom. Stokes parameters can be measured in a spatially resolved manner with appropriate detectors, although they still require projection onto the polarisation basis states. However, polarisation is a finite-dimensional space, and so this does not lead to experimental restrictions.

Concurrence is a measure of entanglement used for bipartite states and can be thought of as a measure that tells us how much information we could gain about a quantity by measuring a different, correlated quantity. Here the correlation is between spatial structure and polarisation, for which we shall define the concurrence. The following proof is originally due to Dr Sarah Croke at the Quantum Theory group of the University of Glasgow, and it is gratefully acknowledged. Later, we compare the results of this basis-independent method with a projection measurement method which was developed at the University of Witwatersrand by Dr Carmelo Rosales-Guzmán and Prof Andrew Forbes, who kindly agreed to the reproduction of some of their results.

3.2.1 Concurrence

Let us start by considering the kinds of beams we study here. They are composed by two orthogonal polarisation components, each with an arbitrary amplitude and phase,

$$\begin{aligned} |\Psi_{vec}\rangle &= |H\rangle|\tilde{\psi}_H\rangle + |V\rangle|\tilde{\psi}_V\rangle \\ &\text{with} \\ |\tilde{\psi}_{H,V}\rangle(x,y) &= \int \mathrm{d}x\mathrm{d}y A_{H,V}(x,y)e^{i\phi_{H,V}(x,y)}|x,y\rangle. \end{aligned} \tag{3.1}$$

In this description $|x,y\rangle$ are position eigenstates, and the beam is defined by electric fields with spatially varying amplitudes A and phases ϕ in two orthogonal polarisation directions, which are $|H\rangle$ and $|V\rangle$ here. Such beams can be written in the form

$$|\Psi_{vec}\rangle = a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle, \tag{3.2}$$

where $|p,s\rangle$ contain polarisation (p) and spatial (s) modes. The orthogonal polarisation states p can be e.g. $|H\rangle = 0$ and $|V\rangle = 1$, and for the spatial modes s any orthogonal pair of modes can be represented by 0 and 1. We can transform equation 3.1 into the form of equation 3.2 by making a number of substitutions, defined as

$$\begin{aligned} p_H &= |\langle\tilde{\psi}_H|\tilde{\psi}_H\rangle|, \\ |\psi_H\rangle &= \frac{|\tilde{\psi}_H\rangle}{\sqrt{\langle\tilde{\psi}_H|\tilde{\psi}_H\rangle}} = \frac{|\tilde{\psi}_H\rangle}{\sqrt{p_H}}, \\ p_V &= |\langle\tilde{\psi}_V|\tilde{\psi}_V\rangle|, \\ |\psi_V\rangle &= \frac{|\tilde{\psi}_V\rangle}{\sqrt{p_V}}. \end{aligned} \tag{3.3}$$

Stokes measurements are insensitive to global phase, that is, phase that is common to both $|H\rangle$ and $|V\rangle$ polarisation components. Luckily, global phase does not matter for our purposes here. Thus

we can define, without loss of generality, the global phase of $|\psi_V\rangle$ such that $\langle\psi_H|\psi_V\rangle = \langle\psi_H|\psi_V\rangle^*$. And so we can use the following orthonormal qubit basis:

$$\begin{aligned} |0\rangle &= \frac{1}{\sqrt{2c_+}}(|\psi_H\rangle + |\psi_V\rangle), \\ |1\rangle &= \frac{1}{\sqrt{2c_-}}(|\psi_H\rangle - |\psi_V\rangle) \\ \text{with} \quad c_\pm &= 1 \pm |\langle\psi_H|\psi_V\rangle|. \end{aligned} \tag{3.4}$$

Choosing such a basis is going to be useful for defining a concurrence, since we can now explicitly write $|\Psi_{vec}\rangle$ in the form of equation (3.2), since

$$\begin{aligned} |\psi_H\rangle &= \frac{1}{2}\sqrt{2c_+}|0\rangle + \frac{1}{2}\sqrt{2c_-}|1\rangle, \\ |\psi_V\rangle &= \frac{1}{2}\sqrt{2c_+}|0\rangle - \frac{1}{2}\sqrt{2c_-}|1\rangle \end{aligned} \tag{3.5}$$

and thus

$$|\Psi_{vec}\rangle = \sqrt{p_H}\frac{1}{2}\sqrt{2c_+}|H\rangle|0\rangle + \sqrt{p_H}\frac{1}{2}c\sqrt{2c_-}|H\rangle|1\rangle + \sqrt{p_V}\frac{1}{2}\sqrt{2c_+}|V\rangle|0\rangle - \sqrt{p_V}\frac{1}{2}\sqrt{2c_-}|V\rangle|1\rangle, \tag{3.6}$$

where $|0,1\rangle$ refer to two orthogonal spatial modes. We then find the expression for the concurrence in terms of the parameters from equation (3.2) in the form suggested by [53],

$$C(|\Psi\rangle) = 2|ad - bc|, \tag{3.7}$$

which we can do by comparison with the expression in equation (3.6). This yields

$$\begin{aligned} C(|\Psi_{vec}\rangle) &= 2\left|-\sqrt{p_H p_V}\frac{1}{2}\sqrt{c_+c_-} - \sqrt{p_H p_V}\frac{1}{2}\sqrt{c_+c_-}\right| = 2\sqrt{p_H p_V}\sqrt{1 - |\langle\psi_H|\psi_V\rangle|^2} \\ &= 2\sqrt{\langle\tilde{\psi}_H|\tilde{\psi}_H\rangle\langle\tilde{\psi}_V|\tilde{\psi}_V\rangle - |\langle\tilde{\psi}_H|\tilde{\psi}_V\rangle|^2}, \end{aligned} \tag{3.8}$$

where we have defined the polarisation basis as $|H\rangle = |0\rangle$ and $|V\rangle = |1\rangle$. Stokes parameter measurements can provide all information required here. Spatially resolved Stokes parameters are obtained from spatially resolved intensity measurements of the state Ψ_{vec} projected onto the 6 polarisation basis vectors:

$$\begin{aligned} I_H(x,y) &= |\langle H|\langle x,y|\Psi_{vec}\rangle|^2, \\ I_V(x,y) &= |\langle V|\langle x,y|\Psi_{vec}\rangle|^2, \\ I_D(x,y) &= |\langle D|\langle x,y|\Psi_{vec}\rangle|^2, \\ I_A(x,y) &= |\langle A|\langle x,y|\Psi_{vec}\rangle|^2, \\ I_R(x,y) &= |\langle R|\langle x,y|\Psi_{vec}\rangle|^2, \\ I_L(x,y) &= |\langle L|\langle x,y|\Psi_{vec}\rangle|^2. \end{aligned} \tag{3.9}$$

Using the definitions of the polarisation basis states from section 1.2 and equation (3.1), and dropping the explicit spatial dependence of $A_{H,V}$ and $\phi_{H,V}$ for conciseness, we can calculate these

as

$$\begin{aligned} I_H(x,y) &= |\langle x,y|\psi_H\rangle|^2 = |A_H|^2, \\ I_V(x,y) &= |\langle x,y|\psi_V\rangle|^2 = |A_V|^2, \\ I_{D/A}(x,y) &= \frac{1}{2}|\langle x,y|\left(|\psi_H\rangle \pm |\psi_V\rangle\right)|^2 = \frac{1}{2}\left(A_H^2 + A_V^2 \pm 2A_H A_V \cos\left(\phi_H - \phi_V\right)\right), \\ I_{R/L}(x,y) &= \frac{1}{2}|\langle x,y|\left(|\psi_H\rangle \pm i|\psi_V\rangle\right)|^2 = \frac{1}{2}\left(A_H^2 + A_V^2 \pm 2A_H A_V \sin\left(\phi_H - \phi_V\right)\right). \end{aligned} \tag{3.10}$$

It can be shown that the concurrence in (3.8) can be expressed in terms of these as

$$C(|\Psi_{vec}\rangle) = \sqrt{(I_H + I_V)^2 - ((I_H - I_V)^2 + (I_D - I_A)^2 + (I_R - I_L)^2)}, \tag{3.11}$$

where the Stokes parameters are integrated over space. This is a particularly simple formula that is very easy to use with experimentally available data. In the following section we outline the experimental procedure where we use this expression to quantify the correlations in several vector beams, and briefly compare the results with an alternative approach based on projection measurement.

3.3 The Glasgow experiment

Based on the above description we expect that we can use concurrence to quantify correlations within light fields with spatially varying polarisation. Therefore we need to generate and measure such beams. Any of the dynamic polarisation-shaping setups described in section 2.4 would be capable of doing this; we opted to use the DMD and Wollaston prism-based device (see figure 2.14) because of its simplicity and the lack of need for high optical powers. The laser beam from a He-Ne source polarised using a PBS, and it is spatially filtered by placing a 50 μm pinhole in the far field of a telescope (lenses: 40 mm to 75 mm). It is then rotated to $|A\rangle$ polarisation using a Fresnel rhomb, and enters the Wollaston prism. The imaging telescope between the prism and the DMD is comprised of a 100 mm and a 125 mm lens. The DMD is imaged onto a camera (Allied Vision GC660) using another telescope (both lenses 200 mm), with a far field filter in the form of a variable aperture (Thorlabs ID36/M) placed in the appropriate plane.

At the two extreme ends of the correlation spectrum are beams with uniform polarisation and some classes of vector beams in which the spatial modes for the two orthogonal polarisation components are also orthogonal. In this section we will call beams in the second category maximally correlated beams. A beam with uniform polarisation carries no correlation between spatial structure and polarisation. For such beams the concurrence is 0. For example, consider any $|H\rangle$-polarised beam: $C = \sqrt{(I_H + 0)^2 - ((I_H - 0)^2 + (1/2I_H - 1/2I_H)^2 + (1/2I_H - 1/2I_H)^2)} = 0$. On the other hand, for maximally correlated beams the concurrence evaluates as 1 if the Stokes parameters are normalised, although this is more involved to show, and we will not do so here.

While it is possible to analytically evaluate the concurrence for all beams discussed here, we opt to do the calculation numerically. We can then pass the measured Stokes parameters and those calculated numerically through the same analysis pipeline. The analysis and numerical simulations are carried out in LabVIEW, using parts of the WaveTrace package.

We generate beams where the two orthogonal polarisation components are $|H\rangle$ and $|V\rangle$ and the corresponding spatial profiles are those of Laguerre-Gaussian modes. We include a parameter θ the value of which allows us to tune the degree of correlation. The beams are then described in

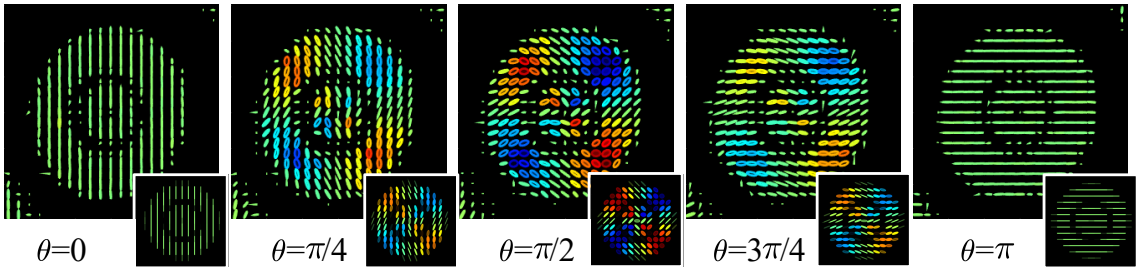

Figure 3.1: Evolution of the polarisation structure of a vector beam $(\cos(\theta/2)LG_1^1\hat{h} + \sin(\theta/2)LG_1^{-1}\hat{v})$. Insets show theoretical structures. Note that the only information missing from these measurements is the spatially dependent phases of the $\hat{h}$ and $\hat{v}$ components, but their relative phases are manifested in the polarisation structures.

the form

$$\mathbf{E} = \cos(\theta/2)LG_{p_H}^{\ell_H}\hat{h} + \sin(\theta/2)LG_{p_V}^{\ell_V}\hat{v}. \tag{3.12}$$

Here varying θ from 0 to π the concurrence should go from 0 through 1 to 0. At $\theta = 0$ $\mathbf{E} = LG_{p_H}^{\ell_H}\hat{h}$ so $C = 0$ and at $\theta = \pi$ $\mathbf{E} = LG_{p_V}^{\ell_V}\hat{v}$, for which $C = 0$ also. At $\theta = \pi/2$ we have an equal weighting of the two polarisation components, and $C = 1$. Some examples of such beams are shown in figure 3.1.

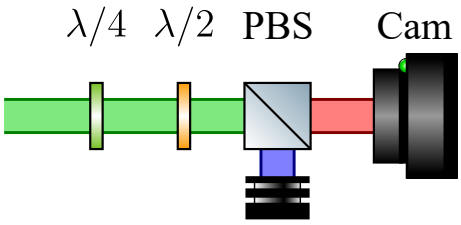

Figure 3.2: Polarimeter setup. Details in main text. PBS: polarising beam splitter.

The Stokes parameters are measured using a simple automated polarimeter setup shown in figure 3.2. The vector beam passes through a quarter- and a half-wave plate followed by a polariser (in the form of a polarising beam splitter) and finally arrives at a detector in the image plane of the DMD. Different orientations of the waveplates allow the polariser to project the incoming beam onto all the polarisation basis states according to table 3.1. This is the most efficient way of measuring all six projections, requiring one to turn only one waveplate for each projection. It is possible to calculate Stokes parameters from only four measurements, however, by using the relation $I_H + I_V = I_D + I_A = I_R + I_L$. Later in this section we will discuss why we have not done this.

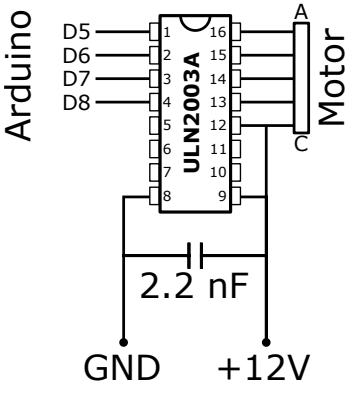

Figure 3.3: Stepper motor driver circuit diagram. For the Arduino connection different ports can also be used.

The automation comes in with the waveplates being placed in 3D-printed rotation stages driven by stepper motors and controlled by Arduino boards connected to the experiment control computer. The Arduino boards need to be connected to the stepper motor through a small control board that we built ourselves. The circuit diagram for this is shown in figure 3.3. Depending on the reduction ratio of the stepper motors this system can measure all Stokes parameters in about 10 seconds. Of course a higher reduction ratio results in increased accuracy but reduced speed. The components of the 3D printed rotation stage, designed by Ermes Toninelli at the University of Glasgow, are shown in figure 3.4.

In principle one could also replace the half-wave plate with a sheet polariser, remove the beam splitter, and rotate the polariser by double the half-wave plate angles listed in table 3.1. The reason we choose not to do this is because the sheet polarisers available to us were mounted between

Table 3.1: Waveplate angles with respect to the horizontal for projection onto the polarisation basis states for Stokes measurements.

Polarisation state	$\lambda/4$	$\lambda/2$
$\|H\rangle$	0°	0°
$\|R\rangle$	0°	22.5°
$\|V\rangle$	0°	45°
$\|L\rangle$	0°	67.5°
$\|A\rangle$	45°	67.5°
$\|D\rangle$	45°	22.5°

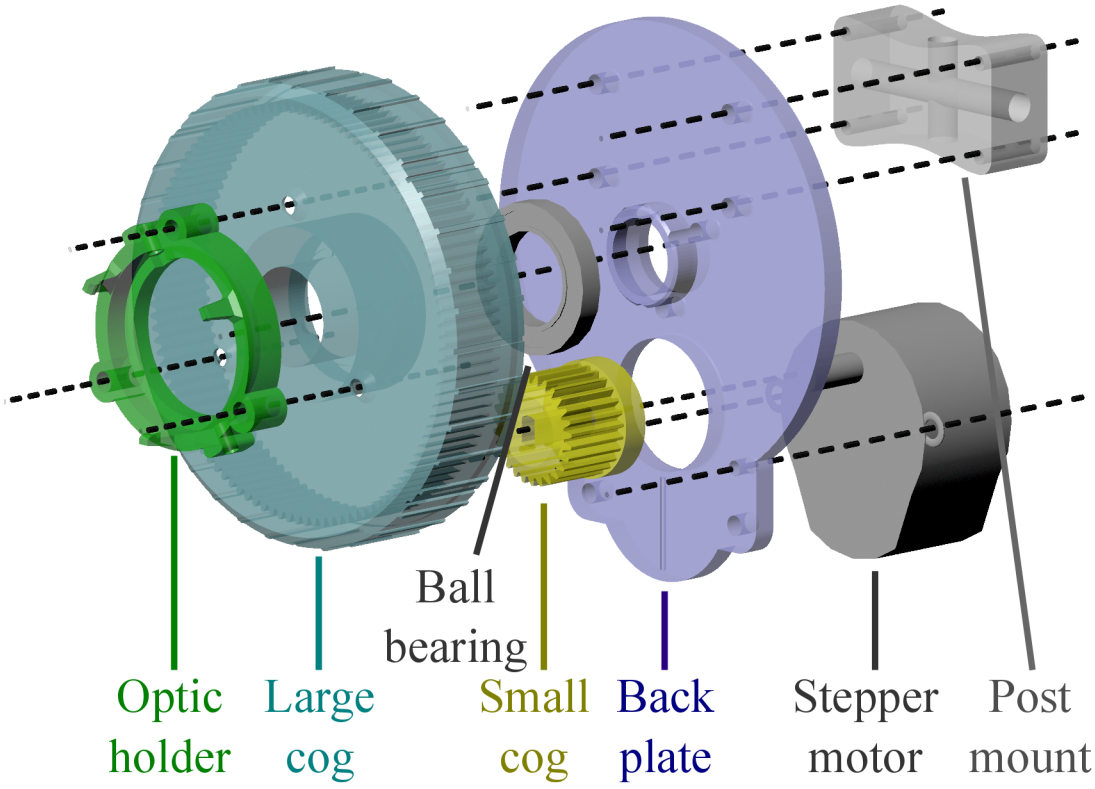

Figure 3.4: Components of the 3D printed rotation stages used to measure Stokes parameters. The waveplates in the experiment were mounted into the 1 inch internal diameter optic holder (green), which was attached to the large cog (cyan). The large cog had teeth on its inner surface, connecting to a small cog (yellow) driven by a stepper motor (dark gray). The stepper motor was attached to a back plate (light purple), which connected to the large cog with a ball bearing (gray) and also to a post mount (light gray) to which standard Thorlabs posts could be attached via M6 grub screws. The assembly provides a roughly 1 cm diameter through path for optical access. The components were held together by M3 screws at the connections indicated by the dashed black lines.

two relatively thick plates of acrylic. This is not, in itself, a problem, but if they are not placed in the beam path exactly perpendicular to the propagation direction then the beam is displaced due to refraction. This is still not much of an issue, but if the rotation stage wobbles, then the displacement due to refraction changes with polariser angle, which leads to the beam moving on the detector. This is undesirable, because it would make the calculation of polarisation patterns difficult, since it would make it necessary to align the recorded Stokes images in software. Since the Stokes parameters used in determining concurrence are integrated over the beam profile one could focus the output of the polarimeter setup onto a photodiode instead of imaging onto a camera.

Noise is a significant issue for Stokes parameter measurements because of normalisation. The expression for the concurrence, including the normalisation chosen for the experiment, is a form of equation 3.11,

$$C = \sqrt{1 - \left(\left(\frac{I_H - I_V}{I_H + I_V} \right)^2 + \left(\frac{I_D - I_A}{I_D + I_A} \right)^2 + \left(\frac{I_R - I_L}{I_R + I_L} \right)^2 \right)}. \tag{3.13}$$

Of course in theory $I_H + I_V = I_D + I_A = I_R + I_L$, but it is possible that in the experiment we have optical elements that absorb in different ways for different polarisation states, in which case a slightly more reliable normalisation could be obtained using the expression in (3.13). Issues arise in areas of the images where the signal-to-noise ration (SNR) is low, i.e. in areas where there is little light in the beams. Because of the division in the normalisation these areas are prone to causing large fluctuations in the Stokes parameters. Recall that the point of the concurrence measurement is to discern whether there is spatial information correlated with polarisation. If there is any kind of spatial fluctuation due to noise this will cause us to overestimate concurrence.

In a way this is not surprising at all. The background noise does not carry polarisation structure, but neither is it associated with any particular polarisation. The main sources for this noise is background light scattered from room lights, LEDs and screens in the lab, and CCD noise. Scattered light is unpolarised, and CCD noise is not related to polarisation at all. A characteristic of unpolarised light is that the normalised Bloch vector, the vector comprised of the last three normalised Stokes parameters $S_{1,2,3}$, associated with it has a magnitude less than 1. But looking at the definition of concurrence, this property is shared with vector beams, if we integrate over the beam profile! The Bloch vector of vector beams integrated over their profile is indeed shorter than 1, and it has a length 0 for maximally correlated beams. So unpolarised light, like noise, is indistinguishable from vector beams using this measure. This is why we pay careful attention to reducing noise in the recorded images. The way we do this is by low-pass filtering the images. We Fourier transform the images and remove a few of the highest frequency components. In practice we chose to remove the two highest frequency components as this gave very good results without an overly harsh reduction of detail in the beams. This, of course, limits the spatial resolution that we can detect in the beams, but for the smooth beam profiles we were interested in this was not an issue.

Results of the measurements are shown in figure 3.5. It can be seen that there is excellent agreement between the numerical simulations and the measured data. We obtain the error bars by performing the same measurement 21 times and taking the standard deviation of the calculated concurrences.

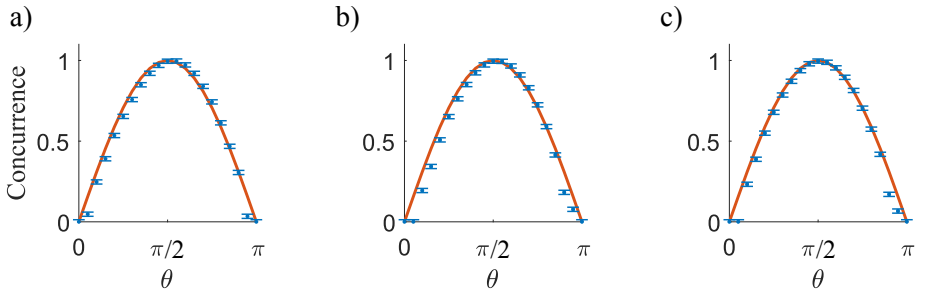

Figure 3.5: Concurrence measurements calculated from data (blue) and numerical simulations (orange) from Stokes measurements at the University of Glasgow. The polarisation basis is $\hat{h}, \hat{v}$. The LG modes used are, in the format $[\ell_0, p_0; \ell_1, p_1]$, a) [1,0;-1,0], b) [1,1;-1,1], c) [3,0;1,1].

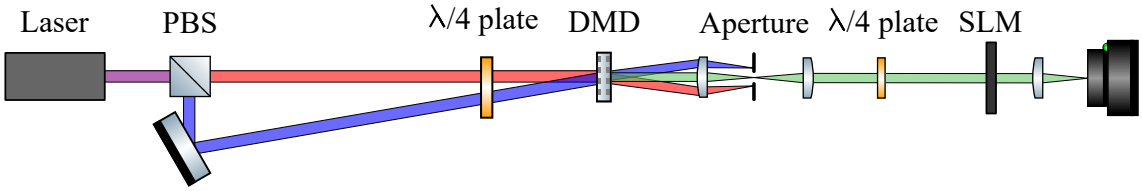

Figure 3.6: Vector beam generating and measurement setup built at the University of Witwatersrand.

3.4 A related experiment at the University of Witwatersrand

Our collaborators at the University of Witwatersrand carried out a different experiment intended to measure the concurrence of vector beams. Their experimental setup for generating vector beams, shown in figure 3.6 was conceptually identical and practically similar to ours as well. They have opted to use a polarising beam splitter instead of a Wollaston prism and overlapped the two orthogonal polarsation components on the DMD by directing one of the outputs from the beam splitter using a mirror. They have included a quarter-wave plate before the DMD to transform from the $|H\rangle - |V\rangle$ basis to the circular one.

Of more interest is the difference in concurrence measurement setups. Our collaborators projected the vector beams onto not only polarisation basis states but also spatial modes using an SLM in the reverse operation compared to how it would be used for beam shaping, followed by on-axis intensity measurement. This exploits that beam propagation is symmetric in z; in other words, if we know the optical field in any plane we can determine it in any other plane. Thus if we propagate for example a LG_p^ℓ beam backwards through a hologram that is designed to generate a LG_p^ℓ beam from a Gaussian, the output field will be a Gaussian. This is a form of projection; the more the output of such a hologram is like a Gaussian the more the projected beam was like the one the hologram was designed to generate if a Gaussian propagated through it. In fact, one can measure the on-axis intensity of the output, which is related to the projection $\langle \Psi_{in} | \Psi_{test} \rangle$.

Naturally, for this to work one needs to know which subspace of spatial modes to test, since the whole space of spatial modes is infinite. With this restriction, though, the method works very well, as can be seen in figure 3.7 showing measurement results kindly provided by our collaborators. The uncertainty in these measurements depends on the particular beam chosen. The measurement was performed using a camera, and uncertainty was calculated by taking the standard deviation of the 64 pixels around the on-axis pixel as the uncertainty of the on-axis pixel value, and propagating this through equation 2 in [52]. Interestingly these measurements have a tendency to overestimate

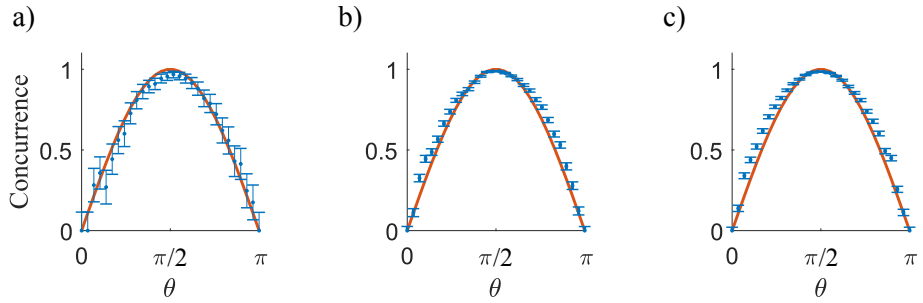

Figure 3.7: Concurrence measurements calculated from data (blue) and numerical simulations (orange) from projection measurements at the University of Witwatersrand. The polarisation basis is $\hat{r}, \hat{l}$. The LG modes used are, in the format $[\ell_0, p_0; \ell_1, p_1]$, a) [1,3;-1,3], b) [3,1;-3,1] and c) [4,0;-4,0].

concurrence, which suggests that it is also susceptible to noise issues. Our collaborators generated beams of a form similar to ours in 3.5 but using the $|R\rangle - |L\rangle$ polarisation basis,

$$E = \cos(\theta/2) LG_{p_R}^{\ell_R} \hat{r} + \sin(\theta/2) LG_{p_L}^{\ell_L} \hat{l}. \tag{3.14}$$

It can be seen that their results also agree very well with the theoretical predictions. However, to achieve these results a subset of spatial modes that contain the ones that make up the vector beams needed to be known *a priori*, which would be difficult to guarantee if the beams to be characterised would be unknown.

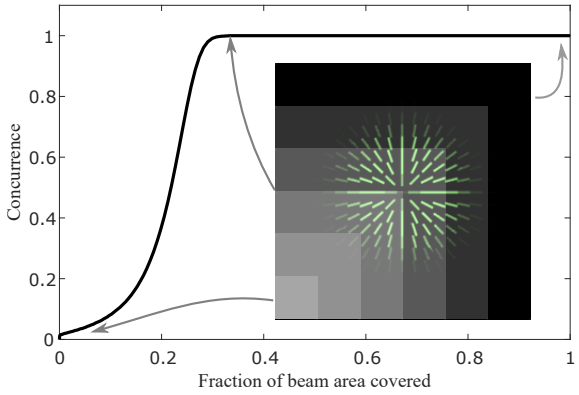

Figure 3.8: Concurrence evaluated for sub-areas of a maximally correlated beam. The inset shows some example sub-areas. Notice how concurrence reaches its maximum value as soon as all polarisation directions are included in the sub-area.

Let us return briefly to the idea that based on concurrence measurement we cannot tell the difference between structured polarisation and unpolarised light. The difference between the two cases is, of course, that for structured polarisation the polarisation state at any point in space, and more critically any transverse plane, is well-defined, whereas for unpolarised light it is not. So we can tell the difference between the two cases if we test the polarisation at a single point. Of course for such a scenario concurrence would be zero, since there would be no spatial information to support a correlation. This leads to a question: if we look at a single pixel in our spatially resolved Stokes measurements for a maximally correlated beams, we measure zero concurrence, but if we integrate over the entire beam we measure maximal concurrence, so how exactly does the observed area influence our measured concurrence? We investigated this question numerically with one example of a maximally correlated beam that was radially polarised. We evaluated the concurrence for growing sub-areas of the beam, as shown in figure 3.8. While the exact shape of the curve in the figure depends on the polarisation structure of the beam in question, the key observation is that the concurrence reaches its maximal value as soon as all polarisation directions are included in the sub-area. This suggests that for beams with certain symmetries it is not necessary to integrate over the entire beam, but it is sufficient to interrogate a small area which contains all polarisations. This should be the case for vector vortex beams, for example, but not for full Poincaré beams [54].

3.5 In summary

We have shown that Stokes measurements can be used to calculate concurrence for the internal degrees of freedom (polarisation and spatial modes) of vector beams. We have experimentally demonstrated the effectiveness of this approach for quantifying correlations in vector beams. We also compared our method to the estabilished technique of projection measurements on both degrees of freedom.

Part II

Matter

Chapter 4

Fluorescence and the optical Bloch equations

4.1 Introduction

So far in this thesis we have exclusively dealt with optics, in manipulating light using dynamic and static elements, and saw how we can extract information about correlation between degrees of freedom of photons using only such things. It is time to turn our attention to the effects of light on the material world, and look into how light-matter interactions can be used to manipulate matter, especially through using properties of structured light.

The system we will be concerned with is rubidium atoms under the influence of resonant, and later near resonant, coherent laser radiation. Atoms are nice for the study of light-matter interactions because they interact rather strongly with electromagnetic radiation near resonance to an atomic transition (and in some cases even far from resonance the effects can be surprisingly strong [55]). Atoms of the same species and isotope are absolutely identical so they also interact with a light field in the same way,[1] enhancing or suppressing interaction when working in concert [56, 57, 58].

Rubidium is a particularly experiment-friendly species of atoms for two reasons. First, rubidium has only one electron in its outermost shell, which leads to a rather simple, almost hydrogen-like electronic level structure. Despite being simple, there are enough levels to provide a rich system with many opportunities to investigate new physics even today, when atomic physics has been at the centre of significant research efforts for a century.

Second, some very useful rubidium transitions are in the near infrared (IR), extremely close in frequency to that of laser diodes found in DVD drives, which used to make laser sources extremely cheap until DVDs went out of style in favour of online streaming services[2]. Coincidentally, detectors such as CCDs and photodiodes are usually very sensitive to near-IR light, which makes the detection and analysis of radiation scattered/transmitted by rubidium atoms easy. Liquid crystal devices developed for visible light are also an excellent base technology on which to build spatial light modulators and Q-plates for near-IR.

[1] Provided that they do not experience different external environmental parameters, such as magnetic fields.

[2] Isn't it interesting how two such seemingly unconnected fields as atomic physics and the entertainment industry can have connections? The history of technology is, I'm sure, full of such fascinating connections.

In this part of this thesis we will first look at scattering phenomena, using the scattering properties and electronic level structure of rubidium to learn something about the intensity structure of shaped laser beams, and in turn dictate the electronic state of atoms in three dimensions. Then, we will describe methods of cooling atoms to extremely low temperatures, and exploit the very slow thermal motion of cold atoms and the magnetic coupling of electronic levels to gain information about magnetic fields.

4.2 Rubidium level structure

Figure 4.1 summarises the structure of the electronic states of both relevant isotopes of rubidium. Throughout Part II of this thesis we will need to consider different structures of these levels, so it makes sense to discuss all of the details here. We will look at energy level shifts due to various effects in order of magnitude.

In a very rudimentary picture of an atom one can consider electrons to orbit the nucleus[3]. For our purposes here it is sufficient to consider two orbital angular momentum states of the outermost electron of the rubidium atom, or two orbitals. These are the S and P orbitals, with angular momentum quantum numbers $L = 0$ and $L = 1$, respectively. The S orbital is the ground state for this electron, and is separated by the P orbital by roughly 380 THz[4] [59], or in units of wavelength, 790 nm.

The electron, being a spin 1/2 particle, carries its own spin angular momentum as well as orbital angular momentum. The spin angular momentum quantum number can take the values $S = \pm 1/2$. The two angular momenta can interact, an effect that is known as spin-orbit coupling. The total electronic angular momentum can then be expressed as $\mathbf{J} = \mathbf{L} + \mathbf{S}$ (since angular momenta are vectors) [60]. $|\mathbf{J}|$ can take values between $|L - S| \leq J \leq L + S$, which means that for the S state $J = 1/2$. However, for the P state, with $L = 1$, J can take two values: $J = 1/2$ and $J = 3/2$. This is known the fine splitting, indicated in figure 4.1 by the spectroscopic notation L_J. This fine splitting is of the order of 7 THz, or, expressed in more conventional units for this range, 15 nm.

The electron is not the only particle with angular momentum in the atom, however. The nucleus carries its own angular momentum $\mathbf{I}$, which also couples to the electronic angular momentum, leading to a total atomic angular momentum $\mathbf{F} = \mathbf{J} + \mathbf{I}$, the magnitude of which again takes values between $|J - I| \leq F \leq J + I$. Therefore the fine structure levels also contain sublevels labelled by F, an effect known as the hyperfine splitting. This is where we encounter the first major difference between the two relevant isotopes of rubidium. The nuclear angular momentum quantum number is different for the isotopes: $I_{85} = 5/2$, while $I_{87} = 3/2$. This means that the F quantum numbers for ^{85}Rb are higher by 1 when compared to ^{87}Rb. There are too many hyperfine levels to list here, but they are clearly labelled in figure 4.1. In any dipole-allowed transition $|\Delta F| \leq 1$. The hyperfine levels are separated by a few gigahertz in the ground state S, and by energies from a few tens to a few hundreds of megahertz for the excited P levels.

The last structure we need to consider is the magnetic structure of the hyperfine levels. Each hyperfine level contains $2F + 1$ magnetic sublevels, labelled by the quantum number m_F, which can take values between $-F \leq m_F \leq F$. The magnetic sublevels are degenerate in the absence of external magnetic fields, but they undergo Zeeman shift in the presence of one. This shift is given

[3] It is well known that this picture is inaccurate, but for this discussion this simple and intuitive picture will work.

[4] When we refer to some energy E in units of frequency ω, we mean that $E = \hbar\omega$. Wavelength is meant in a similar way.

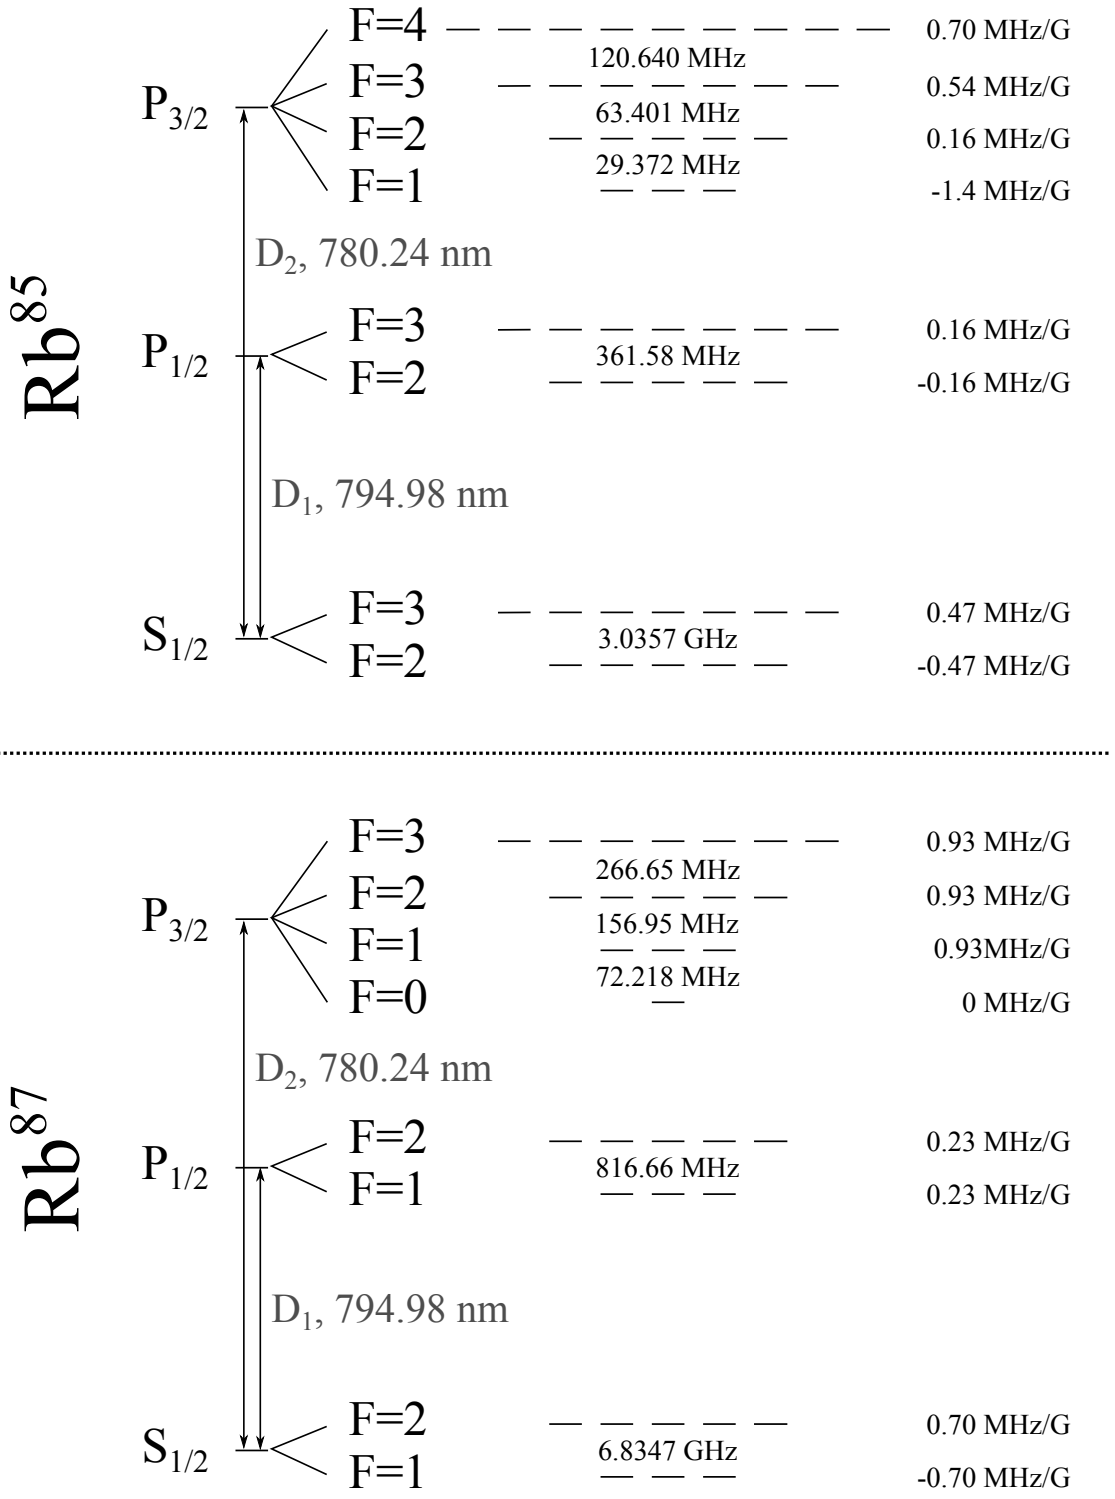

Figure 4.1: Electronic level structure of the outermost electron (5s) of two isotopes of rubidium, ^{85}Rb and ^{87}Rb. Energy increases from the bottom to the top, although level spacings are not to scale. Fine and hyperfine structure are shown, with splitting energies expressed in nanometres for fine structure and units of frequency for hyperfine splitting. Magnetic (m_F) sublevels are also shown, with the Zeeman shift per m_F indicated on the right.

by

$$\Delta E_{m_F} = \mu_B g_F m_F B_z, \tag{4.1}$$

where μ_B is the Bohr magneton, as usual, B_z is the z-component of the magnetic field, with z being the atomic quantisation axis, and g_F is the hyperfine Landé g-factor, given by

$$\begin{aligned} g_F &= g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} \\ \text{with} \quad g_J &= 1 + \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)}, \end{aligned} \tag{4.2}$$

where on the first line we omitted a factor proportional to the nuclear g-factor due to its three orders of magnitude smaller effect, and on the second line we took the approximate values for the spin and orbital g-factors $g_S \approx 2$ and $g_L \approx 1$. Because of the weakness of magnetic interaction the energy shift due to the Zeeman effect is usually rather small. Neighbouring m_F states tend to shift by around 0.5 MHz per Gauss external field, and even the difference between the two furthest magnetic sublevels $P_{3/2}$ manifold of ^{85}Rb is, at 1 G, less than 7 MHz, which is on the scale of the linewidth of these transitions. The scales of the energy separations between adjacent levels in orbital, fine, hyperfine and Zeeman levels in rubidium is shown for interest in figure 4.2.

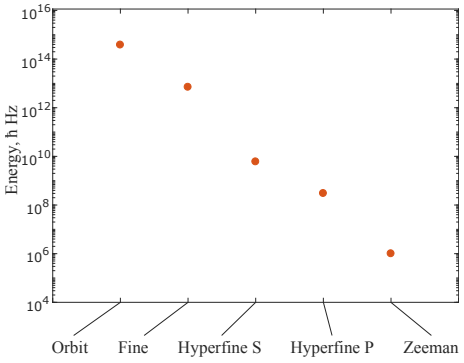

Figure 4.2: Scales of energy separations between adjacent levels in orbitals, fine, hyperfine and Zeeman levels in rubidium.

Fine and hyperfine levels are separated in energy enough to be easily resolved even by cheap, home-made lasers, which will be described in the following experimental chapter. In fact, it is not the laser linewidth that prevents us from resolving the magnetic sublevels, but rather the natural transition linewidth. However, just because the m_F levels are not resolvable, it does not mean that they cannot be addressed. They are, in fact, crucial to our work on cold atoms.

4.3 Absorption and emission

As we have mentioned a few times already, the electronic structure of atoms can be influenced by the presence of light[5]. Perhaps the most obvious interaction between atoms and light is the absorption of light. In everyday life we are familiar with opaque materials (which are, of course, made of atoms) that absorb light; this is usually a solid state phenomenon, however, and has quite different properties compared to atomic interactions.

An atom can only absorb photons from a light field whose energy is equal to the energy difference between two electronic energy levels of the atom, $\hbar\omega = \Delta E$. Even this can only happen when the atom is in the lower energy state. In such a scenario the photon is removed from the external light field, and the atom is excited to the higher energy state[6]. Excited states have a finite lifetime, meaning that atoms decay to lower energy states over some time with the emission of a photon of the same energy. This is known as spontaneous emission. The finite lifetime τ leads to a broadening of the emission spectrum due to the uncertainty relationship between the energy and lifetime of the state, $\Delta E\tau \approx \hbar$. The naturally broadened emission line has the shape of a Lorentzian [60] with a width related to the inverse of the lifetime of the excited state $\Gamma = 1/\tau$, which is the radiative decay rate. Γ gives the rate of spontaneous emission from an atom if there are enough photons in

[5]This should come as no surprise, since the electrons are bound to the nucleus through electromagnetic interaction, which is, in a way, not very different from light.

[6]In a more accurate description the photon and atom become a single "thing", but this makes little difference here.

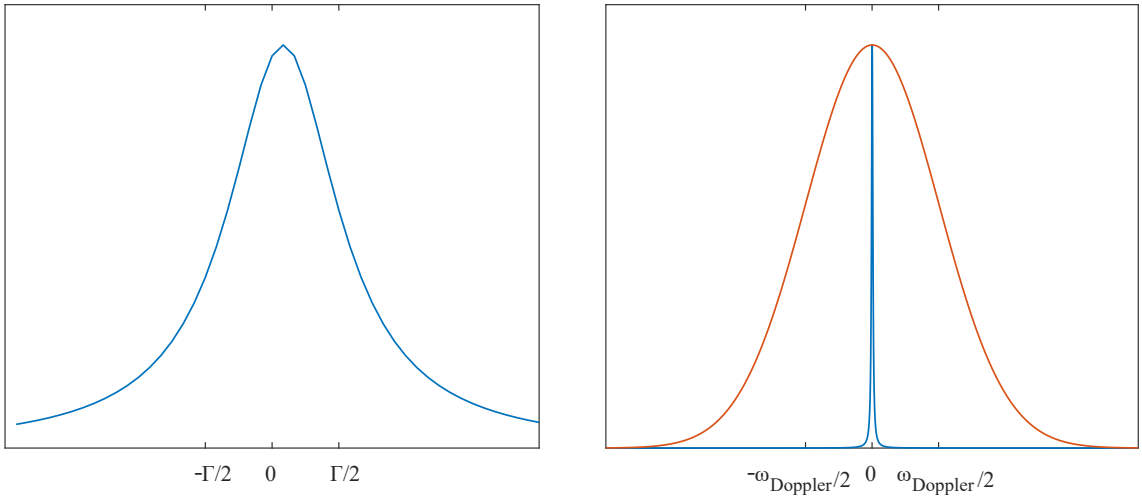

Figure 4.3: Lorentzian shape of natural broadening (left) and Gaussian shape of Doppler broadening (right). Their relative width for the rubidium transitions considered here is shown on the right.

the external light field to excite it from the lower state over a time that is much shorter than the excited state lifetime. We also ignore the time it takes for an atom to enter the excited state, since this is extremely short for atoms[7] It is also possible to de-excite an atom by using an external light field. This process is called stimulated emission. This process is crucial to the operation of lasers, but we will not consider it here as it is not relevant to the processes we study.

During absorption and emission processes a number of quantities are conserved. In subsequent chapters the conservation of linear momentum during absorption (and emission) of radiation will be of great importance. It can be expressed as

$$\mathbf{p_{final}} = \mathbf{p_{initial}} + \hbar\mathbf{k_{photon}}, \tag{4.3}$$

where the $\mathbf{p}$ is the linear momentum of the atom. Naturally, energy is also conserved, so

$$\begin{aligned} \Delta E_{total} &= \left(E_e + \frac{\mathbf{p_{final}} \cdot \mathbf{p_{final}}}{2M}\right) - \left(\hbar\omega + E_g + \frac{\mathbf{p_{initial}} \cdot \mathbf{p_{initial}}}{2M}\right), \\ \text{so} \quad \hbar\omega &= \hbar\omega_0 + \frac{\hbar\mathbf{k_{photon}} \cdot \mathbf{p_{initial}}}{M} + \frac{\hbar^2\mathbf{k_{photon}} \cdot \mathbf{k_{photon}}}{2M}, \end{aligned} \tag{4.4}$$

where $E_e - E_g = \hbar\omega_0$ is the energy difference between the atomic states, $\hbar\omega$ is the photon energy, M is the total atomic mass. The second and third terms on the right hand side in the second line of equations 4.4 correspond to the Doppler shift and recoil energy, respectively. Since the recoil energy varies strongly with photon energy, it is significant only for high frequency transitions, and for the optical transitions considered here it is insignificant and so we ignore it.

In a thermal gas of temperature T the atomic momentum distribution follows the Maxwell-Boltzmann law, so it has a Gaussian profile with a width of $\omega_{\mathbf{Doppler}} \approx k_{\mathbf{photon}}\sqrt{k_B T/M}$. The absorption and emission spectra are broadened by this momentum distribution, and they also have a Gaussian shape[8] with a width $\hbar\omega_{\mathbf{Doppler}}$. For room temperature rubidium this broadening is on the order of 500 MHz. The shapes and relative widths of natural and Doppler broadening are shown in figure 4.3.

[7] But the transition can be observed in other analogue systems, and it has been shown that the evolution from one state to another is a continuous, coherent process that can even be reversed with appropriately high temporal resolution in the experiment control [61].

[8] Technically speaking their profile is the convolution of the Doppler Gaussian with the natural Lorentzian known as the Voigt profile, but the natural width is so narrow compared to the Doppler width that this hardly matters.

Polarisation	Δm_F
σ_+	+1
π	0
σ_-	-1

Table 4.1: Changes in magnetic quantum numbers under the influence of different light polarisations.

We should note that for absorption the photon momentum **k** is dictated by the laser, both in magnitude and direction. For spontaneous emission, however, there is no constraint on its direction and so spontaneous emission is isotropic.

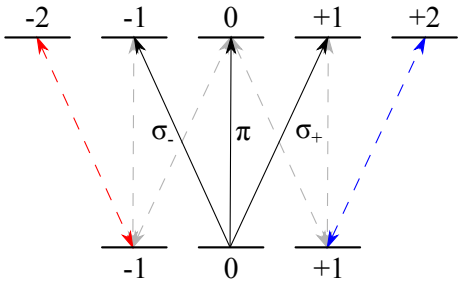

Figure 4.4: Polarisation selection rules for dipole-allowed transitions, illustrated here for the $S_{1/2}F = 1 \rightarrow P_{1/2}F' = 2$ transition of ^{87}Rb. All possible decay and excitation channels are shown. Notice that the $m'_F = \pm 2$ can only be reached by (and decay through) $\sigma_\mp$ respectively.

Angular momentum is also conserved during absorption and emission. This leads to selection rules on the possible transitions for different polarisations of the light field [62]. Here we only consider dipole-allowed transitions. In this situation we need to consider three types of light polarisations: σ_+, σ_- and π. Depending on the angle between the atomic quantisation axis[9] and the light propagation direction these map to different light polarisations. In the case when the two axes are aligned $\sigma_+ = \hat{l}$, $\sigma_- = \hat{r}$ and π corresponds to polarisation along the z-direction (or, alternatively, linear polarisation emitted orthogonally to the atomic quantisation axis, where the polarisation direction is along the quantisation axis).

Interaction with these polarisations affects the magnetic m_F sublevels differently. For all of the transitions considered here the electronic orbital angular momentum changes by 1, $\Delta L = 1$. During such interactions $\Delta m_F = 0, \pm 1$. How these correspond to the three light polarisations is summarised in table 4.1. These selection rules are extremely important as they make it possible to manipulate the internal structure of atoms simply by an appropriate choice of light polarisation, which we have excellent control over (see chapters 1 and 2).

Figure 4.4 illustrates the polarisation selection rules for the $S_{1/2}F = 1 \rightarrow P_{1/2}F' = 2$ transition of ^{87}Rb. Note that most m_F levels can be excited, and can decay to, a number of other levels. Usually the laser field, in conjunction with the atomic quantisation axis, dictates which excitation channels are utilised in any experiment. The polarisation of spontaneous emission depends on the specific decay channel. The probability of following any specific transition is proportional to the overlap of the angular momenta of the start and end states, which is quantified by square of the Clebsch-Gordan coefficients. These numbers can be found in data tables.

We would like to draw attention to a configuration of magnetic sublevels that can be found in $F \rightarrow F' = F + 1$ transitions. In such transitions there exist two pairs of magnetic sublevels, $|m_F| = F \rightarrow |m'_F| = F + 1$, which are connected by only one possible excitation and decay channel, through $\sigma_\pm$ polarisation. An example is shown in colour in figure 4.4. Such $m_F = \pm F'$ states are often called stretched states, and the transitions to them are known as cycling transitions, because if an atom can undergo one such transition once, it can only decay to a level from which it must undergo the same transition, provided that the polarisation of the driving light does not

[9]The quantisation axis is the direction onto which angular momenta are projected to find values of quantum numbers.

contain polarisation that can drive the other transitions. The redistribution of atomic populations between states by interaction with light is called optical pumping.

4.4 Fluorescence and the optical Bloch equations

The discussion in this chapter follows the work outlined in [63] and [59].

When illuminated by a continuous stream of coherent photons, in other words, CW laser radiation near resonance, atoms continuously absorb and spontaneously emit light. As we have said before, the emission is isotropic, and is called fluorescence. We are interested in the rate at which fluorescence photons are emitted by the atoms. This, of course, depends on the intensity of the incoming light field as well as properties of the states. To find this rate we need to introduce the optical Bloch equations. Here we will consider only two levels, which requires a bit of justification. All levels apart from the magnetic sublevels are easily resolved by the narrow linewidth lasers we use, because their separation is large compared to the atomic linewidth.

This is not the case for the magnetic sublevels. However, there are several scenarios in which we do not have to worry about treating them as separate levels. If the laser field carries polarisation that contains σ_+, σ_- and π polarisation simultaneously,[10] then there is no selectivity in the interaction between the magnetic sublevels and they can be all addressed and the interaction is averaged over them. In the case of atoms optically pumped into a stretched state the interaction can only happen between two Zeeman sublevels, so the 2-level approximation makes sense again.

We begin by introducing the quantum mechanical density operator for a state of a two-level atom $|\Psi\rangle = c_g|g\rangle + c_e|e\rangle$, where the ground $|g\rangle$ and excited $|e\rangle$ states have complex probability amplitudes. The density operator is then defined as $\hat{\rho} = |\Psi\rangle\langle\Psi|$. It is, however, more commonly written in the form of a matrix, the well-known density matrix:

$$\hat{\rho} = \begin{pmatrix} \rho_{ee} & \rho_{ge} \\ \rho_{eg} & \rho_{gg} \end{pmatrix} = \begin{pmatrix} c_e c_e^* & c_g c_e^* \\ c_e c_g^* & c_g c_g^* \end{pmatrix}. \tag{4.5}$$

The diagonal elements of the density matrix, ρ_{gg} and ρ_{ee}, are known as populations of the states $|g\rangle$ and $|e\rangle$ respectively. They represent to the probability of finding the atom in the corresponding state, and for ensembles of atoms they can represent the occupation of the states in the ensemble. The off-diagonal elements are called coherences, and depend on the relative phases between the complex amplitudes associated with the states. The density matrix formalism can be readily extended to multiple levels, where the elements contain pairwise coherences between levels in the off-diagonal elements. The diagonal always contains the populations of the levels.

Of course absorption and spontaneous emission processes are all time-dependent, so we need to be able to express the evolution of the density matrix under the influence of an external radiation field. The von Neumann equation gives this time evolution, if the interaction can be written as a Hamiltonian operator $\hat{H}$:

$$\frac{d\hat{\rho}}{dt} = \frac{i}{\hbar}\left[\hat{\rho}, \hat{H}\right]. \tag{4.6}$$

Unfortunately, spontaneous emission cannot be described in terms of a Hamiltonian. However, the power of the density matrix approach is that it allows us to incorporate even such processes. In the previous section (4.3) we stated that excited states decay at a rate given by the inverse lifetime of the state, Γ. After this decay a two-level atom must be in the ground state, so it fills up at the

[10] This is achievable with an appropriate angle between the atomic quantisation axis and the laser direction.

same rate. So we can write, for the populations, that

$$\frac{d\rho_{ee}}{dt} = -\frac{d\rho_{gg}}{dt} = -\Gamma\rho_{ee}. \tag{4.7}$$

We also need to describe the evolution of the coherences. It turns out that they also depend on the decay rate,

$$\frac{d\rho_{eg}}{dt} = -\frac{\Gamma}{2}\rho_{eg} \quad \text{and} \quad \frac{d\rho_{ge}}{dt} = -\frac{\Gamma}{2}\rho_{ge}. \tag{4.8}$$

We can thus write down the entire time evolution of the density matrix as

$$\frac{d\rho}{dt} = \frac{i}{\hbar}\left[\hat{\rho}, \hat{H}\right] + \begin{pmatrix} \Gamma\rho_{ee} & -\frac{\Gamma}{2}\rho_{ge} \\ -\frac{\Gamma}{2}\rho_{eg} & -\Gamma\rho_{ee} \end{pmatrix}. \tag{4.9}$$

All we have to do now is obtain an expression for the interaction Hamiltonian, and we can write down the complete time-evolution of the populations. To do this, we start with the Schrödinger equation,

$$i\hbar\frac{\partial\Psi(\mathbf{r},t)}{\partial t} = \hat{H}\Psi(\mathbf{r},t), \tag{4.10}$$

with the atomic wavefunction

$$|\Psi\rangle(\mathbf{r},t) = c_g(t)|g\rangle + c_e(t)|e\rangle e^{-i\omega_0 t}. \tag{4.11}$$

The Hamiltonian in equation 4.10 is constructed from two parts, one corresponding to the bare atom (with eigenstates $|g\rangle$ and $|e\rangle$), and one to the external field, which is due to the atomic dipole interacting with the external field: $\hat{d}\cdot\hat{E}$. Using these with equations 4.10 and 4.11, we obtain time-evolution equations for the complex amplitudes $c_{g,e}$,

$$\begin{aligned} i\hbar\frac{dc_g}{dt} &= c_e\langle g|\hat{d}\cdot\hat{E}|e\rangle e^{-i\omega_0 t} \\ i\hbar\frac{dc_e}{dt} &= c_g\langle e|\hat{d}\cdot\hat{E}|g\rangle e^{+i\omega_0 t}. \end{aligned} \tag{4.12}$$

The laser field that we consider here is a monochromatic plane wave, which we can write as the expression at the beginning of this thesis, equation 1.3, recast to the form $\mathbf{E} = E\exp[i(\mathbf{k}\cdot\mathbf{r}-\omega t)]\hat{e}$. Then, we can write equations 4.12 as

$$\begin{aligned} i\hbar\frac{dc_g}{dt} &= c_e\hbar\Omega^*\left(\frac{e^{i(\omega-\omega_0)t} + e^{-i(\omega+\omega_0)t}}{2}\right) \\ i\hbar\frac{dc_e}{dt} &= c_g\hbar\Omega\left(\frac{e^{i(\omega+\omega_0)t} + e^{-i(\omega-\omega_0)t}}{2}\right), \end{aligned} \tag{4.13}$$

where we have defined the Rabi frequency $\Omega = E/\hbar\langle e|\hat{d}\cdot\hat{e}|g\rangle$, which charaterises the strength of the interaction between the light field and the atom. It depends on the light intensity[11] and the dipole matrix element $\hat{d}\cdot\hat{e}$, but because the latter is fixed for each transition we often consider the Rabi frequency analogous to intensity.

Note that on the right-hand side of equations 4.13 there are oscillations with two frequencies, $\omega-\omega_0$ and $\omega+\omega_0$. Close to resonance the former is small whereas the latter is very large. Because of this we can make the approximation that over the slow evolution of the complex amplitudes terms with the large frequency undergo many oscillations and their value averages to 0. This is known

[11] Recall that $I = |E|^2$.

the Rotating Wave Approximation. We can then write equations 4.13 in the much simpler form

$$\begin{aligned} i\hbar\frac{dc_g}{dt} &= c_e\hbar\Omega^*\frac{e^{i\Delta t}}{2} \\ i\hbar\frac{dc_e}{dt} &= c_g\hbar\Omega\frac{e^{-i\Delta t}}{2} \end{aligned} \tag{4.14}$$

by defining the detuning $\Delta = \omega - \omega_0$, the difference between the frequency of the driving field and that of the atomic transition.

Comparing these expressions with 4.10 we could read off the matrix form of the Hamiltonian, but at this point it would contain time-dependent terms. We can move the time-dependence to the complex amplitudes by moving into a new, time-dependent basis (in a reference frame co-rotating with the light field). This can be done by making the new definitions $c'_g = c_g$ and $c'_e = c_e e^{i\Delta t}$, which allows us to rewrite equations 4.14 as

$$\begin{aligned} i\hbar\frac{dc'_g}{dt} &= c'_e\frac{\hbar\Omega}{2} \\ i\hbar\frac{dc'_e}{dt} &= c'_g\frac{\hbar\Omega}{2} - c'_e\hbar\Delta. \end{aligned} \tag{4.15}$$

We can thus write down the time-independent Hamiltonian:

$$\hat{H} = \frac{\hbar}{2}\begin{pmatrix} -2\Delta & \Omega \\ \Omega & 0 \end{pmatrix}. \tag{4.16}$$

We are finally at the stage where we can substitute $\hat{H}$ into equation 4.9 and write down the optical Bloch equations in full glory,

$$\begin{aligned} \dot{\tilde{\rho}}_{ee} &= -\frac{i\Omega}{2}\left(\tilde{\rho}_{ge} - \tilde{\rho}_{eg}\right) - \Gamma\tilde{\rho}_{ee} \\ \dot{\tilde{\rho}}_{gg} &= \frac{i\Omega}{2}\left(\tilde{\rho}_{ge} - \tilde{\rho}_{eg}\right) + \Gamma\tilde{\rho}_{ee} \\ \dot{\tilde{\rho}}_{ge} &= -\frac{i\Omega}{2}\left(\tilde{\rho}_{ee} - \tilde{\rho}_{gg}\right) - i\Delta\tilde{\rho}_{ge}\frac{\Gamma}{2}\tilde{\rho}_{ge} \\ \dot{\tilde{\rho}}_{eg} &= \frac{i\Omega}{2}\left(\tilde{\rho}_{ee} - \tilde{\rho}_{gg}\right) + i\Delta\tilde{\rho}_{eg}\frac{\Gamma}{2}\tilde{\rho}_{eg}, \end{aligned} \tag{4.17}$$

where we transformed the density matrix to the co-rotating reference frame as well, indicated by the tilde above the matrix elements, e.g. $\tilde{\rho}_{ge} = \rho_{ge}e^{-i\Delta t}$. Note that the trace of the density matrix is 1, so $\rho_{ee} + \rho_{gg} = 1$ at all times, because there are no other states that the atom could occupy. The same logic can be used to extend these equations for any number of levels, although they do get quite messy very quickly. Solving the optical Bloch equations is another matter entirely; analytical solutions are known for a few special cases, but usually we solve them numerically. We also note that it is possible to incorporate the decohering effect of collisions into these equations by an additional term, $-i\tilde{\rho}_{ge/eg}\gamma_c$, in the equations for the coherences, where the collisional decay rate γ_c can be obtained from various models for collisions. For the densities and temperatures that we work with collisions between atoms are negligible, so we need not worry about this.

Luckily, in order to obtain an expression for the steady-state fluorescence rate all we need to know is the population in the excited state ρ_{ee}, since only atoms in the excited state can fluoresce, so the fluorescence (or scattering) rate is just $\Gamma\tilde{\rho}_{ee}$. Setting $\dot{\tilde{\rho}}_{ee} = 0$ (steady state) we find that the fluorescence rate is

$$R_f = \Gamma\tilde{\rho}_{ee} = \frac{\Gamma}{2}\frac{(I/I_{sat})}{1 + 4(\Delta/\Gamma)^2 + (I/I_{sat})}, \tag{4.18}$$

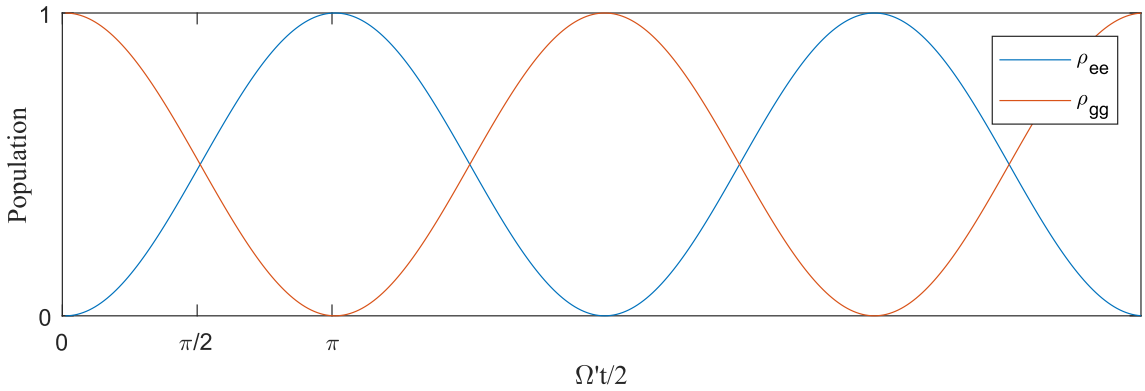

Figure 4.5: Rabi oscillations of the ground and excited state populations using a CW laser.

where we have defined the saturation intensity

$$I_{sat} = \frac{c\epsilon_0 \Gamma^2 \hbar^2}{4|\hat{e} \cdot \hat{d}|^2} \tag{4.19}$$

such that $I/I_{sat} = 2(\Omega/\Gamma)^2$. We will revisit the expression for the fluorescence rate in a subsequent chapter on controlling the populations in a warm vapour in a spatially resolved manner.

4.5 Rabi flops

Let us return briefly to equations 4.14. It is possible to find solutions to these equations by differentiation, which gives

$$\begin{aligned} \frac{d^2 c_g}{dt^2} - i\Delta \frac{dc_g}{dt} + \frac{\Omega^2}{4} c_g = 0 \\ \frac{d^2 c_e}{dt^2} + i\Delta \frac{dc_e}{dt} + \frac{\Omega^2}{4} c_e = 0. \end{aligned} \tag{4.20}$$

It is interesting to look at how the populations, $|c_e|^2$ and $|c_g|^2$, vary with time in the presence of the external radiation field. Taking the initial condition that at $t = 0$ $c_g = 1$ (and consequently $c_e = 0$) yields the solution

$$\begin{aligned} |c_g(t)|^2 &= \frac{\Omega^2}{\Omega'^2} \cos^2\left(\frac{\Omega' t}{2}\right) \\ |c_e(t)|^2 &= \frac{\Omega^2}{\Omega'^2} \sin^2\left(\frac{\Omega' t}{2}\right), \end{aligned} \tag{4.21}$$

where we have defined $\Omega' = \sqrt{\Omega^2 + \Delta^2}$. figure 4.5 shows the time-evolution of the two states. The probability to find the atom in either state varies sinusoidally, an effect known as Rabi oscillation, or Rabi flops.

Note that if the radiation field was turned on for a time $t_\pi = 2\pi/\Omega'$ the population of the ground state would be completely and coherently transferred to the excited state. A pulse of this length is known as a π-pulse. The other notable pulse length is $t_{\pi/2} = \pi/\Omega'$ for a $\pi/2$-pulse, which transfers the population to an equal superposition of the ground and excited states. Sequences of such pulses are extremely important in optical clocks and quantum information processing, but they are less critical for our work and so will not be considered further[12].

[12]Despite the rich possibilities... It's a shame that there is finite time for PhDs.

4.6 In summary

We have introduced enough concepts of atom-light interactions to allow us to understand the more advanced techniques of subsequent chapters. We looked at the electronic and magnetic structure of energy levels of the 5s electron in two isotopes of rubidium, ^{85}Rb and ^{87}Rb. We have described the absorption and spontaneous emission of (near-)resonant light by atoms. We have obtained the optical Bloch equations and used them to find an expression for the scattering rate. In the next chapter we will look at ways of obtaining resonant laser light for use in atom optics experiments.

Chapter 5

Resonant laser light for working with atoms

5.1 Introduction

In the previous chapter we saw that the internal state of atoms can be manipulated by (near-)resonant electromagnetic fields. Through the conservation of linear and angular momentum external, motional states are also controllable by such fields, which leads to various cooling and trapping schemes, some of which are described later in this thesis. All such manipulations require monochromatic coherent radiation with tunable frequency. Sometimes internal state manipulation can be done by microwave radiation (a good example is in the operation of caesium clocks), but usually useful atomic transitions require optical frequencies ranging from the ultraviolet to the infrared. For example, the rubidium transitions we work with are all in the near IR. At these wavelengths we use lasers to obtain coherent radiation.

There are many different tunable laser designs, and the experimental requirements[1] should determine which one is to be used. On the most versatile end of the scale we have Ti:Sapphire lasers with extremely large frequency ranges (typically ~ 300 nm [64] that can be increased by frequency conversion [65]), low linewidths (a few dozen kilohertz) and high output powers in the range of a few watts. However, for many important applications much simpler laser systems are sufficient. In our experiments we use external cavity diode lasers (ECDLs), which still provide wavelength tunability over a few nanometres, with a couple of hundred kilohertz linewidth and around 100 mW of output power [66]. Let us look into how they work in more detail.

5.2 External cavity diode lasers

Bare laser diodes are not very useful for atomic physics applications because of their broad linewidth, on the scale of 100s of megahertz, and the difficulty of precisely controlling their frequency. One can introduce a frequency-selective feedback to reduce the linewidth of diode lasers by placing the diode in an external cavity. The frequency selectivity is provided by a reflective grating completing the cavity. The grating splits the relatively broadband laser light into a spectrum. By carefully tuning the angle of the grating with respect to the laser propagation direction one can

[1] and available funding...

direct a specific frequency component back to the diode, completing the cavity for that wavelength and thus forcing the gain medium in the diode to amplify the selected wavelength.

It can then be seen that for a given grating period greater cavity lengths lead to narrower linewidths. This is because the grating dispersion angle is fixed by the period, and so for shorter cavities more wavelengths are allowed to circulate in the cavity. This is illustrated in figure 5.1. A competing effect is that the free spectral range of a cavity is inversely proportional to its length, $\Delta\nu = c/2L$. This translates to a reduced mode-hop-free tuning range for ECDLs, essentially limiting the scannable frequency range. Longer cavities are also more sensitive to angle noise on the grating. As a result in practice one needs to find an optimal external cavity length that yields an acceptable linewidth and tuning range.

The ECDL frequency can obviously be tuned by changing the grating angle, selecting a different frequency for feedback. There are, however, other parameters that need to be considered. The diode gain properties are temperature-dependent, which means that their output spectra vary with temperature. The gain medium is electrically pumped, which makes the spectrum depend on the driving current as well. As a result stable control of these three parameters (grating angle, diode current and temperature) is desirable if one wishes to have good control over the output frequency of an ECDL.

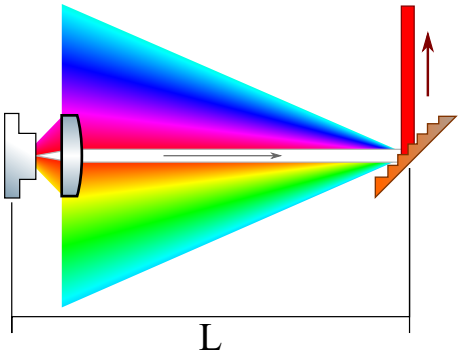

Figure 5.1: Frequency selective feedback in an ECDL. Initially the laser is relatively broadband (white beam), which is dispersed by the grating, and diffracted back towards the diode on the left. Only a limited wavelength range, dependent on the external cavity length L, is captured by the collimating lens and are amplified by the gain medium in the diode. The output is then narrow-band (red beam).

5.2.1 Older ECDL design

The older ECDLs used in our cold atom experiments were constructed from collimator tubes and gratings fixed to kinematic mirror mounts attached to large aluminium blocks. The laser diode and its collimating lens were contained in the collimator tube, which was attached to the kinematic mirror mount back plate by screws, with thermal paste between them ensuring good thermal contact. A holographic diffraction grating (1800 lines/mm) was glued to an aluminium spacer which was screwed to the kinematic mount front plate. The external cavity length was chosen to be close to 20 mm, which gives a sub-MHz linewidth and reasonably large (> 7 GHz) mode-hop-free tuning range. A hole was drilled into the mirror mount front plate opposite to the horizontal angle adjustment screw, and a piezoelectric transducer (PZT[2]) was fixed into this hole for electronic fine control of the grating angle. The kinematic mount assembly was screwed to large (approximately $10 \times 10 \times 5$ cm) aluminium blocks used as heat sinks, with a Peltier cooler between the mount and aluminium block for active temperature control. A thermistor was attached to the mirror mount for feedback to the temperature controller. This construction is shown in figure 5.2. The electronic control of diode current and temperature was handled by commercial Thorlabs diode laser driver boards.

These older ECDLs are in the so-called Littrow configuration, where the grating angle provides the frequency control, but unfortunately the beam output direction also depends on the grating angle. This results in unwanted beam deflection as the frequency is tuned. To be fair, the deflection is

[2]We realise that the abbreviation is often used to refer to lead zirconate titanate, a specific piezoelectric material. Apologies for the confusion.

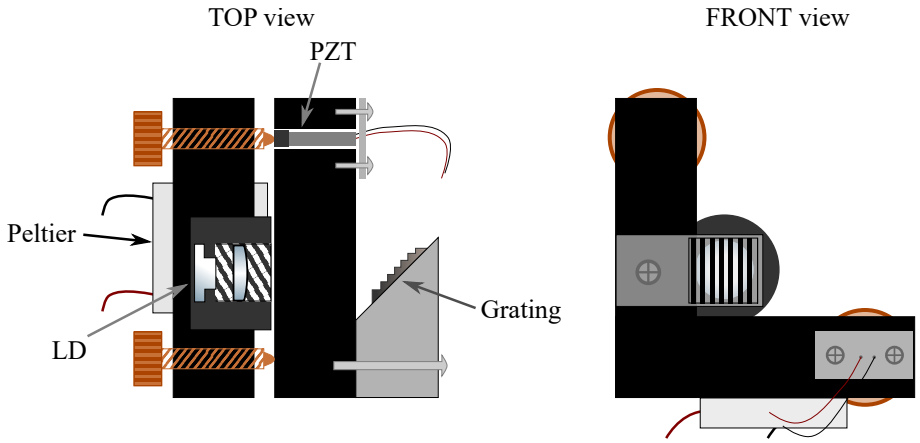

Figure 5.2: Schematic ECDL design for the older models used in our experiments. LD: laser diode; PZT: piezo-electric transducer.

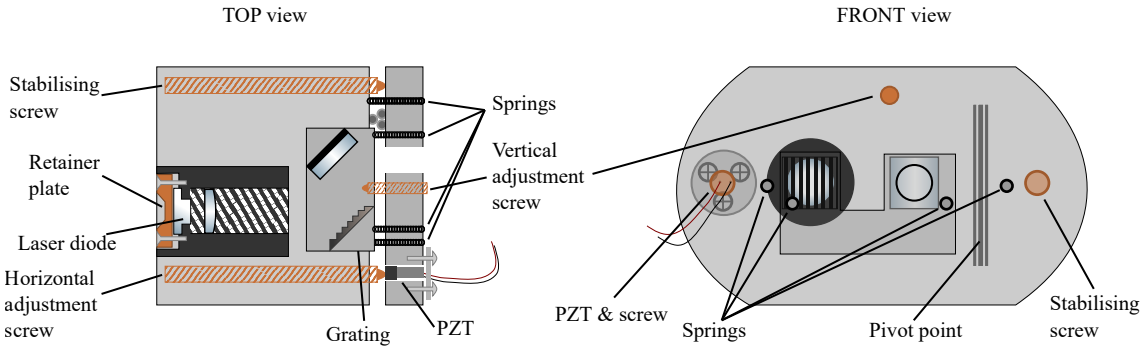

Figure 5.3: Schematic diagram of the new generation ECDLs, excluding the Peltier unit and the adjustable height platform.

extremely small once the laser is set up to emit close to the desired frequency, and is usually barely detectable even after metres of propagation.

5.2.2 New generation ECDLs

Recently a new generation of ECDLs were constructed and used for our warm atom work. The design is based on work by Robert Wylie, Ewan MacLagan, Paul Griffin and Erling Riis. The operating principle of this new design is very similar to that of the older design detailed above, with one modification. A mirror is inserted opposite and parallel to the grating, and its angle is tuned in unison with the grating (they are mounted together). This ensures that the beam output direction is independent of the laser frequency [67] because the additional deflection caused by a change in the grating angle is exactly compensated by the rotation of the parallel mirror. Note that while the beam output direction is unchanged under tuning, there is a very small lateral displacement. This is negligible for all our experiments.

Another change for the new designs is the lack of commercial parts. Most parts of the extended cavity were machined from aluminium at the local workshop. The design comprises 5 main parts: a collimating tube, a housing for the collimating tube, a grating and mirror mount, a front plate, and an adjustable height mounting platform. These parts are shown in figure 5.3. Collimation tube contained a threaded mounted 6.24 mm focal length aspheric collimating lens (Thorlabs C110TME-B). The laser diode was secured to the aluminium tube by a brass back plate screwed to the tube with 3 screws. The collimation tube tightly fit into the aluminium housing block, which provides adequate thermal contact. It was secured by a plastic-ended grub screw. The grating and

the mirror were glued to the machined mount, which was attached to the aluminium front plate via two tight springs. The front plate and grating mount assembly was secured to the collimation tube housing via another two springs.

The front plate assembly was allowed to pivot around a vertically aligned metal cylinder glued to the housing held between two identical cylinders glued to the front plate. The horizontal angle of the front plate assembly was controlled by a long ball-ended brass screw going through the housing. The ball end rested against a sapphire pad attached to a PZT, which pressed against a screwed-on retainer plate on the front plate. On the other side of the housing, near the horizontal pivot point, a spring-loaded ball-ended screw rested against the front plate to increase the plate movement stability at high frequencies. The vertical tilt of the front plate assembly was adjusted by a small ball-ended brass screw going through the front plate above the laser output port, resting between two small metal cylinders glued to the housing. The front plate can be tilted through less than 10° in both horizontal and vertical directions.

The ECDL assembly was mounted on a bespoke adjustable height mounting platform by four screws. A Peltier thermal control unit was inserted between the mounting platform and the laser housing, with the thermistor supplying temperature feedback inserted into a small drilled hole in the laser housing immediately above the Peltier unit, surrounded by ample thermal paste to provide reliable information on the housing temperature.

5.2.3 Setting up an ECDL for spectroscopy

Once basic construction is complete, setting up an ECDL is a multi-stage process taking a few hours for a practiced operator. The first step is collimation, followed by setting up the external cavity for feedback at the desired operating temperature, then a rough frequency calibration takes place. After this the laser can usually be operated with only occasional minimal adjustments for many years, unless something breaks. We will now give a detailed walkthrough for the above stages. We refer to the new generation ECDLs in details, but the principles translate straightforwardly to the older versions. Most of the observations should be equivalent for other ECDL designs as well.

For collimation the collimation tube is removed from the housing and secured in a position that allows the output beam to propagate several metres, either in a straight line, or using mirrors to fold the beam path. We usually clamp the collimation tube in a V-mount, for example Thorlabs VC3C/M. Straight propagation is preferred because adjusting the collimating lens causes it to shift laterally slightly in the tube, which changes the beam output direction significantly, potentially causing it to miss the mirrors. Because of this it is advisable to place signs outside the area in use warning of possible stray light. The laser diode should be driven at currents such that the output power is low, between 2-5 mW should do. Because of the beam direction changing during collimation we usually compare the beam size in two planes, on an infrared detector card, by eye[3] first. The collimating lens position is adjusted[4] such that the beam size is the same in two plane separated by as much distance as practically achievable. Since the size of the beam changes less close to the diode, matching the far plane size to the close plane size is the way to go. We make sure that there is no focal point between the two planes by scanning the detector card along the entire propagation. The beam size of diode lasers is usually difficult to determine with accuracy because they change shape on propagation (they are highly elliptical, with axis ratios changing

[3]This should not need saying, but just in case, NEVER look directly into a laser beam or a specular reflection of it! This is a bad idea even with safety goggles on.

[4]Usually this is done by poking the lens tube with a small precision screwdriver because we find that in practice this tool gives the best control, provided that the operator has decent steady hands.

significantly over long distances). One can use beam profilers, either commercial devices or home made ones built of a camera and a software that fits gaussians to the recorded beam profile, to get more quantitative metrics of beam size along different axes in the two planes. However, we find that in practice very good results can be achieved just by eye; often, intuition about equivalent beam sizes is better than trying to compare the widths of differently shaped beams.

Once satisfactory collimation has been achieved, the collimation tube is placed back in the housing and secured in such a way that the laser polarisation is vertical (the elliptical beam should have its major axis horizontal to sample most of the grating for good frequency selectivity). This can be checked with a polarising beam splitter, by maximising the reflected (not diffracted) power from the grating, or by noting that the major axis of the elliptical beam profile corresponds to the direction of polarisation for most diodes. Vertical polarisation is chosen to minimise the grating efficiency. Only a small amount of diffracted light is necessary in the external cavity, because the laser gain will quickly suppress other wavelengths, and we usually want as much output power in the 0th order of the grating as we can get. This is also the time to set the external cavity length. We set this by aligning the brass retainer plate of the laser diode to be co-planar with the back plane of the housing. This gives an approximately 2 cm long external cavity, which seems to be optimal for tunability and stability. Once the collimation tube is secure we bring the laser assembly to operating temperature using the Peltier unit.

At this point the external cavity is very likely to be misaligned. To align it, the diode current is lowered to just above lasing threshold. For our diodes[5] this is around 40-45 mA, where the ECDL output power is around 1 mW. The output power is then observed on a power meter or an oscilloscope using a photodiode. The goal is to maximise the output power. A jump to at least twice, but likely 3 or 4 times the output power is a hallmark of good feedback and hence cavity alignment. This is because the external cavity enhances the diode laser operation, pushing it further past threshold. First, the horizontal direction is adjusted to maximise laser power. We observe a wide power curve with a slow change in output power. A scan of a few degrees is usually necessary to find this maximum. Then, the vertical direction is optimised. This is a much harder process, and may need repeating for a couple of horizontal alignments. The power curve along the vertical direction is mostly flat with a few jumps to much higher laser powers, and it is *extremely* sensitive to the vertical angle. Sometimes it can be difficult to find one peak in power that is obviously larger than the rest; this indicates a need to change the horizontal direction a bit. A correctly aligned cavity produces dramatically more power than the other configurations for which power peaks are seen.

When feedback is achieved, or in doubt, the power meter is replaced with a spectrometer. The cavity is aligned correctly if when the horizontal adjustment screw is *carefully* rotated the laser spectral line moves smoothly and continuously, without jumps and secondary spectral lines appearing, over several nanometres. A good diode, cavity length and alignment allows scanning at least 3-5 nm. If the tuning range is narrower and the spectral line jumps, sometimes this can be fixed by adjusting the vertical screw just after a jump has occurred, such that the spectral line jumps back to the last position along the smooth scan. If this is doable, one can try to scan back and forth again and see if the scanning range is extended. This is difficult to do and does not always work. If it does not, the alignment should be started from the beginning again, looking at the output power near threshold.

This process assumes that the vertical and horizontal screws tilt the grating in orthogonal direc-

[5] Unfortunately their make and model is unknown, nobody remembers what they are and they are in an unlabeled box.

tions. This is not the case if the grating was glued on tilted in such a way so as to couple these directions. If this is the case it might be impossible to get very good tunability, but even then it should be possible to achieve 1-2 nm tuning range, unless the tilt is excessive. If the tuning range does not reach 1 nm the grating needs re-mounting. When satisfactory tunability is achieved at low powers the driving current should be increased smoothly to operating current, while monitoring the laser spectrum. We frequently see small secondary spectral lines appearing around 100 mA driving current, and sometimes all power concentrates in these at operating currents of > 130 mA. Do not exceed 150 mA for long periods of time with these diodes, it is likely that their lifetime shortens considerably above those currents. We observe that sometimes in this power range previously good tunability deteriorates, or completely disappears. When this is observed the feedback needs to be optimised at low powers again, before ramping up to operating current again. If the ECDL tunes well at operating power we tune it within a few gigahertz of its intended frequency. In our experiments this is done by picking off a small amount (< 10 mW) of the laser light and direct it through a rubidium reference cell. We observe the cell, either through an IR viewer or by an IR camera, while slowly scanning the laser frequency over the expected resonance frequency using the horizontal screw. When fluorescence is detected the laser is ready for stabilisation to a rubidium transition via Doppler-free spectroscopy, a process described in the following sections.

We built three copies of the new generation ECDLs. Only two were ever stabilised, but both of those showed identical performance. As far as deterioration is concerned, we have observed laser diodes ceasing operation only after several years; also, on one occasion a PZT needed replacing as well, after around 10 years of operation. We have observed glued gratings becoming loose early in the life of the new generation ECDLs, but after re-gluing three years ago they seem to be stable. At the time we observed that the gratings had an oval pattern on their surface, likely dust or glue fumes burned on by the laser. However, this does not seem to affect operation in any way. The new generation ECDLs were built in 2015, with one laser diode last replaced in 2017 because of the then high temperature operation. The older designs had their diodes replaced in 2018. We expect that under normal operating conditions all lasers should function adequately until at least 2022.

5.3 Doppler-free Spectroscopy

The ECDLs described above provide narrow linewidth monochromatic coherent light to interact with the atoms. The task that is left is to tune their frequencies to match that of atomic transitions. One advantage of working with neutral atoms is that it is easy to get a frequency reference. We use glass cells filled with low pressure rubidium vapour at room temperature (e.g. Thorlabs GC25075-RB). As described in section 4.3, the natural linewidth of the transitions of interest is $\sim 2\pi \times 6$ MHz and the hyperfine separation for the excited states is a couple dozen megahertz, so the few hundred kilohertz laser linewidth should allow us to resolve all of them. Unfortunately the Doppler broadening at room temperature of around 500 MHz makes it impossible to resolve the excited hyperfine states without doing something more clever than just pointing a laser beam through the vapour cell and looking at the absorption signal with changing laser wavelength.

The standard technique to get around the Doppler broadening issue is saturated absorption spectroscopy [68]. This requires two laser beams propagating in opposite directions, called the pump and probe beams, and can be easily achieved by placing a mirror behind the reference cell, as shown in figure 5.4. We can look at the returning probe beam by placing a non-polarising beam splitter before the reference cell and placing a photodiode into its reflected arm.

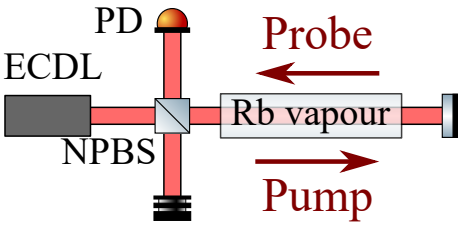

Figure 5.4: Doppler-free spectroscopy setup.

The role of the pump beam is to excite atoms, reducing the ground state population, when the laser frequency is resonant with a transition. This condition is satisfied for atoms with velocity $v_1 = (\omega - \omega_0)/k$, where ω is the laser frequency, ω_0 is the transition frequency and k is laser wavenumber. The probe beam interacts with atoms in a similar manner, but because of the inverted propagation direction it is absorbed by atoms with velocity $v_2 = -(\omega - \omega_0)/k$. An atom can interact with both beams, but in this case $v_1 = v_2$[6], which implies that $\omega = \omega_0$ (or $\omega = 0$, which is not the case because light has a finite frequency). This means that for atoms interact with both beams the Doppler shift is eliminated, essentially they travel in a transverse direction to the beam axes. The probe beam sees less such atoms in the ground state because they have been excited by the pump beam, and so the absorption of the probe beam is reduced. Thus we observe narrow dips in the absorption spectrum, or peaks in the oscilloscope trace (since the photodiode detects transmission, not absorption), at the hyperfine resonances, which have a width of only the natural linewidth.

If there are multiple excited states within the laser scan range, it is possible that the probe beam sees a depleted ground state population not just for stationary atoms, but also atoms moving with a velocity such that the pump beam depleted them on a transition different from what the probe beam is probing [69]. This translates to the condition that $\omega + kv = \omega_{0,1}$ and $\omega - kv = \omega_{0,2}$, in the case where an atom is moving at such a speed that $2kv = \omega_{0,1} - \omega_{0,2}$. Thus we observe another transmission peak at a frequency $\omega = (\omega_{0,1} + \omega_{0,2})/2$, halfway between the two actual transition frequencies. These crossover peaks are usually much stronger than the peaks corresponding to the actual resonances.

In saturated absorption spectroscopy the pump beam is much more intense than the probe beam, which could be achieved in a setup like figure 5.4 by placing a neutral density (ND) filter between the cell and mirror, but in practice we find that this does not improve the signal, so we keep both beams at comparable intensity. To avoid confusion we simply call this setup Doppler-free spectroscopy.

In the experiments we scan the laser frequency by supplying a time-dependent voltage signal to the PZTs in the ECDLs. In the older ECDLs the PZTs were driven with a triangle wave at 30 Hz with an amplitude of 5-10 V, while the new ones were driven with a sawtooth wave at 19 Hz with an amplitude up to 100 V. The scan range could be adjusted by changing the amplitude of the voltage signal, and the centre of the scan range could be chosen by applying a DC offset to the signal. Example oscilloscope traces for all relevant transitions are shown in figure 5.5, with hyperfine transitions labelled by their F quantum number.

Let us give some more practical remarks on obtaining the absorption spectra shown in figure 5.5. In the following we assume that the laser temperature is stabilised to a reasonable value. For the D_2 lines (780 nm) room temperature, or slightly below room temperature, is satisfactory, whereas for the D_1 lines (795 nm) a higher temperature of around 50°C is required. The first step is always[7] to use the ECDL wavelength tuning screw while monitoring fluorescence from the reference cell, while the piezo scan is on and scanning a wide range. This helps the operator to spot fluorescence. The laser should be roughly tuned to produce some fluorescence, at which point the oscilloscope trace should show one or more absorption dips. If this is seen, but hyperfine peaks cannot be seen,

[6] Since it is the same atom, it has only one velocity.
[7] Unless absorption spectra have already been observed recently.

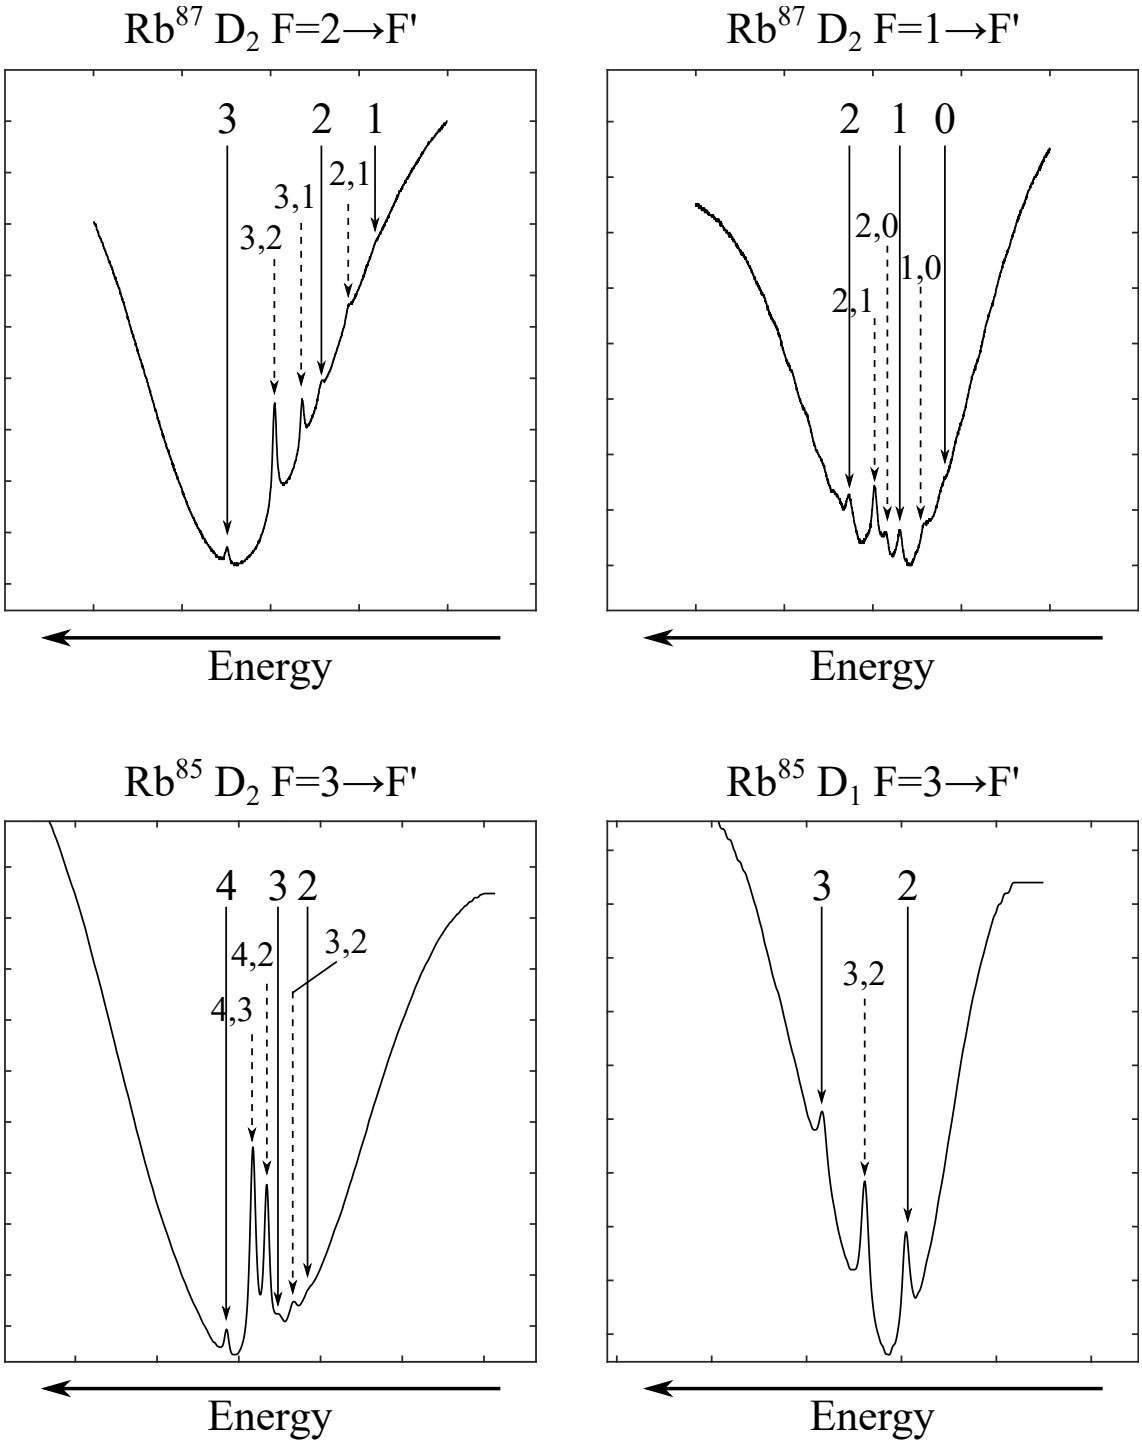

Figure 5.5: Measured oscilloscope traces showing hyperfine transition and crossover peaks within a Doppler-broadened envelope. The top row (^{87}Rb) shows relevant transitions for our cold atom experiments, and the bottom row (^{85}Rb) is relevant for our warm atom work.

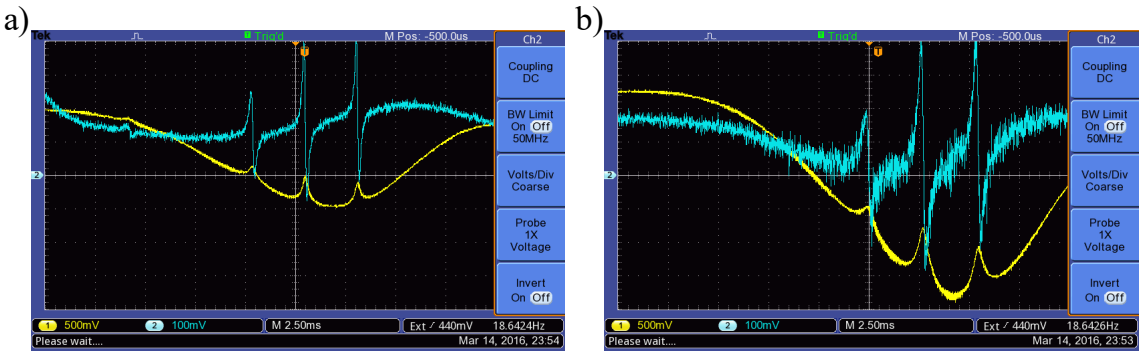

Figure 5.6: Oscilloscope traces showing absorption spectra (yellow) in the case when (a) a laser side-mode or (b) the main mode is resonant with atomic transitions. The two spectra were taken one minute apart, only the laser diode current was changed by about 5 mA.

two things can be done. The laser current can be slowly adjusted in a small range. This changes the diode gain characteristics and changes the spectral composition of its output, so can help if the gain properties were not favourable to lasing on the required frequencies. If this does not help, the retroreflector mirror in figure 5.4 needs to be fine-tuned. To do this, set the oscilloscope to AC coupled mode, and adjust the mirror angle to get maximum peak visibility.

We observe that at high operating currents (130-150 mA) the diodes are often slightly multimode, lasing on usually a few distinct frequency modes that are fairly close together, so our spectrometers have a hard time resolving them. Changing the diode gain characteristics a little can have a dramatic effect on the power distribution in these modes, but the key effect is that it is possible to observe a decent-looking absorption spectrum while most of the output power is in the wrong mode, which makes experimental results rather poor. We observe that in these cases there usually is another absorption spectrum identical in shape but with much better contrast less than a milliamp (or few mA) of diode current away. Slowly adjusting the current, often towards higher currents, results in a sort of 'front' sweeping across the absorption spectrum. On either sides of this 'front' the contrasts are different. An example of two spectra, one on a side-mode and one on the main mode, are shown in figure 5.6. In this specific case the good spectrum was found at about 5 mA higher current, which is an unusually large difference.

5.4 Laser locking

With our ECDLs we clearly have the ability to address single hyperfine transitions of rubidium, since the laser linewidth is narrow enough to resolve them. We can also tune the lasers onto resonance with individual hyperfine transitions, as shown by the absorption spectra in figure 5.5. It is still a challenge to keep the laser frequency tuned to a hyperfine transition, because the laser frequency drifts more than the atomic linewidth on the scale of seconds due to small temperature and diode current fluctuations in the ECDL. However, on a good day the amplitude of these drifts is relatively small, below 100 MHz[8], which is well within the mode-hop-free tuning range of the ECDL, which means that it should be possible to correct the drift by active feedback via the PZT. In the following we describe how we lock our lasers to hyperfine transitions.

The idea is to supply a drift-dependent voltage to the PZT, which would act to compensate the frequency drift. Therefore this so-called error signal should be zero when the laser is resonant with a transition, add a negative voltage if the laser frequency is too low and a positive voltage

[8] And even on bad days, when the air conditioning system is acting up, the drift is not more than a few gigahertz.

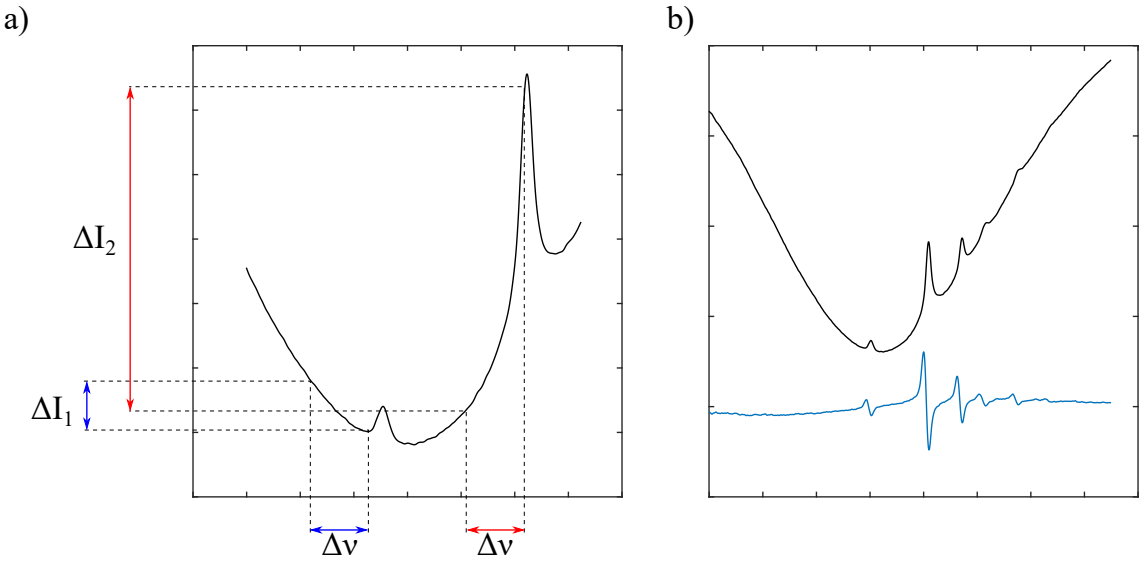

Figure 5.7: (a) Generating a derivative signal by adding a small frequency scan with amplitude $\Delta\nu$ to the laser. Scanning a relatively flat part of the spectrum leads to a small intensity amplitude (ΔI_1), whereas scanning over a steep feature gives a large intensity amplitude (ΔI_2). The scan ranges are vastly exaggerated for illustrative effect. (b) An example absorption spectrum with the corresponding derivative signal calculated numerically using this method.

if it is too high. This is because a positive voltage decreases the laser frequency with our PZT configuration.

The simplest idea would be to use the absorption spectrum directly as the error signal. A DC offset can be added to the absorption signals such that the chosen hyperfine peak crosses zero. Then the laser can be locked to the side of the absorption feature, specifically the right side of the peaks in figure 5.5 since these have the correct slope. There are two problems with this approach. First, the laser frequency is some 3 MHz detuned from the actual transition, because the laser is locked to the side of the absorption feature. This can be a problem if the light frequency is not manipulated further in the experiment, since the laser is not resonant with the atomic transition. Second, by looking at the oscilloscope traces in figure 5.5, we can see that the peaks corresponding to the actual transitions are small, with a relatively small side slope, so the lock stability is expected to be rather weak, especially if the measured intensity at the photodiode varies, which can happen due to air currents, for example.

A much better idea is to obtain the derivative signal of the absorption spectrum. This has the advantages of crossing zero at the centre of the absorption features and large slopes even for small peaks. A derivative signal can be obtained by adding a small amplitude, high frequency voltage signal to the piezo drive. The frequency of this signal should be high compared to the scan frequency, but low compared to the atomic linewidth. In the older ECDLs we use 20.3 kHz, while the newer ones were supplied with 250 kHz. This fast scan essentially probes the amplitude range of the absorption signal. If the gradient of the absorption feature is high, a small scan in laser frequency will result in large changes in the photodiode signal, and for small gradients the photodiode signal becomes flat. So the amplitude of the photodiode signal at the high frequency scan corresponds to the slope, or derivative, of the absorption signal. This concept is illustrated in figure 5.7. The amplitude of the photodiode signal at the scan frequency can be extracted using a lock-in amplifier, which has the added benefit of rejecting detector noise at frequencies other than that the modulation frequency.

We operated the two sets of ECDLs (the older and newer designs) with different electronics. The

current to the laser diodes in the older ECDLs was supplied by Thorlabs laser diode drivers (ITC-102 with ITC-100D). These drivers supplied temperature control as well, except in one case where the temperature control circuit has been burned out. For this laser temperature was stabilised by an ILX Lightwave LDT-5412. The low frequency triangle wave scan signal was supplied to all old ECDLs by a TTI TG1010 programmable function generator[9]. The high frequency scan signal was supplied from the internal oscillator of an EG&G 5208 lock-in amplifier. This lock-in, along with an EG&G 5210 and a Femto LIA-MV-150, also performed the extraction of the derivative signal from the photodiode signal. Two of the older ECDLs were locked to their error signals by home-built integrators, the circuit diagrams of which can be found in [70, 71]. The third laser of that design was locked using a Sacher LB2001 servo controller[10]. The new generation ECDLs were controlled entirely by MOGLabs DLC202s (MOGBoxes). These units provided current control, temperature stabilisation, error signal generation and PZT feedback, which made experimental control a bit easier[11].

In an attempt to abstract away the different electronics for all the different ECDLs, wires to the ECDL components (laser diode, Peltier unit, thermistor and piezoelectric transducer) were bundled together into a convenient serial port connector. For historical reasons the older ECDLs' PZT control was separate to the serial connectors. In the case of the MOGBoxes the output was already bundled into a DVI connector, which we built a converter for. The pinout of the serial connectors is shown in figure 5.8. These common connectors turned out to be a very good idea when the EG&G 5210 lock-in amplifier was destroyed by overcurrent, and one of the home-built integrators failed in 2018 and two of the older ECDLs had to be controlled by the MOGBoxes.

5.5 Frequency control after the laser: acousto-optic modulators

We should mention one more way of controlling light frequency that is completely external to the laser, using acousto-optic modulators (AOMs). These devices contain an acousto-optic crystal, tellurium dioxide or quartz in the ones we use, a piezoelectric transducer on one side of the crystal and an acoustic absorber on the other. The PZT is driven with some radio frequency (RF) signal, typically ranging from a few dozen to a few hundred megahertz. The PZT strikes the crystal, and sets up a standing acoustic wave of changing density and hence refractive index. This acts as a diffraction grating due to Brillouin scattering.

As light is diffracted to the mth diffraction order from this acoustic grating its frequency is also shifted by $m \cdot f$, where f is the PZT driving frequency. This can be thought of as a consequence of a three-wave mixing process between laser photons and acoustic phonons in which the energy and momentum of m phonons are transferred to one photon, which accounts for the deflection as well. We can use this sum-frequency generation to tune the laser frequency. Positive frequency shifted orders are always deflected away from the PZT, and negative ones towards the PZT. The amount of light diffracted into the $m \neq 0$ orders depends on the amplitude of the PZT drive signal, for the same reason as a phase grating efficiency depends on its grating depth. This makes AOMs useful for switching beams on and off fast as well [72], without the need to disturb the laser diode operation or having mechanical shutters which lead to large mechanical vibrations. They can also be used for fine electronic control of laser intensities by controlling the power in the RF

[9] With a broken display, so the settings could not be checked. Fun.
[10] Although to my knowledge nobody at Glasgow ever really understood how that thing worked.
[11] Here's a general rule for experimental science: you can pay either money for complete control solutions, or pay time in developing/building your own.

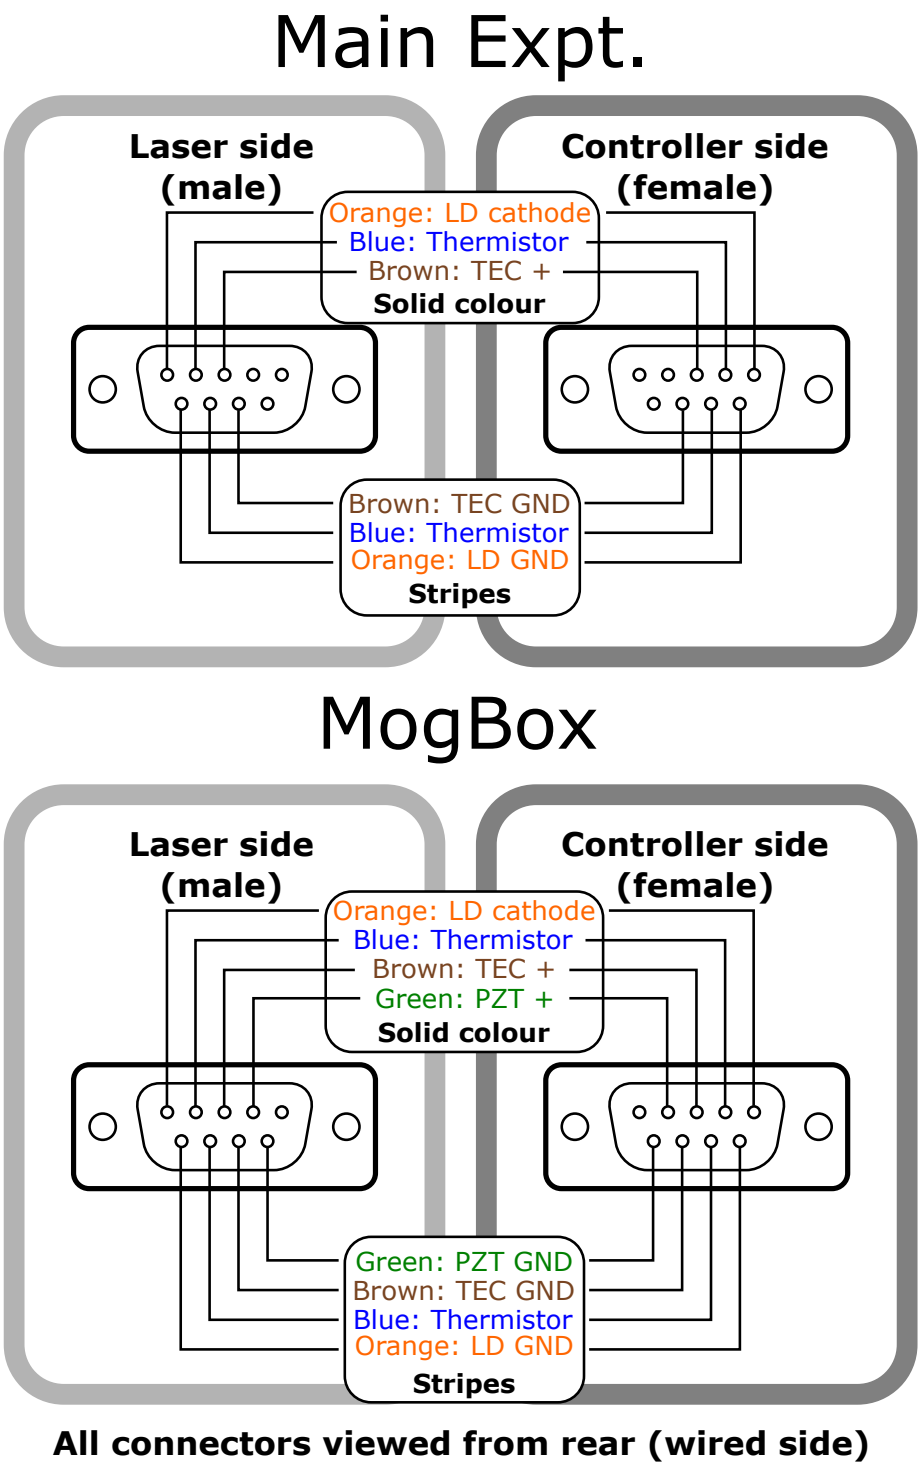

Figure 5.8: Laser driver serial port pinouts.

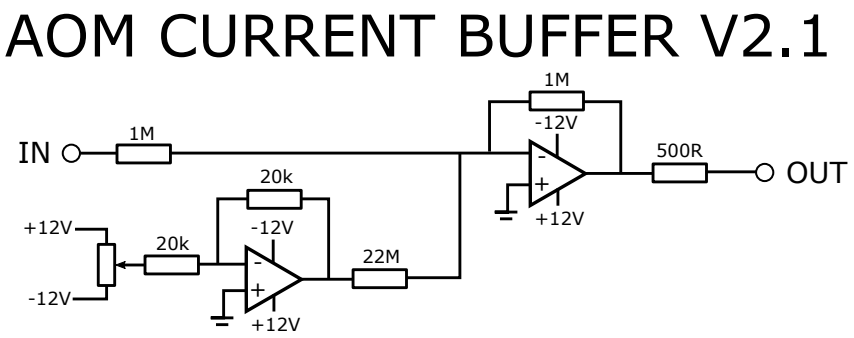

Figure 5.9: Circuit diagram for current buffers to provide non-negligble power TTL signals to mixers. The 500 Ω resistor determines the amount of current (and hence power) this circuit supplies. If this buffer does not allow one to zero the output signal when the TTL is off, the 22 MΩ resistor can be exchanged for a lower resistance one to give a larger range of voltages to be added to the TTL.

drive signal. Typically the maximum diffraction efficiency of AOMs in the first order is 80% at their central frequency, but diminishes when driving the PZT with different frequencies. The exact curves depend on the AOM design, their manuals should contain this information.

Typically, the AOM is supplied with an RF signal generated by a voltage controlled oscillator (VCO) passed through an amplifier. The frequency of the signal can be controlled by a[12] control voltage that can be supplied by simple voltage divider circuits using a potentiometer, or by analogue outputs of computer controlled input/output (I/O) cards. Additionally, the amplitude of the signal, and hence the diffraction efficiency of the AOM, can be controlled in multiple ways. A common strategy is to pass the VCO signal through a voltage variable attenuator (VVA), which can be controlled the same way as a VCO. We have used VVAs in such a way before, but we found them to behave strangely in some cases, introducing delays on the order of milliseconds to the AOM response (which is normally on the scale of fractions of microseconds) and changing the polarisation of the light leaving the AOM. We could not explain these behaviours, but changing to a different strategy for amplitude control resolved the issues.

We can turn the AOM grating on and off by mixing the VCO signal with a TTL signal from an I/O card. This allows us to switch the diffracted beams on or off with precise timing, although to have fine, non-binary, amplitude control we would need an analogue voltage signal from the I/O card, which is also possible. Unfortunately the I/O cards output negligible current, so there is no power in the TTL signal, which means that the mixers output no signal, irrespective of the state of the TTL voltage. For this reason we built a current buffer for each TTL channel. The circuit diagram for these current buffers is shown in figure 5.9. The purpose of these was to supply the same voltage as their input (the TTL signal), but provide a non-negligible current. They also allow the fine-tuning of the zero voltage of their output by adding a small, positive or negative, voltage to the input TTL, allowing us to make sure that when we want to turn the AOMs off we can actually supply zero voltage to the mixers.

Of course changing the AOM frequency changes the laser deflection angle, which makes AOMs difficult to use in experiments that require a range of frequencies and stable beam alignments. A solution to this is to use them in a double pass configuration, shown in figure 5.10. The laser beam is directed through a polarising beam splitter and focused[13] onto the AOM aperture. The

[12] You guessed it...

[13] There is an alternative approach, which can produce a much smaller footprint in an experiment. The interested reader is directed to [73] for more information.

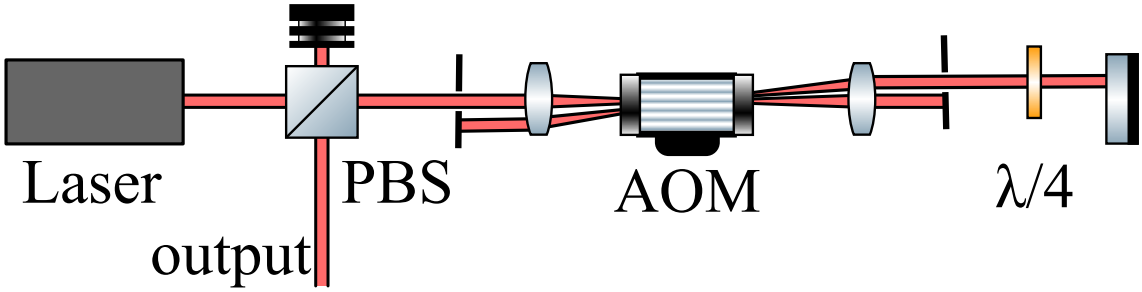

Figure 5.10: AOM in double-pass configuration. The first order of the AOM is reflected straight back, and the laser frequency is shifted twice (once on passing through in either direction). The two passes through the $\lambda/4$ plate rotate the polarisation by π, so the output is reflected from the polarising beam splitter rather than transmitted. Due to the geometry of the two diffractions the beam output direction is independent of the frequency shift.

AOM produces diffraction orders which are collimated by a second lens. The diffracted orders are parallel after the lens, but spatially separated, so we can place an aperture around the first order, blocking the others. This first order is reflected straight back by a mirror, and passes through a $\lambda/4$ plate twice. This rotates its polarisation to the orthogonal polarisation. The retroreflected beam enters the AOM at an angle, and is diffracted again. Because the deflection direction of a diffraction order depends on the frequency shift applied to that order, the first order on the way back is deflected to the same path as the original input. We can place an aperture around this order as well after the input side lens collimates the beams. The orthogonally polarised, doubly frequency shifted beam leaves the polarising beam splitter at the other output instead of being directed back towards the laser. The two deflections due to the two frequency shifts cancel out, and so the output direction of such a setup is independent of the frequency shift.

Aligning AOMs in such a setup can be deceptively difficult. After the PBS the first focusing lens should be inserted, centered on the beam such that it does not change the beam propagation direction. The AOM should be turned on and placed at the focus of the first lens. The task at this point is to align it for maximum diffraction efficiency by translating it in the plane transverse to the beam propagation, and rotating it around two axes (horizontal and vertical) while monitoring the output initially on a screen (piece of paper or detector card). Because of the many degrees of freedom the initial placement can be a bit tricky. To make this easier, the AOM should be kept horizontal. Then aligning the laser with the AOM aperture should not be too difficult, which leaves rotation around the vertical axis as a free parameter. This should be adjusted, along with small adjustments to the transverse position to keep the beam centred on the AOM aperture, until multiple orders are observed on the screen. At this point all degrees of freedom should be slowly optimised while monitoring power in the first order[14] using a power meter. A particularly good indicator of the quality of the initial placement is the AOM tilt in the $z-y$ plane. If it needs to be rotated more than a degree to find maximum diffraction efficiency the height of the AOM is likely to be wrong, and should be adjusted instead of the tilt.

Once the AOM is placed the second lens should be inserted, ideally such that it is centered on the zeroth order, but if this configuration would place the first order too close to the edge of the lens it is more important to keep this order. An aperture should be placed around the first order after the lens, and the waveplate and mirror should be inserted. The quarter-wave plate should be rotated to 45° with respect to the laser polarisation. The mirror should then be aligned such that the beam is propagating along the same path back through the AOM. The second lens position should be optimised to minimise the deflection of the output beam with changing AOM frequency.

[14] Or negative first order, if a negative frequency shift is desired.

The power in the output of the PBS should be monitored and optimised last by small changes to the mirror alignment and, to a smaller extent, the waveplate angle. The AOM angle can also be adjusted slightly to optimise for double-pass efficiency over single-pass efficiency.

5.6 In summary

In this chapter we explained the operation principle of external cavity diode lasers, described the two ECDL designs we used in our experiments and looked in detail at the procedure for setting them up. We also examined Doppler-free spectroscopy as a method for obtaining reference frequencies for atomic transitions to tune the ECDLs to. We described ways of using the absorption spectrum of a reference vapour to provide an error signal for laser locking, either by using the absorption spectrum itself, or its derivative. We briefly described how a derivative signal can be obtained. We also described how to use AOMs to switch lasers on and off and shift their frequency without the need to change the ECDL lock. In the next chapter we will combine our understanding of light shaping from Part 1, atomic structure from Chapter 4 and laser control from this chapter to generate and measure three-dimensional population structures in a room temperature rubidium vapour.

Chapter 6

Experiment: 3D population patterns

6.1 Introduction

The idea of this experiment came about to follow on from a previous work by Radwell, *et al.* [74]. In that paper the 3D structure of a shaped light field was reconstructed using fluorescence from room temperature rubidium vapour. This worked because at intensities below the saturation intensity ($I/I_{sat} << 1$) the fluorescence rate is linear with intensity, i.e. equation 4.18 becomes, on resonance,

$$R_f = \frac{\Gamma}{2} I/I_{sat}. \tag{6.1}$$

Then, in the homogeneous vapour in a reference cell the fluorescence from each point is proportional to the laser intensity at that point. This assumes that the atoms act as passive scatterers, which is a reasonable assumption in the case when there is only one low intensity beam present.

Here we show that simultaneously, the internal state of the atoms can change due to the interaction with the laser. Specifically, in the electronic level structure of rubidium there are two hyperfine S ground states (see figure 4.1) separated far enough in energy that it is possible to drive atoms from one state to the other via excitation to a P state from which spontaneous emission to both ground states is possible. In such a situation fluorescence from the driving laser, exciting atoms from the upper ground state, is not proportional to its local intensity, because eventually the atom will decay to a lower energy ground state which is not accessible for the laser and fluorescence stops. At the same time, in a thermal vapour the upper ground state is repopulated by atoms entering from outside the laser beam, and the volume outside the laser acts as a reservoir of atoms distributed between the ground states. Fluorescence from a separate, spatially uniform laser driving a cycling transition from the same ground state as the first laser can be used to interrogate the internal state of the atoms. In this chapter we investigate the atomic physics in this two-laser system, and describe the tomographic technique we used for reconstructing 3D atomic state structures. This work has been recently published [75].

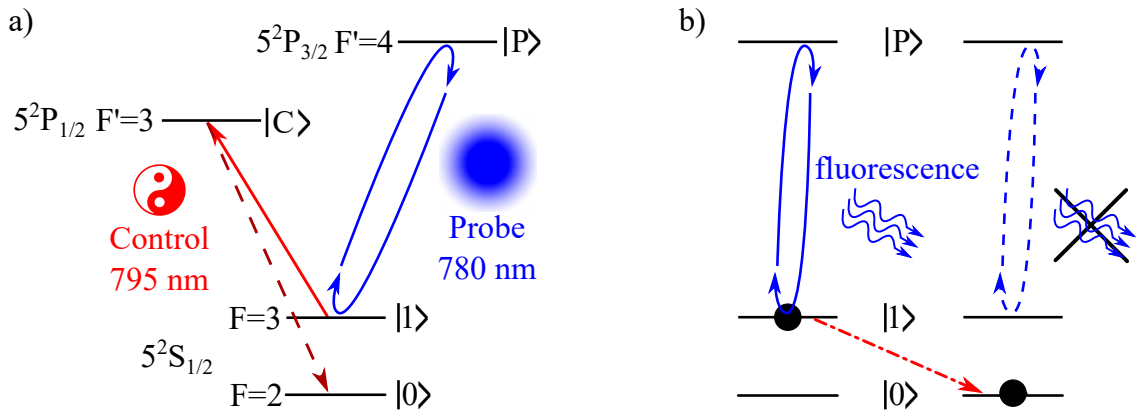

Figure 6.1: a) Simplified energy level scheme of ^{85}Rb showing spectroscopic notation. D_1 and D_2 transitions are shown in red and blue, respectively. b) Interpretation in terms of electron shelving. Detection of fluorescence of the probe light indicates that an atom was in $|1\rangle$, and excitation by the control beam and subsequent spontaneous decay transfers atoms from $|1\rangle$ to $|0\rangle$.

6.2 Atomic states in motion: structured light in warm atoms

We will study the interaction of room-temperature rubidium with two laser beams addressing two transitions in a configuration reminiscent to electron shelving. A ubiquitous technique in ion trapping, electron shelving is used to probe the quantum state of an atom [76, 77]. Typically two lasers are used, one for the control of the state and one for detection. The strong control laser is resonant with a transition of the atom between a stable ground state and a metastable (long-lived) excited state. The probe laser, usually much weaker, is resonant with a transition to a short-lived excited state. When the control laser is on, the atom produces fluorescence from the probe beam, indicating that it is in the ground state. Shining the control beam onto the particle induces Rabi oscillation between the metastable and ground states, so there is no interaction with the probe and fluorescence stops.

Our analogues to the control and probe lasers drive the D_1 and D_2 transitions of rubidium-85 at 795 nm and 780 nm respectively as shown in figure 6.1, coupling to the excited states $5sP_{1/2}$ F'=3 (denoted as $|C\rangle$) and $5sP_{3/2}$ F'=4 (denoted as $|P\rangle$). These two transitions share the ground state $5sS_{1/2}$ F=3, which we refer to as $|1\rangle$. The control beam is shaped by an SLM and the probe beam retains a simple Gaussian profile, expanded to a waist that is larger than the control beam. Selection rules permit $|C\rangle$ to decay to either of the ground states $|1\rangle$ and $|0\rangle$ ($5sS_{1/2}$ F=2), with branching ratio 2:1. This means that where the control laser is bright atoms quickly decay into $|0\rangle$, which is a dark state of both lasers. As a result fluorescence from the probe ceases. Conversely, in volumes of zero intensity in the control beam the probe beam produces fluorescence. By monitoring the probe fluorescence in 3D we can, in principle, characterise darkness structures within the control beam. Crucially, this is done at a wavelength different from that of the control beam, which enables the use of this technique in cases where light from the control laser is difficult to detect.

Inside a room temperature vapour things are slightly more complicated since rubidium atoms move at $\sim$170 ms^{-1} under such conditions. They collide with the walls of the vapour cell, which randomly resets their ground state before they reenter the interaction regions. This process leads to an effective repopulation of $|1\rangle$, so the probe produces some fluorescence even in regions of high intensity volumes in the control beam. Before we get into the details we would like to thank our internship student Sylvain Fayard who derived and calculated equation 6.3 during his stay in Glasgow during the summer of 2016.

Initially, in the absence of light, atoms are distributed between states $|0\rangle$ and $|1\rangle$ with a ratio of 5:7. This ratio comes from the fact that atoms in $|0\rangle$ can be in any one of the 5 Zeeman sublevels, and similarly in one of 7 for $|1\rangle$, and the energy difference between $|0\rangle$ and $|1\rangle$ is too small to lead to a noticable difference in thermal population. As we said before, the control beam rapidly drives atoms into the dark state $|0\rangle$, in regions of high intensity in the control beam. Of course where the control beam has no intensity atoms remain in $|1\rangle$. In the presence of the probe beam the populations in the various states could be calculated using the optical Bloch equations, in a spatially dependent fashion, but this is rather complicated. It turns out that for our purposes here, however, it is sufficient to consider rate equations, because we assume that coherences do not matter much in this system. This is a reasonable assumption because at room temperatures decoherence times are very short.

The rates of change of populations can be written as

$$\begin{aligned}
\dot{P_0} &= \Gamma_{C0} P_C, \\
\dot{P_1} &= -(R_{1C} + R_{1P}) P_1 + \Gamma_{C1} P_C + \Gamma_P P_P, \\
\dot{P_C} &= R_{1C} P_1 - (\Gamma_{C0} + \Gamma_{C1}) P_C, \\
\dot{P_P} &= R_{1P} P_1 - \Gamma_P P_P, \\
1 &= P_0 + P_1 + P_C + P_P,
\end{aligned} \tag{6.2}$$

where R_{ij} is the transition rate from state i to j, caused by excitation due to a laser and $\Gamma_i(j)$ refers to spontaneous emission from state i (to state j).

We can make a number of simplifying assumptions here. First, the lifetimes of the two excited states are short compared to the lifetimes of the ground states, so they are essentially never populated. This means that we can forget about P_P completely. P_C is still relevant, since atoms are transferred to $|0\rangle$ through $|C\rangle$. Because of the assumption that $|C\rangle$ is populated for negligible times, which comes from the short lifetime and hence fast decay of $|C\rangle$ ($\Gamma_C >> R_C$, which is true for low control intensities), we can re-interpret the system as a two-level system, where the excitation rate from $|1\rangle$ to $|0\rangle$ incorporates the excitation rate from $|1\rangle$ to $|C\rangle$ and from $|C\rangle$ to $|0\rangle$.

In this case we have a system that is essentially identical to the populations in equation 4.17. Then we can find the spatially dependent depletion rate of $|1\rangle$ as $\frac{d}{dt}P_0 = -\frac{d}{dt}P_1 = R_d P_1$. The depletion rate is essentially given by the fluorescence rate on the control transition, modified because of the fact that from the excited state $|C\rangle$ the atom can decay back to $|1\rangle$ as well as $|0\rangle$. We have already derived the fluorescence rate in equation 4.18. Here we need to find the detuning Δ in terms of Doppler shifts due to the atoms in motion. This can be done by integrating the Maxwell-Boltzmann velocity distribution in 3D. Then R_d is given by

$$R_d(\mathbf{r}) = \frac{1}{3} \frac{\Gamma\lambda}{4\sqrt{3\pi k_b T / m_{Rb85}}} \frac{\Gamma}{2} \frac{I(\mathbf{r})/I_S}{\sqrt{1 + I(\mathbf{r})/I_S}}, \tag{6.3}$$

where $I(\mathbf{r})$ is the intensity of the control beam at a position $\mathbf{r}$, $T \simeq 293$ K the temperature in the cell, m_{Rb85} is the mass of a rubidium 85 atom, $\lambda = 795$ nm the wavelength, $I_S = 4.49$ mW/cm^2 the saturation intensity and the decay rate is $\Gamma = 2\pi \times 5.75$ MHz [78]. The leading factor of 1/3 comes from the branching ratio and relative pumping rate between the probe and control beams, and represents the probability of the atom decaying to $|0\rangle$ rather than back into $|1\rangle$.

As we said before, without the two laser beams, the atoms within the cell are distributed between the hyperfine ground states according to their degeneracy, $P_0 = 5/12$ and $P_1 = 7/12$. This

distribution is maintained through collisions with the cell wall, which resets their state between the ground states randomly. The probe beam drives a transition to an excited state that cannot decay to any other state than the one the same beam excites, so this beam cannot change the populations between $|0\rangle$ and $|1\rangle$.

In the presence of resonant light, the local population distribution is determined by two competing processes. Consider an observation region near the middle of the control beam. The $|1\rangle$ population is depleted at a rate R_d given by equation 6.3, which depends on the local control beam intensity. At the same time it is repopulated by atoms that have collided with cell walls and drifted back into the observation region. However, atoms in $|1\rangle$ may not reach the observation region, because they may have had to pass through high intensity regions of the control beam and so may be depleted. This depends on the exact spatial profile of the control beam and the path atoms take through the cell. The vapour density in unheated vapour cells is low (on the order of 10^{-8} Torr), which means that the mean free path between atom-atom collisions is on the order of kilometres so we can assume that atoms travel along straight lines.

It is also sufficient to consider atoms travelling in a plane transverse to the beam propagation. This is because we use laser light resonant with stationary atoms because we lock them to Doppler-free spectroscopy features. Therefore the atoms which have a component of their velocity along the beam propagation direction see the laser Doppler-shifted out of resonance. Both the control and probe beams need to address the same velocity class of atoms in order for us to be able to measure P_1. This means that the combined detuning of the two lasers needs to be within a narrow frequency band which we measured to be approximately 40 MHz wide, corresponding roughly to the width of the convolution of the probe and pump transition line profiles.

Considering a 1D slice in a transverse plane these arguments can be combined to show that the variation of $|1\rangle$ population of atoms passing through the cell can be written as

$$\frac{\mathrm{d}P_1}{\mathrm{d}r} = \frac{\mathrm{d}P_1}{\mathrm{d}t} / \frac{\mathrm{d}r}{\mathrm{d}t} = \frac{-P_1 R_d}{v_r}, \tag{6.4}$$

where $v_r = \sqrt{2k_bT/m_{Rb85}}$ is the most probable transverse speed of atoms. Technically we could obtain a more accurate expression by integrating the Maxwell-Boltzmann distribution in 2D, but we find that using this value for v_r works very well. Integrating this equation, considering all the directions an atom can travel along in 2D we find that the $|1\rangle$ population at a position (x, y) in a transverse plane can be expressed as

$$\begin{aligned} P_1(x,y) &\propto \frac{7}{12} \int_0^{2\pi} \mathrm{d}\theta P_{1,\theta}(\theta), \quad \text{with} \\ P_{1,\theta}(\theta) &= \exp\left(\int_0^{wall} \mathrm{d}r \frac{-R_d(r\cos\theta + x, r\sin\theta + y)}{v_r} \right), \end{aligned} \tag{6.5}$$

where θ and r are polar coordinates with the origin at (x, y) and the prefactor of 7/12 is the initial value of P_1, as described above.

The population in $|1\rangle$ is detected from the fluorescence of the probe beam, so it should have a uniform intensity throughout the cell, or at least in any transverse plane. In practice it is hard to make a flat-top beam that propagates a reasonable distance, so we use a beam with a truncated Gaussian profile with a width exceeding the control beam diameter. This is obtained by passing a wide Gaussian through a variable aperture. The probe fluorescence is related to the population

in $|1\rangle$ through a rate R_p identical in form to R_d, of course with the appropriate values for the saturation intensity $I_{sat} = 3.90$ mW/cm^2, decay rate $\Gamma = 2\pi \times 6.07$ MHz and the prefactor (equal to 1) due to the branching ratio for this cycling transition. Then the probe fluorescence rate is proportional to $P_1 R_p$. For a spatially uniform R_p the populations can be inferred directly from the fluorescence, but even for a non-uniform R_p we can recover the population by dividing by the spatial structure of the probe beam.

We note that this discussion has implications for techniques that rely on measuring the absorption of laser light due to thermal atoms, such as Doppler-free spectroscopy discussed in chapter 6. As long as the laser beams used in those techniques have a small[1] cross-section and intensity, the repopulation due to new atoms can be considered homogeneous throughout the beam and leads to a systematic increase in absorption. However, for extended beams or very high intensities atoms in the centre of the beam will be depleted while atoms in the outer regions of the beams still experience repopulation. The result is that one observes different levels of absorption when measuring different parts of the beam.

In subsequent sections we will experimentally generate light structures in the control beam and measure the populations these structures imprint on the atoms. We will first look at the experimental methods used.

6.2.1 Tomographic reconstruction

Tomography is the reconstruction of 3D structures from 2D images. We reconstruct the spatial structure of the fluorescence using the tomographic technique described in [74]. The fluorescence in the cell is imaged from the side, giving a projection of the intensity onto the yz plane. While a single image of the fluorescence contains information about only a two-dimensional projection of the populations, we can also obtain three-dimensional information.

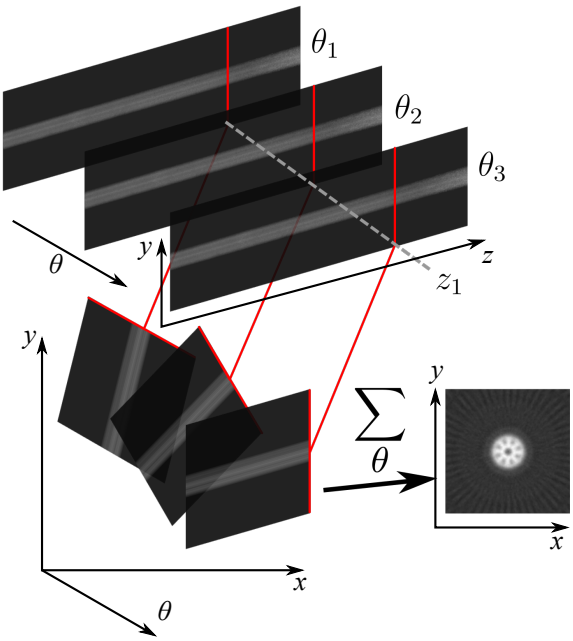

Figure 6.2: Illustration of the tomographic reconstruction process for 3 rotation angles. A column of pixels is taken from a series of 3 fluorescence images recorded at different rotation angles of the control beam. These pixel columns are stretched into 2D, the resulting squares are rotated to the appropriate angle and they are summed to produce a reconstructed cross-section.

In simple terms, tomographic reconstruction is the process of obtaining two-dimensional information from a series of one-dimensional data[2]. Our series of 1D data are columns of pixels from the images taken from the side of the cell, corresponding to a particular distance along the propagation axis z. The multiple images are recorded, each after the control beam is rotated by an angle θ around its propagation axis z. This is equivalent to using cameras viewing the cell from different angles, where the spatial offset between multiple cameras is translated to a temporal offset between multiple rotation angles of the beam. Each recorded image correspond to a different projection of the 3D structure of

[1] What constitutes 'small' depends on the temperature and size of vapour cell.

[2] Although mathematically the same ideas can be used to generate N-dimensional data from a series of (N-1)-dimensional data.

the fluorescence onto the detector plane. We need to undo this projection. We use a process called filtered back projection, based on an inverse Radon transform [79, 80]. The Radon transform is the gives the projection of a 2D structure onto a 1D line by integrating along the direction transverse to the 1D line, and its inverse allows us to undo such a projection. The process is illustrated in figure 6.2.

Each n pixel long column from the images is Fourier transformed, and a ramp filter is applied to the frequency space data in order to compensate for the oversampling of regions closer to the rotation axis. The inverse Fourier transform is then applied, and the n pixel long 1D column is 'stretched' to an $n \times n$ square in 2D. These squares are rotated by the angle that the control beam was rotated at in the given image. The rotated squares are summed together to produce cross-sections at the corresponding z-coordinate. The full 3D structure within the cell is built up from all the cross-sections.

6.3 Experimental setup

The experimental setup is shown in figure 6.3. In this experiment two new generation external cavity diode lasers were used. To achieve the required frequency for the control laser the diode was heated to close to 48°C. This required a small modification to one of the MOGBoxes, a variable resistor needed to be changed to allow setting the temperature controller target so high. The probe laser was locked to the D_2 F=3 F'=2,3 crossover and tuned to the F=3 F'=4 transition using an acousto-optic modulator (AOM) for fine frequency and power control. Both lasers were linearly polarised, which was ensured by Faraday isolators placed immediately after the ECDL exit port.

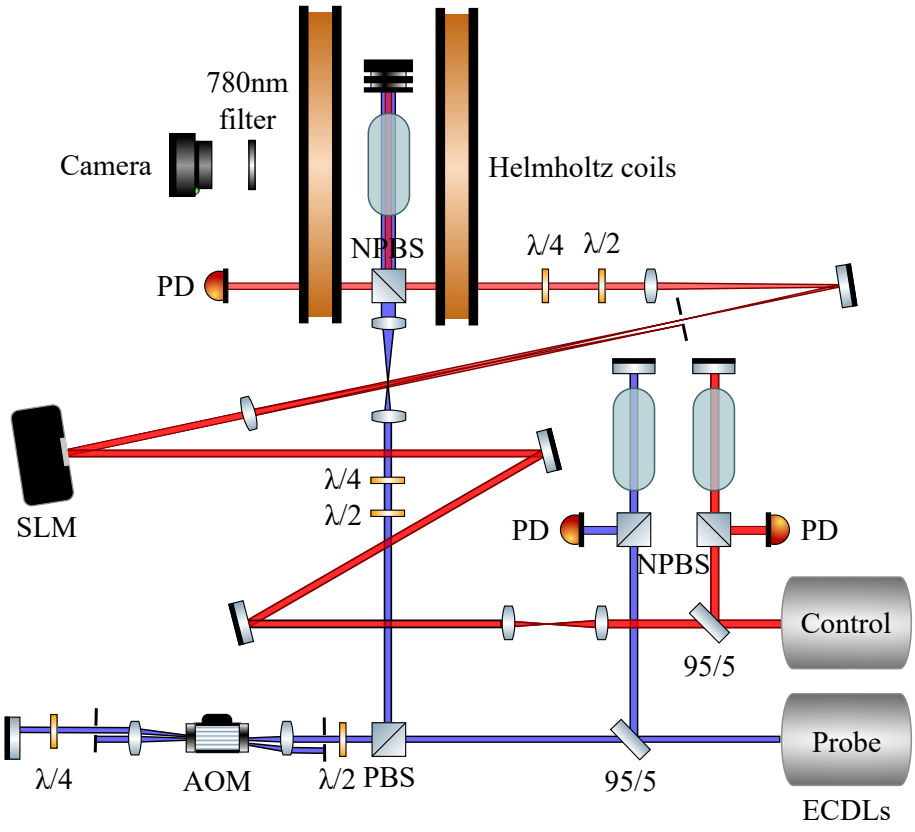

Figure 6.3: Experimental setup as described in the main text.

The control beam was cleaned through a pinhole and was then directed onto an SLM (Hamamatsu

LCOS) where it interacted with an amplitude modulated holographic grating to acquire a chosen shape, for details see chapter 2. All orders except from the first order of the grating were filtered out by an aperture placed in the Fourier plane of the SLM. After collimation the control beam passed through a $\lambda/2$ and a $\lambda/4$ plate for polarisation control, and it entered a non-polarising beamsplitter used to combine the two beams.

The probe beam was sent through the AOM in double-pass configuration. The AOM frequency was kept at 76.15 MHz to offset the laser frequency from the crossover to the F=3 F'=4 transition (the frequency difference between the crossover and the transition is $2 \times 76.15 = 152.3$ MHz). This beam was sent through a $\lambda/2$ and a $\lambda/4$ plate, it was magnified such that its cross-section was larger than that of the shaped beam, and then the two beams were combined.

After the beamsplitter the now co-propagating beams passed through a rubidium vapour cell containing both ^{85}Rb and ^{87}Rb in their natural isotopic abundance of 72:28, to produce fluorescence. The fluorescence from the cell was imaged onto a camera using a camera-mounted imaging lens (16 mm focal length), chosen for its decent depth of field. The lens was adjusted such that the centre of the beams were in focus. A narrow-band filter centered on 780 nm (bandwidth ∼10 nm) was inserted before the imaging lens so that only fluorescence from the probe beam was recorded; example images are shown in the next section.

6.3.1 Experiment control software

The software used to run the experiment was written in LabVIEW, progressing through many iterations. The final version[3], used to take all data below, was constructed for maximum memory efficiency, since during the experiments large amounts of image data was recorded, which had a tendency to fill the available random-access memory. The software had to perform four major functions: designing of light structures for the control beam and calculating holograms for the SLM; displaying the holograms, while controlling the frequency and intensity of the probe laser via an AOM; recording fluorescence images while the control beam is rotating; and finally reconstructing fluorescence cross-sections.

The holograms were constructed using the techniques in chapter 2. The intensity and phase structures could be generated in a very flexible manner in the software. There were options for generating superpositions of Laguerre-Gaussian or Hermite-Gaussian modes, beams with flat phases, only their amplitude shaped, or beams with completely arbitrary phase and amplitude structures defined as PNG files, generated in other software. The experiment control software took care of amplitude and phase correction, although the corresponding profiles of the input beam were measured with other software written specifically for that purpose.

For each beam several holograms were generated, corresponding to different rotation angles of the control beam. The software was designed to take into account rotational symmetries of the desired beams. For example, a HG_0^1 has a 2-fold symmetry under rotation by 360°, in other words, rotation by 180° returns to the beam at 0° (up to an unimportant global phase). Therefore it is sufficient to record data for half the full 360° rotation, the second half would contain no new information, since the holograms would be identical. We can exploit this to either perform experiments faster (using less rotation angles) or get higher spatial frequencies in the reconstruction (by using a large number of rotation angles in a smaller overall rotation). These were relevant concerns because during most of the measurements presented here laser locking stability was rather low, with ECDLs coming

[3]Access at `http://dx.doi.org/10.5525/gla.researchdata.630`

off lock over periods on the scale of 30 minutes, which limited the available time for recording fluorescence images from the same control beam under the same conditions.

Each hologram was saved to hard disk to save memory space. During the experiment they were individually loaded by the control software and displayed on the SLM. Because in this experiment we had no access to the SLM refresh timings we waited 100 ms after the display request to give ample time to the SLM to display the required hologram. We have seen in tests that waiting less would occasionally result in the previous hologram still being displayed. Then the camera exposure would begin, and on completion of the exposure the fluorescence image was saved to hard disk.

For the reconstruction the images were loaded one by one, and only the columns from a chosen subset were kept in memory. A bandpass filter was applied to the columns in order to remove high frequency detector noise. In practice we used quite aggressive smoothing (removed the top 20% of frequency components) because the beams we used had reasonably low spatial frequencies. Then the ramp filter was applied to the smoothed pixel columns, the 2D squares were calculated and they were summed together to produce fluorescence cross-sections.

It was important that the centre of rotation of the beams be identified accurately for the reconstruction to work. A small displacement (less than 5 pixels) of the reconstruction axis resulted in completely wrong results. We could identify the rotation axis by generating a known pattern, measuring the fluorescence of the control beam alone and comparing the reconstruction with the desired pattern. Great care was taken to make sure that both the control and probe beams propagate horizontally and coaxially, parallel to the camera horizontal, so that the rotation axis would not change throughout the camera images. This was done by monitoring the fluorescence from either beam in real time, and fitting Gaussians to two columns of pixels, separated significantly. We could then match the centres of the Gaussian fits by careful alignment of the beams using the mirrors and beamsplitters. This has worked very well, giving us an estimated precision of ± 0.5 pixels when determining the beam axes.

6.4 Results

First we attempted to determine the ideal combination of beam powers (hence pumping rates) that produces the best reconstructed images, with the aim of maximising contrast. In order to quantify the quality of data we introduced different measures for the contrast. First, let us describe the measurements available to us. We first propagated only the Gaussian-shaped probe beam through the cell, blocking the control beam, and recorded the fluorescence at some value of the propagation distance z. This cross-section serves as a reference to which we can compare the following measurements. We then unblocked the control beam, gave it a simple Gaussian profile and recorded the fluorescence. Finally we 'threw away' a part of the shaped beam, making a hole in it, and recorded the fluorescence. This cycle was repeated for different combinations of powers in both the probe and shaped beam controlled by the AOM and the grating depth, respectively. We performed this 2D power scan for a number of different polarisations in the beams, and at different magnetic fields and directions. Sample cross-sections of these fluorescence measurements are shown in figure 6.4a.

We then calculate the difference between the first and second fluorescence profiles, shown in figure 6.4b. Here we have a number of choices that we can make in order to find a single value for the contrast. We can either fit the difference profile with a Gaussian and find its peak, or we can integrate it. Whichever method we choose we call the result 'bright contrast', and it measures how

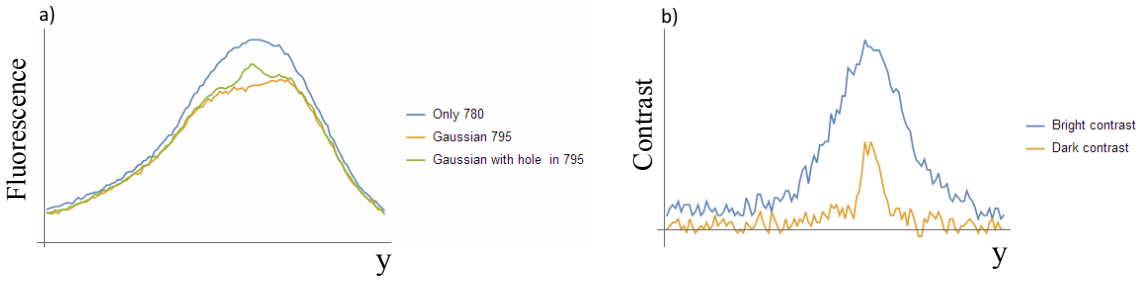

Figure 6.4: a) Sample cross-sections of fluorescence measurements. Blue: fluorescence from the probe beam only. Yellow: fluorescence from the probe beam with the Gaussian-shaped beam turned on. Green: same as yellow, with a imprinting a hole in the shaped beam. b) Sample contrast profiles. Blue: bright contrast, yellow: dark contrast.

well the depletion process works.

Of more interest is another quantity we call 'dark contrast', which measures how well we can differentiate between the shaped beam having a hole or not based on the reconstruction. This is calculated by first taking the difference between the first and the third profiles, and subtracting this from the bright contrast profile. Both the bright and dark contrast profiles, calculated from the data shown in figure 6.4a, are shown in figure 6.4b. Here we again have the choice of finding the peak of the profile or integrating. The result of both of these is shown in figure 6.5.

From figure 6.5 it can be seen that in general higher power in the shaped beam gives better contrast. It can be seen that contrast ramps up for low probe powers and then flattens, however this is not related to the atom-light interactions. It arises simply because the exposure time for the camera was too short for low probe powers and the profiles were very small. Once there was enough light so that the CCD would have saturated the exposure was lowered to prevent this, so all profiles had roughly the same peak values and the contrast flattens. The length of the initial ramp can be changed significantly by choosing different initial settings for the exposure. These measurements suggest that overall high powers in both beams are desirable for a good contrast.

Also of note are a set of dispersive features seen in the middle of figure 6.5b (and to a lesser extent, figure 6.5a). These features are present for the same powers irrespective of relative beam polarisations (not shown in figures), but shift with changing magnetic fields. This can be seen from figure 6.5d, taken at a different magnetic field, where the features shifted to higher powers in both beams. We believe that these features are caused by the difference in saturation intensities for the two transitions, but the exact mechanism is not understood and was not investigated further.

In order to show the complementarity of the control and probe fluorescences we generated a control beam with a Yin-yang symbol cross-section show in figure 6.6a. We have used our reconstruction technique to produce the cross-sections in 6.6b and d. As can be seen in the figure, control and probe fluorescence cross-sections are indeed visually complementary to each other, with some notable differences. The the probe cross-section has less sharp edges. From figure 6.6 it also seems that the technique is less accurate at reproducing population patterns due to dark features surrounded by brightness than bright regions surrounded by darkness, leading to a loss of contrast. This contrast loss can be explained by the fact that atoms need to cross high intensity regions in the control beam to get to dark holes.

We investigated the variation of contrast measured at a dark hole in a beam as a function of the amount of surrounding control light power. We numerically simulated the populations in a beam

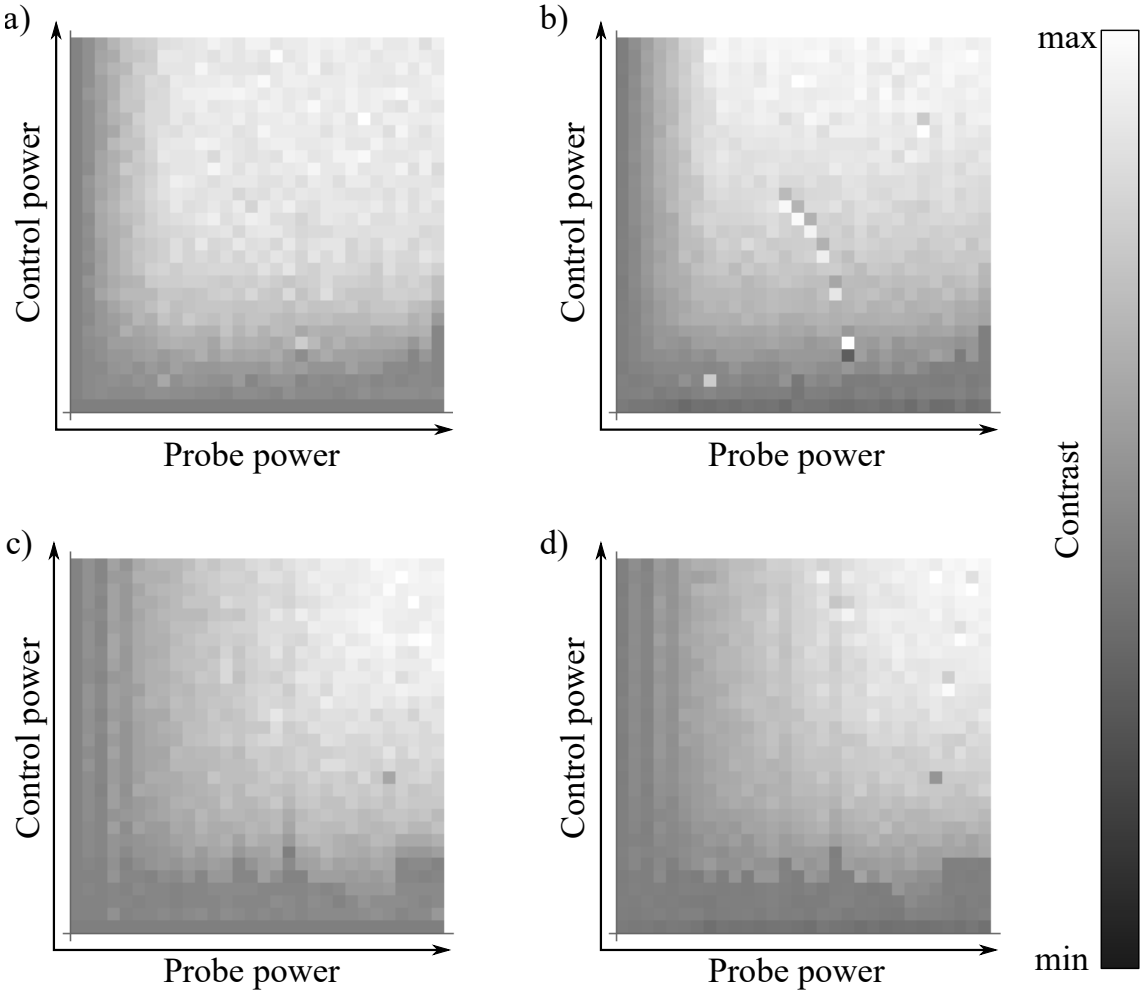

Figure 6.5: Dark contrast measurement results. Left: contrast measured in terms of fluorescence peak, right: contrast measured in terms of integrated fluorescence. Top: no applied magnetic field, bottom: magnetic field of approximately 0.75 G aligned with beam propagation direction

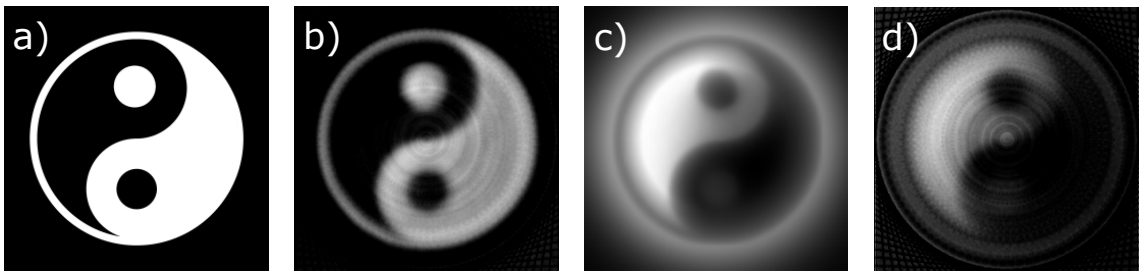

Figure 6.6: Simulated and reconstructed cross-sections from a control beam with a yin-yang symbol profile, measured in the image plane of the SLM. a) Desired control intensity profile. b) Light reconstruction from fluorescence of the control beam. c) Simulation of the reconstruction of probe fluorescence indicating population in $|1\rangle$. d) Measured probe fluorescence reconstruction.

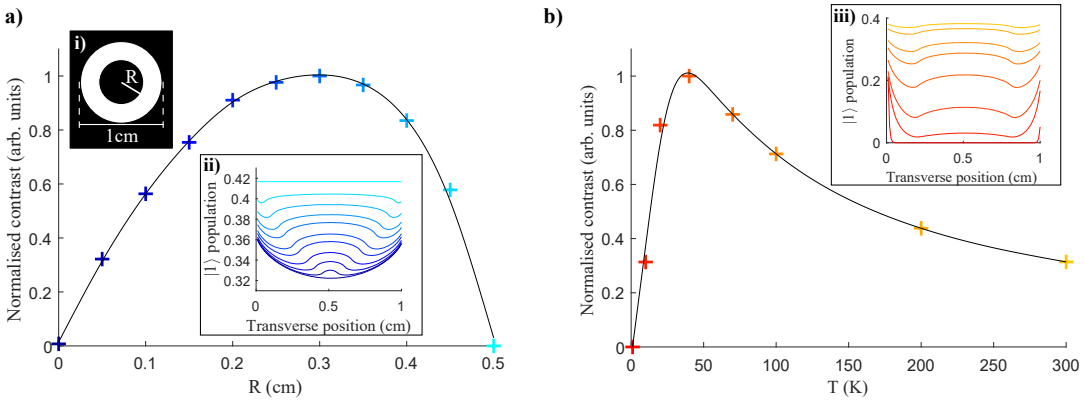

Figure 6.7: Normalised population contrast as a function of varying a) dark core radii R in a flat-top control beam and b) temperature. In a) radii R are ranging from 0 cm (shown in dark blue) to the full beam radius of 0.5 cm (cyan) at a temperature of 300 K. In b) the temperature is ranging from 1 K (shown in red) to 300 K (yellow), for R=0.3. Fits show as black lines are smoothed spline interpolations calculated from 15 data points in both figures. Inset i) shows the control beam shape, indicating the definition of R. Insets ii) and iii) show population pattern cross-sections through the centre of the control beam.

with a central dark core of varying radius R. The outer beam radius was fixed at 0.5 cm and the local intensity in each pixel of the simulated beam was $0.1I_S$. To evaluate the population contrast we defined it as the difference between the minimum $|1\rangle$ population within the pattern and population in the beam centre. The results are shown in figure 6.7. If the hole radius is 0 (no hole), there is of course no contrast, and we see that the population of $|1\rangle$ in the centre is reduced compared to the outer regions. With increasing dark core radii, the contrast peaks at a radius of approximately 0.3 cm. For larger dark core radii the contrast decreases because the control beam no longer has enough power to effectively pump atoms into $|0\rangle$ during the transit across the annulus of brightness. For beams of complicated shapes the exact population contrast structure can become hard to predict without simulations, but these observations should provide some intuition.

Naturally, the temperature of the atoms in the vapour makes a significant impact on the population structures we can inscribe. Figure 6.7 b) shows simulations using the same control beam as in the previous paragraph, but this time with a fixed core radius of R= 0.3 cm and evaluating the contrast for a range of temperatures. It can be seen that at very low temperatures, below approximately 20 K, the contrast is low. This is because the slow-moving cold atoms traverse the bright control beam annulus slowly and have a long time to be efficiently transferred to $|0\rangle$, and so atoms in $|1\rangle$ cannot reach the dark core. For higher temperatures, above around 100 K, contrast becomes low again because of the opposite reason to the cold atom case: hot, fast atoms do not spend sufficient time in the control beam to be transferred out of $|1\rangle$. So for the parameters used here there is an optimal temperature range, close to 50 K, for maximising contrast. Again, this generalises to complicated beam structures, but the exact optimal temperature can be hard to determine without a lot of numerical simulation.

In order to further investigate the relationship between control beam shape and $|1\rangle$ populations we performed the experiment using a control beam with a linear intensity ramp in the azimuthal direction. These results are shown in figure 6.8. It is particularly instructive to look at the measured control and probe azimuthal intensity profile. We expect the control profile to be linear and increasing with azimuthal angle ϕ. This is confirmed by the results in figure 6.8. It can be seen that both control and probe reconstructions possess reasonably linear profiles away from sharp edges, as expected. However, the probe reconstruction is significantly less sharp near high

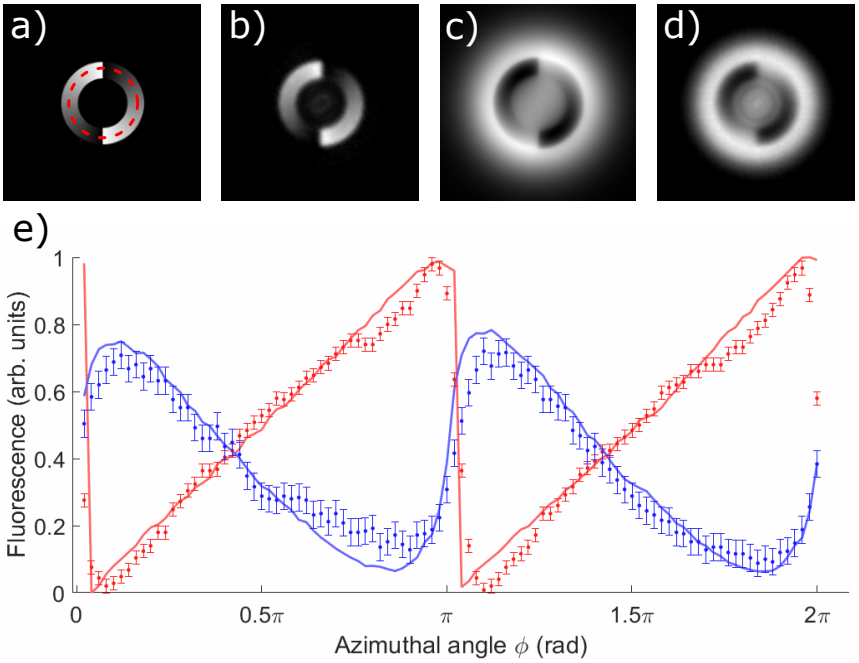

Figure 6.8: Reconstructed cross-sections from a control beam with a linear azimuthal intensity ramp. a) Desired control intensity profile. b) Control fluorescence reconstruction. c) Simulated $|1\rangle$ populations. d) Reconstruction of state $|1\rangle$ populations. e) Unwrapped control (blue) and probe (red) fluorescence profiles at a fixed beam radius corresponding to the red dashed circle in a). Solid lines are simulations. Error bars represent the standard deviation calculated from 10 cross-sections.

contrast edges as can be seen from the smooth curving of the red curve around $\phi = \pi$. A naive way of modeling the expected dark intensity profile would be subtracting the control profile from a constant value, since the intensity of the probe is azimuthally constant. It is apparent from the measured intensity curve that the probe images, and hence the populations, are blurred. This blurring is in fact a direct consequence of the fact that at points within the control beam atoms sample areas of varying intensity from different directions, as expressed in equation 6.5.

Figure 6.9 shows full 3D control and probe reconstructions. The control beams used here were a 3-by-3 array of discs, and a superposition of Laguerre-Gaussian modes with $l_1 = 3$ and $l_2 = 11$. The waists and relative amplitudes corresponding to the two modes are chosen such that the superposition gives a so-called optical ferris wheel [81]. Both the control and probe beams were focused by a lens with focal length 75 mm placed immediately after the recombining beam splitter. The cross-sections used to build up the 3D structures were peak normalised. The halo due to the larger diameter of the probe beam, seen in the probe cross-sections in 6.9b), was removed from the corresponding 3D reconstruction where possible for clarity.

It can be seen that the control reconstruction retains good resolution throughout most of the propagation. The dark lobes are still clearly resolved 2 cm from the focus. In contrast, the probe reconstruction does not show the corresponding bright lobes at the same distance. As can be seen in the 3D image the ability to resolve these lobes within the bright halo is lost near a distance of 4 cm from the focus. Furthermore, the overall contrast between the bright lobes and the background is also lower than the contrast observed in the control cross-sections. This illustrates the resolution limit at which we can manipulate atomic population structures in a warm vapour. In our case the transverse resolution limit is around 200 μm. We made this estimate by convolving a sawtooth by a Gaussian, and fit the result to the measured population shown in figure 6.8. Using a Gaussian with a width of 200 μm gave a very good fit.

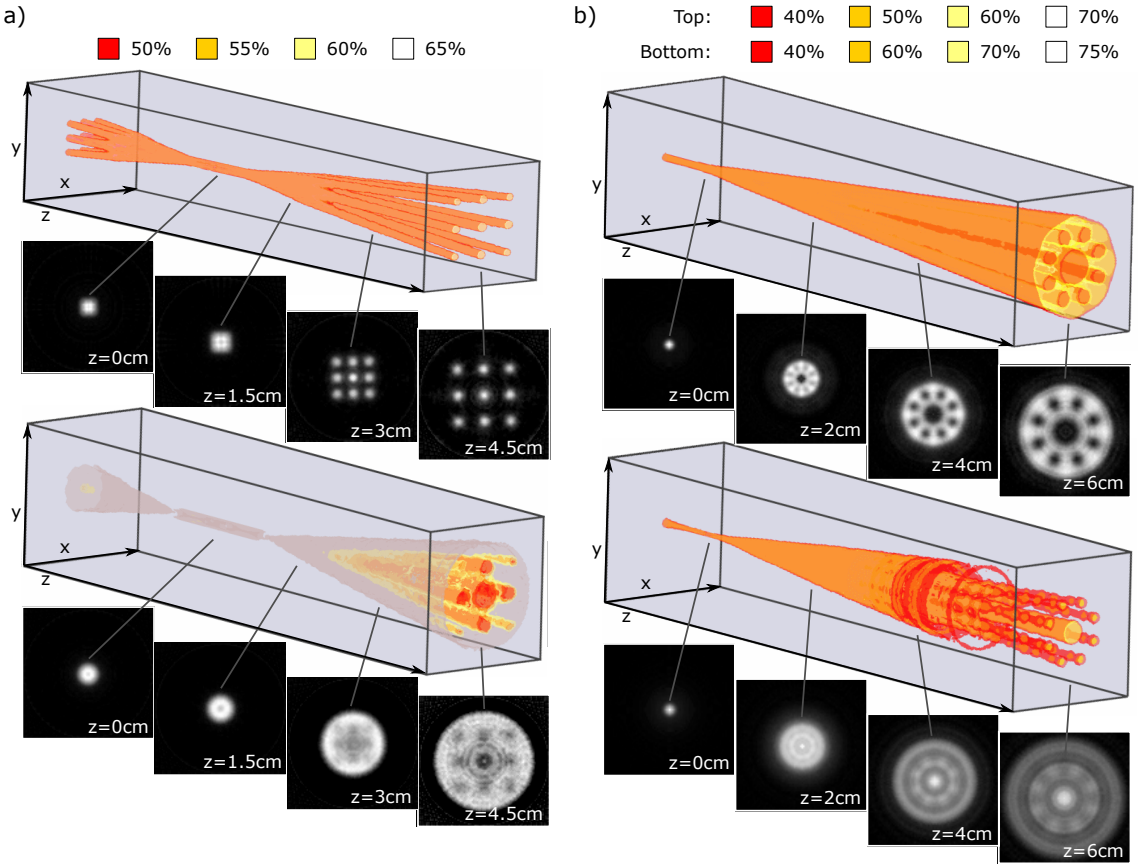

Figure 6.9: Full 3D reconstructions from focused control (top) and probe (bottom) beams. The control beams had a cross-section of a) a 3-by-3 array of discs and b) an optical Ferris wheel. Sample cross-sections are also shown. The dimensions of the boxes are approximately 1 cm x 1 cm x 7 cm. Propagation distances z are measured from the focus. For clarity, isosurfaces were plotted at intensities indicated in the colourboxes. In these figures the beams propagate in the $-z$ direction.

6.5 In summary

We have described an experiment that uses all our expertise with beam shaping and atomic physics to inscribe 3D population patterns into a room temperature rubidium vapour. We reconstructed the inscribed population patterns using fluorescence imaging and tomographic techniques. We paid particular attention the the limitations of resolution and contrast in the achievable atomic populations due to the competition between local population depletion and globally sampled re-population, developing numerical simulations based on rate equations that match our experimental results extremely well. We note that all the work here, both theoretical and experimental was done using CW lasers, and it might be interesting to consider the possibilities of a pulsed control beam, which would allow us to effectively generate a short-lived 3D image memory. In any case, in the next chapter we look at a completely different regime of atomic physics in setting up a cold atom trap.

Chapter 7

Cold atom physics: optical traps

7.1 Introduction

Widely accessible warm atomic vapours already offer an interesting and useful platform to study atom-light interactions, and they can be used very practically for example in high-sensitivity magnetometry [82] or quantum information storage [83]. The main issue with using warm vapours for such applications used to be the very short coherence times of the atoms at these temperatures, essentially because they frequently collide with the cell walls which causes decoherence. While recently this issue has been somewhat resolved by the use of anti-relaxation coatings [84, 85] and/or buffer gases [86], the classic approach of cooling and trapping atoms in a small volume is widely used, for example in high-stability optical clocks [87]. We need cold atoms for their slow movement, to avoid smearing out spatial structures when atoms interact with shaped light. Slow atoms also enhance atom-light interaction because the lack of Doppler shift allows the light to be resonant with more atoms. Here we look at the basic theory of magneto-optical traps, their construction and operation, and we describe a more advanced trapping configuration known as a dark spontaneous force optical trap.

In section 7.2 I describe the theory of magneto-optical traps, and in the subsequent sections in this chapter I describe my work on building and optimising such a trap.

7.2 Magneto-optical trap

Consider a two-level atom at rest in its electronic ground state, with a laser beam shining on it. The atom will only interact with photons in the beam with the appropriate wavelength to match the difference in energy between the two atomic levels (on resonance). Such a photon will be absorbed, and in line with the conservation of momentum cause the atom to move in the direction of travel of the photon. During the absorption process the atom will be promoted to its excited state and after a short time decay back to the ground state through the spontaneous emission of a photon of identical wavelength in a random direction. If the spontaneously emitted photon is emitted in the same direction as the absorbed photon the atom will end up at rest. However, in all other cases its kinetic energy will be higher compared to that before the interaction.

The same process can be used to cool the atom if the laser beam is red-detuned[1] and the atoms

[1] The laser wavelength is tuned slightly off-resonant towards longer wavelengths.

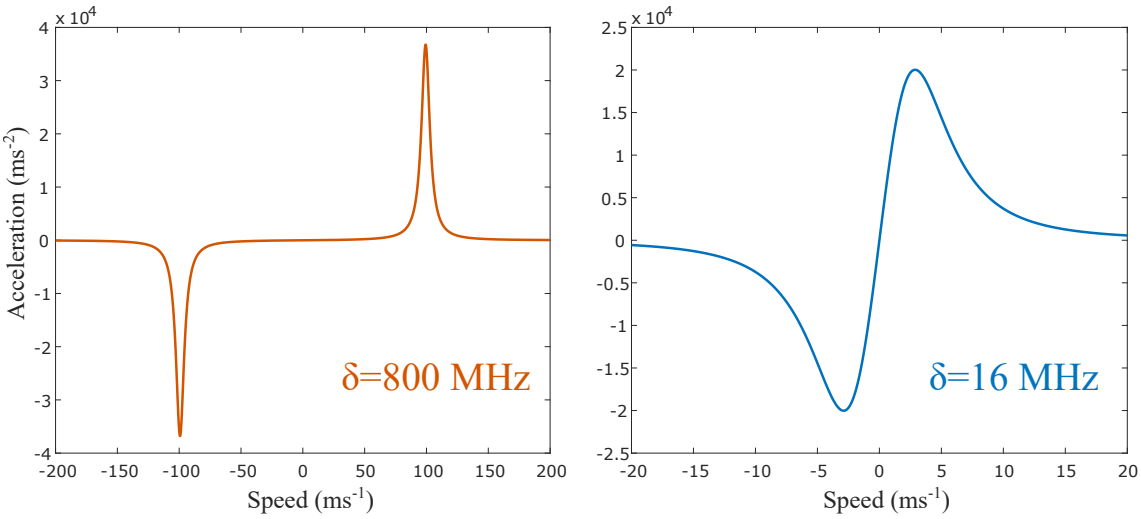

Figure 7.1: Velocity-dependent acceleration in Doppler cooling for rubidium-87 and different detunings δ and I=I_{sat}.

are in thermal motion [88, 89]. In this case in the frame of atoms moving *towards* the laser source photons will be blue-shifted to resonance with the internal transition and be readily absorbed. Due to the geometry of the situation this reduces the momentum of the absorbing atom. Since the effect of spontaneous emission is, averaged over a large number of interaction cycles, zero, the net effect is the slowing of atoms moving towards the laser source. On the other hand in the frame of atoms *not* moving towards the laser there will not be resonant photons to be absorbed provided that the laser spectrum is narrow enough, and hence will not be affected. This process 'burns a hole' in the velocity distribution, but does not slow all atoms if the laser linewidth is narrow. One can choose the speeds affected by selecting the laser detuning. In the case of two counter-propagating laser beams the acceleration experienced by atoms is given by [90]

$$a(v) = \frac{\hbar k}{m}\frac{\Gamma}{2}\left(\frac{I/I_{\text{sat}}}{1+4(\delta-kv)^2/\Gamma^2+2I/I_{\text{sat}}} - \frac{I/I_{\text{sat}}}{1+4(\delta+kv)^2/\Gamma^2+2I/I_{\text{sat}}}\right), \tag{7.1}$$

where $\hbar k$ is the photon momentum and m is the atomic mass. This expression essentially amounts to the statement that absorbed photons change the momentum of the atom, and photons are absorbed at the scattering rate.

As we said before, this velocity-dependent acceleration burns a hole in the velocity distribution around the speed where the Doppler shift brings the atoms exactly to resonance. For room temperature rubidium the detuning would be close to 800 MHz if we wanted to slow the most atoms[2]. However, in that case atoms with lower speeds would not be affected, and the velocity distribution would not be compressed appropriately. At low detunings the acceleration is pretty much linear around 0 m/s, up to a peak at some speed v_c known as the capture velocity. This is shown in figure 7.1. In this regime cooling can be modelled as linear damping, and it allows us to efficiently achieve low temperatures. We chose a detuning of around 16 MHz because we find that it is optimal for trapping, even though the highest capture velocity with proper cooling would occur at around 10 MHz higher detuning. As we will see in the next paragraph, in the presence of a magnetic field gradient needed for trapping there is an extra detuning due to energy level shifts, which optimises the capture velocity for the lower laser detuning we set here. The downside of this is that the lasers are talking to a relatively small fraction of all velocity classes, but for rubidium this is still sufficient to capture a good number of atoms.

[2]The most probable speed is around 100 m/s.

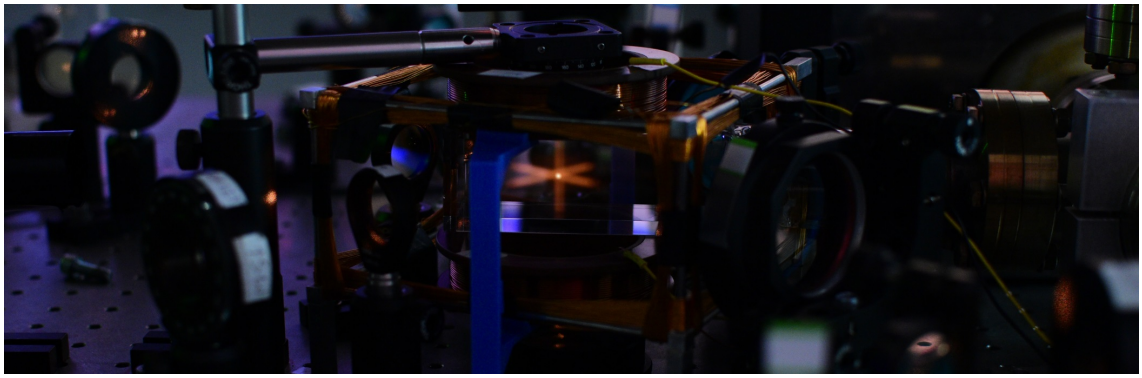

Figure 7.3: X marks the MOT: fluorescence from the trap beams is visible when the vapour pressure is high.

Doppler-cooling allows us to slow atoms down to a few 100 μK, and other effects can cool even further [91, 92, 93], but they do not readily allow us to collect atoms in a small volume. In magneto-optical traps (MOTs) this is achieved by a quadrupole magnetic field generated by a pair of anti-Helmholtz coils [94], see figure 7.2. The quadrupole field results in a spatially-dependent force seen by the atoms because of the Zeeman effect [95]. The magnetic sublevels of the hyperfine structure of the atomic energy levels are not degenerate in the presence of an external magnetic field. In fact the magnetic sublevels are split proportionally to the supplied magnetic field strength and direction. Atoms at different positions within the quadrupole field experience different laser detunings because their energy levels are shifted closer to resonance, provided that the lasers are circularly polarised. This effectively increases the capture velocity and cooling efficiency, but slow atoms are also pushed towards the trap centre. Since the quadrupole magnetic field gets stronger away from the trap centre the Zeeman shifted levels of slow atoms get closer to resonance, and so such atoms will scatter more light farther out, experiencing a restoring force pushing them towards the middle. In the case of neutral atom traps one normally needs counter-propagating lasers in three orthogonal directions with appropriately balanced intensities to slow all atoms, compressing their velocity distribution and so cooling them. This is shown in figure 7.3. Several other configurations also exist [96, 97]. The intensities of lasers in different configurations can be calculated by following [98], for example.

Another important point to consider is that to efficiently cool atoms a suitable set of energy levels needs to exist in the atoms. This is because after scattering events atoms need to return to their original ground state in order to interact with the laser again in a closed loop transition; otherwise the energy difference between the new ground state and the upper state does not match that of the photons, and interaction stops. If the atoms do fall into such a dark state, another 'repump' laser can be used to excite them from the dark state of the trap laser into another excited state, from which they can again decay back into the original ground state. In general the more complicated the (hyper)fine structure of an atomic species is the more repump lasers are necessary

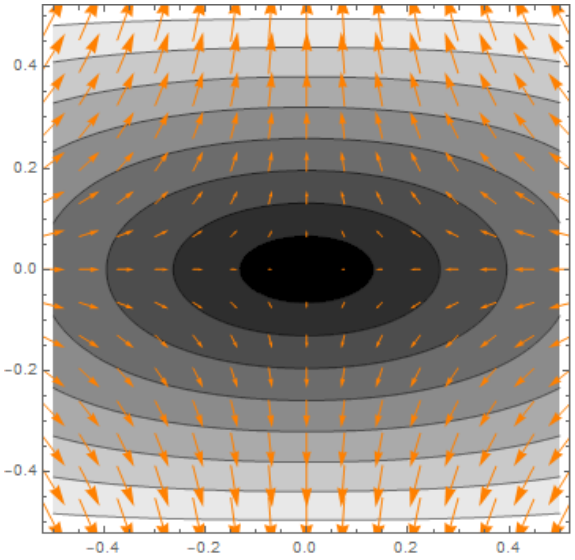

Figure 7.2: Quadrupole magnetic field centered on the trapping region. Represented area is 1 cm×1 cm, and the maximum field magnitude in this region is close to 2 G and it is 0 G in the centre.

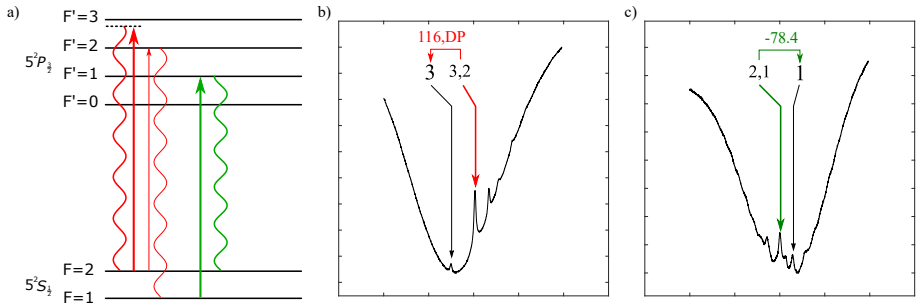

Figure 7.4: Simplified level scheme of ^{87}Rb D_2 line showing a) MOT transitions. Red: trap laser, green: repump. Curved lines indicate decays and thin solid lines indicate off-resonant pumping. ECDL lock points are shown for b) trap (F=2→F'=3) and c) repump (F=1→F'=1) transitions, indicating frequency shifts in MHz from the lock point to the transition. DP indicates that the AOM is in double pass configuration, so its frequency shift should be half of the indicated number.

and the less efficient cooling becomes; for rubidium 87 used in our experiments one repump beam is sufficient. This is because we address the F=2 F'=3 D_2 transition with our trap laser, which can only decay back to the F=2 ground state; however, because of the finite line width of the atomic levels and the laser, occasionally atoms are off-resonantly excited to the F'=2 excited state, from which they can decay into the F=1 ground state as well. A repump tuned to the F=1 F'=1 transition can return these atoms to the cooling cycle. This level scheme is illustrated in figure 7.4a.

7.3 Aligning the MOT

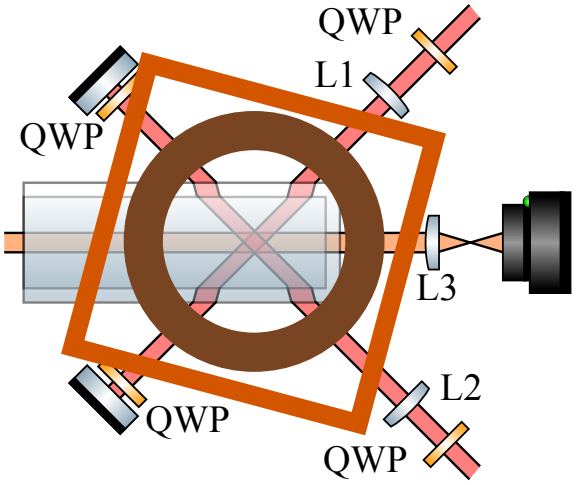

Figure 7.5: Schematic diagram of the current MOT setup, top view. Red lasers are responsible for forming the trap and the orange beam is used for probing the atoms. Quarter-wave plates placed in front of mirrors ensure the correct polarisation of the returning beams. Note the highly exaggerated refraction through the vacuum cell walls. Lenses L1 and L2 in the horizontal beams before they enter the cell are responsible for correcting the returning beam intensity. Quadrupole (brown) and compensation (copper) magnetic coils are also shown.

The correct alignment of the trapping beams is essential to the good operation of the trap. A schematic model of the key components of the trap setup are shown in figure 7.5. Previously, the alignment was done simply by aligning the horizontal beams to pass directly over a set of holes in the optical table in a 90° cross pattern, and aligning the vertical beam with the center hole. The anti-Helmholtz coils were placed such that the magnetic field zero would be at the beam crossing point. Then, the trap quality was iteratively improved by making small changes to the beam alignment. This process does not explicitly take into account refraction at the vacuum chamber walls, nor power loss through reflection, so it was a slow, difficult and non-rigorous process. Since the experiment was recently disassembled for cleaning, we performed calculations of the above effects to speed up alignment and potentially improve trap quality.

Firstly, we calculated the effect of refraction at the vacuum cell walls on the beam travel paths. We found that inside the vacuum chamber the horizontal beams, propagating at 45° to the cell walls, are displaced by about 2.7 mm on a single refraction. This means that after exiting the

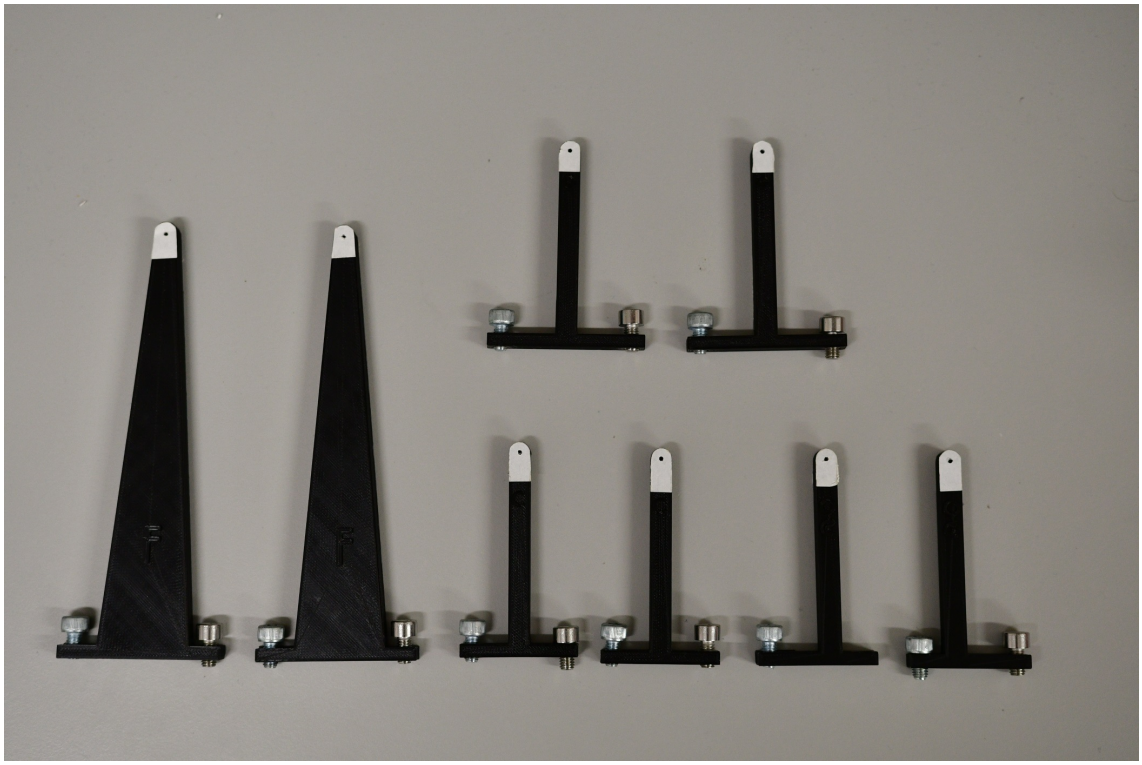

Figure 7.6: 3D printed alignment aids. The tall ones were used for the vertical MOT beam. The various smaller ones were used for aligning the horizontal arms and other beam paths.

cell they are displaced by approximately 5.4 mm compared to going straight through (without any glass present). As a result the magnetic field zero needs to be displaced from its previous position by 2.7 mm[3]. To guarantee this we 3D printed support structures for the coils that ensure the correct magnetic field zero position. The vertical beam also needed to be displaced by the same amount, to do this we again 3D printed a small indicator plate that keys into holes in the optical table, with the correct beam position marked on it. The vertical beam was then aligned with this indicator.

We also need to be able to accurately align the horizontal beams with the magnetic field zero. For this purpose we again 3D printed alignment markers that could be screwed to the optical table in various positions. They had holes at the correct height and transverse location at the corresponding position on the optical table. The trap and repump beams could be aligned through these holes by maximising transmission monitored on a power meter. This allowed very quick alignment of these beams. Because the beams were not perfectly symmetric Gaussians, however, it was difficult to judge when they were aligned best through the holes. As a result the beam alignments all could be optimised after this initial alignment. However, we estimate that the improvements to the initial alignment technique reduced alignment time by about 75%, to about one hour. Figure 7.6 shows the 3D printed alignment aids used in our experiments.

We then turned our attention to the effect of power loss through reflection from the vacuum chamber walls. We found that approximately 5% of light is lost per surface in the horizontal beams. This means that beam intensities between the incident and returning beams are imbalanced, resulting a shift of trap position and change in the trap shape. Since this shift moves the optical trapping region away from the magnetic field zero it reduces trap quality. This effect can be minimised if we reduce the returning beam size such that intensity is the same in the retroreflected beam at

[3]Since it is within the cell.

the trap location as it was on the first pass of the beam. This was achieved by placing lenses in front of the cell (see figure 7.5).

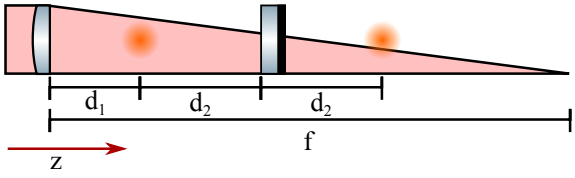

Figure 7.7: Compensating intensity losses on reflection in the MOT vacuum chamber by using lenses, decreasing beam area to compensate for lost power.

The beams pass through the glass walls four times after the MOT region, once on exiting the vacuum cell and once entering again after retroreflection. Thus the returning power is $P_2 = 0.95^4 P_1$. This means that the radius of the returning beam needs to be $R_2 = 0.95^2 R_1$. The lens we need is going to be of a very long focal length, so we can treat the focusing of the beam as linear. Then, using similar triangles, the distance between the lens and the trap position d_1 and the distance from the trap to the mirror d_2 are related to the lens focal length f as

$$\begin{aligned} (f - d_1) - 2d_2 &= 0.95^2(f - d_1), \quad \text{so} \\ 2d_2 + 0.05d_1 &= 0.05f, \end{aligned} \tag{7.2}$$

see figure 7.7.

The parameters are constrained by the placement of magnetic field compensation coils (described later), which imply that we need a focal length that is as long as possible. The longest focal length commercial lens we could find has $f = 2500$ mm. In this case $2d_2 + 0.05d_1 = 125$ mm. The closest we can place the lens to the cell is around 60 mm, in which case $2d_2 \approx 122$ mm so $d_2 \approx 61$ mm. Unfortunately the mirrors could not be placed in these exact locations, so we placed them as close to the ideal locations as possible. The error in the placement was no more than 10 mm, which leads to an intensity error of less than 1%.

After implementing these improvements we measured the number of atoms trapped without any further optimisation. We determined atom numbers from absorption measurements. Using a camera we measure the intensity of a probe beam, tuned to resonance with the trap transition, in the absence of a trapped cloud (I_{in}) and in its presence (I_{out}). The atom number N was found by integrating the two-dimensional projection of the atomic density distribution, assuming low optical density. The density distribution is related to the measured images via

$$N = \frac{1 + I/I_S + 4\Delta^2/\Gamma^2}{A} \int\int \ln\left(\frac{I_{in}(x,y)}{I_{out}(x,y)}\right) dxdy, \tag{7.3}$$

where I is the average intensity of the probe beam, Δ is the probe detuning, which in our experiment was zero, Γ is the natural linewidth of the probe transition, and the resonant scattering cross-section is $A = 2.90 \times 10^{-13}$ m^{-2} and the saturation intensity of the transition is $I_S = 1.67$ mW/cm^2 [59]. The absorption imaging process is illustrated in figure 7.8. After the initial alignment of the trap with the new technique, without further optimisation, we find atom numbers of approximately $2.5{\times}10^8$, which are consistent with the best results that were previously achieved after lengthy alignment. The shape of the cloud, shown in figure 7.9, is also qualitatively better, more symmetric.

We also measured the MOT temperature. This is done by loading a MOT, turning the trapping beams and quadrupole field off and letting the cloud expand for some time δt, then turning the probe beam on and recording a absorption image. The size of the expanded cloud is estimated by fitting it as a 3D Gaussian. This process is repeated with a number of different δt, and from the series of cloud sizes after different times of expansion the most probable atomic speed is estimated.

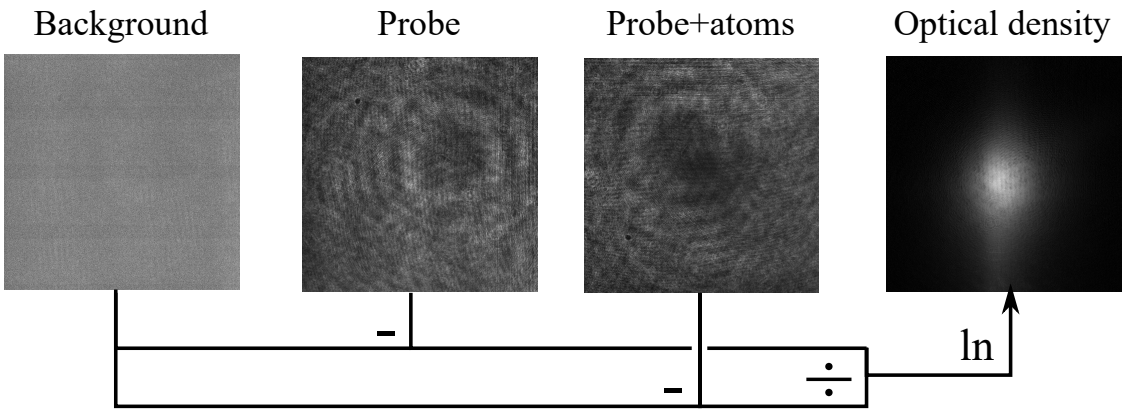

Figure 7.8: Absorption imaging proces. The natural logarithm of the ratio of the background-subtracted probe beam images without and with the atom cloud present is the optical density, or absorption image, up to a constant factor. The greyscale images were scaled to show detail; the background is usually barely perceptible when view on the same scale as the probe beam.

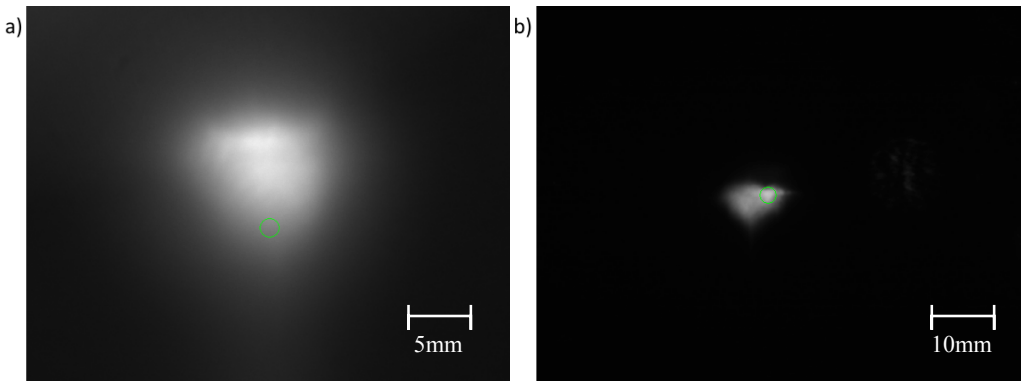

Figure 7.9: Fluorescence images of a MOT when viewed a) along the probe direction and b) perpendicular to the probe direction. The two images are not to the same scale. The green circles were meant to help alignment in the experiment, and are not intended to indicate the same volume.

This can be related to temperature by

$$T_j = \frac{m}{k_B}\frac{\sigma_{j2}^2 - \sigma_{j1}^2}{\delta t_2^2 - \delta t_1^2}, \quad j \in \{x, y\} \tag{7.4}$$

where m is the atomic mass, and σs are the cloud widths along the j direction at times δt. We find initial temperatures of $T_x \approx 220\mu$K and $T_y \approx 385\mu$K, which are on the warm side, but similar to previously measured results.

7.4 Magnetic field control

There are numerous sources of magnetic fields affecting the atomic clouds used in the experiments. The largest contribution likely comes from the magnetic field of the Earth, but the control electronics, and especially the vacuum pump located in close proximity to the experiment volume, provide significant contributions. These must be accounted for, and compensated, in order to perform all experiments presented in this thesis.

Magnetic fields in and around the trap region are controlled by 4 sets of coils. A pair of circular coils in anti-Helmholtz configuration provide the quadrupole field necessary for the MOT, and three pairs of rectangular coils in Helmholtz configuration are responsible for cancelling the Earth's

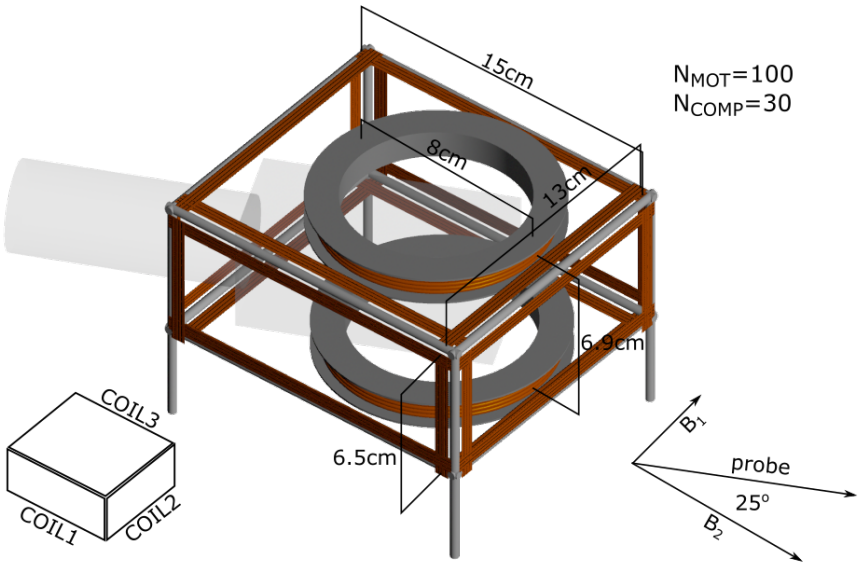

Figure 7.10: Location and orientation of the four pairs of coils around the vacuum cell of the trap. Coil dimensions and the number of turns N in each coil are indicated. Inset shows the labelling of the compensation coils and the directions of the magnetic fields generated by them.

residual magnetic field and providing a small bias field determining the quantisation axis of the atoms. The location and orientation of these coils are shown in figure 7.10.

The magnetic fields generated by the coils depends on the current driven through them. In practice we do this by setting a voltage, which sets a current through Ohm's law, $V = IR$. The problem is that the resistance of the wires that make up the coil is temperature dependent, and the coils heat up considerably when driven at the high currents we use. This means that if the set voltage is constant the current is temperature dependent, and stabilises to different values depending on external factors such as the room temperature and air currents. The current also takes a long time to stabilise, because the coils thermalise on the order of several minutes. This is not acceptable if we need fast switching of fields, which we do.

This issue is solved by a current driver shown in figure 7.11. The current in the coil is monitored by the voltage across a 1Ω resistance. If a TTL signal is on, switching the coils on, this voltage is compared to a reference set either by a potentiometer or an analogue voltage from an I/O card connected to the experiment control computer. The difference between the reference and monitor voltages passes through a low-pass filter to reduce ringing and is used to control the resistance of a FET in series with the coil, thus feeding back to the current through the coil.

7.4.1 Field strengths and coil currents

It is useful to be able to determine the transverse and longitudinal components of the bias field in the frame of the probe direction. Assuming that the Earth's field is properly compensated, coil 3 does not contribute to either of these since the bias field is applied only by coils 1 and 2. Then, the two components of interest, longitudinal and transverse, can be found in a straightforward fashion:

$$\begin{aligned} B_{\text{long}} &= B_1 \sin 25^\circ + B_2 \cos 25^\circ \\ B_{\text{trans}} &= B_1 \cos 25^\circ - B_2 \sin 25^\circ. \end{aligned} \tag{7.5}$$

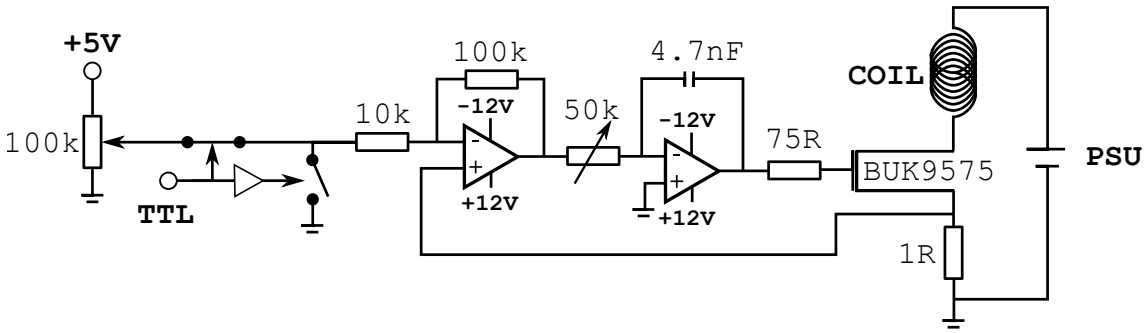

Figure 7.11: Coil current driver circuit diagram. The PSU supplies a constant voltage, and the current in the coil-side circuit is adjusted by feeding back to a BUK9575 FET in series with the coil.

However, in the experiment we can directly control B_1 and B_2. Equations 7.5 can be inverted to find how the experimentally controllable parameters depend on the ones of theoretical interest:

$$\begin{aligned} B_1 &= \frac{B_{\text{long}} \tan 25^\circ + B_{\text{trans}}}{\cos 25^\circ + \sin 25^\circ \tan 25^\circ} = B_{\text{long}} \sin 25^\circ + B_{\text{trans}} \cos 25^\circ \\ B_2 &= \frac{B_{\text{long}} \cot 25^\circ - B_{\text{trans}}}{\sin 25^\circ + \cos 25^\circ \cot 25^\circ} = B_{\text{long}} \cos 25^\circ - B_{\text{trans}} \sin 25^\circ. \end{aligned} \tag{7.6}$$

Experimentally these fields are generated by pairs of rectangular coils. For such coils with N turns, dimensions a, b and separation h the current I required to generate a desired magnetic field is given by

$$I = \frac{5B}{8N} \frac{\left(a^2+h^2\right)\left(b^2+h^2\right)\sqrt{a^2+b^2+h^2}}{ab\left(a^2+b^2+2h^2\right)}. \tag{7.7}$$

The direction of the current in the coils required to give the desired magnetic field direction can easily be determined using the right-hand rule. In the experiment the compensation coils have been wired to produce magnetic fields in the directions indicated in figure 7.10. Coil 3 provides a magnetic field pointing upwards.

7.4.2 Magnetic field rise-time

Because of the non-zero inductance of the coils used to generate magnetic fields in the experiment it takes a finite time for the current in the coils, and hence the magnetic fields, to reach their intended value after supplying the coils with a voltage. The current at time t after switching a voltage V_0 on is given by equation 7.8,

$$I = \frac{V_0}{R}\left(1 - e^{-t\frac{R}{L_{coil}}}\right), \tag{7.8}$$

where R is the coil resistance and L is the inductance of the coil. These can be determined from experimentally measurable quantities, listed in

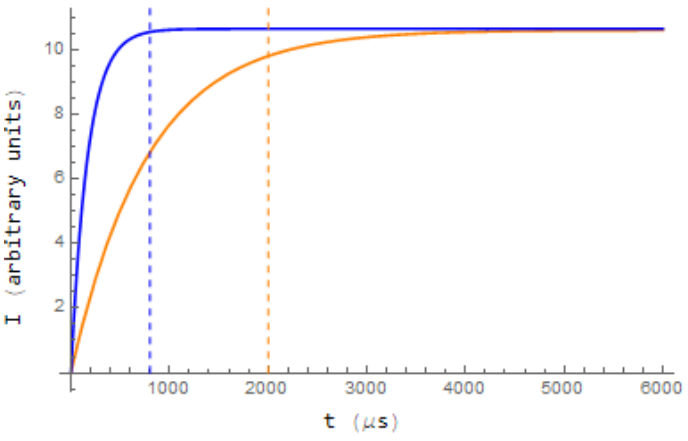

Figure 7.12: Time-dependence of the currents driven in the different coils used for the MOT. Blue: compensation coils, orange: quadrupole field coils. Vertical dashed lines indicate experimentally measured times by which magnetic fields reach their intended value.

Quantity	Symbol	Value	Unit
Number of loops per layer	n	20	-
Number of layers	m	4	-
Coil radius	r	0.04	[m]
Coil length	l	0.015	[m]
Wire diameter	d	0.0075	[m]
Material resistivity	ρ	1.724×10^{-8}	[Ωm]

Table 7.1: Physical parameters of the quadrupole field coils.

Quantity	Symbol	Value	Unit
Number of loops per layer	n	15	-
Number of layers	m	2	-
Coil radius[4]	r	0.0575	[m]
Coil length	l	0.005	[m]
Wire diameter	d	0.0045	[m]
Material resistivity	ρ	1.724×10^{-8}	[Ωm]

Table 7.2: Physical parameters of the compensation field coils.

tables 7.1 and 7.2, as described in equations 7.9 and 7.10.

$$R = \frac{2N\rho r}{(d/2)^2}. \tag{7.9}$$

To find the total inductance L_{coil} of a multi-layered cylindrical coil with n current loops in m layers we need to consider the self-inductances L of each circular current loop as well as the mutual inductance $M_{j,k}$ of all pairs j, k of circular loops,

$$L_{coil} = \sum_{i=1}^{n\times m} L_i + \sum_{j=1}^{n\times m} \sum_{k=1,k\neq j}^{n\times m} M_{j,k}. \tag{7.10}$$

The self-inductance of a current loop can be expressed as [99]

$$L = \frac{8\mu_0 r^3}{3d^2}\left(\frac{2k_s^2-1}{k_s^3}E(k_s) + \frac{1-k_s^2}{k_s^3}K(k_s) - 1\right) \text{ with } k_s = \sqrt{\frac{4r^2}{4r^2+d^2}}, \tag{7.11}$$

where r is the loop radius, d is the wire diameter and $K(k)$ and $E(k)$ are the complete elliptic integrals of the first and second kind, respectively. The mutual inductance of a pair of coaxial circular loops is given by [100]

$$M_{j,k} = \frac{8\pi\sqrt{r_j r_k}}{\sqrt{k_m}}\left(K(k_m) - E(k_m)\right) \text{ with } k_m = \frac{\sqrt{(r_j+r_k)^2+x^2} - \sqrt{(r_j-r_k)^2+x^2}}{\sqrt{(r_j+r_k)^2+x^2} + \sqrt{(r_j-r_k)^2+x^2}}, \tag{7.12}$$

where r_j and r_k are the radii of the two current loops and x is the separation of their centres. These formulae hold for a single coil and do not take into account the fact that in the experiment coils come in pairs and the pairs are near other coil pairs. However, experimentally measured rise-times agree nicely with the results obtained using these expressions, see figure 7.12.

[4]Since the compensation coils are rectangular, and each has different dimensions, it is incorrect to use the equations for cylindrical coils. However, it turns out that one can use an approximate radius, calculated as the mean of the side lengths, and get a rise-time that agrees well with the experimentally measured rise-time.

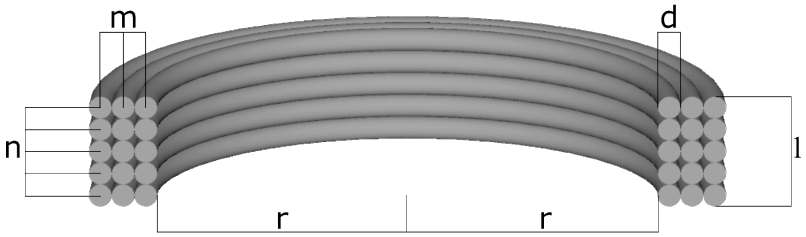

Figure 7.13: Schematic diagram showing magnetic coil parameters, labeled as in tables 7.1 and 7.2, relevant for calculating rise-times.

7.5 Background field compensation

The quadrupole magnetic field due to a pair of cylindrical coils in anti-Helmholtz configuration, Taylor-expanded to 3rd order around the origin, is given by

$$\mathbf{B_{quad}}(r,y) = (B_r(r,y), B_y(r,y)) = B_1 I \left(\frac{-1}{2}, y\right) + B_3 I \left(\frac{3r^3}{8} - \frac{3r^2 y}{2}, y^3 - \frac{3r^2 y}{2}\right), \tag{7.13}$$

with

$$B_1 = \frac{24\pi N d^2 s}{5(d^2+s^2)^{5/2}}, B_3 = \frac{96\pi N d^2 s(4s^2-3d^2)}{(d^2+s^2)^{9/2}}, \tag{7.14}$$

where N is the number of turns of the coils, d is their diameter, s is their separation and I is the current throught them.

Figure 7.2 a few pages ago shows this quadrupole magnetic field in the $x-y$ plane centered on the trap location in a region roughly corresponding to the area imaged by a camera in the actual experiment for a typical trapping current $I = 2.5$A. In this region for this current the field is very close to being linear, with a gradient of approximately 4 Gcm^{-1} in the y-direction and 2 Gcm^{-1} in the x-direction (and equivalently in the z-direction).

Since a MOT forms around the absolute zero of the magnetic field, in the absence of any background fields the trap would be located at the origin of the quadrupole field. However, any background field will offset the MOT, as it forms where the quadrupole field exactly cancels the background field. Naturally, for different quadrupole field gradients this occurs at different positions. This effect can be used to determine the background field which in turn can be used to cancel it. In the regime where the linear part of the quardupole field dominates (which we define as $B_{quad,linear} > 10B_{quad,non-linear}$, which corresponds to a transverse area with dimensions 14.6cm x 0.55cm for the lowest field, 0.5G, that can be supplied by the quadupole coils and still obtain a decent MOT), the MOT position in the radial r-directions and vertical y direction respectively are given simply by rearranging equation 7.13, omitting the higher order terms:

$$\begin{aligned} r &= 2B_{bg,r}/B_1 I \\ y &= B_{bg,y}/B_1 I. \end{aligned} \tag{7.15}$$

So the MOT position offset is inversely proportional to the current supplied to the coils. In principle this makes the measurement of the background field rather straightforward, all one has to do is record the MOT position recorded by two orthogonal cameras for a known coil current, and the background field can be directly obtained. However in reality this is slightly complicated by the fact that the location of the quadrupole field zero, the origin of the above system, is not known to begin with, and as a result the offset position cannot be immediately determined. Happily, this is

not an insurmountable problem, as equations 7.16 are only slightly modified by this situation:

$$\begin{aligned} r &= 2B_{bg,r}/B_1 I - r_0 \\ y &= B_{bg,y}/B_1 I - y_0, \end{aligned} \tag{7.16}$$

where r_0 and y_0 are the unknown coordinates of the quadrupole field zero in the cameras' frame of reference. Loading a series of MOTs for a series of different coil currents and monitoring their position by centre-of-mass estimation of their fluorescence allows one to produce r and y plots agains I, and fitting these for I^{-1} it is straightforward to obtain r_0 and y_0 and hence calculate the background field.

Figure 7.14a) shows a set of measurements used to determine the background field. The data for these graphs was obtained with an arbitrary known level of compensation field applied, because without such a field the MOT formed outside the field of view (FOV) of the cameras observing the trapping region. The arbitrary additional field applied to keep the MOT visible even for low quadrupole coil currents was $\mathbf{B}(x, y, z) = (1.4\text{ G}, 4.5\text{ G}, 2\text{ G})$. Using the fits shown in the figure, we determined that the additional compensation field required to cancel all background magnetic fields was (0.1 G, 0.69 G, 0.33 G), giving a total compensation field of (1.5 G, 5.19 G, 2.33 G). It can be seen that the measured positions in the x and y directions match the expected curve, but the z-displacements are not a good match. This was because the MOT shape changed as a function of atoms in the trap, and for larger quadrupole field gradients more atoms were trapped. This made it very difficult to estimate the MOT position when the shape was not symmetric. To reduce this effect we ran a similar experiment with a very short MOT loading time of 2 s, such that the trap would contain a small, and hopefully consistent, number of atoms for each quadrupole field gradient. The downside of this approach was that because of the low atom numbers background light threw off the centre-of-mass estimations. Figure 7.14b) shows these results, which gave a compensation field estimate of (1.48 G, 5.15 G, 2.48 G).

In practice a good approximation to zero background field can also be achieved by simply observing the MOT position using a high coil current, and adjusting the compensation coil fields until the MOT forms in the same place for low quadrupole coil currents. We needed to do this, because the results obtained above were not good enough to achieve satisfactory compensation. The reason for this is that in practice it was exceedingly difficult to determine the position at which the MOT forms, as the trapped cloud was asymmetrical in our case, being highly elongated in the z direction with bulges in both the x and y direction. As a result the position measurements were not quite as precise as we would have needed them to be for a good compensation. However, the measured compensation field was very close to the one deemed good enough, which was (1.6 G, 5.35 G, 2.57 G). In the end these values were confirmed by the most sensitive test we have access to, which is based on spatially dependent electromagnetically induced transparency, which will be detailed in the next chapter. It is unclear whether the background magnetic field changes on the scale we measure it over periods longer than a few weeks, but it certainly appears to be constant for at least two weeks, which was the length of the longest continuous experimental run.

7.6 Dark spontaneous force optical trap

In our experiment we have the ability to holographically shape our laser beams using an SLM. This allows us to use different trapping configurations other than the standard MOT. One particularly useful configuration is the so-called dark spontaneous-force optical trap (SpOT) [101, 102, 103]. In a standard MOT temperatures and densities are limited by recoil heating and radiation trapping [104]. Ideally, hot atoms should be exposed to a high light intensity so that they cool quickly.

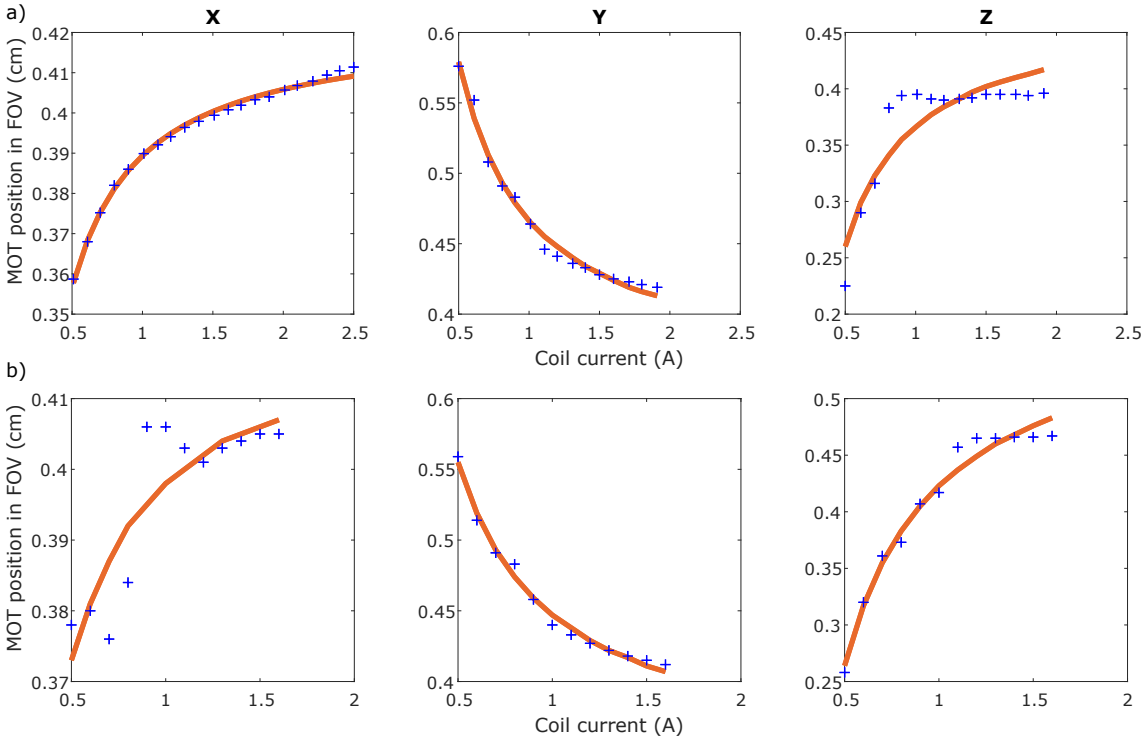

Figure 7.14: MOT displacement as a function of quadrupole field gradient, written here in terms of current in the quadrupole coils, with uncompensated background magnetic field. The top row (a) was measured using a long MOT loading time of 10 s, which allowed us to collect a large number of atoms. The bottom row (b) was measured using a short loading time, 2 s, so that the MOT was as symmetric as possible.

However, cold atoms should not encounter large photon numbers because at temperatures where the kinetic energy of atoms is comparable to that provided by the absorption-emission cycles laser cooling turns into heating. At the same time the scattered trapping light from the trapped, cold cloud produces radiation pressure pushing the atoms apart, limiting density. The dark SpOT approach is one way to combat this. In such a trap the repump laser has a dark core in its middle, where atoms slowly decay to the F=1 ground state via off-resonant excitation by the trap laser to the F'=2 excited state. These atoms no longer interact with any light as the trap beam does not drive transitions from this level and the repump beam is dark. In practice this decay in rubidium is too slow, especially since the cooling beam is off-resonant for cold atoms. This is helpful for a MOT, but undesirable for the SpOT. We introduce another laser called the 'depump' to encourage atoms to decay to the F=1 ground state, shown in figure 7.15a. The depump beam drives the F=2 F'=2 transition, and it has the opposite purpose to that of the repump laser. The depump is shaped to fit exactly in the dark core of the repump beam. Using the dark SpOT trap has several benefits, such as increased density and a different electronic state for the trapped atoms, which is the main benefit for us.

There is a downside to using a SpOT, though, in that loading the trap from background pressure is not very efficient, in contrast with the MOT. This is because the lack of repump light severely limits the number of atoms that we can load into the trap. This is where using an SLM gives us a remarkable advantage. We can first load a standard MOT with high atom numbers at an already low temperature, then holographically shape the repump beam into one with a dark core filled with depump light. Such a loading sequence allows us to retain most of the atoms from the MOT with all the advantages of the SpOT. This advantage comes with its own challenges, however. Aligning the MOT well is already not trivial as discussed above, but optimising a SpOT requires precise control over three new degrees of freedom: the size of the hole in the repump, the intensity of the

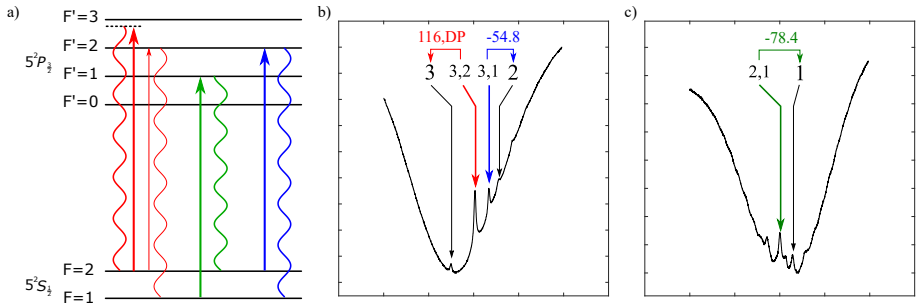

Figure 7.15: Simplified level scheme of ^{87}Rb D_2 line showing a) SpOT transitions. Red: trap laser, green: repump, blue: depump. Curved lines indicate decays and thin solid lines indicate off-resonant pumping. ECDL lock points are shown for b) trap and depump (F=2→F'=3 and F'=2) and c) repump (F=1→F'=1) transitions, indicating frequency shifts in MHz from the lock point to the transition. DP indicates that the AOM is in double pass configuration, so its frequency shift should be half of the indicated number.

depump, and the SpOT loading time (the time for which the trap is in the SpOT configuration after loading a MOT but before any probing). We find that the SpOT is much more sensitive to beam alignment accuracy compared to the MOT. It is also harder to characterise the SpOT, since the atoms in this trap are in the F=1 ground state from which there are no cycling transitions, because any $\sigma_\pm$ or π polarised light on the F=1→F'=0 transition would drive the atoms into a dark state, so we cannot use absorption imaging to determine atom numbers. This problem can be alleviated somewhat by either mixing different polarisations in the probe beam or applying a magnetic field that mixes the Zeeman sublevels. In either of these cases, however, it is more difficult to calculate the saturation intensity with any accuracy.

A simple solution is to transfer the atoms back to the F=2 ground state and measure their numbers there. The issue with this approach is that there are MOT atoms in the F=2 ground state at all times because of the presence of the repump light around the edges. To illustrate this, consider a very small dark core in the repump. In this case we effectively still have a MOT, with a little bit of a SpOT. So we need to somehow determine how many atoms are still in the MOT to be able to determine how many are in the SpOT. The procedure for this is as follows. We first load a MOT for 10 seconds, then transfer to a SpOT, and measure atom numbers using absorption imaging on the usual F=2→F'=3 probe transition. This gives us the number of atoms left in the MOT after changing over to the SpOT configuration. Second, we load a MOT again, transfer to the SpOT, but this time we apply repump light for a short time to get all atoms back into the F=2 ground state, and perform absorption imaging to determine the total number of atoms in both the MOT and the SpOT. The difference of the two numbers gives us the number of atoms in the SpOT. Of course all of these measurements need to be done several times and the results averaged, because there is a small but noticeable difference in numbers from one trap to the next. The SpOT density could be measured by using a repump beam along the probe path for absorption imaging. This is a good approach for recording the shape and size of the SpOT cloud, but cannot be used to measure atom numbers. The SpOT size can then be calculated by fitting a 2D Gaussian to the recorded SpOT image, and assuming that the depth is similar to the width of one of the transverse Gaussians, e.g. $\sigma_z = \sigma_x$, we can estimate the volume as $V = (2\pi)^{3/2}\sigma_x\sigma_y\sigma_z$, where σ_j is the j direction. The SpOT density is then

$$n = \frac{N_{total} - N_{MOT}}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_z}. \tag{7.17}$$

Now that we have some diagnostics for the SpOT, we can proceed to discuss how to optimise

the trap. A well-aligned MOT with high atom numbers is paramount to getting a good SpOT. The MOT should be built such that the repump beam comes from the first order of a blazed grating displayed on the SLM, without any amplitude shaping to maximise power. The shape of the repump is not critical for a good MOT, a large size, at least covering the trap beam volume, is much more important. In our traps the repump beams are only included in the horizontal arms, but not in the vertical. This has purely practical reasons; for the SpOT we need to image the SLM plane onto the atoms, which we did with a single lens, but in this case the path lengths for all arms containing repump light needed to be the same. This was manageable for the horizontal arms, but difficult to include the vertical arm as well, and since a tiny amount of repump light is enough for effective trapping, it was not critical to have repump light in the vertical direction anyway. The depump laser was then directed onto the same area of the SLM as the repump, but under a slightly different angle. This angle was chosen such that the zeroth order reflection of the grating for the depump coincided with the first order diffraction for the repump. This is shown in figure 7.16.

The co-linear alignment of the repump and depump beams was important for the effective operation of the SpOT, requiring precision beyond what could be achieved by looking at the beams on a detector card. Instead, we constructed a telescope to image the MOT centre onto a camera. The camera could also be placed into the far field of the telescope. We then iteratively aligned the depump into the hole of the repump using the combining beamsplitter and a mirrors before it, observing the alignment of the two beams in the image and far field planes. We could use the same camera setup to precisely align the repump-depump combination with the trap beam, although this was much harder since in the image plane both the trap and repump were very large, making it hard to estimate the position of their centres[5].

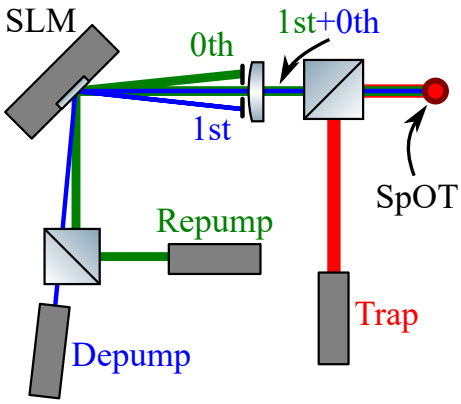

Figure 7.16: Schematic diagram of SpOT beam alignments, with emphasis on the SLM orders. The arrangement is a bit different in the actual experiment, see figure 7.19.

Confident in the alignment we turned our attention to the new degrees of freedom. Our strategy was to perform a rough scan for each property in order, keeping the rough optimal value from each scan to the next, then optimised around the best values found this way. We decided that for our experiments atom number in the SpOT was more important than density, so we optimised for atom number. The experiment timing sequence is shown in figure 7.17. In all experiments a MOT is loaded first for about 10 seconds (trap and repump beams and quadrupole coils on, everything else off). Then the SLM is commanded to switch to the SpOT configuration, cutting a hole in the repump. Unfortunately we do not have direct access to the internal clock of the SLM, so it is challenging to synchronise the experiments to the SLM refresh. We manage this by utilising a small section of the SLM as a separate grating from the repump shaping grating. A small beam (derived from the first order of the depump AOM) is diffracted off this part of the SLM, whose first order is detected by a photodiode. When the SLM is in MOT mode (unshaped repump) there is no grating on this 'trigger' part of the SLM and the photodiode signal is low. The SpOT hologram contains a grating in the trigger area, so when the SLM is displaying this the photodiode signal is high. The photodiode signal is sent through a Schmitt trigger, connected to the experiment control electronics and used as a trigger signal for the rest of the experiment.

[5]And their messy, non-Gaussian shapes did not help either.

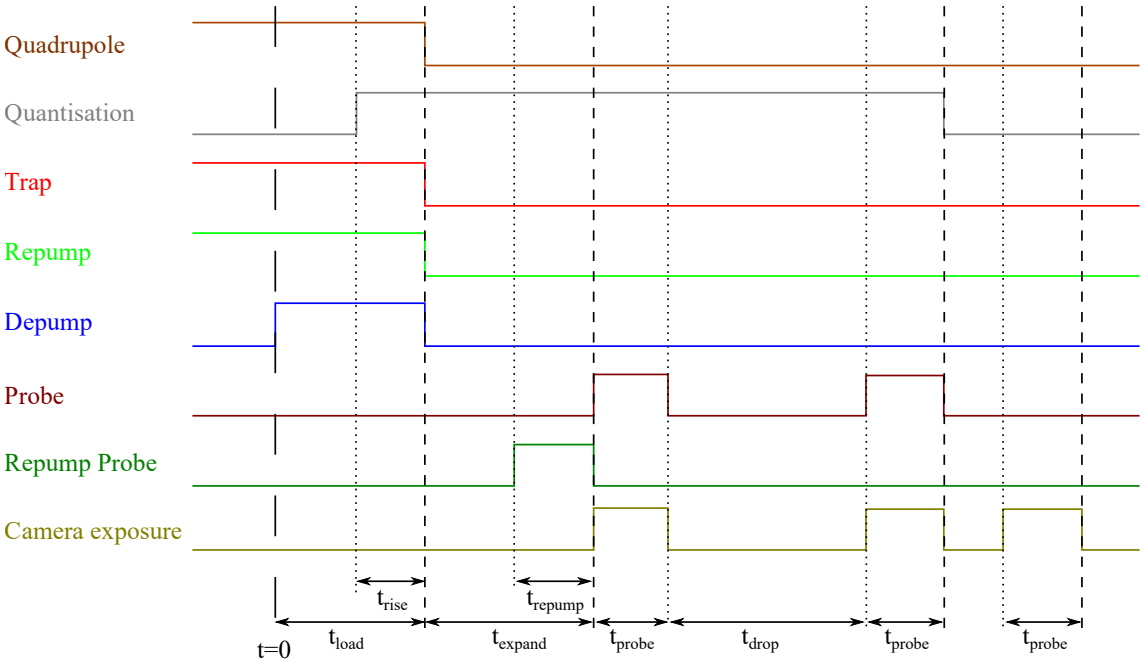

Figure 7.17: Experiment timings for loading and characterising a SpOT. The graphs indicate when the respective beams, magnetic fields, or cameras are on (high) or off (low). Time axis is not uniform or to scale.

Table 7.3: Typical timings for loading and characterising a SpOT.

Pulse	t_{load}	t_{rise}	t_{expand}	t_{repump}	t_{probe}	t_{drop}
Duration (μs)	250000	800	1500	1000	2000	150000

When the SLM switched over to the SpOT configuration and triggered the timing electronics, the SpOT loading begins (t=0 on figure 7.17). The depump is turned on at this time, and remains on while the SpOT is loading, for t_{load}. After this time all trapping beams and the quadrupole field are turned off. The quantisation field, supplied by the compensation coils, needs to be on at this time, so it is turned on slightly earlier as determined by the coil risetime t_{rise} (see figure 7.12). The cloud of atoms is allowed to expand and fall under gravity for a time t_{expand} to reduce density to a point where the cloud does not absorb all photons from the probe beam[6]. In this time the repump probe beam may be turned on for t_{repump} to repump all atoms to the F=2 ground state when required (see above about measuring SpOT numbers). Then the probe beam is turned on for t_{probe} and the camera is exposed to record an absorption image. Afterwards all beams are turned off and the cloud is allowed to drop out of the detection region over t_{drop}. The probe beam is then turned back on and the camera is exposed for another t_{probe} to record an image of the probe beam in the absence of atoms. After another delay (typically t_{drop}) the camera is exposed again without any beams on for t_{probe} to measure the background illumination. Typical durations of these 'pulses' are shown in table 7.3.

For the optimisation, first we set the depump power and SpOT loading times to some reasonable value based on [103] and changed the radius of the dark core in the repump at the SLM in 10 pixel increments from 100 to 200 pixels. The SLM is imaged onto the atoms with a magnifaction dictated by the lens focal length and the distance between the trap and the SLM. 100 SLM pixels are magnified to approximately 4 mm at the trap (which corresponds to a 2× magnification). The optimal was usually found between a 150 to 180 pixel radius, depending on the alignment of the beams. We then kept the depump hole size at the rough optimal value and scanned the depump

[6]If it was allowed to absorb all photons we would not be able to measure accurate numbers, but an estimate of the lowest possible number.

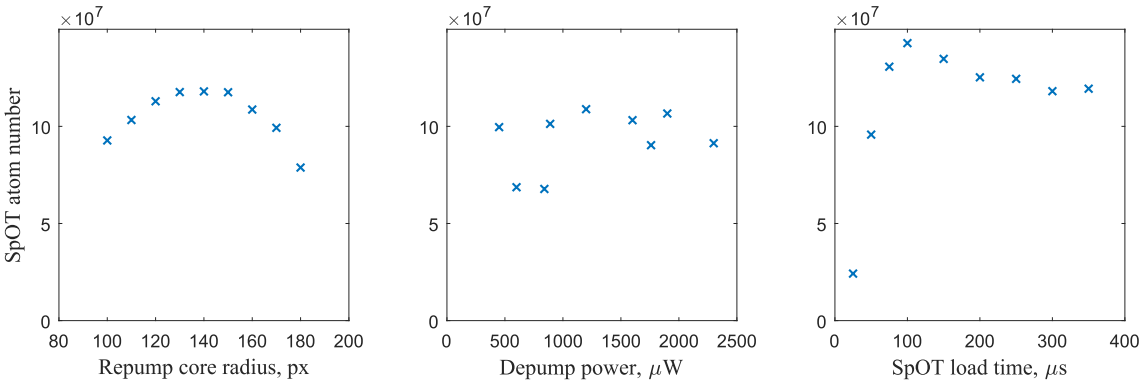

Figure 7.18: Sample data from the SpOT optimisation process from May 2017.

power using a variable ND filter[7]. Too little depump power would lead to too small atom numbers in the SpOT because atoms would not decay to the MOT dark state (F=1) efficiently. At too high depump powers we again find very few atoms in the SpOT because then the depump overwhelms even the small amount of repump light and we are left with no trap. We typically found that very small amounts of depump power were sufficient, but it was strongly dependent on alignment and the optimum was within a broad range, between 45 and 1600 μW measured before the SLM. Lastly, we adjusted the SpOT loading time $\mathbf{t_{load}}$ from 50000 μs to 350000 μs in 50000 μs steps. We expected behaviour similar to the depump power scan for similar reasons: at too short loading times we don't have enough atoms, and at too long times atoms diffuse out of the trap. We found the optimal load time to be around 250000 μs. Before further experiments these parameters were optimised around these best values, although we almost always found that there was no noticeable benefit to departing from the rough values. We found beam alignment (influencing the dark core radius and depump power required) to be the most critical parameter for a good SpOT. Sample data from such scans is shown in figure 7.18.

7.7 In summary

In this chapter we have looked at the basics of the theory of neutral atom trapping, describing the operation of a standard magneto-optical trap. We paid a lot of attention to the practical operation of our MOT, including the alignment of the setup and absorption imaging as the primary means of measuring atom numbers, density and temperature in the trap. We have looked at the control of magnetic fields with the available coils in great detail, which, we will see, is extremely important for experiments in the next chapter. We also described our improvement on the standard MOT in the form of a MOT-loaded dynamic dark spontaneous force optical trap, and examined our method for optimising atom numbers in the SpOT. In the next chapter we will put this trap to use in our investigation of the interaction of atomic states and magnetic fields, driven by laser light with structured polarisation.

[7]Originally a stack of fixed ND filters were used, and the increments were highly discretised by the available ND filters.

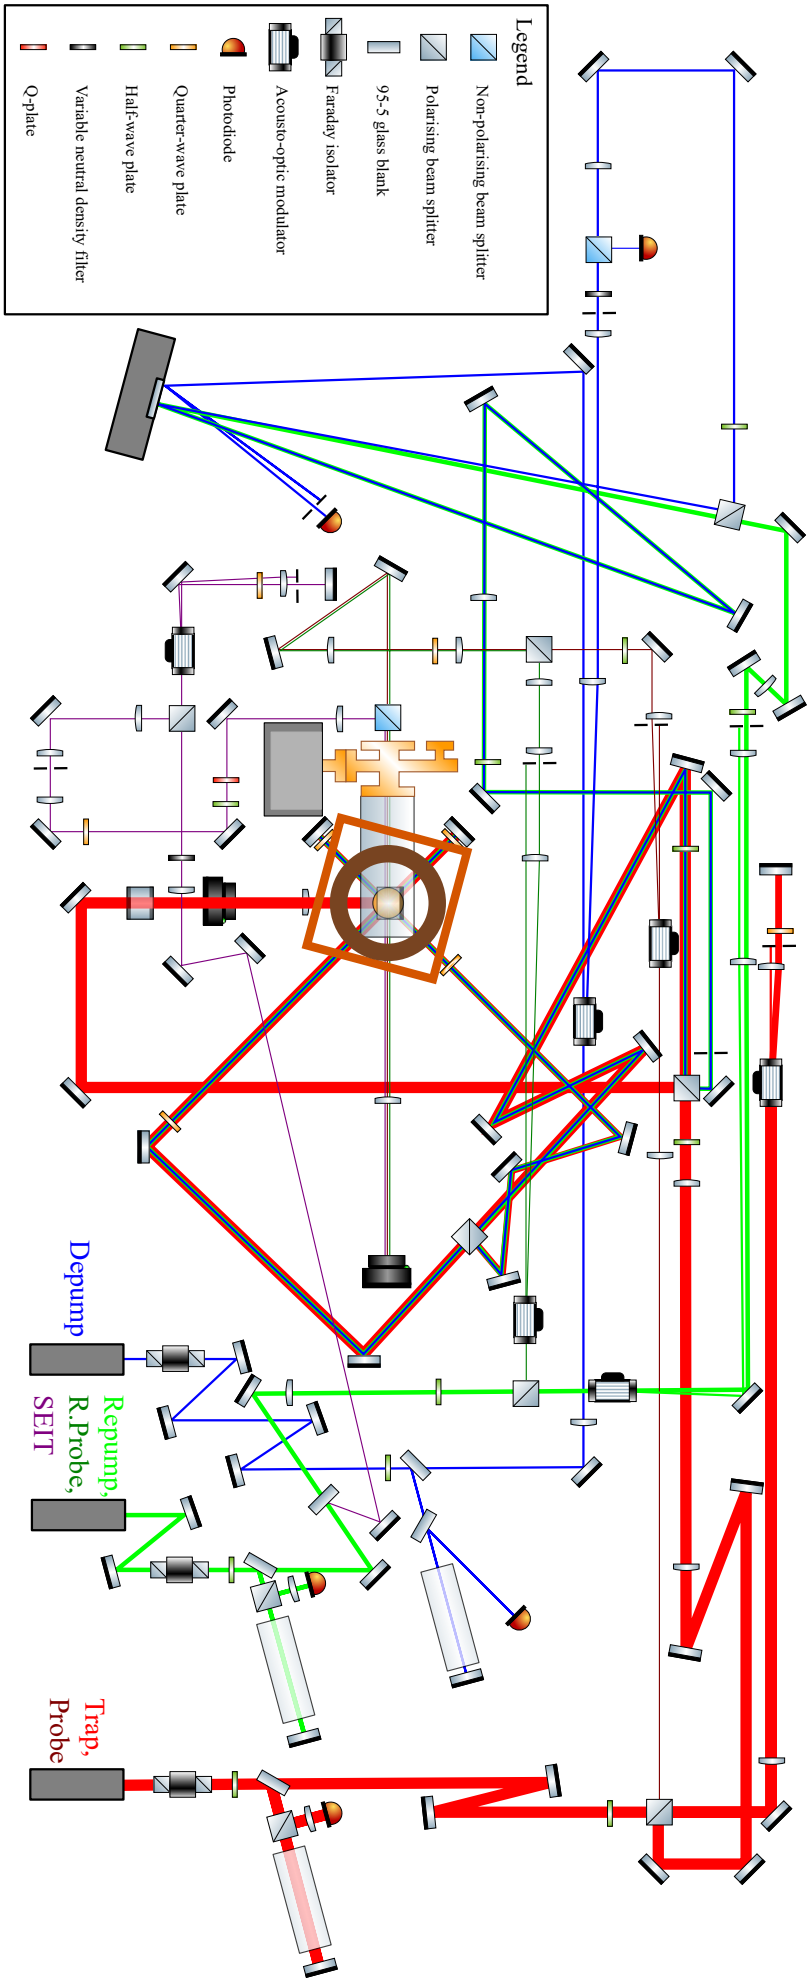

Figure 7.19: Detailed diagram of the cold atoms experiment, showing all optical elements at the time of writing.

Chapter 8

Experiment: atomic compass

8.1 Introduction

In this last chapter on our work with atoms we finally gather all our knowledge of structured light and atoms to investigate the interaction of atomic states in a modified Λ level configuration with light with structured polarisation. In particular, we are interested in the absorption (or lack of absorption) of resonant laser light in electromagnetically induced transparency (EIT). EIT is a quantum effect arising in atoms when two electronic states are coupled to a third state by two exciting light fields. Usually studied in terms of the detuning of one of these beams while the other is kept on resonance, under certain conditions the excitation amplitudes on the two transitions cancel such that the atoms become transparent to resonant radiation they would absorb if it was on its own, hence the name. In our system this picture is slightly modified because of the presence of a fourth atomic level. It turns out that we can incorporate this level to manipulate the EIT condition, making it dependent on the relative phase between the two driving laser fields. It has been shown [105] that EIT dependent on the local polarisation direction can be observed in such a system. It was observed that small magnetic fields transverse to the propagation direction of the laser light is required for this polarisation dependence to manifest. Here we investigate in detail the role of this magnetic field, and show that its direction influences the observed pattern of absorption produced by the atom-light interaction. The magnetic field dependence can be exploited to deduce the direction of a magnetic field from a single absorption image. In this sense the atoms act as a sort of compass needle.

Since we develop a technique for measuring the direction of magnetic fields it is sensible to give an overview of current technologies used for this. There are a wide variety of techniques used to measure magnetic fields that can be used as a compass, with various sensitivities and complexities, such as Hall effect, search coil, anisotropic magnetoresistance, SQUID and optically pumped magnetometers to name only a few. For a recent review of magnetic sensors see [106]. It is worth discussing the operation of optically pumped magnetometers (OPMs) since they exploit similar physics to what we will look at in section 8.4.

OPMs use warm alkali vapours to probe magnetic fields. A circularly polarised laser beam with a frequency matching a F $\rightarrow$ F'=F transition is directed through the vapour cell. Provided that any background magnetic field is not exactly perpendicular to the beam propagation direction this light drives σ_+ (or σ_-) transitions in the atoms, which drives the atoms into a dark state (similar

to pumping into a stretched state, as discussed in section 4.3). This is called a dark state because atoms in this state can no longer absorb photons from the laser because there is no m_F state in the excited state reachable by a σ_+ (or σ_-) transition. If a large population of atoms in the vapour can be maintained in such a state, for example by suppressing spin-exchange relaxation [107], the atomic medium becomes highly sensitive to magnetic fields because the atomic dipole precesses around the magnetic field vector, redistributing populations in the m_F states, destroying the dark state. The absorption of the laser beam through the cell in this a regime is then related to the strength of the magnetic field. Using laser amplitude modulation and lock-in techniques (e.g. Bell-Bloom method [108], analogous to the error signal generation in the dither laser lock in chapter 5) this dependence can be made linear within a narrow range of magnetic field strengths [109]. OPMs have exquisite sensitivity, achieving 1 pT/$\sqrt{\text{Hz}}$ [110]. Our system is not dissimilar to that used in OPMs, although we interpret it in different ways, considering EIT.

In this chapter we will first look at why EIT arises in the first place, studied in the common Λ system in terms of detuning. Then we introduce a fourth level present in the real rubidium system we use, and investigate how its presence affects absorption. We first do this by following the arguments of Thomas Clark[1] [111], then from a different perspective that is, in our opinion, simpler. We then show that recently obtained data agrees with the predictions of both theoretical approaches. Lastly we describe how we can deduce the direction of a magnetic field from an absorption image, and also how we can use multiple absorption images to very accurately cancel residual magnetic fields.

8.2 Electromagnetically induced transparency

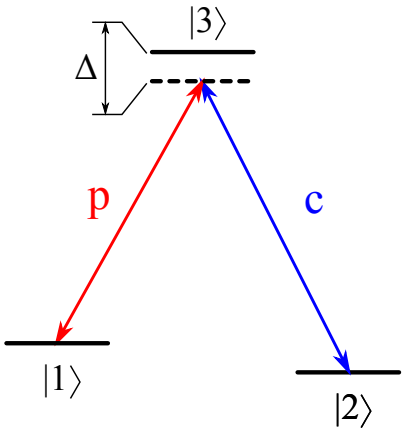

Figure 8.1: Λ-type EIT level scheme. The ground states are coupled to the excited state via optical transitions.

The following discussion roughly follows [112] and [113]. Consider a 3-level atom in a Λ configuration shown in figure 8.1. In this system there are two ground states ($|1\rangle$ and $|2\rangle$) coupled to a common excited state ($|3\rangle$) by two optical transitions (c for coupling and p for probe), but crucially there is no dipole-allowed transition between the ground states. Let us have two lasers driving these transitions with Rabi frequencies Ω_c and Ω_p respectively, and the probe beam detuned from resonance by a frequency Δ. We keep the control beam on resonance. We can describe the evolution of the states using a Hamiltonian, but of course in the presence of (near-)resonant radiation we need to augment the bare atomic Hamiltonian H_0 with a Hamiltonian describing the interaction, H_{int}: $H = H_0 + H_{int}$. Unfortunately the eigenstates of the bare atomic Hamiltonian, $|1\rangle$, $|2\rangle$ and $|3\rangle$, are not, in general, eigenstates of H. It makes sense then to find at the eigenstates of the interaction Hamiltonian. It is actually fairly straightforward to write down the H [114] in the basis $(|1\rangle, |2\rangle, |3\rangle)$,

$$H = -\frac{\hbar}{2}\begin{pmatrix} 0 & 0 & \Omega_p \\ 0 & 0 & \Omega_c \\ \Omega_p & \Omega_c & -2\Delta \end{pmatrix}, \tag{8.1}$$

where we have taken the rotating wave approximation as usual.

[1] Previously PhD student at Glasgow, at the time of writing postdoc at the Wigner Institute, Budapest.

The eigenstates of H are then

$$\begin{aligned} |+\rangle &= \sin\theta\sin\phi|1\rangle + \cos\theta\sin\phi|2\rangle + \cos\phi|3\rangle, \\ |-\rangle &= \sin\theta\cos\phi|1\rangle + \cos\theta\cos\phi|2\rangle - \sin\phi|3\rangle, \\ |0\rangle &= \cos\theta|1\rangle - \sin\theta|2\rangle, \end{aligned} \tag{8.2}$$

where we have defined the so-called mixing angles θ and ϕ, given by

$$\tan\theta = \frac{\Omega_p}{\Omega_c} \quad \text{and} \quad \tan\phi = \frac{\sqrt{\Omega_p^2 + \Omega_c^2}}{\sqrt{\Omega_p^2 + \Omega_c^2 + \Delta^2} + \Delta}. \tag{8.3}$$

These dressed states correspond to viewing the atomic states in a basis more appropriate in the presence of light, and we could express the bare states in terms of them. The interesting thing about these new states is that one of them, $|0\rangle$, does not couple to the atomic excited state $|3\rangle$. Since there are no optical transitions from $|0\rangle$ it is a dark state. The idea is that the coupling beam transfers the atoms into this dark state via spontaneous emission from $|3\rangle$, at which point the atoms cannot interact with the beams anymore, and they become transparent. Of course the presence of the probe beam is entirely necessary for this process. The definition of state $|0\rangle$ intimately involves the intensity of the probe beam through its Rabi frequency. The downside of the dressed state approach is that often it is quite difficult to intuitively interpret what the dressed states and transitions between them mean, since the dressed states mix the bare atomic states that are easy to understand.

The most intriguing aspects of EIT are revealed when we consider the phenomenon in terms of the detuning Δ. For the proper treatment we need to include spontaneous emission, which the above discussion does not include. We do this the same way as we did in chapter 4, through the density matrix approach. As always, the time evolution of the density matrix is given by equation 4.6. We have all the necessary components already with the Hamiltonian given in equation 8.1, so we can write down the optical Bloch equations for our 3-level system, adding in the decay terms γ as before,

$$\begin{aligned} \dot{\tilde{\rho}}_{11} &= -i\frac{\Omega_p}{2}(\tilde{\rho}_{13} - \tilde{\rho}_{31}) + \gamma_1\tilde{\rho}_{33}, \\ \dot{\tilde{\rho}}_{22} &= -i\frac{\Omega_c}{2}(\tilde{\rho}_{23} - \tilde{\rho}_{32}) + \gamma_{32}\tilde{\rho}_{33}, \\ \dot{\tilde{\rho}}_{33} &= i\frac{\Omega_c}{2}(\tilde{\rho}_{23} - \tilde{\rho}_{32}) + i\frac{\Omega_p}{2}(\tilde{\rho}_{13} - \tilde{\rho}_{31}) - (\gamma_1 + \gamma_2)\rho_{33}, \\ \dot{\tilde{\rho}}_{12} &= \dot{\tilde{\rho}}_{21}^* = (-\gamma_{12} + i\Delta)\tilde{\rho}_{12} - i\frac{\Omega_c}{2}\tilde{\rho}_{13} + i\frac{\Omega_p}{2}\tilde{\rho}_{32}, \\ \dot{\tilde{\rho}}_{13} &= \dot{\tilde{\rho}}_{31}^* = (-\gamma_{13} + i\Delta)\tilde{\rho}_{13} - i\frac{\Omega_c}{2}\tilde{\rho}_{12} + i\frac{\Omega_p}{2}(\tilde{\rho}_{33} - \tilde{\rho}_{11}), \\ \dot{\tilde{\rho}}_{23} &= \dot{\tilde{\rho}}_{32}^* = -\gamma_{23}\tilde{\rho}_{23} + i\frac{\Omega_c}{2}(\tilde{\rho}_{33} - \tilde{\rho}_{22}) - i\frac{\Omega_p}{2}\tilde{\rho}_{21}, \end{aligned} \tag{8.4}$$

where the off-diagonal decay rates γ_{ij} are given by $\gamma_{ij} = (\gamma_i + \gamma_j)/2$. These equations would be rather cumbersome to work with as is, but thankfully under certain conditions we can simplify things quite a bit. Looking at equations 8.2 and 8.3, we can see that if the probe intensity is much lower than that of the coupling beam then $\sin\theta \approx 0$, $\cos\theta \approx 1$ and so $|0\rangle \to |1\rangle$. So in such a configuration the bare atomic state $|1\rangle$ becomes the dark state and we expect that over time most of the population accumulates there. Thus $\rho_{11} \approx 1$ and $\rho_{22} = \rho_{33} \approx 0$. This allows us to write

equation 8.4 as

$$\begin{aligned}
\dot{\tilde{\rho}}_{12} &= (-\gamma_{12} + i\Delta)\tilde{\rho}_{12} - i\frac{\Omega_c}{2}\tilde{\rho}_{13} + i\frac{\Omega_p}{2}\tilde{\rho}_{32}, \\
\dot{\tilde{\rho}}_{13} &= (-\gamma_{13} + i\Delta)\tilde{\rho}_{13} - i\frac{\Omega_c}{2}\tilde{\rho}_{12} - i\frac{\Omega_p}{2}, \\
\dot{\tilde{\rho}}_{23} &= -\gamma_{23}\tilde{\rho}_{23} - i\frac{\Omega_p}{2}\tilde{\rho}_{21},
\end{aligned} \tag{8.5}$$

where we do not show the complex conjugates for simplicity. The advantage of these expressions is that now $\dot{\tilde{\rho}}_{23}$ depends on Ω_p linearly, so the term $i\frac{\Omega_p}{2}\tilde{\rho}_{32}$ in $\dot{\tilde{\rho}}_{12}$ is quadratic in Ω_p. This is great, because the low probe intensity assumption allows us to drop this very small term. This decouples $\dot{\tilde{\rho}}_{12}$ and $\dot{\tilde{\rho}}_{23}$, leaving us with a much simpler pair of coupled equations (since $\dot{\tilde{\rho}}_{23}$ is now independent of the other two equations),

$$\begin{aligned}
\dot{\tilde{\rho}}_{12} &= (-\gamma_{12} + i\Delta)\tilde{\rho}_{12} - i\frac{\Omega_c}{2}\tilde{\rho}_{13}, \\
\dot{\tilde{\rho}}_{13} &= (-\gamma_{13} + i\Delta)\tilde{\rho}_{13} - i\frac{\Omega_c}{2}\tilde{\rho}_{12} - i\frac{\Omega_p}{2}.
\end{aligned} \tag{8.6}$$

A symbolic algebra package like Mathematica has no trouble solving these in steady state, giving

$$\begin{aligned}
\tilde{\rho}_{12} &= -\frac{\Omega_c\Omega_p}{\Omega_c^2 - 4(i\gamma_{12} + \Delta)(i\gamma_{13} + \Delta)}, \\
\tilde{\rho}_{13} &= \frac{2(i\gamma_{12} + \Delta)\Omega_p}{-\Omega_c^2 + 4(i\gamma_{12} + \Delta)(i\gamma_{13} + \Delta)}.
\end{aligned} \tag{8.7}$$

This is all well and good, but density matrix elements are not directly measurable in an experiment. We need to relate them to observable parameters, such as absorption. We do this by noting that the complex linear susceptibility of an ensemble of atoms as perceived by a light field driving the $j \to k$ transition is proportional to $\tilde{\rho}_{jk}$. The imaginary part of the linear susceptibility is related to absorption and its real part to the refractive index of the ensemble at the driving frequency. The imaginary and real parts of $\tilde{\rho}_{13}$ are plotted figure 8.2 in the presence and absence of coupling light. The absorption and dispersion profiles follow these curves. Notice that absorption goes to zero on resonance. At the same time the large negative gradient of the dispersion leads to a dramatically increased refractive index of the atoms. This can be exploited to produce slow, or even stopped light[2] [115, 116, 117].

8.3 Spatially dependent EIT

All right, let us complicate things a bit more by adding one extra level into the mix. In our SpOT we have atoms trapped in the F=1 ground state, from which we can drive transitions to the F'=0 excited state. If we use circularly polarised light we can have a system very much like the one shown in figure 8.1, if we consider the $m_f = \mp 1$ levels to be $|1\rangle$ and $|2\rangle$, respectively. The $m_f = 0$ excited state corresponds to $|3\rangle$, shown in figure 8.3. The complication is that there is one more ground state, $m_f = 0$. Let us relabel these states as follows: F=1, $m_f = \mp 1, 0 \to |g_\mp\rangle, |g_0\rangle$ and F'=0, $m_f = 0 \to |e\rangle$.

The original reason for considering this system was to see if EIT can be made dependent on the optical phase difference between the two driving fields. In this section we will follow this logic, summarising the thinking outlined in [111]. An easy way to see this phase difference dependence in an experiment would be to give one or both of the light fields a spatially varying phase. Then EIT should vary in space, which would be visible in an absorption image, hence the name spatially

[2] Although in a sense on interaction with the atoms the light and atoms stop being separate things.

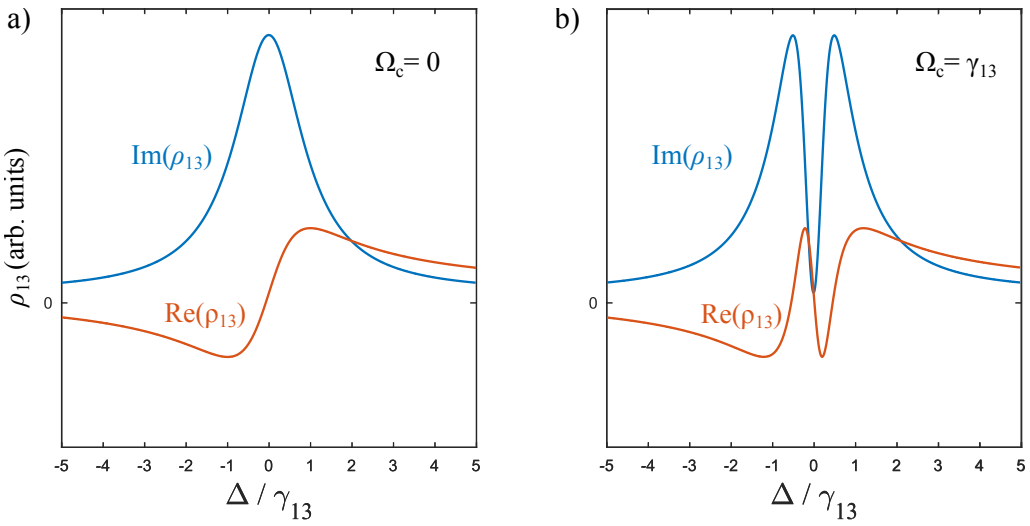

Figure 8.2: Absorption (proportional to $\mathrm{Im}(\rho_{13})$) and dispersion (proportional to $\mathrm{Re}(\rho_{13})$) curves (a) in the absence of coupling light and (b) in EIT.

dependent EIT (SEIT). But for the phase dependence to manifest in this atomic system the transitions involved should form a closed loop [118]. Therefore we need to provide a coupling between $|g_{\mp}\rangle$, which can be achieved using magnetic fields. Specifically, any magnetic field transverse to the quantisation axis set by the two beams' propagation direction will introduce magnetic coupling between all ground states $|g\rangle$. This is shown by green arrows in figure 8.3.

This complicates the interaction Hamiltonian a bit. It now has two components, corresponding to the electric and magnetic interactions:

$$\hat{H}_{int} = \hat{H}_E + \hat{H}_B. \tag{8.8}$$

Let us define the magnetic field in spherical coordinates, pointing in some direction relative to the quantisation axis $\hat{z}$, $\mathbf{B} = B_0(\cos\theta_B\hat{z} + \sin\theta_B\cos\phi_B\hat{x} + \sin\theta_B\sin\phi_B\hat{y})$. This magnetic field then contributes to the total interaction Hamiltonian via

$$H_B = \frac{\hbar}{2}\begin{pmatrix} \Omega_L\cos\theta_B & -e^{+i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 & 0 \\ e^{-i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 & -e^{+i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 \\ 0 & -e^{-i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & -\Omega_L\cos\theta_B & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{8.9}$$

where we have defined the Larmor frequency $\Omega_L = g_F\mu_B B_0$, with $g_F = -1/2$ the hyperfine Landé g-factor and μ_B the Bohr magneton. The basis here is $(|g_-\rangle, |g_0\rangle, |g_+\rangle, |e\rangle)$. This matrix incorporates two effects of the magnetic field. The energies of $|g_{\mp}\rangle$ are Zeeman shifted due to the magnetic field component along the quantisation axis, captured in the diagonal elements. The off-diagonal elements indicate a new coupling between $|g_0\rangle$ and $|g_{\mp}\rangle$ due to the transverse component of the magnetic field. The direction of the transverse component breaks the cylindrical symmetry set up by the $\hat{z}$ quantisation direction, defining an azimuthal direction $\phi_B = 0$.

The light fields are also a bit different from the ones considered so far in this chapter, in that they acquire a spatially varying phase. We use light fields described by

$$\mathbf{E}(r,\phi) = \frac{E_0(r)}{\sqrt{2}}\left(\hat{l}e^{-i\ell\phi} + \hat{r}e^{+i\ell\phi}\right), \tag{8.10}$$

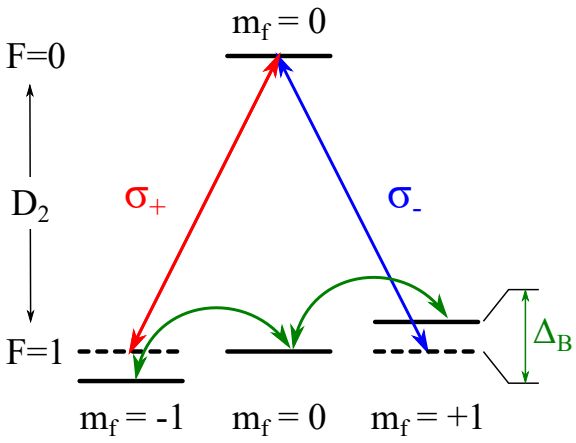

Figure 8.3: EIT level scheme on the ^{87}Rb D_2 line, with the quantisation axis along the beam direction. The $m_f = \mp 1$ ground states are coupled to the $m_f = 0$ excited state via optical transitions. In the presence of a transverse magnetic field they are also coupled to the $m_f = 0$ ground state magnetically. The central ground state is not optically coupled.

which can easily be generated using Q-plates (see chapter 1). Such beams comprise two orthogonal circularly polarised components corresponding to $\sigma_\pm$, each with $\pm\ell\hbar$ orbital angular momentum, represented by the azimuthal phase $\exp \pm i\ell\phi$. An alternative way of interpreting these beams is that they are linearly polarised, but the polarisation direction rotates azimuthally. For example, the $\ell = 1$ Q-plate beam is radially polarised. The intensity of these beams is allowed to vary radially, and in reality it does vary quite a bit, for simplicity we will not take this into account in the following, restricting our analysis to some $r = r_0$. With this in mind, the Hamiltonian for the electric interaction can be written, similarly to before but including the azimuthal phase, as

$$H_E = \frac{\hbar}{2}\begin{pmatrix} 0 & 0 & 0 & \frac{1}{\sqrt{6}}e^{i\ell\phi}\Omega \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{\sqrt{6}}e^{-i\ell\phi}\Omega \\ \frac{1}{\sqrt{6}}e^{-i\ell\phi}\Omega & 0 & \frac{1}{\sqrt{6}}e^{i\ell\phi}\Omega & 0 \end{pmatrix}, \tag{8.11}$$

where we have assumed that both beams are resonant and have the same intensity (hence Rabi frequency Ω). The factors of $1/\sqrt{6} = 1/\sqrt{2F_{\text{initial}}+1}$ are geometrical coefficient coming from the Wigner-Eckart theorem, arising from the orientation of the atomic dipole. We can now write down the total interaction Hamiltonian,

$$H_{int} = \frac{\hbar}{2}\begin{pmatrix} \Omega_L\cos\theta_B & -e^{+i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 & \frac{1}{\sqrt{6}}e^{i\ell\phi}\Omega \\ e^{-i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 & -e^{+i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & 0 \\ 0 & -e^{-i\phi_B}\frac{1}{\sqrt{2}}\Omega_L\sin\theta_B & -\Omega_L\cos\theta_B & \frac{1}{\sqrt{6}}e^{-i\ell\phi}\Omega \\ \frac{1}{\sqrt{6}}e^{-i\ell\phi}\Omega & 0 & \frac{1}{\sqrt{6}}e^{i\ell\phi}\Omega & 0 \end{pmatrix}. \tag{8.12}$$

This is a rather complicated Hamiltonian, which is difficult to manipulate. The associated Bloch equations can be solved numerically relatively painlessly, which we did using the AtomicDensityMatrix Mathematica package[3] and we show some results from such a solution at the end of this section. But it is possible to simplify the Hamiltonian by a clever choice of basis, allowing some level of analytic solution. We can construct a system in which new levels, built from the bare atomic states, are coupled by the optical and magnetic interactions in a sort of cascade. Let us define the following new levels:

$$\begin{aligned} |g_{C,NC}\rangle &= \frac{1}{\sqrt{2}}\left(e^{i\ell\phi}|g_-\rangle \mp e^{-i\ell\phi}|g_+\rangle\right), \\ |\Psi_0\rangle &= \frac{1}{N(\phi)}\left(\cos\theta_B|g_{NC}\rangle + \sin\theta_B\sin(\ell\phi-\phi_B)|g_0\rangle\right), \\ |\Psi_{NC}\rangle &= \frac{1}{N(\phi)}\left(i\sin\theta_B\sin(\ell\phi-\phi_B)|g_{NC}\rangle - \cos\theta_B|g_0\rangle\right), \end{aligned} \tag{8.13}$$

with the normalisation done by $N(\phi) = \sqrt{\cos^2\theta_B - \sin^2\theta_B\cos^2(\ell\phi-\phi_B)}$. Using this basis, the

[3] http://rochesterscientific.com/ADM/

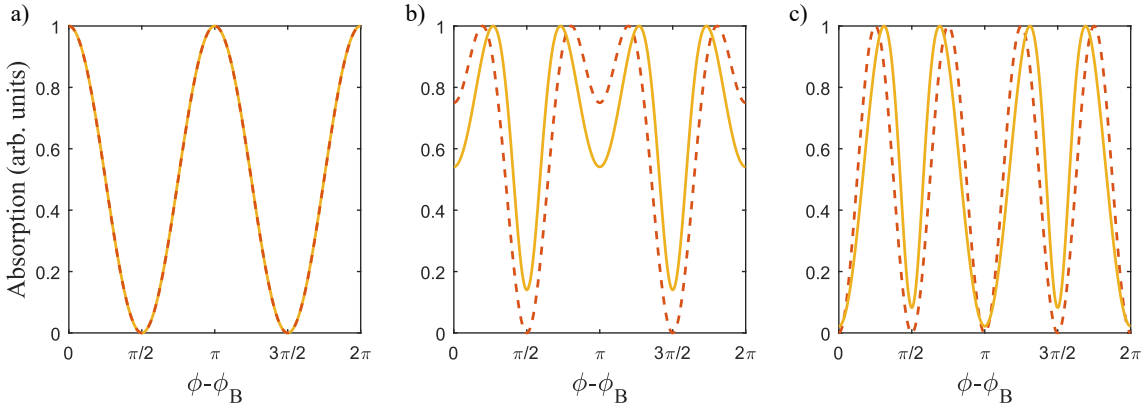

Figure 8.5: Comparison of SEIT patterns predicted by numerical solutions of the optical Bloch equations (yellow) and Fermi's golden rule (orange, dashed) for $\ell = 1$ at magnetic field inclination angle (a) $\theta_B = \pi/36$, (b) $\theta_B = \pi/3$ and (c) $\theta_B = \pi/2$. The Bloch equations were evaluated with experimentally relevant parameters, $B_0 = 0.2$ G, $I_{tot} = 200$ mW/cm^2. Curves are individually peak-normalised.

interaction Hamiltonian can be written in a simpler form in the new basis $(|\Psi_{NC}\rangle, |\Psi_0\rangle, |g_C\rangle, |e\rangle)$,

$$H_E = \frac{\hbar}{2}\begin{pmatrix} 0 & i\Omega_L \sin\theta_B \cos(\ell\phi - \phi_B) & 0 & 0 \\ -i\Omega_L \sin\theta_B \cos(\ell\phi - \phi_B) & 0 & -\Omega_L N(\phi) & 0 \\ 0 & -\Omega_L N(\phi) & 0 & \frac{1}{2\sqrt{3}}\Omega \\ 0 & 0 & \frac{1}{2\sqrt{3}}\Omega & 0 \end{pmatrix}. \tag{8.14}$$

In this basis the states are pairwise coupled. The excited state is optically coupled to the coupling state $|g_C\rangle$. This in turn is magnetically coupled to the intermediate state $|\Psi_0\rangle$, which is also magnetically coupled to the (sometimes) non-coupling state $|\Psi_{NC}\rangle$, forming the ladder system shown in figure 8.4. Notice how the coupling between $|\Psi_0\rangle$ and $|\Psi_{NC}\rangle$ turns off at certain angles $\phi - \phi_B = n\pi/\ell$, $n \in \mathbb{N}$. At these angles $|\Psi_{NC}\rangle$ is a dark state which allows EIT, and we can see that the phase dependence between the two circular polarisation components of the driving light sets whether EIT occurs or not. Since this relative phase is space dependent, EIT also follows this.

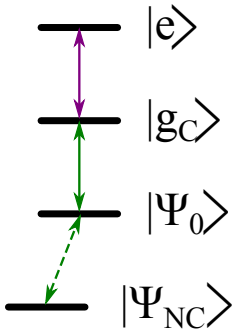

Figure 8.4: Ladder basis for SEIT.

For the weak driving fields considered here we can estimate the absorption using Fermi's golden rule, since the excited state population then remains negligible and so the absorption is proportional to the probability of excitation from the sometimes non-coupling state,

$$T_{\Psi_{NC}\to e} \propto \left(\frac{2\pi}{\hbar}\right)^3 |\frac{\hbar}{4\sqrt{3}}\Omega|^2 |\frac{1}{2}\Omega_L|^4 |\sin\theta_B \cos(\ell\phi - \phi_B)|^2 |\frac{1}{2}\Omega_L N(\phi)|^2, \tag{8.15}$$

which can be written, after expanding $N(\phi)$, as

$$T_{\Psi_{NC}\to e} \propto \frac{\hbar\pi^3}{96}\Omega_L^4 \Omega_R^2 \sin^2\theta_B \cos^2(\ell\phi - \phi_B)\left(\cos^2\theta_B + \sin^2\theta_B \sin^2(\ell\phi - \phi_B)\right). \tag{8.16}$$

The azimuthal variation of absorption is shown in figure 8.5 at different magnetic field inclination angles θ_B. At small θ_B we see a (squared) sinusoidal variation in transmission, with 2ℓ lobes. Strikingly, however, at larger θ_B the lobes of EIT split, transitioning to 4ℓ lobes at $\theta_B = \pi/2$. We will get a more intuitive picture in the next section as to why this happens.

The absorption pattern also shifts with ϕ_B, rotating from one lobe to the next in $\Delta\phi_B = \pi$. This is very easy to see from equation 8.16, where ϕ_B only appears alongside $\ell\phi$, and can be understood as the azimuthal angle of the magnetic field being the origin of the azimuthal component of the coordinate system. This is because in this picture it is the magnetic field that breaks the cylindrical symmetry.

Figure 8.5 also shows the absorption patterns predicted by a numerical solution of the optical Bloch equations following from equation 8.12. In this picture we consider the absorption to be proportional to the excited state population, since the only way to populate the excited state is via the absorption of photons. We can see that the general, qualitative behaviour is very similar: 2ℓ lobes at small θ_B and 4ℓ lobes at $\theta_B = \pi/2$. At intermediate values of θ_B the absorption patterns are similar in character, but the depth of the lobe splitting is slightly different. This effect is sensitive to beam intensity; the figure shows results calculated for the intensity measured in the experiment. This result highlights the fact the Fermi's golden rule is only applicable when certain assumptions, mainly low beam intensity, are met.

8.4 SEIT: a simpler model

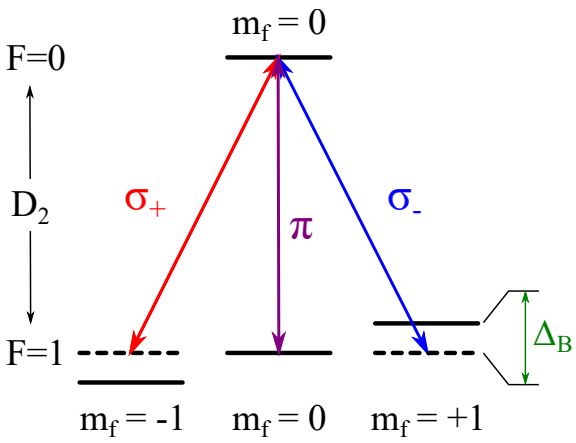

Figure 8.6: EIT level scheme on the ^{87}Rb D_2 line, with the quantisation axis along the magnetic field. The $m_f = \mp 1$ ground states are coupled to the $m_f = 0$ excited state via $\sigma_+ + \sigma_-$ transitions and the $m_f = 0$ ground state is coupled to the excited state via a π transition.

Another way to look at the SEIT process involves a different choice of quantisation axis. Instead of the probe direction $\hat{z}$, we can also choose the magnetic field direction[4]. Again, the magnetic field is given by

$$\mathbf{B} = B_0 \begin{pmatrix} \sin\theta_B \cos\phi_B \\ \sin\theta_B \sin\phi_B \\ \cos\theta_B \end{pmatrix}. \tag{8.17}$$

The electric field vector of the vector vortex beams we are considering here point in a direction given by

$$\mathbf{E}(r,\phi,z) = E_0 \begin{pmatrix} \cos\ell\phi \\ \sin\ell\phi \\ 0 \end{pmatrix}, \tag{8.18}$$

which is all in a plane transverse to the propagation of the beam. If our quantisation axis follows the magnetic field, however, the electric field does not necessarily oscillate in a transverse plane to the quantisation axis. It is well-known that linear polarisation transverse to the quantisation axis drives superpositions of σ_+ and σ_- transitions, just like we had in the previous section. The component of the electric field along the quantisation axis, however, drives a π transition, for which $\Delta m_f = 0$. The intensities of the π and σ_+ and σ_- superposition driving fields can be calculated from the parallel and perpendicular components of the electric field to the magnetic field,

$$\begin{aligned} I_\parallel &= \frac{1}{2}c\epsilon_0 \left(\frac{\mathbf{E}\cdot\mathbf{B}}{B_0}\right)^2 = \frac{1}{2}c\epsilon_0 E_0^2 \cos^2(\ell\phi - \phi_B)\sin^2\theta_B, \\ I_\perp &= \frac{1}{2}c\epsilon_0 \left(E_0^2 - \left(\frac{\mathbf{E}\cdot\mathbf{B}}{B_0}\right)^2\right) = \frac{1}{2}c\epsilon_0 E_0^2 \left(1 - \cos^2(\ell\phi - \phi_B)\sin^2\theta_B\right). \end{aligned} \tag{8.19}$$

[4] In fact we could choose any direction we like, but these two are probably the most obvious options.

The atomic states are very simple when viewed along this quantisation axis, because the ground states are not coupled by magnetic fields, shown in figure 8.6. We can use the bare states, although it makes sense to make a small modification. Since the magnetic field is now along the quantisation axis, it does not mix the Zeeman sublevels, but it does introduce a Zeeman shift. However, if the laser is tuned to resonance with the un-shifted levels, then the Zeeman shifts are equal and opposite for the $m_f = \pm 1$ states, so the effective laser detuning is the same for both of them, although with a different sign, but that will not matter for our discussion. Then it is perhaps more sensible to consider their superposition states, $|+\rangle = (1/\sqrt{2})(|m_f = 1\rangle + |m_f = -1\rangle)$ and $|-\rangle = (1/\sqrt{2})(|m_f = 1\rangle - |m_f = -1\rangle)$. $|+\rangle$ is driven by the $\sigma_+ + \sigma_-$ light, whereas $|-\rangle$ is not accessible by spontaneous emission, which we consider to be by far the dominant decay process from the excited state, so we assume that its population is 0. These states should be familiar from the previous section, where they were called the coupling and non-coupling ground states, $|g_{C,NC}\rangle$.

Then we can write down rate equations for the populations in the atomic states. We calculate the Rabi frequency as [59]

$$\begin{aligned} \Omega_{\parallel} &= \frac{\Gamma}{2}\sqrt{\frac{I_{\parallel}}{I_{sat}}} = \frac{\mathbf{E}\cdot\mathbf{d}}{\hbar}\cos^2(\ell\phi - \phi_B)\sin^2\theta_B \\ &\text{and} \\ \Omega_{\perp} &= \frac{\Gamma}{2}\sqrt{\frac{I_{\perp}}{I_{sat}}} = \frac{\mathbf{E}\cdot\mathbf{d}}{\hbar}(1 - \cos^2(\ell\phi - \phi_B)\sin^2\theta_B), \end{aligned} \tag{8.20}$$

with $I_{sat} = c\epsilon_0\Gamma^2\hbar^2/(4|\mathbf{E}\cdot\mathbf{d}/E_0|^2)$. The rate equations for the populations P_i in this system are

$$\begin{aligned} \dot{P}_E &= \Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B P_0 + (\Omega_0(1 - \cos^2(\ell\phi - \phi_B)\sin^2\theta_B) + \Omega_L)P_+ - \Gamma P_E, \\ \dot{P}_0 &= -\Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B P_0 + \frac{1}{3}\Gamma P_E, \\ \dot{P}_+ &= -\left(\Omega_0(1 - \cos^2(\ell\phi - \phi_B)\sin^2\theta_B) + \Omega_L\right)P_+ + \frac{2}{3}\Gamma P_E, \\ 1 &= P_E + P_0 + P_+ \end{aligned} \tag{8.21}$$

where we have defined $\Omega_0 = \mathbf{E}\cdot\mathbf{d}/\hbar$, and the Larmor frequency Ω_L represents the detuning due to Zeeman shift for $|+\rangle$. The steady-state solutions for the populations are rather involved:

$$\begin{aligned} P_0 &= \left(1 + \frac{2\Omega_0}{-\Omega_0 + (\Omega_0 + \Omega_L)\csc^2\theta_B\sec^2(\ell\phi - \phi_B)} + \frac{3\Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B}{\Gamma}\right)^{-1}, \\ P_+ &= \frac{2\Omega_0\Gamma}{(\Omega_0 + \Omega_L)\Gamma\csc^2\theta_B\sec^2(\ell\phi - \phi_B) + \Omega_0(3(\Omega_0 + \Omega_L) + \Gamma - 3\Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B}, \\ P_E &= \frac{-3\Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B(-\Omega_L - \Omega_0 + \Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B)}{(\Omega_0 + \Omega_L)\Gamma + \Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B(3(\Omega_0 + \Omega_L) + \Gamma - 3\Omega_0\cos^2(\ell\phi - \phi_B)\sin^2\theta_B)}. \end{aligned} \tag{8.22}$$

We can relate the absorption to these populations through noting that the only way to get from either ground state to the excited state is via the absorption of a photon. Hence the absorption is proportional to the excited state population P_E, although the excited state population itself is very small; the maximum excited state population is approximately 0.02 for a reasonable choice of parameters ($\Gamma \approx 10\Omega_0 \approx 100\Omega_L$). Figure 8.7 compares the predictions of this result with the absorption obtained using Fermi's Golden Rule in the previous section. The two results are extremely similar in appearance, and they both produce the expected behaviour with changing magnetic field direction. This is shown explicitly only for the peak splitting, since the rotation is trivial.

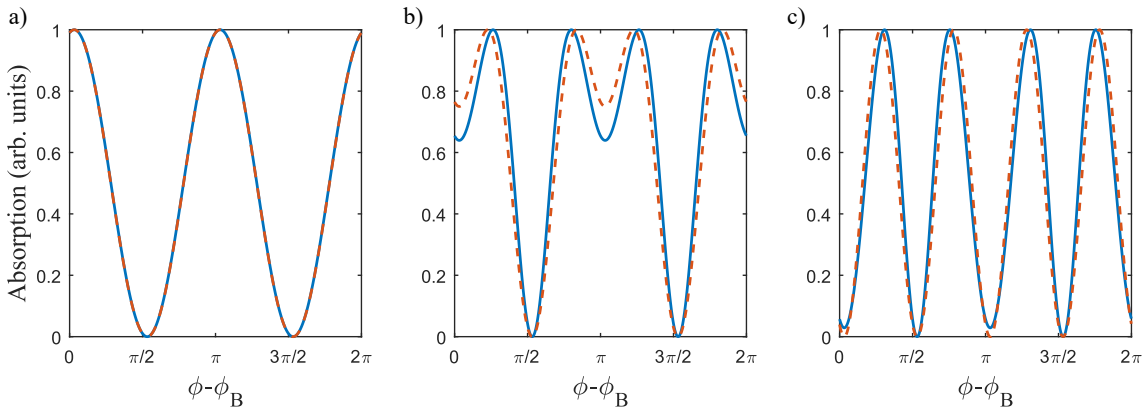

Figure 8.7: Comparison of SEIT patterns predicted by rate equations (blue) and Fermi's Golden Rule (orange, dashed) for $\ell = 1$ at magnetic field inclination angle (a) $\theta_B = \pi/36$, (b) $\theta_B = \pi/3$ and (c) $\theta_B = \pi/2$. In these simulations the magnetic field strength was small, $\Omega_L = 0.1\Omega_0$, and the decay was fast, $\Gamma = 100\Omega_0$.

The two models are in nearly perfect agreement at small and large magnetic field inclination angles θ. The slightly different splitting depth seen in figure 8.7 b) for intermediate θs is not entirely unexpected, since the very same behaviour can be seen if we numerically solve the full optical Bloch equations for this system at intermediate probe intensities ($I < I_{sat}$ but not $I \ll I_{sat}$). A comparison of such numerical solutions with the absorption obtained from Fermi's Golden Rule was shown in figure 8.5. We see that the absorption pattern obtained from the excited state population in the numerical solutions is again very similar to the pattern obtained from the rate equation model.

This investigation suggests that all of these methods are valid ways of analysing SEIT. They are perhaps useful in different situations. I think it is easiest to think about the situation when the quantisation axis follows the magnetic field, since in this case we can think about populations in the bare atomic states. It is easy to see how the dark states arise, and how they are populated via optical pumping. In this picture it is easier to think of the SEIT as polarisation dependent absorption. The populations in this picture are analytic, so can be used in fast calculations.

The Fermi's golden rule picture is also analytic, and its predictions are remarkable in their simplicity. It is worth noting that the Fourier transform of the absorption profile predicted by Fermi's golden rule is analytic. As we will see in the next section, this gives very simple and reliable predictions for the experiments so it is preferable for investigating the experimental results. While the absorption profile obtained from rate equations can also be Fourier transformed analytically, the result is too complicated to be of any practical use[5]. I would not use the full optical Bloch equations, because its solution is cumbersome, but in our investigations it, too, gave similar results to the other two methods.

8.5 Experimental results

As we discussed above, the four-level tripod shown in figures 8.3 and 8.6 is realised in rubidium 87 on the F=1→F'=0 transition on the D_2 line. Our SpOT collects cold atoms in the ground state of this transition, which is excellent since we can ignore the Doppler effect for the probing, because its effects are minimal. We generate optical fields much like those in equation 8.18 using Q-plates introduced in chapter 2. The Q-plate and a lens are placed such that the atoms are in the far field

[5]It took Mathematica about an hour to come up with the transform, and it is far too complicated for me to interpret in any meaningful way.

of the Q-plate, because in this case the beam from the plate becomes a good approximation of a superpositon of circularly polarised Laguerre-Gaussian beams with $\ell = \pm 2q$ due to diffraction. This helps clean up the phase as well. This far field plane is imaged onto a camera by a single lens. The complete experiment setup can be seen in figure 7.19 at the end of the previous chapter.

The magnetic field used to set the quantisation direction is controlled by the coils responsible for compensating background magnetic fields. On their own these coils only allow the quantisation axis to point in $\pi/2$ sr solid angle defined by their axes and the direction of current flow through them. In the experiment we had no way of reversing the current flow, which causes this limitation. However, because the coils were designed to work against the background magnetic field, if we turn them off we get a field pointing in a direction opposite to what the coils can generate. Therefore as long as we use small enough fields[6] we can exploit the background field to give us access to full 4π sr of pointing direction. This gives us all the tools we need to experimentally investigate the models derived in sections 8.3 and 8.4.

We performed absorption imaging (see chapter 7) using the Q-plate beam as the probe at different quantisation field angles, using a magnetic field of 0.2 G. The atom cloud was allowed to expand for 2 ms to grow to a large enough size to encompass the Q-plate beam. This also reduced the density of the cloud to a point where the absorption would not be saturated by the atoms. If the cloud was too dense parts of the beam which had a low probability to be absorbed by a single atom would still be absorbed, which would change the observed absorption pattern compared to what we would expect from theory. A sample set of absorption images are shown in figure 8.8, showing the splitting and rotation of the EIT lobes. These absorption images are extremely noisy where the beam has no appreciable intensity because of the division in equation 7.3. This is not a problem for the analysis of the absorption.

During the analysis absorption images were azimuthally unwrapped around the beam centre. We found the beam centre by assuming that the outer edge of the beam was circular due to diffraction from a pinhole in the beam path. The central vortex of the beams was not useful for determining the beam centre, because the Q-plates did not work very well near their centre so the central dark region was rather asymmetric. We then took the Fourier transform of the azimuthal variation of absorption at a radius of choice, usually at the peak intensity. The 2ℓ and 4ℓ components of the complex Fourier transform contain information about the quantisation magnetic field direction, and we can relate them easily to the absorption function given by Fermi's golden rule. The Fourier transform is calculated along ϕ, $\mathcal{F}(T(\phi)) = t(\Phi)$. The argument of the 2ℓ is immediately related to ϕ_B via

$$\arg(\mathcal{F}(\Phi = 2\ell)) = 2\phi_B, \tag{8.23}$$

where the factor of 2 comes from the symmetry of the specific beams we used, i.e. the fact that all specific linear polarisations appear an even number of times in the beam cross section.

The splitting of the EIT lobes can be seen in the Fourier transform as the appearance of a 4ℓ component as a function of the inclination angle θ_B,

$$|\mathcal{F}(\Phi = 4\ell)| = \frac{1}{4}\sqrt{\frac{\pi}{2}}\sin^4\theta_B. \tag{8.24}$$

It could be interesting to look at the variation of the 2ℓ component with θ_B as well. This is given

[6]The magnitude of the desired magnetic field is less than that of the background magnetic field.

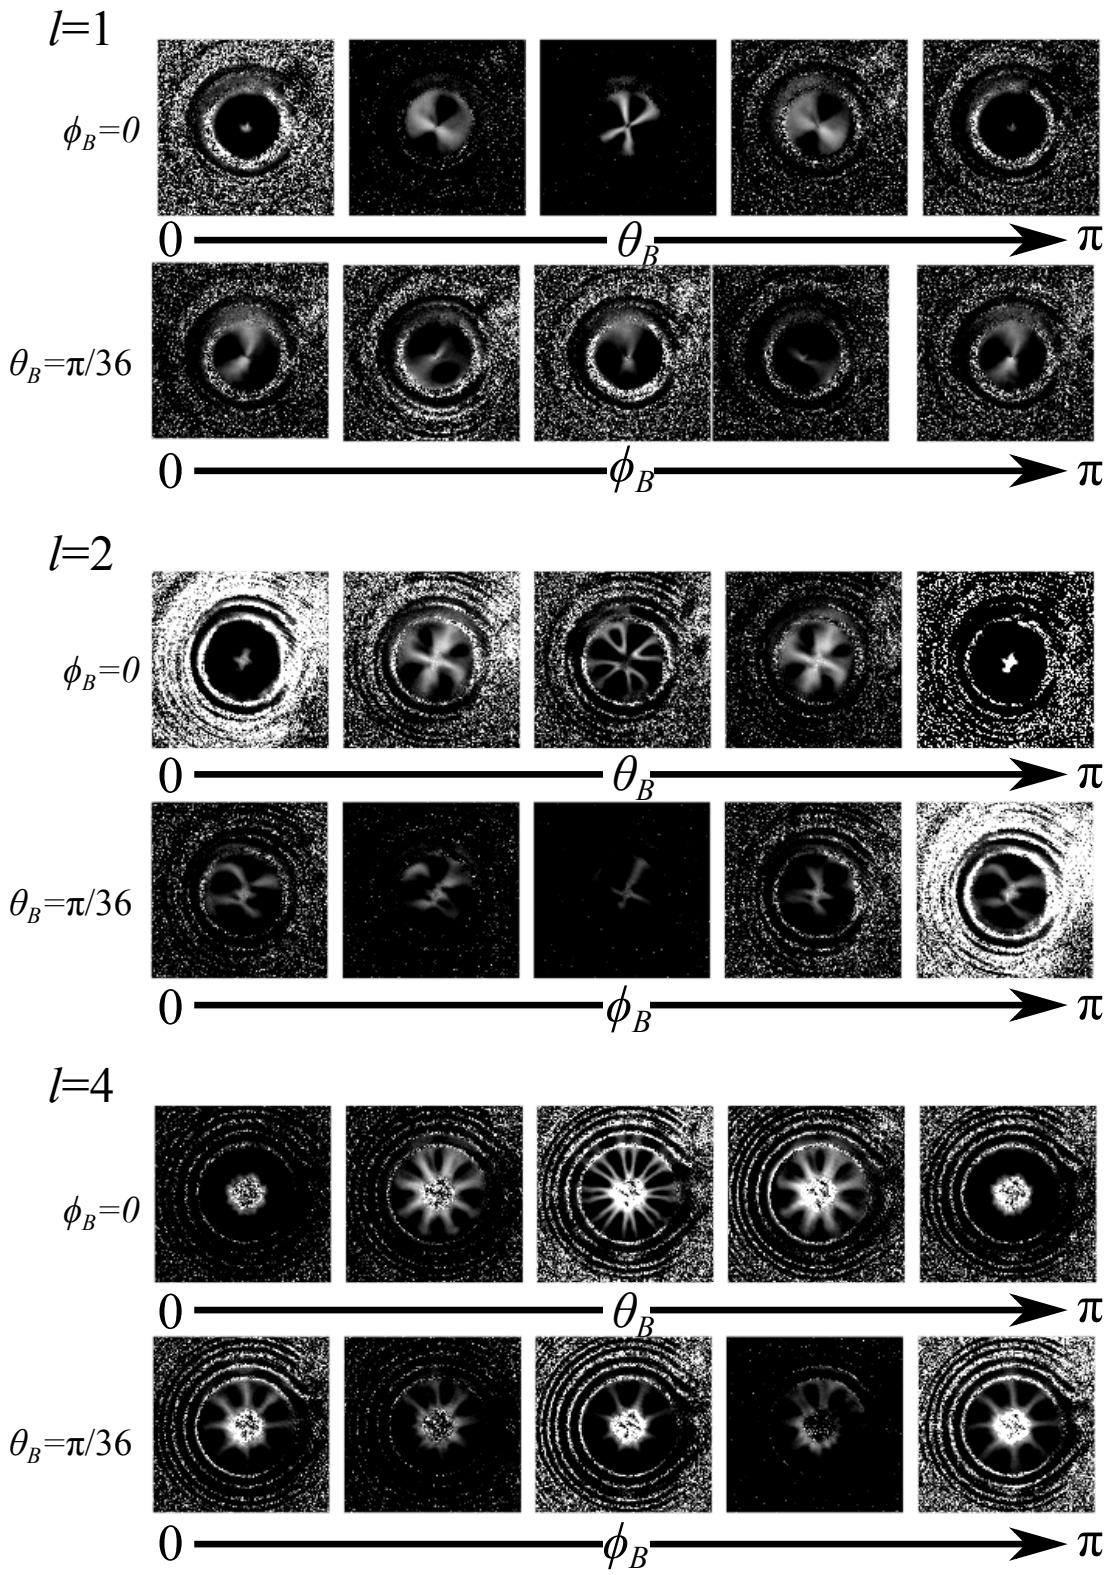

Figure 8.8: SEIT absorption images, demonstrating the rotation and splitting of the EIT lobes with magnetic field angles θ_B and ϕ_B. The images have been thresholded to enhance the visibility of the absorption lobes.

by a very similar expression,

$$|\mathcal{F}(\Phi = 2\ell)| = \sqrt{\frac{\pi}{2}} \cos^2\theta_B \sin^2\theta_B. \tag{8.25}$$

This has twice the frequency of the 4ℓ component amplitude, being zero for $\theta_B = 0$ and $\theta_B = \pi/2$ and maximum for $\theta_B = \pi/4$ and $\theta_B = 3\pi/4$.

Figure 8.9 shows the argument of the 2ℓ and the amplitude of the 4ℓ Fourier components, both for data and the above expressions. The data presented in figure 8.9b) were all normalised to the $\ell = 1$ data set. Error bars are only shown for the $\ell = 1$ case, and we assume the uncertainty for the other beams is comparable. We calculated the uncertainty by repeating each experiment 10 times and calculating the standard deviation of the calculated Fourier components. Note that the ϕ_B-dependence was negative because of a mistake in the definition of the coordinate system used to generate magnetic fields in the experiment; the x-axis was reversed. We take this into account when calculating the model dependence.

In any case, the data and theory are generally in very good agreement. A small discrepancy can be seen in both the 2ℓ and the 4ℓ components. The arguments of the 2ℓ component were consistently too large for $\phi_B < \pi/2$ and too small (although by a smaller amount) for $\phi_B > \pi/2$. The amplitude of the 4ℓ component changed slightly with ϕ_B as shown in the inset in figure 8.9a), which was not expected. We attribute these phenomena to either a slightly incorrect background field compensation, or imprecisely generated magnetic fields. A small offset magnetic field would displace the origin of the coordinate system for the quantisation axis rotation, which can explain the systematic discrepancies between theory and data. The magnetic field generated was never measured, we only measured the coil currents from which we could calculate the magnetic field generated. However, given experimental uncertainties in the dimensions and alignment of the coils it is possible that the actually generated fields differed somewhat from the intended fields, but this was not measured.

It is possible to introduce further uncertainty to the Fourier components when unwrapping the absorption images. If the centre around which the unwrapping is performed is chosen slightly off the actual beam axis, or the pattern is not circularly symmetric, then the unwrapped profile will have additional low frequency components not due to the absorption pattern. However, we expect that these effects would have a diminishing effect with spatial frequency. In any case, care was taken to find the centre for unwrapping by iteratively choosing a range of centre points and making sure that the unwrapped profile was as azimuthally symmetric as possible (which assumes that the beams themselves were azimuthally symmetric). Image noise can be neglected as a source of uncertainty, because it is at very high spatial frequencies and so does not affect the low Fourier components we are concerned with here.

We can also see that for larger values of ℓ the results were somewhat less in agreement with theory, especially in the case of the amplitudes of the 4ℓ component. There were a couple of reasons for this. Firstly, beams with higher ℓ are larger in the far field, meaning that they overlapped with fewer atoms at the edge of the cloud leading to lower signal-to-noise ratio. To counteract this we could expand the cloud for longer, but this lead to other issues. During a longer expansion the cloud slightly shifts under gravity, meaning that if the beam was centered on the trap location the absorption images would be slightly less symmetric, which leads to undesired variations in the Fourier components, distorting the results. A more expanded cloud is also less dense, although it is at least more uniformly so. Technically these effects could all be compensated with careful beam

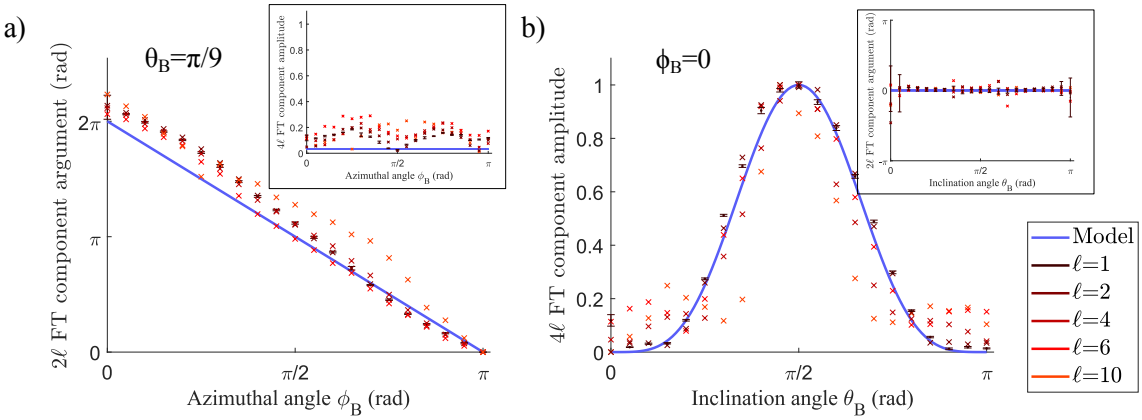

Figure 8.9: Fourier transform components of SEIT as a function of magnetic field angles. Model curves were obtained from Fermi's golden rule.

alignment and by adjusting the beam power to account for fewer atoms in a larger cloud or at the edge of the cloud. However, we decided that such adjustments between data sets would make their comparison more difficult because of the new free parameters, so we performed all experiments with the same expansion and beam power. Note that keeping beam power constant did change the intensity of the beams, as larger beams had slightly less intensity.

Another reason for the data for different ℓ to look somewhat different is that the Q-plates we used were, at this point, quite old, over 4 years old. Q-plates degrade over time, their phase retardation changes. As a result, some of the beams we used were better representations of equation 8.18 than others. This can be seen in figure 8.8. For example, the $\ell = 2$ beam looks markedly worse than the others in the figure, with a strange radial variation in absorption, and non-uniform azimuthal variation. Even so, the corresponding Fourier components behaved remarkably similar to theory, which demonstrates that the Fourier method of analysis is rather robust. This is likely because we specifically analyse spatial frequency components of interest, and so the method is insensitive to noise at other frequencies.

8.6 An atomic compass

The results in the previous sections show that the absorption patterns that appear when a vector vortex beam interacts with our rubidium SpOT in the presence of a magnetic field depend on the direction of the magnetic field. This means that we can deduce the direction of the magnetic field just by looking at an absorption image, or, more precisely, the Fourier components of the unwrapped absorption images. In this sense the atoms can be used more or less as a compass. We can calculate the azimuthal and inclination angles of the magnetic field from the Fourier components by inverting equations 8.23 and 8.24. The results of this calculation are shown in figure 8.10. The error bars here are calculated from the uncertainties shown in figure 8.9 by adding and subtracting the uncertainty from the corresponding Fourier component before performing the inverse of equations 8.23 and 8.24.

The average error in our measured angles is $\pm 2.4°$. It is difficult to compare this value to other currently available devices that can be used to measure the direction of magnetic fields, because their specifications are stated in terms of sensitivity. As an example we take a commercially available OPM from QuSpin[7]. They quote the heading error, the error in the measured field due to the alignment between the device axis and the total magnetic field, as 3 nT in the 1-100 μT

[7] `https://quspin.com/qtfm/`

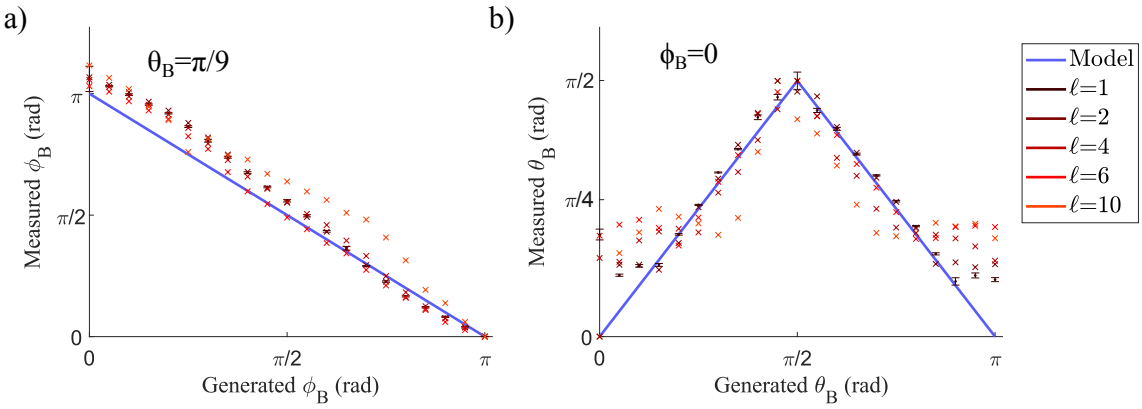

Figure 8.10: Measuring magnetic field angles using SEIT. Model curves were obtained from Fermi's golden rule.

range for their device. Since we used a 20 μT field we can compare our results to this OPM. Let us assume an unfavourable configuration for the OPM, where we use two devices at right angles to measure the direction of a 20μT magnetic field in the plane of the two devices, aligned with the magnetic field such that the heading error is large for both, i.e. at 45° to both. In this case both devices would measure a $20/\sqrt{2} \pm 0.003 = 14.142 \pm 0.003$ μT field, where we have taken the maximum heading error for worst case scenario. Then the uncertainty in the measured field angle is $45° - \tan^{-1}((14.142 - 0.003)/(14.142 + 0.003)) = 0.01°$. This shows that even in an scenario unfavourable to commercial devices our method of measuring magnetic field angle cannot yet compete with their sensitivity.

In fact cheap commercially available orienteering sensors are available at a direction accuracy comparable to our atomic sensor, with a far lower complexity [106]. Having said that, our sensor has very low maturity with room for improvement, discussed below. We also note that we exploit a new approach that, to our knowledge, has not been used to measure magnetic fields before. Because of this we believe there is inherent scientific interest in our technique.

There are a couple of things to notice about the deduced angles, some of which indicate sources of our uncertainty. First, the measured values of ϕ_B have twice the periodicity of the generated ϕ_B because of the even symmetry of the polarisation patterns in the beams we used. This means that with such beams this technique cannot tell the difference between a rotation by Φ and $\Phi + n\pi$, $n \in \mathbb{N}$. This limitation is quite easy to remove, however. One only needs to use beams where the polarisation has an odd azimuthal symmetry. These can be generated by specially designed Q-plates, but are easy to make using any of the dynamic polarisation shaping setups discussed in chapter 2.

A similar issue exists for determining the inclination angle θ_B. We can see from equation 8.24, and figures 8.9 and 8.10, that the technique cannot differentiate between θ_B and $\pi/2 -$ mod $(\theta_B, \pi/2)$. This is because to determine θ_B we are essentially looking at how split the absorption lobes are, which monitors how transverse the magnetic field is to the beam propagation direction. Unfortunately there is no easy way around this.

There was a suggestion by Thomas Clark[8] that detuning the beam from resonance with the $m_f = 0 \rightarrow m_f = 0$ transition would be able to lift this degeneracy. The logic was that the Zeeman shifted $m_f = \pm 1$ states would lead to different absorption on the σ_+ and σ_- transitions if the beam is detuned, because one or the other would be closer to resonance. This should be detectable as a

[8] Personal communication.

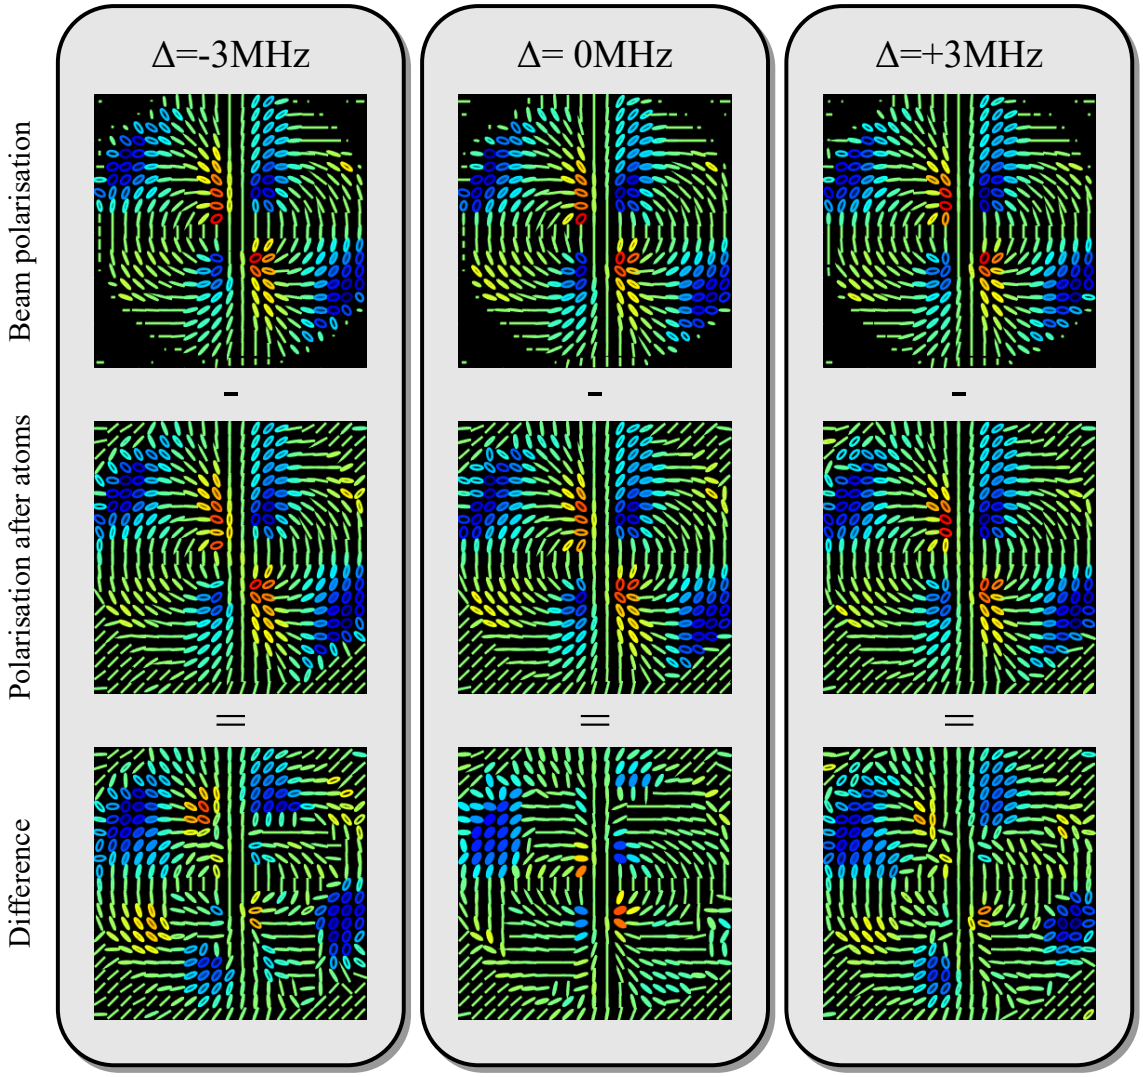

Figure 8.11: Polarisation of vector vortex beams before and after interacting with the SpOT.

change of ellipticity of the polarisation of the beam as it passes through the atoms. We investigated this by using spatially resolved Stokes polarimetry as described in chapter 3. Sample plots of the recorded polarisation patterns are shown in figure 8.11.

Because the atom numbers in the SpOT changed somewhat between experiments we performed the experiment with the same settings 14 times, and averaged the absorption images for each polarisation projection ($\hat{h}, \hat{v}, \hat{a}, \hat{d}, \hat{r}$ and $\hat{l}$). Unfortunately the results were still inconclusive. We saw no obvious pattern in the way polarisation patterns change as a function of detuning. It would perhaps be more useful to compare the absorption of the two circular polarisation components, since a detuning of the laser beam would bring one of the $\sigma_{\pm}$ components closer to resonance when the ground state levels are Zeeman shifted, so we expect to see a difference in the absorption of the circular polarisation components. This could give us an experimental improvement, too: we could use a quarter wave plate and a Wollaston prism to simultaneously record absorption images for the two orthogonal circular polarisation components using the same atom cloud. We expect that such a method should show some difference between the components. We attempted this experiment, but, at the time of writing, there are no results to speak of due to unexpected experimental issues. The Wollaston prism we used has a large (20°) beam separation angle, which lead to one of the beams being compressed in one dimension, see figure 8.12. New, small angle Wollaston prisms have since been purchased, but the experiment has not yet been done with them. Of course, the stretched images could be rescaled in post-processing, but this could introduce further unnecessary

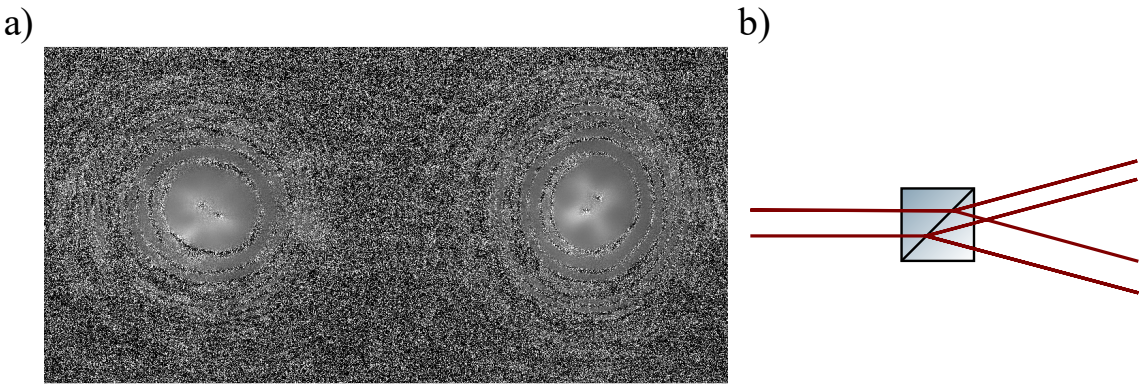

Figure 8.12: Compressing effect of a large separation angle Wollaston prism, shown in (a) the absorption image of two orthogonal polarisation components, where the shape of the two beams appear different and (b) conceptually, showing how two rays are separated by a Wollaston prism.

complications.

Returning to figure 8.10, we note that our observations about figure 8.9 still hold, in that the deduced ϕ_B values are not linear, and the measured θ_B values have a 'shoulder', we never measured values less than $\pi/8$. We attribute these effects to incorrect background field compensation. In fact it seems like there was a small remaining field lying mostly in the $x-y$ plane. We suspect that it had basically no component along the z axis, because such a component would make measured θ_B values asymmetric around $\pi/2$, which we do not see. It is also possible that Fermi's golden rule is simply not an accurate enough approximation to the real system with the parameters we used. Looking at the absorption images in figure 8.8, we can see that at $\theta_B = \pi/2$ the absorption patterns are not completely azimuthally symmetric. This feature appears in both the numerical solutions to the optical Bloch equations and the solutions to the rate equations, but is missing from Fermi's golden rule. Since in the two more complete solutions the splitting (and so the 4ℓ component) appears for lower θ_B using Fermi's golden rule would overestimate this angle for low angles. Figure 8.13 shows an example where shoulders appear in the angles measured when θ_B is calculated from the numerical fast Fourier transform of the absorption pattern using the rate equations solution. The graph was obtained by violating the low intensity assumption ($\Omega_0 = 50\Gamma$) and with an uncompensated transverse background magnetic field with strength of $0.5B_0$.

Nevertheless, our observations suggest that SEIT can also be used to very accurately compensate background magnetic fields, as long as we have precise control over the magnetic field generated by the coils. In our experiments we attempted to cancel the background magnetic field by the following process. First, we applied no quantisation field, and iteratively adjusted the compensation field (in the x-y plane) until the azimuthal absorption contrast (given by the difference between the maximum and minimum absorption at a chosen radius) was minimised. This way we set the compensation such that the remaining field was along the z direction. It can be seen from equations 8.22 that the excited state population, and so the absorption,

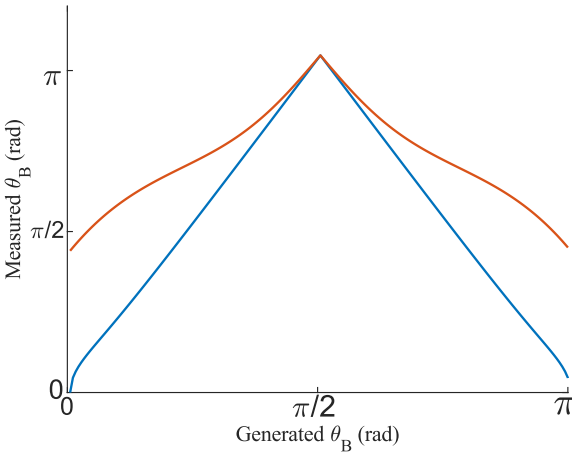

Figure 8.13: Measured θ_B (orange) can be overestimated if beam intensities are too high and/or there is an uncompensated background magnetic field. The blue curve is calculated from Fermi's golden rule. Compare with figure 8.10b.

is zero irrespective of ϕ if $\theta_B = m\pi$, $m \in \mathbb{N}$, i.e. the magnetic field is axial. Then, we applied a quantisation field such that $\theta_{quant} = \pi/2$. We then maximised the azimuthal absorption contrast by iteratively adjusting the axial compensation field strength.

This process can actually be made much more efficient. Consider a controlled quantisation field of strength B_q. Let us set the direction of this field to $\phi_{q1,q2} = 0$, $\theta_{q1} = \pi/4$ and $\theta_{q2} = 3\pi/4$, and measure $\theta_{B1,2}$ of the total magnetic field using the 4ℓ Fourier component according to equation 8.24. The geometry of the situation is shown in figure 8.14. The x and y components of the total magnetic field are the same for both cases, so we can write down

$$\tan\theta_{B1} = \frac{\sqrt{2}B_q + B_0\sin\theta_0\sin\phi_0}{\sqrt{2}B_q + B_0\cos\theta_0}$$
$$\text{and} \tag{8.26}$$
$$\tan\theta_{B2} = \frac{\sqrt{2}B_q + B_0\sin\theta_0\sin\phi_0}{-\sqrt{2}B_q + B_0\cos\theta_0}.$$

We can calculate the ratio

$$\frac{\tan\theta_{B1}}{\tan\theta_{B2}} = \frac{-\sqrt{2}B_q + B_0\cos\theta_0}{\sqrt{2}B_q + B_0\cos\theta_0}, \tag{8.27}$$

which we can rearrange to find the z-component of the background field $\mathbf{B}_0$,

$$B_{0,z} = B_0\cos\theta_0 = \frac{2\sqrt{2}B_q}{1 + \frac{\tan\theta_{B1}}{\tan\theta_{B2}}} - \sqrt{2}B_q. \tag{8.28}$$

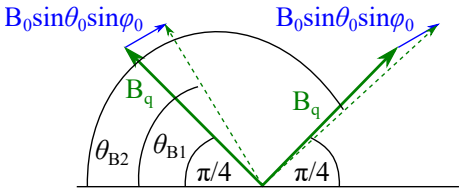

Figure 8.14: Measuring the z-component of a background magnetic field (blue) using SEIT by setting a controlled quantisation field (green) to two inclination angles $\theta_{q1} = \pi/4$ and $\theta_{q2} = 3\pi/4 = \pi - \pi/4$.

This is then easily compensated by adjusting the z-component of the compensation field by $-B_{0,z}$. To determine the azimuthal angle of the remaining background field we generate a series of controlled quantisation fields with $0 \leq \phi_q < 2\pi$ and retrieve the measured ϕ_B using equation 8.23. ϕ_B lie on a curve that crosses the $\phi_B = \phi_q$ line when $\phi_q = \phi_0$ or $\phi_q = -\phi_0$, because in these cases the background field is collinear with the generated quantisation field, so the measured direction is the same as the one set by the quantisation field. This way we can directly read off $|\phi_0|$. The transverse part of the background field can be compensated by adding a component in the $|\phi_0|$ direction to the compensation field, and varying its strength until the absorption contrast goes to zero. At that point the background field will be exactly compensated.

Of course we could measure the background field direction simply by turning any compensation or quantisation field off. In this case we would be uncertain about the θ_0 direction as discussed above, and we would have no information about the strength of the field. We could try setting the compensation field against one of the possible $\mathbf{B}_0$ directions and vary the strength of the compensation field, attempting to minimise the absorption contrast. Unfortunately if we choose the wrong θ_0 we would never achieve zero contrast. We could monitor the rate of change of the splitting with positive and negative compensation fields, since the total magnetic field direction θ_B would change slower with the compensation field pointing in a direction close to $\mathbf{B}_0$ than away from it if they are not collinear.

It is perhaps worthwhile to point out that the system presented here cannot compete in sensitivity

with existing magnetometers, whether based on atoms [82, 119] or other technologies [120]. We think that SEIT is interesting from a fundamental scientific point of view, allowing us to deduce (with some ambiguity) the direction of a magnetic field from an absorption image. This can be done because our cloud of atoms maps the projection of an electric field vector (in the light field) onto the magnetic field vector onto the transmitted intensity of the probe light. Having said that, the technique for compensating stray magnetic fields outlined in the preceding few paragraphs could be of practical use in atom traps, where it is impractical or impossible to place magnetometers of other kinds close to the trap volume.

8.7 In summary

In this chapter we have described electromagnetically induced transparency, and developed theoretical descriptions for investigating EIT in a 4-level tripod system. This system is realised in our experiments comprising the Zeeman sublevels in the F=1→F'=0 D_2 manifold of rubidium 87. We found that in this system absorption depends on the angle between the polarisation direction and the quantisation axis set by an external magnetic field. We can exploit this, in conjunction with beams with structured polarisation, such as the vector vortex beams used here, to deduce the direction of the magnetic field. Here we were only sensitive to magnetic field direction within π sr solid angle because of the even azimuthal symmetry of the beams and the degeneracy of the system with magnetic field inclination angle. By using beams with an odd azimuthal symmetry we can extend the detection range to 2π sr, which is equivalent to a compass with an added orthogonal alignment measurement. We have shown data obtained from vector vortex beams generated from Q-plates. We also described how to use this system to measure and compensate an unknown background magnetic field.

For further study it would be interesting to investigate dispersion in SEIT. We showed in figure 8.2 that the reduced absorption is accompanied by an increased dispersion. This effect reflects the Kramers-Kronig relations for the real and imaginary parts of complex functions. The same relations must hold for our spatial EIT, in which case we should be able to engineer the dispersion by structuring the light field driving SEIT. If so, we can perhaps engineer the dispersion profile in a spatially structured manner, in a scheme that could be used as a spatial quantum memory. We think that we can investigate the dispersion by looking at the rotation of beam polarisation, since polarisation components experiencing EIT ought to experience a different refractive index compared to polarisation components that are more absorbed, and so should be phase shifted.

The end

The previous chapter marked the end of the core chapters of this thesis. In Part I we learned about light as electromagnetic radiation, and learned to control the amplitude, phase and polarisation nearly arbitrarily in any plane transverse to optical beam propagation. These light structures extend into three dimensions via propagation, and we described a simple way of calculating this in the paraxial limit. We explored in detail methods of controlling light using digital devices, spatial light modulators and digital micromirror devices, their advantages and disadvantages, and their optimal use through calibration and aberration correction. In the last chapter of the first part we used our knowledge of beam shaping and analysis to quantify correlations between spatial and polarisation degrees of freedom in vector vortex beams. These correlations mimic the mathematics of quantum entanglement, although they are physically distinct in that entanglement can involve non-local effects whereas classical correlations cannot. Nevertheless, quantum entanglement measures can be adapted to these correlations. In particular we related concurrence to Stokes parameter measurements. Quantifying correlations this way may be useful, apart from their pure scientific interest, for quantum algorithms that require correlations, but not non-locality.

We began Part II by introducing basic atomic physics, including absorption and scattering of (near-)resonant coherent light. We introduced the optical Bloch equations for a 2-level atom. Then we described in quite a lot of detail how to generate the resonant coherent light using external cavity diode lasers. We detailed the construction of the ECDLs used in the experiments in the subsequent chapters, and explained how to stabilise the laser frequency to atomic transitions of rubidium, the atomic species of choice for our experiments. We described an experiment in which we combined the knowledge gained in Part I about structured light with the atomic state control via interaction between atoms and light. We used structured light to drive atoms to dark states in specific volumes in a warm rubidium vapour, and probed the populations of atoms remaining in bright states using a separate probe beam. We reconstructed the population structures tomographically from the fluorescence of the probe beam. We also modelled the populations using rate equations, taking into account the transverse velocity of atoms in the warm cell, which yielded good predictions matched by the experimental results.

We then turned up the complexity another notch, exploring cold atom traps. We described how the Doppler effect can be exploited to cool atoms with red-detuned laser light. We introduced magneto-optical traps, in which cold atoms are collected around the zero of a quadrupole magnetic field. We discussed in detail experimental considerations in our trap, mostly considering alignment and magnetic field control. We also introduced a more advanced trap configuration, where atoms are trapped in a dark state, leading to higher densities because of the elimination of re-radiation pressure from scattered photons within the cold atom cloud. This spontaneous-force optical trap formed the backbone of our next, most complicated experiment, where we once again used structured light and atoms, this time of the cold variety. We developed a theory of

polarisation-dependent EIT, where the absorption of the cold atom cloud is controlled by the angle between a magnetic field setting the quantisation axis for the atoms and the polarisation of a vector vortex beam. We showed that this system allows us to deduce the magnetic field direction, acting like a 3D compass, although with some restrictions that lead to sensitivity to only 2π sr solid angle.

In the remainder of this thesis we report on short projects we took part in that are tangentially related to the main matter of this thesis. I feel they are worth including because they helped build my intuition about topics that form the core of this work[9]. There will be some more talk of atoms when we look at single ion trapping in an experiment to measure the longitudinal component of a focused radially polarised beam. We will also briefly investigate the diffraction of vector vortex beams.

[9] And because I think they are interesting.

Part III

Short projects

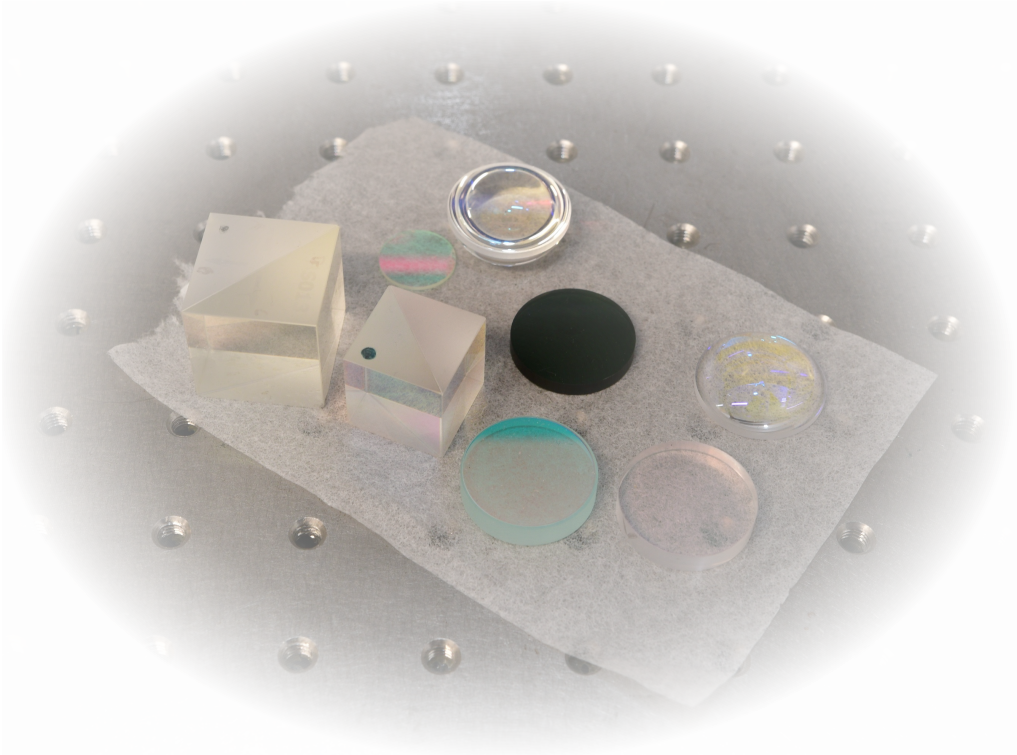

Chapter 9

Short projects

9.1 Introduction

In this short chapter we describe two projects we worked on that are tangentially related to the preceding core chapters of this thesis. The first of these projects is related to our work on spatial polarisation control (chapter 2), here put to use in the context of single ion trapping. The physics of ion trapping is similar to atom trapping (chapter 7), although it is possible to manipulate the motion of ions with higher precision than the motion of neutral atoms in a magneto-optic trap[1]. In the project described in section 9.2.2 we combined our expertise in beam shaping with the ion trapping expertise in the Schmidt-Kaler group at the Johannes Gutenberg University, Mainz, Germany. I visited the Schmidt-Kaler group in March 2017, and the work described here was carried out then. I also present a detailed discussion of ion trapping theory; while I only did minor work on the construction and operation of ion traps, I find the concept exciting, which is why the theory is included here.

Our work in the project described in section 9.3 involved simulations of double-slit diffraction of vector vortex beams, using methods described in chapter 1. This work complemented student projects carried out in the experiments we built to generate arbitrary polarisation structures, as described in chapter 2. We include the simulation results here out of interest, as well as future reference. I built the experiment and wrote the simulations during the summer of 2018, alongside the work presented in chapter 8.

9.2 Ion traps

In chapter 1 we briefly touched on the fact that in the paraxial limit the electric (and magnetic) fields of a propagating light field oscillate transverse to the propagation direction. In strongly focusing, high numerical aperture systems this is no longer necessarily valid. For example, at a tight focus radially polarised beams have a strong longitudinal electric field component oscillating in the propagation direction. Due to the extremely small focal spots in such situations the electric field distributions are difficult to measure. Previously, gold nanoparticles were used to perform such measurements [122, 123]. In collaboration with the Schmidt-Kaler group at the Johannes Gutenberg University of Mainz we intended to measure the distribution of the various polarisation

[1] Although it is possible to exert comparable level of control over neutral atoms, for example in optical lattices [121]

components of tightly focused radially and azimuthally polarised beams using single ytterbium ions.

Today laser cooling is routinely used to prepare experiments investigating the properties of single atoms or ions, and recently using both individual ions and ion strings for quantum computing and quantum simulation. However, just sixty years ago such experiments were considered to never be feasible; this is well illustrated by the famous statement by Rutherford in 1952: "... we never experiment with just one electron or atom or (small) molecule. In thought experiments, we sometimes assume that we do; this invariably entails ridiculous consequences..." Indeed, the difficulty of isolating single particles and performing measurements on it would be insurmountable without tunable lasers. The first tunable dye lasers were invented in the 1960s, and two decades after Rutherford's statement the first narrow-band tunable laser was reported[124]. This development led to the first ion trapping experiments performed by Neuhauser *et al.* [125] and Wineland *et al.* [126] in 1978. Soon after experiments with single trapped ions were realised. Today experiments with single trapped ions are used as accurate time and frequency standards, for precision spectroscopy and quantum computing and simulation.

The early experiments used tunable lasers to cool the ions in conjunction with static quadrupole electromagnetic fields to form the trap. This configuration is known as a Penning trap, proposed in 1936 [127]. In this essay we will focus on a different setup proposed by Paul in 1990 [128]. Paul traps use a dynamic electric field, oscillating at a radio frequency, to generate the quadrupole trapping potential. We will first describe the quantum-mechanical treatment of this potential. We will then look at the interactions between the trapped ions and laser fields. Finally, we apply the concepts developed in these sections to describe laser cooling in this context. This section is largely based on the reviews of Leibfried, *et al.* [129] and Eschner, *et al.* [130], and the PhD thesis of Cornelius Hempel [131].

9.2.1 Cooling to the motional ground state in a linear Paul trap

Let us first look at the form of the potential in a Paul trap. It is easy to see that in order to trap a particle we require a restoring force that is proportional to the distance from the trap centre. This condition is satisfied by the quadrupole potential in a Paul trap, which has the form

$$\Phi = U\frac{1}{2}\left(\alpha x^2 + \beta y^2 + \gamma z^2\right) + U_{RF}\cos(\omega_{RF}t + \phi_{RF})\frac{1}{2}\left(\alpha' x^2 + \beta' y^2 + \gamma' z^2\right) \tag{9.1}$$

where U, U_{RF} correspond to voltages applied to the trap electrodes, ω_{RF} is the trap drive frequency and coefficients α, β, γ determine the shape of the potential. Note that the potential is a sum of separate time-dependent and time-independent parts. To constrain the coefficients one has to solve the Laplace equation $\nabla^2\Phi = 0$. Doing so yields the relations $\alpha + \beta + \gamma = 0$ and $\alpha' + \beta' + \gamma' = 0$. A common choice for these coefficients is

$$-(\alpha + \beta) = \gamma,\ \alpha' = -\beta' \tag{9.2}$$

which results in a potential that generates confinement in the *x*-*y* plane via the oscillating field and by a static field in the z direction; this configuration is known as a linear Paul trap, shown in Fig. 9.1. Such traps are attractive because they can be used to trap strings of ions, for example for applications in quantum computing. Note that the sign difference between the different coefficients indicates that there is no static minimum of the potential in all three dimensions, hence ions will be trapped dynamically.

We now transform into a quantum mechanical picture by replacing the coordinates with their respective operators. Since the potential in different dimensions can be decoupled it is sufficient to look at the part of the potential in the x direction; the other directions can be treated analogously. So the potential $V(t)$ can be written as

$$V(t) = \frac{1}{2}mW^2\hat{x}^2 + \frac{1}{2}mW_{RF}^2\hat{x}^2 \tag{9.3}$$

where m is the ion mass and we have made the substitutions

$$W = \frac{U\alpha}{m}$$
$$W_{RF} = \frac{U_{RF}\alpha' \cos(\omega_{RF}t + \phi_{RF})}{m} \tag{9.4}$$

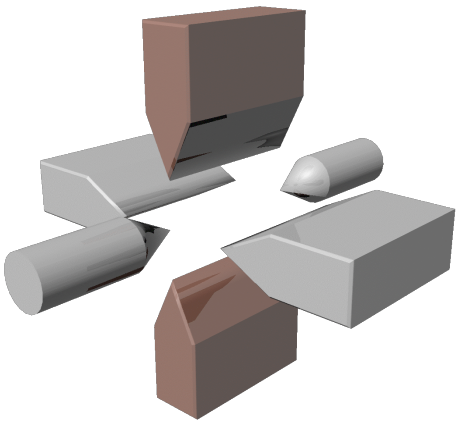

Figure 9.1: Diagram of a linear Paul trap. The trap consists of four blade electrodes and two end-cap electrodes. RF voltage is applied to two opposite blade electrodes (shown in copper) while the rest of the electrodes are held at a constant voltage (grey).

First, let us look at the time-independent part. It is straightforward to write down the Hamiltonian for this part of the ion motion

$$\hat{H}_{h.o.} = \frac{\hat{p}}{2m} + \frac{m}{2}W\hat{x}^2 \tag{9.5}$$

We recognise this Hamiltonian as that of the familiar harmonic oscillator. In fact we can introduce the creation and annihilation operators

$$\hat{a}^\dagger = \sqrt{\frac{mW}{2\hbar}}\left(\hat{x} - \frac{i}{mW}\hat{p}\right) \text{ and } \hat{a} = \sqrt{\frac{mW}{2\hbar}}\left(\hat{x} + \frac{i}{mW}\hat{p}\right) \tag{9.6}$$

Using these operators we can rewrite the time-independent part of equation 9.3 as

$$\hat{H}_{h.o.} = \hbar W\left(\hat{a}^\dagger\hat{a} + \frac{1}{2}\right) \equiv \hbar W\left(\hat{N} + \frac{1}{2}\right) \tag{9.7}$$

The simultaneous eigenstates of $\hat{H}_{h.o.}$ and $\hat{N}$ are known as number states or Fock states, and they are evenly spaced with a ground state $|0\rangle$ with energy $E_0 = \hbar W/2$. The repeated application of the creation operator to the ground state yields all possible eigenstates. The eigenvalues n of $\hat{N}$ corresponding to $|n\rangle$ give the number of motional states or phonons. The motion described by these states is known as secular motion.

We can now consider the time-dependent part of the potential. The total Hamiltonian is given by

$$\hat{H} = \hat{H}_{h.o.} + \hat{H}_{RF} \tag{9.8}$$

where $\hat{H}_{RF}$ is the perturbing Hamiltonian corresponding to the time-dependent potential. This induces a driven oscillation at the trap drive frequency; the resultant motion is called micromotion. The magnitude of micromotion is much smaller than that of secular motion, and its frequency is normally much higher, which is why it is often neglected. However, because of its large frequency it does contribute to the energy in the system[132]. It is easy to find the effect of the micromotion Hamiltonian by performing the transformation

$$\hat{H}_{int} = \hat{U}_0^\dagger \hat{H}_{RF}\hat{U}_0, \tag{9.9}$$

where $\hat{U}_0 = e^{-i\hat{H}_{h.o.}t/\hbar}$ is simply the time-evolution of the unperturbed system. This transformation to the so-called interaction picture separates the perturbing Hamiltonian from the unperturbed

(harmonic oscillator) system. It is convenient to express the RF Hamiltonian in terms of a time-dependent force:

$$\hat{H}_{RF} = -(\hat{a}+\hat{a}^\dagger)x_0\frac{A}{2i}\left(e^{i(\omega_{RF}t+\phi_{RF})}+e^{-i(\omega_{RF}t+\phi_{RF})}\right) \equiv -(\hat{a}+\hat{a}^\dagger)x_0F, \tag{9.10}$$

where we have expressed $\hat{x}$ in terms of the creation and annihilation operators, with the characteristic length scale of the harmonic oscillator $x_0 = \sqrt{\hbar/(2mW)}$. We have also replaced the time-dependent trigonometric function by its exponential form. Having done this, the full interaction Hamiltonian can be written down:

$$\begin{aligned}\hat{H}_{int} = x_0\frac{A}{2i}\Big(&\hat{a}e^{-i((W+\omega_{RF})t+\phi_{RF})}+\hat{a}e^{-i((W-\omega_{RF})t-\phi_{RF})}\\ &+\hat{a}^\dagger e^{i((W-\omega_{RF})t-\phi_{RF})}+\hat{a}^\dagger e^{i((W+\omega_{RF})t+\phi_{RF})}\Big).\end{aligned} \tag{9.11}$$

To obtain this we have used the identities

$$\begin{aligned}e^{i\theta\hat{a}^\dagger\hat{a}}\hat{a}e^{-i\theta\hat{a}^\dagger\hat{a}} &= e^{-i\theta}\hat{a},\\ e^{i\theta\hat{a}^\dagger\hat{a}}\hat{a}^\dagger e^{-i\theta\hat{a}^\dagger\hat{a}} &= e^{i\theta}\hat{a}^\dagger.\end{aligned}$$

One can now perform the rotating wave approximation neglecting terms with frequency sums to simplify this to

$$\hat{H}_{int} = x_0\frac{A}{2i}\left(\hat{a}e^{-i((W-\omega_{RF})t-\phi_{RF})}+\hat{a}^\dagger e^{i((W-\omega_{RF})t-\phi_{RF})}\right). \tag{9.12}$$

On resonance (i.e. when the trap drive frequency is the same as that of the harmonic oscillator) this further simplifies to

$$\hat{H}_{int} = x_0\frac{A}{2i}\left(\hat{a}e^{i\phi_{RF}}+\hat{a}^\dagger e^{-i\phi_{RF}}\right). \tag{9.13}$$

The time evolution corresponding to this Hamiltonian can be re-written as a displacement operator

$$D(\alpha) = e^{\alpha\hat{a}^\dagger-\alpha^*\hat{a}}, \tag{9.14}$$

where we have introduced the dimensionless amplitude $\alpha = e^{-i\phi_{RF}}Ax_0t/(2\hbar)$. Applying the displacement operator to the ground (vacuum) state creates a coherent state[133]. Coherent states have a Gaussian probability distribution with the same shape as the ground state wave packet in the harmonic oscillator, but their centre follows the trajectory of the ion (including micromotion) and do not spread.

To describe the dynamics of ion-laser interactions we will assume that the internal electronic level structure of the ion is well approximated by that of a two-level atom, with states $|g\rangle$ and $|e\rangle$ the ground and excited states respectively. This simple model is adequate in most experiments, where the energy separation of two electronic levels addressed by the laser is very different compared to other level separations (i.e. the laser detuning is small only for these two levels) and the Rabi frequencies (and consequently laser intensities) are small compared to this difference as well.

Let us revisit the two-level atom introduced in chapter 4. The Hamiltonian describing the internal states of a two-level atom is given by

$$\hat{H}_e = \hbar\frac{\omega_e+\omega_g}{2}\left(|e\rangle\langle e|+|g\rangle\langle g|\right)+\hbar\frac{\omega_e-\omega_g}{2}\left(|e\rangle\langle e|-|g\rangle\langle g|\right), \tag{9.15}$$

where $\hbar\omega_g$ and $\hbar\omega_e$ are the energies of the ground and excited states respectively. It is convenient to re-cast this expression in terms of the familiar Pauli matrices. This can be done because the

algebra of two-level operators is identical to the spin-1/2 algebra if one maps ground and excited states $|g\rangle, |e\rangle$ to the down- and up-spin states $|\downarrow\rangle, |\uparrow\rangle$ respectively. We perform the change of basis

$$|e\rangle\langle e| + |g\rangle\langle g| \mapsto \hat{\sigma_0}, |g\rangle\langle e| + |e\rangle\langle g| \mapsto \hat{\sigma}_x,$$
$$i\left(|g\rangle\langle e| - |e\rangle\langle g|\right) \mapsto \hat{\sigma}_y, |e\rangle\langle e| - |g\rangle\langle g| \mapsto \hat{\sigma}_z,$$

with $\hat{\sigma_0}$ the 2x2 identity matrix to write Eq. 9.15 as

$$\hat{H}_e = \hbar\frac{\omega_0}{2}\hat{\sigma}_z, \tag{9.16}$$

where $\hbar\omega_0$ is the energy difference between the ground and excited states. The first part of equation 9.15 could be removed because it is state-independent and does not influence the analysis that follows (its removal can be considered a choice of a zero energy level halfway between the two levels). In the spin-1/2 formalism one can define the electronic raising and lowering operators $\sigma_\pm = (\sigma_x \pm i\sigma_y)/2$ which will be of use later.

Having described the energy of the two-level atom on its own let us now consider the effect of an incident laser field. Such a light field will induce a flipping between the two states at the Rabi frequency Ω, proportional to the electric field amplitude of the incoming field:

$$\hat{H}_l = \hbar\Omega\hat{\sigma}_x\cos(\omega t + \phi), \tag{9.17}$$

where ω is the frequency of the incident (monochromatic) light field and ϕ is its phase, and t is the duration of the laser pulse. We note that in our definition the Rabi frequency corresponds to the frequency at which populations are exchanged between the ground and excited levels. The total atomic Hamiltonian is the sum of the two Hamiltonians in equations 9.16 and 9.17

$$\hat{H}_t = \hbar\frac{\omega_0}{2}\hat{\sigma}_z + \hbar\Omega\hat{\sigma}_x\cos(\omega t + \phi). \tag{9.18}$$

Once again it is useful to transform to an interaction picture to look at the dynamics induced by the laser pulse. The transformation again takes the form $U_0^\dagger \hat{H}_l U_0$ with $U_0 = e^{-i\hat{H}_e t/\hbar} = e^{-i\omega_0 t\sigma_z/2}$, and expressing $\cos(\omega t + \phi) = \frac{1}{2}\left(e^{i(\omega t+\phi)} + e^{-i(\omega t+\phi)}\right)$ yields, after quite some algebra,

$$\begin{aligned}\hat{H}_{int} = \frac{\hbar\Omega}{4}\Big(&e^{i((\omega-\omega_0)t+\phi)}\sigma_x + e^{-i((\omega-\omega_0)t+\phi)}\sigma_x \\ &+ie^{i((\omega-\omega_0)t+\phi)}\sigma_y - ie^{-i((\omega-\omega_0)t+\phi)}\sigma_y\Big),\end{aligned} \tag{9.19}$$

where we performed a rotating wave approximation to drop terms with frequency sums since, considering the optical frequencies involved, they oscillate too fast to affect the time evolution of the remainder. Finally, we make use of the electronic raising and lowering operators defined above and trigonometric identities to arrive to

$$\hat{H}_{int} = \frac{\hbar\Omega}{2}\left(e^{-i(\Delta t+\phi)}\sigma_+ + e^{i(\Delta t+\phi)}\sigma_-\right), \tag{9.20}$$

where we also introduced the laser detuning $\Delta = \omega - \omega_0$. Unsurprisingly, for a resonant laser beam ($\Delta = 0$) during the time evolution of the interaction Hamiltonian $e^{-i\hat{H}_{int}t/\hbar}$ populations in the ground and excited states are exchanged periodically.

In order to adequately describe the interaction between ions in a trap with incident laser beams we have to combine the motional and electronic Hamiltonians. We start with combining the Hamiltonian corresponding to the motion in a quantum harmonic oscillator (equation 9.7) and the

Hamiltonian corresponding to the energy of the unperturbed two-level atom (equation 9.16) to obtain a description of the unperturbed system

$$\hat{H}_0 = \hbar W \left(\hat{a}^\dagger \hat{a} + \frac{1}{2}\right) + \hbar \frac{\omega_0}{2} \hat{\sigma}_z. \tag{9.21}$$

The laser-ion interaction is described as a perturbation on this system, which takes the form of equation 9.20 modified by the inclusion of a factor coupling the motional and electronic states:

$$\hat{H}_1 = \frac{\hbar\Omega}{2}(\sigma_+ + \sigma_-)\left(e^{i\eta(\hat{a}+\hat{a}^\dagger)}e^{-i(\omega t+\phi)} + e^{-i\eta(\hat{a}+\hat{a}^\dagger)}e^{i(\omega t+\phi)}\right), \tag{9.22}$$

where we have defined the Lamb-Dicke parameter $\eta = \frac{2\pi}{\lambda}\sqrt{\frac{\hbar}{2m\omega_{RF}}}\cos(\theta)$, which includes the laser wavelength λ and the angle of incidence θ relative to the motion of the ion.

To investigate this perturbation we transform to the interaction picture with $U_0 = e^{-i\hat{H}_0 t/\hbar}$:

$$\begin{aligned}\hat{H}_{int} = \frac{\hbar\Omega}{2} &\left(e^{i\frac{\omega_0}{2}t\sigma_z}(\sigma_+ + \sigma_-)e^{-i\frac{\omega_0}{2}t\sigma_z}\right) \\ &\times \left(e^{iWt\hat{a}^\dagger\hat{a}}e^{i\eta(\hat{a}+\hat{a}^\dagger)}e^{-iWt\hat{a}^\dagger\hat{a}}e^{-i(\omega t+\phi)}\right. \\ &\left. + e^{iWt\hat{a}^\dagger\hat{a}}e^{-i\eta(\hat{a}+\hat{a}^\dagger)}e^{-iWt\hat{a}^\dagger\hat{a}}e^{i(\omega t+\phi)}\right).\end{aligned} \tag{9.23}$$

This can be simplified using the following relations:

$$\begin{aligned}e^{i\frac{\theta}{2}\sigma_z}\sigma_\pm e^{-i\frac{\theta}{2}\sigma_z} &= e^{\pm i\theta}\sigma_\pm \\ e^{i\theta\hat{a}^\dagger\hat{a}}e^{i\eta(\hat{a}+\hat{a}^\dagger)}e^{-i\theta\hat{a}^\dagger\hat{a}} &= e^{i\eta\left(\hat{a}e^{-i\theta}+\hat{a}^\dagger e^{i\theta}\right)}.\end{aligned}$$

Performing the simplifications, recognising the laser detuning $\Delta = \omega - \omega_0$ and performing a rotating wave approximation leaves us with

$$\hat{H}_{int} = \frac{\hbar\Omega}{2}\sigma_+ e^{-i(\Delta t-\phi)}e^{i\eta\left(\hat{a}e^{-iWt}+\hat{a}^\dagger e^{iWt}\right)} + H.c. \tag{9.24}$$

where *H.c.* denotes the Hermitian conjugate.

In the so-called *Lamb-Dicke regime* the spatial extent of the wavepacket of the ion is much smaller than the laser wavelength and the following inequality is satisfied:

$$\eta\sqrt{\langle(\hat{a}+\hat{a}^\dagger)^2\rangle} \equiv \eta\sqrt{2\overline{n}+1} \ll 1 \tag{9.25}$$

where $\overline{n}$ is the mean phonon number. We can then Taylor-expand the second exponential in equation 9.24,

$$e^{i\eta\left(\hat{a}e^{-iWt}+\hat{a}^\dagger e^{iWt}\right)} = \mathbf{1} + i\eta\left(\hat{a}e^{-iWt} + \hat{a}^\dagger e^{iWt}\right) + \mathcal{O}\left(\eta^2\right). \tag{9.26}$$

We can justify the omission of terms of order η^2 in the case when $\overline{n} < 20$ which is normally satisfied after Doppler cooling in experiments. We can then rewrite equation 9.24 in its final form

$$\hat{H}_{int} = \frac{\hbar\Omega}{2}\sigma_+ e^{-i(\Delta t-\phi)}\left(\mathbf{1} + i\eta\left(\hat{a}e^{-iWt} + \hat{a}^\dagger e^{iWt}\right)\right) + H.c. \tag{9.27}$$

This Hamiltonian contains three resonances. The first one at $\Delta = 0$ takes the form

$$\hat{H}_{car} = \frac{\hbar\Omega_0}{2}\left(e^{i\phi}\sigma_+ + e^{-i\phi}\sigma_-\right), \tag{9.28}$$

where we have defined the Rabi frequency Ω_0 as the coupling strength at $n = 0$ and where we

assumed the resolved sideband limit $\Omega \ll W$ to neglect terms with W (this is another rotating wave approximation). This carrier transition couples the states $|g,n\rangle \leftrightarrow |e,n\rangle$; this means that it induces (de-)excitation in the internal electronic states but does not affect the motional state of the ion.

The condition $\Delta = -W$ leads to a transition of the form

$$\hat{H}_{rsb} = \frac{i\hbar\Omega_{n,n-1}\eta}{2}\left(\hat{a}e^{i\phi}\sigma_+ + \hat{a}^\dagger e^{-i\phi}\sigma_-\right) \tag{9.29}$$

with Rabi frequency $\Omega_{n,n-1} = \eta\sqrt{n}\Omega_0$. This red sideband transition couples $|g,n\rangle \leftrightarrow |e,n-1\rangle$, i.e. it *lowers* the phonon number of the motional state by one unit.

Finally, the third resonance, at $\Delta = W$ gives rise to

$$\hat{H}_{bsb} = \frac{i\hbar\Omega_{n,n+1}\eta}{2}\left(\hat{a}^\dagger e^{i\phi}\sigma_+ + \hat{a}e^{-i\phi}\sigma_-\right) \tag{9.30}$$

coupling $|g,n\rangle \leftrightarrow |e,n+1\rangle$, with Rabi frequency $\Omega_{n,n-1} = \eta\sqrt{n+1}\Omega_0$. It raises the phonon number by one unit. The level scheme described here is shown in Fig. 9.2.

The red sideband transition is of interest in general because it entangles the motional and electronic states of the ion, and is in fact equivalent in form to the well-known Jaynes-Cummings Hamiltonian describing the interaction between a two-level atom and optical cavity modes, widely used in quantum optics. The blue sideband transition is sometimes known as anti-Jaynes-Cummings coupling, and is not analogous to any cavity QED phenomenon (in fact an analogous interaction would violate energy conservation). However, the red sideband transition can also be used for the purpose of cooling; this shall be our focus for the next few paragraphs.

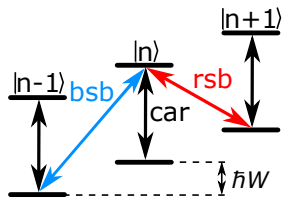

Figure 9.2: Resolved sideband level structure in a harmonic oscillator. The motional states $|m\rangle$ form a ladder with energy offset $\hbar W$, and each motional state has an electronic ground and excited state. The electronic states are coupled by the carrier (car) transition, and the motional states are coupled to the internal states via the blue sideband (bsb) and red sideband (rsb) transitions.

In many ion trapping experiments it is desirable to bring ions to their motional ground state. In most cases this is achieved by a two-stage cooling process. First, after ions are released into the trap volume they are slowed down by Doppler cooling; this is no different from Doppler cooling for neutral atoms described in chapter 7. This allows for an ion (or chain of ions) to be confined in the trap, where they are cooled by the same process until they reach the Lamb-Dicke regime. Then resolved sideband cooling is applied to bring the ion to its motional ground state. Let us look at this procedure in more detail.

The principle of resolved sideband cooling is very similar to that of Doppler cooling. Provided that the linewidth of the decay on the cooling transition is narrower than the frequency of the motional state of the ion the red and blue sidebands are resolved (hence the name resolved sideband cooling). We can neglect the sidebands due to micromotion here because they are at a much higher frequency and so they are far from resonance. It is then possible to tune the cooling laser to the (secular) red sideband transition. In the Lamb-Dicke regime the ion will absorb photons on this transition but spontaneous emission will predominantly occur on the carrier transition which does not change the motional state, so in the cooling cycle the motional state of the ion will be lowered. This process can be repeated until the ion will be in its motional ground state with high probability. This can

be seen relatively easily by looking at rate equations, as we shall do in what follows.

If there were no heating mechanisms present, after a sufficiently large number of cooling cycles the ion would arrive to the motional ground state $|n=0\rangle$, and since this is a *dark state* for red sideband detuning the ion would stay in this state indefinitely. However, there are two obvious heating effects, namely excitation on the carrier and on the blue sideband. Carrier excitation occurs with probability $(\Omega/(2W))^2$ and will decay on the blue sideband with probability $(\Omega/(2W))^2 \tilde{\eta}\Gamma$, where Γ is the decay rate of the excited state in a given motional state. The Lamb-Dicke factor $\tilde{\eta}$ is different from that of the excitation (and the red sideband excitation) because the decay is not restricted in direction to that of the cooling beam. Note that there is a much higher probability that the carrier excitation will decay on the carrier transition, but since that does not change the motional state we ignore it. Also note that we consider the detuning of transitions from the red sideband. The probability of excitation on the blue sideband is $(\eta\Omega/(4W))^2$ and decay on the carrier after this excitation occurs with probability $(\eta\Omega/(4W))^2 \Gamma$. Here we ignore the process of excitation and decay both happening on the blue sideband because the net effect is no change in the motional state.

We can now write down the rate equations for the motional ground state and the first excited state $|0\rangle, |1\rangle$:

$$\begin{aligned} \dot{p}_0 &= p_1 \frac{(\eta\Omega)^2}{\Gamma} - p_0 \left[\left(\frac{\Omega}{2W} \right)^2 \tilde{\eta}^2 \Gamma + \left(\frac{\eta\Omega}{4W} \right)^2 \Gamma \right] \\ \dot{p}_1 &= -\dot{p}_0 \end{aligned} \tag{9.31}$$

In the case of a steady state (i.e. the rates being equal to zero) the probability of finding the ion in the motion ground state is given by

$$p_0 = 1 - \left(\frac{\Gamma}{2W} \right)^2 \left(\frac{1}{4} + \left(\frac{\tilde{\eta}}{\eta} \right)^2 \right) \tag{9.32}$$

where $\Gamma \ll W$ so the first bracketed term is close to zero meaning that the ion reaches the motional ground state with a probability close to unity.

State-of-the-art experiments with trapped ions, used in a wide variety of contexts ranging from high-precision time standards to quantum computing, require precise, high-fidelity state preparation and control. This includes the treatment of both motional and internal electronic states of ions. State preparation is routinely achieved through a multi-stage cooling process, including Doppler cooling and resolved sideband cooling. So far we have described the quantum mechanical treatment of ion motion in the harmonic potential of a linear Paul trap as well as the coupling of the internal states of a two-level ion to its motion via the red and blue sideband transitions. We have also outlined the principles involved in the standard method of cooling via Doppler cooling and sideband cooling in the Lamb-Dicke regime. We note that other cooling processes have been proposed (some also experimentally realised) that were not discussed here; these include cooling via electromagnetically induced transparency (EIT), Sisyphus cooling, and state selection cooling. The interested reader is referred to [130] and references therein for more details.

9.2.2 Experiment: single ion in a vector vortex

In 2017 we proposed an experiment with Christian Schmiegelow from the University of Buenos Aires, Argentina, during his visit at Glasgow. In this experiment we measured the polarisation

components of strongly focused radially and azimuthally polarised vector beams by observing their interaction with a dipole transition. Specifically, a DMD-based vector beam generating device (described in chapter 2) was placed in the repump beam of an ytterbium trap and fluorescence on the Doppler cooling transition was measured. The experimental setup is shown in figure 9.3. Whether the repump transition driven by certain polarisations is allowed or forbidden depends on the Zeemann sublevels. The selection rules for the different magnetic sublevels in turn depend on the projection of the field components on the quantisation direction given by an external magnetic field. This is very similar to our understanding of SEIT in chapter 8. A light field parallel to the magnetic field (a π field) cannot exert torque so only transition with $\Delta m = 0$ are allowed. A linearly polarised field perpendicular to the magnetic field (a field containing equal amounts of σ^+ and σ^- polarisations) exerts torque but in no specific direction so the allowed transitions are both $\Delta m = \pm 1$. A field perpendicular and rotating with respect to the magnetic field (a σ^+ or σ^- field) exerts torque in its sense of rotation so the allowed transition will be either $\Delta m = 1$ or $\Delta m = -1$.

As a result different polarisation components can drive different transitions depending on the angle between the beam propagation direction and the external magnetic field. In simple cases of parallel and perpendicular magnetic fields the following situations can arise:

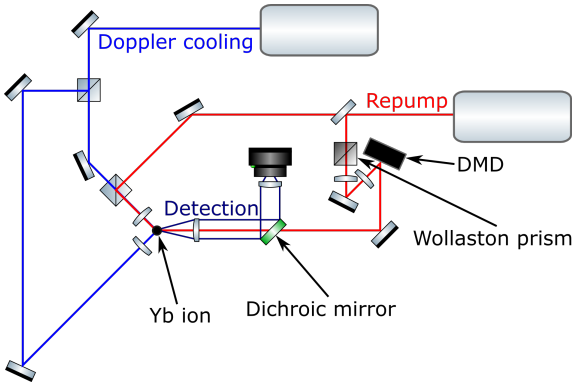

Figure 9.3: Experimental setup used to measure polarisation components of vector beams using single ions. The dichroic mirror transmits 935 nm light and reflects 370 nm.

- Parallel magnetic field and beam propagation direction.
  - circular polarisation is seen as σ^+ or σ^- field respectively.
  - linear polarisation is seen as a $\sigma^+ + \sigma^-$ mixture.
  - longitudinal polarisation is seen as a π field.
- Perpendicular magnetic field and beam propagation direction.
  - circular polarisation is seen as $\sigma^+ + \sigma^- + \pi$ mixture.
  - linear polarisation is seen as a $\sigma^+ + \sigma^-$ mixture or as a π field depending whether it is polarised perpendicular parallel to the plane formed by the magnetic field and the propagation direction.
  - longitudinal polarisation is seen as a $\sigma^+ + \sigma^-$ mixture.

Figure 9.4 shows the energy levels of singly charged Yb^+ ions. In the Zeeman sublevel structure of the levels involved in the repump transition π, σ_- or σ_+ light drives the ion into a dark state where it cannot be repumped so cannot produce fluorescence. Consequently only $\sigma_- + \sigma_+$ light can drive the repump transition. In particular, the interesting longitudinal component of radially polarised beams can be measured by observing the difference of fluorescence for the two magnetic field directions described above. The experiment can be repeated for an azimuthally polarised repump beam, in which case we expect no difference of the fluorescences since azimuthally polarised beams do not have any longitudinal component at a focus. Comparing the radial and azimuthal cases gives us information about the strength of the longitudinal polarisation component at the position

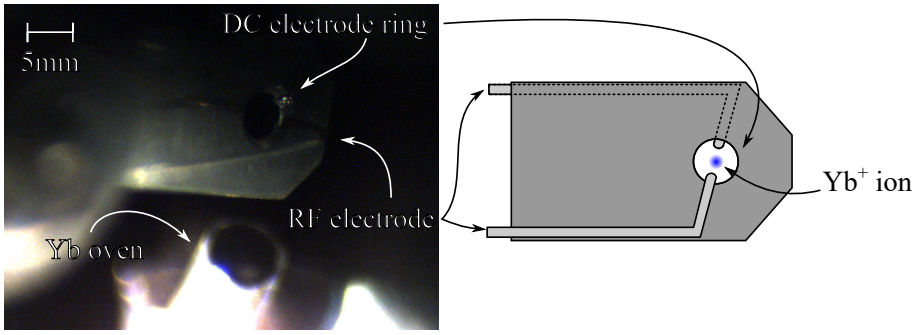

Figure 9.5: Photograph and schematic diagram of a ring ion trap.

of the ion. Either the ion or the beam can be scanned to build up a 3D structure of longitudinal polarisation near the focus.

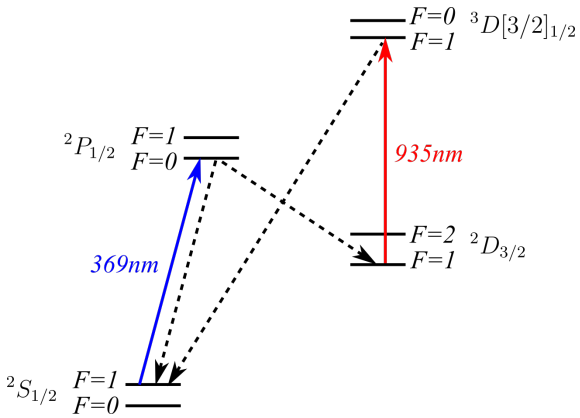

Figure 9.4: Singly charged ytterbium ion level scheme. Blue and red solid lines indicate Doppler cooling and repump transitions, respectively, and dotted lines indicate relevant spontaneous decays.

It is very easy to scan the position of an ion in a segmented linear Paul trap. Unfortunately one was not available for this experiment. Instead we used a much simpler ring trap in which ions are confined well only in a specific spot and cannot be moved. This ring trap, shown in figure 9.5, was built by Christian Schmiegelow, Martin Drechsler and Abasalt Bahrami[2]. We opted to scan the beam instead, which is straightforward using the DMD already in the setup. The trap was constructed in Mainz, but has since been transported to Buenos Aires to become the first functioning ion trap in Latin America. Since it was recently put into operation most components were not optimised yet, and unfortunately to date no conclusive measurements of longitudinal polarisation components were made.

9.3 Double slit diffraction of vector vortex beams

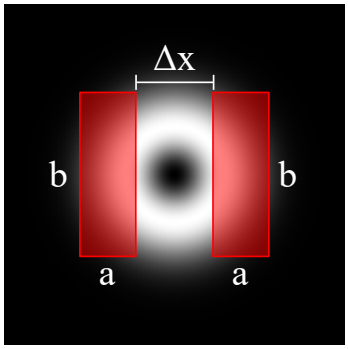

Figure 9.6: Rectangular slits.

The following project was run as a student project on an experiment I built, and I took part in the day-to-day supervision of the students. I produced the simulations shown below to complement the students' work.

In 2012 Li *et al.* investigated the diffraction patterns produced by vector vortex beams passing through double slits [134]. They found that far from the propagation axis along the slit direction the usual double slit diffraction pattern is recovered. However, near the propagation axis there is a more intricate interference pattern. This pattern arises from different, spatially dependent, polarisation components interfering constructively and destructively in different places. Of course this means that the interference pattern depends

[2]... if I remember right. I apologise if this list is inaccurate.

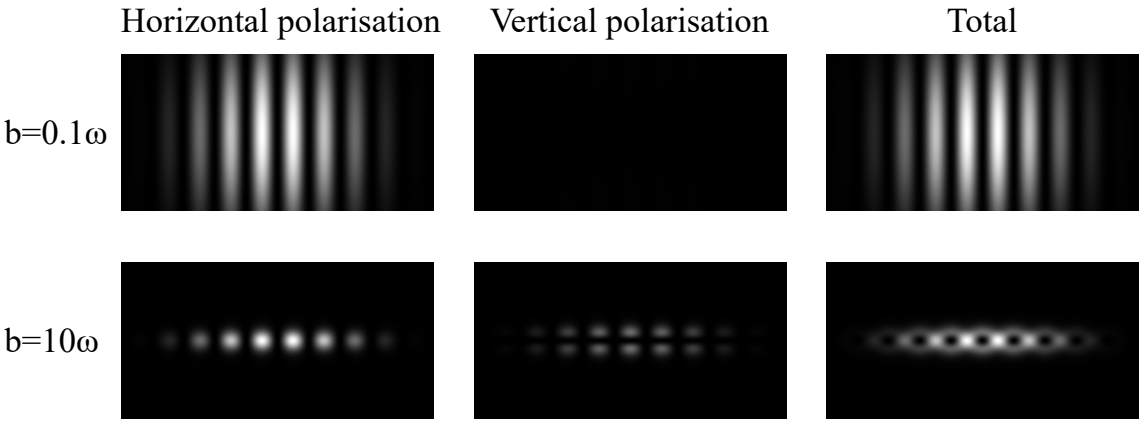

Figure 9.7: Simulated double slit interference of a radially polarised beam, where the slit height b is related to the beam waist ω as shown on the left. The width of the slits was always $a = 0.1\omega$ and their separation was $\Delta x = 0.4\omega$. The horizontally and vertically polarised components of the interference in the far field are also shown.

strongly on how many polarisation directions are transmitted through the slits, which is not discussed in [134]. We had a setup capable of generating any vector field using a DMD (the same setup as in chapter 3), as well as arbitrary slit patterns by multiplying the hologram by programmed slit masks. This prompted us to investigate double slit diffraction.

Consider a radially polarised beam. Placing two slits, shown in figure 9.6 that are small compared to the beam size along the x axis would mostly admit horizontally polarised light, and we expect interference pattern very similar to the standard double slit interference. However, if the slits are larger in the y direction they start to sample more polarisations, and the interference pattern becomes more intricate. This is shown in figure 9.7, where we numerically evaluated the interference in the far field. We can see that when the slits are larger additional interference fringes appear with a different structure. In this case we have a vertically polarised component appearing away from the x axis in two lobes, which is responsible for the dark, horizontally aligned hole intersecting the vertically polarised interference fringes. For higher order vector vortex beams the number of lobes in the vertical direction will depend on the number of polarisation lobes that the slits allow to pass through.

But this discussion suggests that for vector vortex beams rectangular slits are not ideal, because while they allow us to sample different polarisations, they oversample certain polarisations because the intensity of the beams is not uniform. They also never allow us to sample all polarisations. Since the beams are cylindrically symmetric it would perhaps be better to define the slits in cylindrical coordinates, shown in figure 9.8. These are just as easy to generate in our system as the rectangular slits, and numerical simulations are shown in figure 9.9. Of course when the slit opening angle is set to π we just generated an annular aperture, and so the diffraction patterns are just the annular diffraction patterns of the component beams, which, in the case of a radially polarised beam in the horizontal and vertical polarisations shown in the figures, are just Hermite-Gaussian modes.

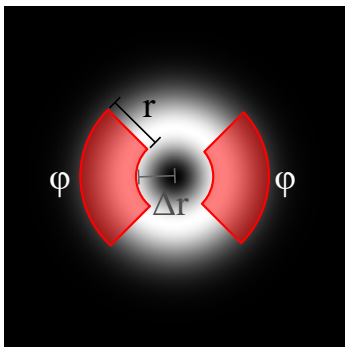

Figure 9.8: Circular slits.

Project students also generated these beams and slit patterns in an experiment, and measured their polarisation structure. The data was in good agreement with our numerical simulations. Interestingly, we never saw the simple double slit interference patterns emerge far from the x axis

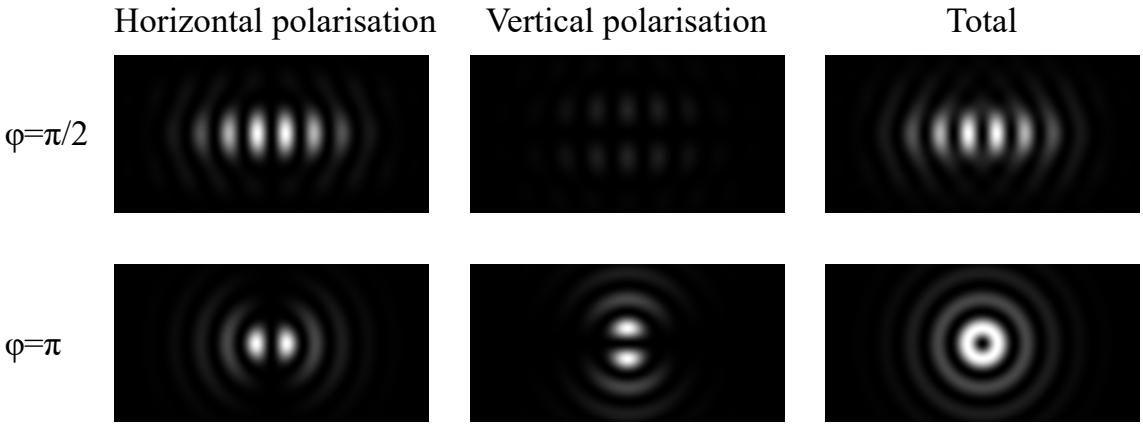

Figure 9.9: Simulated circular slit interference of a radially polarised beam, where the slit opening angle ϕ is shown on the left. The radial width of the slits was always related to the beam waist ω by $r = 0.1\omega$, and their radial separation was $\Delta r = 0.4\omega$. The horizontally and vertically polarised components of the interference in the far field are also shown.

as shown by Li *et al.*, neither in the simulations nor in the experiment. However, a project student developed analytical expressions for calculating the interference patterns, and in her simulations those features appeared. We are not quite sure why this is, but due to lack of time we did not investigate this further.

9.4 In summary

In this chapter we described the theory of cooling a single ion to its motional ground state in a harmonic trap, discussed in the context of a linear Paul trap. We then described an experiment that can use selection rules in atomic transitions to measure longitudinal polarisation components arising from strongly focused radially polarised beams. In this experiment we intended to use a singly charged ytterbium ion in a simple ring trap with a magnetic field setting its quantisation axis to which the polarisation of the illuminating strongly focused repump beam would be compared. This is actually very similar in concept to how we thought about SEIT in chapter 8. Unfortunately to date the experiment yielded no results.

Then we briefly described far field interference of vector vortex beams when they pass through double slits, motivated by a study by Li *et al.* in 2012. We extended study to slits with circular symmetry, and showed some numerical simulations of far field interference for horizontal and vertical polarisation components. These are in qualitative agreement with experimental results, but interestingly enough, are somewhat different from the results of Li *et al.* Unfortunately due to a lack of time at the time of writing this was not investigated further.

Bibliography

[1] J. C. Maxwell, *A Treatise on Electricity and Magnetism*, ser. Cambridge Library Collection - Physical Sciences. Cambridge University Press, 2010, vol. 2.

[2] F. J. Duarte, *Tunable Lasers Handbook*. Academic Press, 01 1995.

[3] D. M. Boroson, "Optical communications: A compendium of signal formats, receiver architectures, analysis mathematics, and performance characteristics," 2005.

[4] G. M. Gibson, B. Sun, M. P. Edgar, D. B. Phillips, N. Hempler, G. T. Maker, G. P. A. Malcolm, and M. J. Padgett, "Real-time imaging of methane gas leaks using a single-pixel camera," *Opt. Express*, vol. 25, no. 4, pp. 2998–3005, Feb 2017. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-25-4-2998

[5] M.-J. Sun, M. P. Edgar, G. M. Gibson, B. Sun, N. Radwell, R. Lamb, and M. J. Padgett, "Single-pixel three-dimensional imaging with time-based depth resolution," *Nature Communications*, vol. 7, pp. 12 010 EP –, Jul 2016, article. [Online]. Available: https://doi.org/10.1038/ncomms12010

[6] N. Radwell, A. Selyem, L. Mertens, M. P. Edgar, and M. J. Padgett, "Hybrid 3d ranging and velocity tracking system combining multi-view cameras and simple lidar," *Scientific Reports*, vol. 9, no. 1, p. 5241, 2019. [Online]. Available: https://doi.org/10.1038/s41598-019-41598-z

[7] A. Weiner, *Ultrafast optics*. John Wiley & Sons, 2011, vol. 72.

[8] S. Rossel and R. Wehner, "Polarization vision in bees," *Nature*, vol. 323, no. 6084, pp. 128–131, 1986. [Online]. Available: https://doi.org/10.1038/323128a0

[9] Z.-Y. Chen, L.-S. Yan, Y. Pan, L. Jiang, A.-L. Yi, W. Pan, and B. Luo, "Use of polarization freedom beyond polarization-division multiplexing to support high-speed and spectral-efficient data transmission," *Light: Science &Amp; Applications*, vol. 6, pp. e16 207 EP –, Feb 2017, original Article. [Online]. Available: https://doi.org/10.1038/lsa.2016.207

[10] W. F. McGrew, X. Zhang, R. J. Fasano, S. A. Schäffer, K. Beloy, D. Nicolodi, R. C. Brown, N. Hinkley, G. Milani, M. Schioppo, T. H. Yoon, and A. D. Ludlow, "Atomic clock performance enabling geodesy below the centimetre level," *Nature*, vol. 564, no. 7734, pp. 87–90, 2018. [Online]. Available: https://doi.org/10.1038/s41586-018-0738-2

[11] A. Hinton, M. Perea-Ortiz, J. Winch, J. Briggs, S. Freer, D. Moustoukas, S. Powell-Gill, C. Squire, A. Lamb, C. Rammeloo, B. Stray, G. Voulazeris, L. Zhu, A. Kaushik, Y.-H. Lien, A. Niggebaum, A. Rodgers, A. Stabrawa, D. Boddice, S. R. Plant, G. W. Tuckwell, K. Bongs, N. Metje, and M. Holynski, "A portable magneto-optical trap with prospects for atom interferometry in civil engineering," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 375, no. 2099, p. 20160238, 2017. [Online]. Available: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2016.0238

[12] D. Budker and M. Romalis, "Optical magnetometry," *Nature Physics*, vol. 3, pp. 227 EP –, Apr 2007, review Article. [Online]. Available: https://doi.org/10.1038/nphys566

[13] H. Häffner, C. Roos, and R. Blatt, "Quantum computing with trapped ions," *Physics Reports*, vol. 469, no. 4, pp. 155 – 203, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0370157308003463

[14] E. Hecht, *Optics*, 6th ed. Addison-Wesley, 2002.

[15] R. D. Guenther, *Modern Optics*. Wiley VCH, Jan. 1990.

[16] J. Hebling, "Derivation of the pulse front tilt caused by angular dispersion," *Optical and Quantum Electronics*, vol. 28, no. 12, pp. 1759–1763, Dec 1996. [Online]. Available: https://doi.org/10.1007/BF00698541

[17] P. Duffieux, *The Fourier Transform and Its Applications to Optics*. UMI, 1983. [Online]. Available: https://books.google.co.uk/books?id=-yEmnQEACAAJ

[18] E. A. Sziklas and A. E. Siegman, "Mode calculations in unstable resonators with flowing saturable gain. 2: Fast fourier transform method," *Appl. Opt.*, vol. 14, no. 8, pp. 1874–1889, Aug 1975. [Online]. Available: http://ao.osa.org/abstract.cfm?URI=ao-14-8-1874

[19] F. Pampaloni and J. Enderlein, "Gaussian, Hermite-Gaussian, and Laguerre-Gaussian beams : A primer," 2004.

[20] 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," *Phys. Rev. A*, vol. 45, pp. 8185–8189, Jun 1992. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.45.8185

[21] L. Allen, S. Barnett, and M. Padgett, *Optical Angular Momentum*, ser. Optics & Optoelectronics. Taylor & Francis, 2003. [Online]. Available: https://books.google.co.uk/books?id=Vf32PZXJ2gMC

[22] N. Radwell, R. F. Offer, A. Selyem, and S. Franke-Arnold, "Optimisation of arbitrary light beam generation with spatial light modulators," *Journal of Optics*, vol. 19, no. 9, p. 095605, aug 2017. [Online]. Available: https://doi.org/10.1088%2F2040-8986%2Faa7f50

[23] T. W. Clark, R. F. Offer, S. Franke-Arnold, A. S. Arnold, and N. Radwell, "Comparison of beam generation techniques using a phase only spatial light modulator," *Opt. Express*, vol. 24, no. 6, pp. 6249–6264, Mar 2016. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-24-6-6249

[24] J. A. Davis, D. M. Cottrell, J. Campos, M. J. Yzuel, and I. Moreno, "Encoding amplitude information onto a phase-only filter," *Appl. Opt*, 1999.

[25] N. Radwell, R. D. Hawley, J. B. Götte, and S. Franke-Arnold, "Achromatic vector vortex beams from a glass cone," *Nat. Commun.*, vol. 7, p. 10564, 2016.

[26] von F. Zernike, "Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode," *Physica*, vol. 1, no. 7, pp. 689 – 704, 1934. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0031891434802595

[27] R. W. Bowman, A. J. Wright, and M. J. Padgett, "An SLM-based shack–hartmann wavefront sensor for aberration correction in optical tweezers," *Journal of Optics*, vol. 12, no. 12, p. 124004, nov 2010. [Online]. Available: https://doi.org/10.1088%2F2040-8978%2F12%2F12%2F124004

[28] S. Wise, V. Quetschke, A. J. Deshpande, G. Mueller, D. H. Reitze, D. B. Tanner, B. F. Whiting, Y. Chen, A. Tünnermann, E. Kley, and T. Clausnitzer, "Phase effects in the diffraction of light: Beyond the grating equation," *Phys. Rev. Lett.*, vol. 95, p. 013901, Jun 2005. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.95.013901

[29] K. J. Mitchell, S. Turtaev, M. J. Padgett, T. Čižmár, and D. B. Phillips, "High-speed spatial control of the intensity, phase and polarisation of vector beams using a digital micro-mirror device," *Opt. Express*, vol. 24, no. 25, pp. 29 269–29 282, 2016.

[30] L. Marrucci, C. Manzo, and D. Paparo, "Optical spin-to-orbital angular momentum conversion in inhomogeneous anisotropic media," *Phys. Rev. Lett.*, vol. 96, p. 163905, 2006.

[31] A. Selyem, C. Rosales-Guzmán, S. Croke, A. Forbes, and S. Franke-Arnold, "Basis independent tomography of complex vectorial light fields by stokes projections," *arXiv preprint arXiv:1902.07988*, 2019.

[32] R. J. C. Spreeuw, "A classical analogy of entanglement," *Foundations of Physics*, vol. 28, no. 3, pp. 361–374, Mar 1998. [Online]. Available: https://doi.org/10.1023/A:1018703709245

[33] A. Holleczek, A. Aiello, C. Gabriel, C. Marquardt, and G. Leuchs, "Classical and quantum properties of cylindrically polarized states of light," *Opt. Express*, vol. 19, p. 97149736, 2011.

[34] X.-F. Qian, B. Little, J. C. Howell, and J. H. Eberly, "Shifting the quantum-classical boundary: theory and experiment for statistically classical optical fields," *Optica*, vol. 2, no. 7, pp. 611–615, Jul 2015. [Online]. Available: http://www.osapublishing.org/optica/abstract.cfm?URI=optica-2-7-611

[35] P. Li, B. Wang, and X. Zhang, "High-dimensional encoding based on classical nonseparability," *Opt. Express*, vol. 24, no. 13, p. 15143, 2016.

[36] B. Perez-Garcia, J. Francis, M. McLaren, R. I. Hernandez-Aranda, A. Forbes, and T. Konrad, "Quantum computation with classical light: The Deutsch Algorithm," *Phys. Lett. A*, vol. 379, pp. 1675–1680, 2015.

[37] S. Berg-Johansen, F. Töppel, B. Stiller, P. Banzer, M. Ornigotti, E. Giacobino, G. Leuchs, A. Aiello, and C. Marquardt, "Classically entangled optical beams for high-speed kinematic sensing," *Optica*, vol. 2, no. 10, pp. 864–868, Oct 2015. [Online]. Available: http://www.osapublishing.org/optica/abstract.cfm?URI=optica-2-10-864

[38] C. Rosales-Guzmán, B. Ndagano, and A. Forbes, "A review of complex vector light fields and their applications," *Journal of Optics*, vol. 20, no. 12, p. 123001, 2018. [Online]. Available: http://stacks.iop.org/2040-8986/20/i=12/a=123001

[39] M. Meier, V. Romano, and T. Feurer, "Material processing with pulsed radially and azimuthally polarized laser radiation," *Appl. Phys. A*, vol. 86, pp. 329–334, 2007.

[40] Q. Zhan, "Cylindrical vector beams: from mathematical concepts to applications," *Adv. Opt. Photonics*, vol. 1, pp. 1–57, 2009.

[41] M. Woerdemann, C. Alpmann, M. Esseling, and C. Denz, "Advanced optical trapping by complex beam shaping," *Laser Photon. Rev.*, vol. 7, no. 6, pp. 839–854, 2013.

[42] R. Chen, K. Agarwal, C. J. R. Sheppard, and X. Chen, "Imaging using cylindrical vector beams in a high-numerical-aperture microscopy system." *Optics letters*, vol. 38, no. 16, pp. 3111–4, 2013.

[43] H. Rubinsztein-Dunlop, A. Forbes, M. V. Berry, M. R. Dennis, D. L. Andrews, M. Mansuripur, C. Denz, C. Alpmann, P. Banzer, and T. Bauer, "Roadmap on structured light," *Journal of Optics*, vol. 19, no. 1, p. 013001, 2017. [Online]. Available: http://stacks.iop.org/2040-8986/19/i=1/a=013001

[44] G. Milione, M. P. Lavery, H. Huang, Y. Ren, G. Xie, T. A. Nguyen, E. Karimi, L. Marrucci, D. A. Nolan, R. R. Alfano *et al.*, "4×20 Gbit/s mode division multiplexing over free space using vector modes and a q-plate mode (de) multiplexer," *Opt. Lett.*, vol. 40, no. 9, pp. 1980–1983, 2015.

[45] V. D'Ambrosio, G. Carvacho, F. Graffitti, C. Vitelli, B. Piccirillo, L. Marrucci, and F. Sciarrino, "Entangled vector vortex beams," *Phys. Rev. A*, vol. 36, p. 030304(R), 2016.

[46] B. Ndagano, B. Perez-Garcia, F. S. Roux, M. McLaren, C. Rosales-Guzmán, Y. Zhang, O. Mouane, R. I. Hernandez-Aranda, T. Konrad, and A. Forbes, "Characterizing quantum channels with non-separable states of classical light," *Nature Phys.*, vol. 13, pp. 397–402, 2017.

[47] B. Ndagano, I. Nape, M. A. Cox, C. Rosales-Guzmán, and A. Forbes, "Creation and detection of vector vortex modes for classical and quantum communication," *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 292–301, Jan 2018.

[48] A. D'Errico, M. Mafei, B. Piccirillo, C. de Lisio, F. Cardano, and L. Marrucci, "Topological features of vector vortex beams perturbed with uniformly polarized light," *Sci. Rep.*, vol. 7, p. 40195, 2017.

[49] J. G. Ureta, B. Khajavi, and E. J. Galvez, "Decoding vortex beams with shear interferometry," in *Latin America Optics and Photonics Conference*. Optical Society of America, 2018, p. W3A.2. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=LAOP-2018-W3A.2

[50] C. V. S. Borges, M. Hor-Meyll, J. A. O. Huguenin, and A. Z. Khoury, "Bell-like inequality for the spin-orbit separability of a laser beam," *Phys. Rev. A*, vol. 82, p. 033833, Sep 2010. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.82.033833

[51] M. McLaren, T. Konrad, and A. Forbes, "Measuring the nonseparability of vector vortex beams," *Phys. Rev. A*, vol. 92, p. 023833, 2015.

[52] B. Ndagano, H. Sroor, M. McLaren, C. Rosales-Guzmán, and A. Forbes, "Beam quality measure for vector beams," *Opt. Lett.*, vol. 41, p. 3407, 2016.

[53] W. Wootters, "Entanglement of formation and concurrence," *Quantum Information and Computation*, vol. 1, no. 1, pp. 27–44, 2001.

[54] A. M. Beckley, T. G. Brown, and M. A. Alonso, "Full Poincaré beams," *Opt. Express*, vol. 18, no. 10, pp. 10 777–10 785, May 2010. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-18-10-10777

[55] S. Lukman Winoto, M. T. DePue, N. Bramall, and D. S. Weiss, "Laser cooling at high density in deep far-detuned optical lattices," *Physical Review A*, vol. 59, 01 1999.

[56] A. V. Andreev, V. I. Emel'yanov, and Y. A. Il'inskiĭ, "Collective spontaneous emission (Dicke superradiance)," *Soviet Physics Uspekhi*, vol. 23, no. 8, pp. 493–514, aug 1980. [Online]. Available: https://doi.org/10.1070%2Fpu1980v023n08abeh005024

[57] W. Guerin, M. O. Araújo, and R. Kaiser, "Subradiance in a large cloud of cold atoms," *Phys. Rev. Lett.*, vol. 116, p. 083601, Feb 2016. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.116.083601

[58] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman, "Observation of Rydberg blockade between two atoms," *Nature Physics*, vol. 5, pp. 110 EP –, Jan 2009. [Online]. Available: https://doi.org/10.1038/nphys1178

[59] D. A. Steck, "Rubidium 85 D Line Data," 2013, available online at http://steck.us/alkalidata (revision 2.1.6, 20 September 2013).

[60] C. Foot and D. Foot, *Atomic Physics*, ser. Oxford Master Series in Physics. OUP Oxford, 2005. [Online]. Available: https://books.google.co.uk/books?id=_CoSDAAAQBAJ

[61] Z. Minev, S. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. Schoelkopf, M. Mirrahimi, H. Carmichael, and M. Devoret, "To catch and reverse a quantum jump mid-flight," *arXiv preprint arXiv:1803.00545*, 2018.

[62] C. Cohen-Tannoudji and D. Gury-Odelin, *Advances in Atomic Physics*. WORLD SCIENTIFIC, 2011. [Online]. Available: https://www.worldscientific.com/doi/abs/10.1142/6631

[63] M. Jones, "Atom-Light Interactions. Archived lecture notes," 2015. [Online]. Available: http://community.dur.ac.uk/thomas.billam/PreviousNotes_MPAJones.pdf

[64] P. F. Moulton, "Spectroscopic and laser characteristics of Ti:Al2O3," *J. Opt. Soc. Am. B*, vol. 3, no. 1, pp. 125–133, Jan 1986. [Online]. Available: http://josab.osa.org/abstract.cfm?URI=josab-3-1-125

[65] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich, "Generation of optical harmonics," *Phys. Rev. Lett.*, vol. 7, pp. 118–119, Aug 1961. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.7.118

[66] A. Arnold, J. S. Wilson, and M. G. Boshier, "A simple extended-cavity diode laser," *Review of Scientific Instruments*, vol. 69, 03 1998.

[67] C. J. Hawthorn, K. P. Weber, and R. E. Scholten, "Littrow configuration tunable external cavity diode laser with fixed direction output beam," *Review of Scientific Instruments*, vol. 72, no. 12, pp. 4477–4479, 2001. [Online]. Available: https://doi.org/10.1063/1.1419217

[68] D. W. Preston, "Doppler-free saturated absorption: Laser spectroscopy," *American Journal of Physics*, vol. 64, no. 11, pp. 1432–1436, 1996.

[69] D. A. Smith and I. G. Hughes, "The role of hyperfine pumping in multilevel systems exhibiting saturated absorption," *American Journal of Physics*, vol. 72, no. 5, pp. 631–637, 2004. [Online]. Available: https://doi.org/10.1119/1.1652039

[70] A. Arnold, "Preparation and Manipulation of an ^{87}Rb Bose-Einstein Condensate," Ph.D. dissertation, University of Sussex, 1999.

[71] A. Vernier, "Phase dependent atom optics," Ph.D. dissertation, University of Glasgow, 2011.

[72] N. Radwell, D. Brickus, T. W. Clark, and S. Franke-Arnold, "High speed switching between arbitrary spatial light profiles," *Opt. Express*, vol. 22, no. 11, pp. 12 845–12 852, Jun 2014. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-22-11-12845

[73] E. A. Donley, T. P. Heavner, F. Levi, M. Tataw, and S. R. Jefferts, "Double-pass acousto-optic modulator system," *Review of Scientific Instruments*, vol. 76, no. 6, p. 063112, 2005.

[74] N. Radwell, M. A. Boukhet, and S. Franke-Arnold, "3d beam reconstruction by fluorescence imaging," *Opt. Express*, vol. 21, no. 19, pp. 22 215–22 220, 2013.

[75] A. Selyem, S. Fayard, T. W. Clark, A. S. Arnold, N. Radwell, and S. Franke-Arnold, "Holographically controlled three-dimensional atomic population patterns," *Opt. Express*, vol. 26, no. 14, pp. 18 513–18 522, Jul 2018. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-26-14-18513

[76] W. Nagourney, J. Sandberg, and H. Dehmelt, "Shelved optical electron amplifier: Observation of quantum jumps," *Phys. Rev. Lett.*, vol. 56, p. 2797, 1986.

[77] T. Sauter, W. Neuhauser, R. Blatt, and P. E. Toschek, "Observation of quantum jumps," *Phys. Rev. Lett.*, vol. 57, p. 1696, 1986.

[78] D. A. Steck, "Rubidium 85 D Line Data," *http://steck.us/alkalidata*, 2013.

[79] J. Radon, "Über die bestimmung von funktionen durch ihre integralwerte lngs gewisser mannigfaltigkeiten," *Ber. Ver. Sächs. Akad. Wiss. Leipzig, Math-Phys. Kl.*, vol. 69, pp. 262–277, 1917.

[80] M. W. Groch and W. D. Erwin, "SPECT in the year 2000: basic principles," *J. Nucl. Med. Technol.*, vol. 28, no. 4, pp. 233–244, 2000.

[81] S. Franke-Arnold, J. Leach, M. J. Padgett, V. E. Lembessis, D. Ellinas, a. J. Wright, J. M. Girkin, P. Ohberg, and a. S. Arnold, "Optical ferris wheel for ultracold atoms." *Optics express*, vol. 15, pp. 8619–8625, 2007.

[82] F. Wolfgramm, A. Cerè, F. A. Beduini, A. Predojević, M. Koschorreck, and M. W. Mitchell, "Squeezed-light optical magnetometry," *Phys. Rev. Lett.*, vol. 105, p. 053601, Jul 2010. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.105.053601

[83] O. Katz and O. Firstenberg, "Light storage for one second in room-temperature alkali vapor," *Nature Communications*, vol. 9, no. 1, p. 2074, 2018. [Online]. Available: https://doi.org/10.1038/s41467-018-04458-4

[84] M. V. Balabas, K. Jensen, W. Wasilewski, H. Krauter, L. S. Madsen, J. H. Müller, T. Fernholz, and E. S. Polzik, "High quality anti-relaxation coating material for alkali atom vapor cells," *Opt. Express*, vol. 18, no. 6, pp. 5825–5830, Mar 2010. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-18-6-5825

[85] M. Bhattarai, V. Bharti, V. Natarajan, A. Sargsyan, and D. Sarkisyan, "Study of EIT resonances in an anti-relaxation coated Rb vapor cell," *Physics Letters A*, vol. 383, no. 1, pp. 91 – 96, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0375960118309964

[86] J. W. Lou and G. A. Cranch, "Characterization of atomic spin polarization lifetime of cesium vapor cells with neon buffer gas," *AIP Advances*, vol. 8, no. 2, p. 025305, 2018. [Online]. Available: https://doi.org/10.1063/1.5010294

[87] R. Wynands and S. Weyers, "Atomic fountain clocks," *Metrologia*, vol. 42, no. 3, pp. S64–S79, jun 2005. [Online]. Available: https://doi.org/10.1088%2F0026-1394%2F42%2F3%2Fs08

[88] T. Hänsch and A. Schawlow, "Cooling of gases by laser radiation," *Optics Communications*, vol. 13, no. 1, pp. 68 – 69, 1975. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0030401875901595

[89] D. J. Wineland and H. Dehmelt, "Proposed 10^{14} Δv¡v Laser Fluorescence Spectroscopy on Tl+ Mono-Ion Oscillator III," *Bulletin of the American Physical Society*, vol. 20, p. 647, 1975.

[90] C. Adams and E. Riis, "Laser cooling and trapping of neutral atoms," *Progress in Quantum Electronics*, vol. 21, no. 1, pp. 1 – 79, 1997. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0079672796000067

[91] J. Dalibard and C. Cohen-Tannoudji, "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models," *J. Opt. Soc. Am. B*, vol. 6, no. 11, pp. 2023–2045, Nov 1989. [Online]. Available: http://josab.osa.org/abstract.cfm?URI=josab-6-11-2023

[92] P. J. Ungar, D. S. Weiss, E. Riis, and S. Chu, "Optical molasses and multilevel atoms: theory," *J. Opt. Soc. Am. B*, vol. 6, no. 11, pp. 2058–2071, Nov 1989. [Online]. Available: http://josab.osa.org/abstract.cfm?URI=josab-6-11-2058

[93] M. Gajda and J. Mostowski, "Three-dimensional theory of the magneto-optical trap: Doppler cooling in the low-intensity limit," *Phys. Rev. A*, vol. 49, pp. 4864–4875, Jun 1994. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.49.4864

[94] H. Youk, "Numerical study of quadrupole magnetic traps for neutral atoms: anti-Helmholtz coils and a U-chip," *Canadian Undergraduate Physics Journal*, 2005.

[95] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," *Phys. Rev. Lett.*, vol. 59, pp. 2631–2634, Dec 1987. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.59.2631

[96] M. Vangeleyn, P. F. Griffin, E. Riis, and A. S. Arnold, "Single-laser, one beam, tetrahedral magneto-optical trap," *Opt. Express*, 2009.

[97] A. S. Arnold, "Extending dark optical trapping geometries," *Opt. Lett.*, vol. 37, p. 2505, 2012.

[98] L. Podlecki, R. D. Glover, J. Martin, and T. Bastin, "Radiation pressure on a two-level atom: an exact analytical approach," *J. Opt. Soc. Am. B*, vol. 35, no. 1, pp. 127–132, Jan 2018. [Online]. Available: http://josab.osa.org/abstract.cfm?URI=josab-35-1-127

[99] E. B. Rosa and F. W. Grover, *Formulas and Tables for the Calculation of Mutual and Self-Inductance, 3rd Edition.* NBS Scientific paper No. 169, 1916.

[100] J. Clerk-Maxwell, *A Treatise on Electricity and Magnetism, Vol. 2, 3rd Edition.* Dover, 1954.

[101] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard, "High densities of cold atoms in a dark spontaneous-force optical trap," *Phys. Rev. Lett.*, vol. 70, pp. 2253–2256, Apr 1993. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.70.2253

[102] S. Muniz, K. M. aes, E. Henn, L. Marcassa, and V. Bagnato, "Creating a self-induced dark spontaneous-force optical trap for neutral atoms," *Optics Communications*, vol. 235, no. 4, pp. 333 – 340, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0030401804002111

[103] N. Radwell, G. Walker, and S. Franke-Arnold, "Cold-atom densities of more than 10^{12} cm^{-3} in a holographically shaped dark spontaneous-force optical trap," *Phys. Rev. A*, vol. 88, no. 4, p. 043409, 2013.

[104] T. Walker, D. Sesko, and C. Wieman, "Collective behavior of optically trapped neutral atoms," *Phys. Rev. Lett.*, vol. 64, pp. 408–411, Jan 1990. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.64.408

[105] N. Radwell, T. W. Clark, B. Piccirillo, S. M. Barnett, and S. Franke-Arnold, "Spatially dependent electromagnetically induced transparency," *Phys. Rev. Lett.*, 2015.

[106] J. Lenz and S. Edelstein, "Magnetic sensors and their applications," *IEEE Sensors Journal*, vol. 6, no. 3, pp. 631–649, June 2006.

[107] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis, "High-sensitivity atomic magnetometer unaffected by spin-exchange relaxation," *Phys. Rev. Lett.*, vol. 89, p. 130801, Sep 2002. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.89.130801

[108] W. E. Bell and A. L. Bloom, "Optically driven spin precession," *Phys. Rev. Lett.*, vol. 6, pp. 280–281, Mar 1961. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.6.280

[109] T. M. Tierney, N. Holmes, S. Mellor, J. D. Lpez, G. Roberts, R. M. Hill, E. Boto, J. Leggett, V. Shah, M. J. Brookes, R. Bowtell, and G. R. Barnes, "Optically pumped magnetometers: From quantum origins to multi-channel magnetoencephalography," *NeuroImage*, vol. 199, pp. 598 – 608, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811919304550

[110] V. Gerginov, S. Krzyzewski, and S. Knappe, "Pulsed operation of a miniature scalar optically pumped magnetometer," *J. Opt. Soc. Am. B*, vol. 34, no. 7, pp. 1429–1434, Jul 2017. [Online]. Available: http://josab.osa.org/abstract.cfm?URI=josab-34-7-1429

[111] T. W. Clark, "Sculpting shadows. On the spatial structuring of fields & atoms: a tale of light and darkness," Ph.D. dissertation, University of Glasgow, 2016.

[112] "Electromagnetically induced transparency. Archived lecture notes," 2014. [Online]. Available: https://www.kip.uni-heidelberg.de/kw/image/f/group/f17/files/F65_EIT_Anleitung.pdf

[113] W. Erickson, "Electromagnetically induced transparency. Bachelor's Thesis," 2012.

[114] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, "Electromagnetically induced transparency: Optics in coherent media," *Rev. Mod. Phys.*, vol. 77, pp. 633–673, Jul 2005. [Online]. Available: https://link.aps.org/doi/10.1103/RevModPhys.77.633

[115] 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*, vol. 397, no. 6720, pp. 594–598, 1999. [Online]. Available: https://doi.org/10.1038/17561

[116] M. Fleischhauer and M. D. Lukin, "Dark-state polaritons in electromagnetically induced transparency," *Phys. Rev. Lett.*, vol. 84, pp. 5094–5097, May 2000. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.84.5094

[117] C. Liu, Z. Dutton, C. H. Behroozi, and L. V. Hau, "Observation of coherent optical information storage in an atomic medium using halted light pulses," *Nature*, vol. 409, pp. 490–493, 2001.

[118] S. J. Buckle, S. M. Barnett, P. L. Knight, M. A. Lauder, and D. T. Pegg, "Atomic Interferometers: Phase-dependence in Multilevel Atomic Transitions," *Optica Acta*, vol. 33, pp. 1129–1140, Sep. 1986.

[119] V. I. Yudin, A. V. Taichenachev, Y. O. Dudin, V. L. Velichansky, A. S. Zibrov, and S. A. Zibrov, "Vector magnetometry based on electromagnetically induced transparency in linearly polarized light," *Phys. Rev. A*, vol. 82, p. 033807, Sep 2010. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.82.033807

[120] J. Gallop and B. Petley, "SQUIDs and their applications," *Journal of Physics E: Scientific Instruments*, vol. 9, no. 6, p. 417, 1976.

[121] C. Schneider, M. Enderlein, T. Huber, and T. Schaetz, "Optical trapping of an ion," *Nature Photonics*, vol. 4, pp. 772 EP –, Oct 2010. [Online]. Available: https://doi.org/10.1038/nphoton.2010.236

[122] P. Banzer, U. Peschel, S. Quabis, and G. Leuchs, "On the experimental investigation of the electric and magnetic response of a single nano-structure," *Opt. Express*, vol. 18, no. 10, pp. 10 905–10 923, 2010.

[123] T. Bauer, P. Banzer, E. Karimi, S. Orlov, A. Rubano, L. Marrucci, E. Santamato, R. W. Boyd, and G. Leuchs, "Observation of optical polarization Möbius strips," *Science*, vol. 347, no. 6225, pp. 964–966, 2015. [Online]. Available: http://science.sciencemag.org/content/347/6225/964

[124] T. W. Hänsch, "Pulsed tunable dye laser for high resolution spectroscopy," *Applied Optics*, vol. 11, pp. 895–898, 1972.

[125] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt, "Optical-sideband cooling of visible atom cloud confined in parabolic well," *Phys. Rev. Lett.*, vol. 41, pp. 233–236, Jul 1978. [Online]. Available: http://link.aps.org/doi/10.1103/PhysRevLett.41.233

[126] D. J. Wineland, R. E. Drullinger, and F. L. Walls, "Radiation-pressure cooling of bound resonant absorbers," *Phys. Rev. Lett.*, vol. 40, pp. 1639–1642, Jun 1978. [Online]. Available: http://link.aps.org/doi/10.1103/PhysRevLett.40.1639

[127] F. M. Penning, "Die glimmentladung bei niedrigem druck zwischen koaxialen zylindern in einem axialen magnetfeld," *Physica (Amsterdam)*, vol. 3, p. 873, 1936.

[128] W. Paul, "Electromagnetic traps for charged and neutral particles," *Reviews of Modern Physics*, vol. 62, no. 3, 1990.

[129] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions," *Reviews of Modern Physics*, vol. 75, no. 1, pp. 281–324, 2003.

[130] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt, "Laser cooling of trapped ions," *Journal of the Optical Society of America B*, vol. 20, no. 5, p. 1003, 2003.

[131] C. Hempel, "Digital quantum simulation, Schrödinger cat state spectroscopy and setting up a linear ion trap," Ph.D. dissertation, Leopold-Franzens University of Innsbruck, August 2014.

[132] J. I. Cirac, L. J. Garay, R. Blatt, A. S. Parkins, and P. Zoller, "Laser cooling of trapped ions: The influence of micromotion," *Physical Review A*, vol. 49, no. 1, pp. 421–432, 1994.

[133] R. J. Glauber, "Coherent and incoherent states of the radiation field," *Physical Review*, vol. 131, no. 638, 1963.

[134] Y. Li, X.-L. Wang, H. Zhao, L.-J. Kong, K. Lou, B. Gu, C. Tu, and H.-T. Wang, "Young's two-slit interference of vector light fields," *Opt. Lett.*, vol. 37, no. 11, pp. 1790–1792, Jun 2012. [Online]. Available: http://ol.osa.org/abstract.cfm?URI=ol-37-11-1790