UNIVERSITY OF SOUTHAMPTON

Coherent two photon excitation within an extended cloud of Rubidium 85 for the purposes of atomic interferometry and cooling

by

Richard Murray

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Science, Engineering and Mathematics

School of Physics and Astronomy

April 2011

"Shall I compare thee to a summer's day?
Thou art more lovely and more temperate.
Rough winds do shake the darling buds of May,
And summer's lease hath all too short a date.
Sometime too hot the eye of heaven shines,
And often is his gold complexion dimmed;
And every fair from fair sometime declines,
By chance, or nature's changing course untrimmed.
But thy eternal summer shall not fade
Nor lose possession of that fair thou ow'st;
Nor shall death brag thou wand'rest in his shade,
When in eternal lines to time thou grow'st,
So long as men can breathe or eyes can see,
So long lives this, and this gives life to thee"

William Shakespeare

UNIVERSITY OF SOUTHAMPTON

Abstract

Faculty of Science, Engineering and Mathematics
School of Physics and Astronomy

Doctor of Philosophy

by Richard Murray

Cold atom samples, at temperatures of the order 100 μK, are useful for a wide reaching array of new and exciting technological and scientific endeavours.

Atoms are conventionally cooled by Doppler cooling, which relies on the continuous absorption and reemission of photons in a closed optical cycle. This requirement is difficult to achieve when there are many allowed decay paths for the excited atom, making Doppler cooling only feasible for a handful atoms with simple energy level structures. More exotic energy level structures, such as those found in molecules, are notoriously difficult to cool.

Coherent cooling schemes in comparison, offer advantages such as insensitivity to frequency detuning or a higher number of photon momenta which can be imparted for each spontaneous emission event making them promising candidates for the optical cooling of particles with more general energy level structures. In order to demonstrate these schemes, we have explored the coherent manipulation of an atomic cloud of Rubidium cooled using a two photon Raman resonance. Despite the long spontaneous decay times of such systems, we find a significant decay in the fidelity of the coherent manipulations, which we have characterised using the techniques of Raman spectroscopy, Rabi oscillations, Ramsey interferometry and spin echo. We have found the minimum time constant for the decay in the decoherence to be 2.1$\pm$0.2 ms, which is a result of non-radiative and partially non-stochastic dephasing mechanisms. Due to a high level of decoherence during the spin-echo experiments further investigation is required to determine the exact ratio of stochastic to non-stochastic dephasing.

Acknowledgements

Running any atomic physics experiment is always a team effort, and for that I am extremely lucky to have worked with a group of people who are not only are very talented, but who are also a pleasure to work with. For this reason, I am hugely indebted to Tim Freegarde, James Bateman, Matt Himsworth, Hamid Ohadi, André Xuereb, Nathan Cooper and Sunil Patel, who between them I think I owe half a brewery in beers for their help.

I would like to thank my internal and external examiners, Hendrik Ulbricht and Peter Barker for their interesting discussion of my work, and for their advice on improving it.

I would not be the person I am today without my parents who I would like to thank for their constant love and support, my Nan, who was never sure what it was that I did, but was pretty damn sure that it was something *important*, my sister, Anna, who showed me how they do things in London and my friends, George, Peta, Mark, Ben, Adam, Andy, Jackie, Cheryl, Mrs Green, Rach, Matt, Jen, Ham, Dr Phil and all those at Welbeck I love you all, thanks for puting up with me. (-:.

Contents

For Nan

Chapter 1

Introduction

1.1 The uses of cold atoms and molecules

Ultra cold atoms, at temperatures of a milli-Kelvin or less, serve a number of diverse and far reaching applications in modern physics and chemistry. The methods currently used to cool atoms to these ultra low temperatures, however, are severely limited to a very small percentage of atoms and molecules. As a result of this, new techniques which extend the scope of the traditional cooling methods, particularly those which can be applied to more diverse species, are desirable.

As a result of their lower entropy, low temperature gases behave very differently from their room temperature counterparts, which makes them easier to trap and manipulate. Standing optical waves can be used to produce periodic lattice potentials as a result of atomic dipole forces, which can be used to model solid state physics, the lattice depth being tunable via the beam intensity. This can, for example, be used to study the transitions between non-interacting atoms in a Mott insulator phase and strongly correlated atoms in a superfluid phase [1, 2]. These studies promise a better understanding of many macroscopic solid state phenomena, the most prolific and notorious examples of this being superfluidity [3] and high temperature superconductivity [4].

In addition to this, cold gases allow the study of reaction dynamics in chemical interactions; for non spherically symmetric reactants, such as molecules, the interaction strength is dependent upon their relative orientation, which for cold particles can be controlled through externally applied electric fields [5]. The reaction products of these aligned gases can be analysed by a number of processes such as velocity map imaging [6] to give a better indication of the reaction dynamics. Furthermore, state preparation can be used to control to electronic state of the reactants and the products of a reaction, further increasing the control over reaction rates.

During cold reactions, it is necessary to use a full quantum description of the particle, rather than hard bodied spheres, to account for phenomenon such as quantum tunneling. Such work has led to the direct control of reactions between ultra-cold mixtures of weakly bound dimer molecules and free atoms, allowing the final products to be tuned via an externally applied magnetic field [7]. In one specific case of controlled dynamics, theory has been used to show that judicious control over the atomic properties can create entangled pairs of atoms [8].

When excitation energy for the first excited state of a harmonic potential becomes comparable to the external energy of an atom, such as in a cooled and trapped boson gas of a high enough phase space density, the spatial extent of the atomic wave packets overlap. In this situation, the atoms condense into what is known as a Bose-Einstein condensate (BEC). In this state, the atoms maintain a coherent, fixed phase relationship with other atoms in the condensate which is useful in the generation of coherent atom sources, such as for atom lasers [9] and in atomic interferometry [10, 11].

Amongst other properties, atoms have a much greater mass and dipole moment than light, which causes gravitational and electric fields to have a much larger purturbative effect on the particle wavefunction; it is because of this that atom interferometers are much more sensitive to many external environmental conditions than light interferometry. This way, accurate measurements of gravitational [12] and electric [13] field strength, rotation[14], the fine structure constant[15] and of the electron dipole moment [16] have been achieved.

In dilute gases such as those produced in magneto-optical trapping, atomic interactions are negligible over long time scales; the long coherence times obtained in these conditions are useful for quantum computation, which is expected to solve specific problems much faster than their conventional counterparts. So far small quantum computers made up of tens of optically cooled and trapped ions [17] and neutral atoms [18] have been implemented. This research is looking to expand in two ways, the first in scalability, increasing the numbers of atoms while reducing the scale of the equipment such as by using 'atom chips', and the second in applying gate operations [19] to the qubits, where the state of one qubit affects an operation applied to another.

Finally, the slow speed of a cold atomic beam is extremely useful in eliminating transit time broadening of spectral lines, which otherwise limits the accuracy of atomic clocks [20]. The clocks also gain from a lower dispersion, due to the low transverse velocities in a cold atom beam, which further improves the accuracy of the clock.

There is significant interest in replacing atoms in many of these experiments with ultra cold molecular to give greater accuracy: for example, although jets of cooled molecules

are harder to achieve, it has been shown that slowed jets of YbF [21] and PbO [22] molecules can increase the accuracy of the electron dipole measurements; these result would benefit from low temperature samples where the number density can be much higher. There have been successes in the generation of molecular BECs, such as with Lithium associated by Feshbach resonances [23], as well as the generation of high density, ground-state molecular gases in optical lattices [24]. New, more generalised, cooling methods would enable the study of new types of quantum gases such as degenerate polar gases.

1.2 Methods of generating cold atoms and the problems of cooling molecules

Of the many methods of generating samples of ultra-cold atoms, the use of a 'magneto optical trap' (MOT) is favoured for its low setup cost, compact design and reliability. The magneto-optical trap uses the continuous absorption and spontaneous-emission of a near resonant laser beam, which, together with an applied magnetic field gradient, gives a velocity and spatially dependent absorption cross section which causes a cooling and trapping force. Spontaneous emission is an important step in the cooling, making the process irreversible and therefore thermodynamically allowed; it also means, however, that the atom must possess a specific energy level structure, known as a closed optical loop, where each absorption-spontaneous emission event brings the atom back to the energy level from which it started. As this energy level structure is found in only a handful of atoms, magneto-optical trapping can only be used for a small fraction of the periodic table.

1.3 Methods of cold molecule production

Due to their high monoatomic to diatomic conversion efficiency of up to 50% [25], producing cold molecules through association is a particularly successful method of achieving high number densities. This association can be driven via photons [26] or Feshbach resonances [27]. The production of molecules in this manner, however, is complicated by de-excitation steps which must preceed all assosiation if the molecules are to remain in a bound state for any length of time. As the diversity of associated molecules is limited to the range Doppler coolable atoms, this technique is still not sufficiently diverse for many of the subject areas mentioned above.

Two alternatives for the cooling of molecules are sympathetic [28] and buffer gas [29] cooling which thermalise the molecules with atoms which have been cooled either cryogenically, or by Doppler cooling. These techniques are limited to low densities of cooled molecules as a result of the required filtration to remove the buffer gases after cooling.

The final two mechanisms to note are Stark deceleration [30] which achieves translational cooling to one Kelvin: too high for many applications, and evaporative cooling which causes a reduction in the atom number when the wings of the Gaussian velocity distribution are sacrificed to reduce the overall temperature of the gas.

While associative techniques can be used to achieve high density samples, they are limited to those species for which it is possible to cool optically, and therefore by the constraints of Doppler cooling. Alternative methods such as buffer gas and sympathetic cooling are able to cool molecules without any restrictions on the molecular structure, but are unable to produce a high enough densities necessary for the study of heavily interacting gases. The holy grail of cooling is to devise a method of cooling and trapping which combines both of these advantages: cooling without requiring any specific electronic structure to produce a high density sample.

1.4 Purely optical coherent interactions for cooling

In this thesis, we investigate coherent atom-light interactions with the intention of using these methods for the cooling of atoms and molecules of a generic energy level structure: the first method, 'coherent amplification of cooling' [31], uses a cycle of excitation and stimulated emission to cool the atoms down before spontaneous emission occurs, while the second technique, 'interferometric cooling' [32] is insensitive to detuning so that broadband excitation which can be used to couple to, and therefore cool on, many transitions simultaneously. These techniques have the potential to generate clouds of ultra-cold molecular gases with unprecedented temperatures and densities, which may offer substantially improved diversity and further miniaturisation over existing cooling methods.

1.4.1 The advantage of untrapped atoms for cooling

In order to implement these novel cooling schemes, we must first characterise these coherent light-matter interactions, specifically Rabi flopping, as a proof of principle, in a spatially-extended, freely falling and expanding cloud of laser cooled Rubidium 85 atoms.

While the study of the coherent interaction of light with alkali metal atoms is not new, experiments involving interactions with large numbers of un-trapped atoms, such as in our setup, is unreported. In attempting to drive large numbers of atoms without trapping, we encounter additional problems, not discussed in previous literature.

There have been recent studies on the feasibility using clouds of many atoms, of finite size, as qubits [33], using the electronic state of the ensemble to determine the value of the qubit. These systems exhibit many of the same challenges as we find in our spatially extended cloud; in this way, our studies are related to studies of large coherent-atom quantum computers [34].

There are a number of advantages to using large numbers of untrapped atoms. By driving these transitions in free atoms, we avoid any disturbances to the momentum distribution from the trap, such as from the introduction of additional vibrational sidebands. This facilitates the study of the momentum changes as a result of light forces alone. The number of atoms we are able to manipulate at any time is also much higher than in equivalent experiments using trapped atoms, giving a much higher signal, allowing smaller collection optics to be used. Along with the possibility of avoiding the need for bulky trapping equipment such as large lasers or magnetic field coils, using free atoms as opposed to trapped atoms opens the possibility of miniaturised atom optics experiments, which may give these techniques a future in portable quantum devices alongside 'MOT on a chip' experiments [35].

1.5 Outlook

This thesis is structured thus: the first chapters (2-3) look at atom-light interactions, and under what circumstances they can be used to construct a magneto-optical trap. We then look at the limits of the temperature and size of the magneto-optical trap, and compare them with the experimental parameters of our Rubidium 85 MOT. After this, we look at coherent optical interactions, the mathematical framework behind Rabi oscillations, population inversion and interferometry and how we propose to use these to achieve cooling. Our experimental setup is then discussed, along with the advantages and disadvantages of the methods we use for pulse synthesis.

As recent improvements in diode lasers, optical fibers, electro-optics and other telecommunications technologies have been important to the construction of our laboratory equipment, a large emphasis is given to describing our experimental apparatus. In particular, we discuss our method of phase modulation, which gives the multi-frequency

Raman light beams necessary for our setup and our fully computer controllable experimental interface, which enables the remote running, data acquisition and analysis of the experiment.

Finally, we discuss our experimental progress in interferometry and coherent interactions.

Chapter 2

Light, atoms and optical selection rules

When presented with the idea of laser cooling, many people on the street are confused having comfortably lived with the assumption that lasers are only able to add to the energy of a material. This assumption is not surprising, given the many literary examples [36] of lasers as devices which burn, vapourise or incinerate whichever hero is unfortunate enough to fall in their path.

The difference between the very different outcomes of laser-atom interactions is a result of the atomic electronic structure, which unlike molecules and solids is decoupled from internal rotational and vibrational states giving very poor coupling between the light field and the translational, vibrational and rotational energies of the particle.

In order to differentiate between these two regimes, it is important for us to understand the structure of the atomic energy levels, and their interactions with light.

2.1 Atomic structure

The discretisation of the atomic state by quantum physics can be described by a string of quantum numbers, n, l, m and s, which between them describe many of the outcomes of a physical observation on the atom. n ($\in \mathbb{Z}^+$), known as the principle quantum number, describes the radial part of the wavefunction, while l ($\in 0, 1, 2...(n-1)$)[1] and $m(\in -l, .., l)$ describe the angular dependence of the wavefunction. The symmetry of

[1]The set that these quantum numbers must belong to arises because of the need for a single valued wavefunction such that $\psi(\theta) = \psi(\theta + 2\pi)$. A full description can be found in any good text on quantum physics [37]

the Coulomb potential means that, to first a approximation, l and m do not affect the energy eigenstates of the Hamiltonian. Because of this, for each state with quantum number n, there are several degenerate states each with their own individual, orbital quantum number, l. The final quantum number, s $(= 1/2)$ represents the electron spin, a more profound property of the electron, which can be loosely described by a spinning charge distribution. Because of the non-spherical symmetry associated with the l and s quantum numbers, we introduce two additional quantum numbers, $m_l(\in -l, -l+1, ..., 0, ..., l-1, +l)$ and $m_s(= \pm 1/2)$ which describe the orientation of the wavefunction.

The alignment of the spin and orbital angular momentum of the electron with respect to each other, to the nuclear magnetic moment, and to an external applied field causes the energy of the states to depend on the product and orientation of these moments. This interaction has the effect of destroying the time independence of the initial quantum numbers, making it necessary to define new, time invariant quantum numbers. We refer to this interaction as 'coupling' of the different quantum states.

To solve the Schrödinger equation including all coupling potentials would be an extremely difficult task; instead we solve the Schrödinger equation for a central Coulomb potential and, assuming that each of the additional potentials represents only weak disturbances to the system, then add the additional coupling perturbations to the system.

Unlike classical physics, where the order in which these couplings are analysed is unimportant, the non-commutativity of quantum operators makes it vital, if this perturbation approach is to give the correct answers, that the perturbations are applied in order of their strength.

2.2 Spin-orbit interactions: the origin of fine splitting

The l quantum number does not affect the radial distance of the electron from the nucleus and so, to first approximation, does not affect the energy eigenstates of the system. The angular momentum of the electron in each orbital, however, affects the interaction between the electron orbital and spin magnetic moments, which we can see by considering the effect of coupling of the spin, $\vec{s}$ and magnetic field, $\vec{B}$ on the atomic Hamiltonian

$$H_{so} = g_s \mu_B \vec{s} \cdot \vec{B}, \tag{2.1}$$

which arises due to the interaction between the spin magnetic moment, a product of the electron spin g-factor, the Bohr magneton and the spin vector, $g_s \mu_B \vec{s}$ and the magnetic

field of the electron orbit,

$$\vec{B} = -\frac{1}{c^2}\vec{v} \times \vec{E}, \tag{2.2}$$

where c is the speed of light, $\vec{v}$ the velocity of the atom, and $\vec{E}$ the electric field strength. For a hydrogen atom, the electric field produced by the nucleus is given by

$$E = \frac{e}{4\pi\epsilon_0|\vec{r}^2|}\hat{r}. \tag{2.3}$$

Equations 2.2 and 2.3 can be substituted into 2.1 to give

$$H_{so} = -\frac{g_s\mu_B e}{4\pi\epsilon_0 c^2|\vec{r}^2|}\vec{s} \cdot \vec{\hat{r}} \times \vec{v} \tag{2.4}$$

This can be simplified by a substitution of the definition of orbital angular momentum

$$\vec{l} = \frac{m_e\vec{r} \times \vec{v}}{\hbar}, \tag{2.5}$$

The change in the Hamiltonian is hence

$$H_{so} = -\frac{-g_s\mu_B\hbar e}{4\pi\epsilon_0|r^3|c^2 m_e}\vec{s} \cdot \vec{l}. \tag{2.6}$$

As a result of the dot product, the relative orientation of $\vec{s}$ and $\vec{l}$ is important. In this regime, **s** and **l** are coupled, and their eigenstates are no longer constant with time. In this situation, we replace $\vec{s}$ and $\vec{l}$ with $\vec{j} = \vec{l} + \vec{s}$ which has a projection along the z axis given by $m_j = m_s + m_l$.

2.2.1 Hyperfine splitting

In addition to the spin and orbit magnetic moments, most atomic nuclei also have a magnetic moment, $\vec{i}$, which interacts with the $\vec{j}$ spin-orbit quantum number in order to produce an additional hyperfine energy splitting of the fine structure states. Although some of the atoms have zero nuclear moments, ^{85}Rb, which is the subject of this thesis, has a nuclear moment of 5/2. In a similar way to the interaction between $\vec{s}$ and $\vec{l}$, $\vec{j}$ and $\vec{i}$ interact to produce eigenstates which no longer commute with the Hamiltonian. This makes it necessary to define two new quantum numbers, $f = j + i$ and the projection along the z axis, $m_f = m_j + m_i$.

2.2.2 Zeeman splitting

If the atom is in a symmetrical environment, the projection axis, also known as the 'quantisation axis', for the m_f states is irrelevant, with no bearing on the energy eigenstates. By applying an external magnetic field, however, an additional shift is given to the atoms, according to the projection of m_f against the direction of the field. Again adopting the classical approach to the interaction of a dipole (this time including the spin, orbit and nuclear moments) with an external magnetic field, we get:

$$H_{Zeeman} = m_f g_F \mu_B \vec{f} \cdot \vec{B}, \tag{2.7}$$

where m_f is the projection of f along the z axis, and g_F is the Landé g-factor. This Zeeman shift will be of vital importance to us in this thesis. The trapping force of the magneto-optical trap is reliant on the Zeeman shift in atomic energy levels with applied field, which bring atoms displaced from the trap centre onto resonance with a restoring force. During our coherent manipulation experiments, we resolve separate resonances from the different m_f levels which we must address separately in order to give coherence. Due to their importance in many areas of experimental physics, the values of the fine

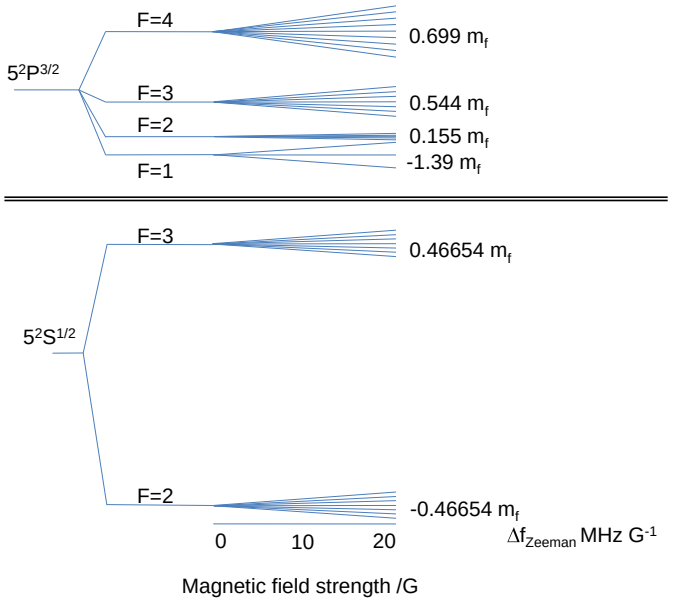

FIGURE 2.1: The Zeeman splitting of the Rubidium $5S^{1/2}$ and $5P^{3/2}$ hyperfine states by an applied static magnetic field. The energy shift is proportional to the applied magnetic field, with the proportionality constant, shown on the right of the diagram, given by the product of g_F and μ_B

and hyperfine splittings in Rubidium 85 have been measured experimentally using laser systems locked to very high quality cavities [38].

Quantum Number	Selection Rule	Polarisation
l	$\Delta l = \pm 1$	
m	$\Delta m = +1$	Circular:σ^+
	$\Delta m = -1$	Circular: σ^-
	$\Delta m = 0$	Linear: $\|\|$ z
	$\Delta m = \pm 1$	Linear: $\|\|$ (x,y)

TABLE 2.1: A table summarising the selection rules governing electric dipole transitions in a hydrogen like atom.

2.3 Light polarisation and the angular momentum of photons

Due to the dipole nature of light, transitions between energy levels have different strengths depending on the transition dipole moment [39]:

$$\vec{\mu}_{if} = \left\langle \vec{i} | \vec{\mu} | \vec{f} \right\rangle, \tag{2.8}$$

where $\vec{\mu} = -e\vec{r}$ is the electric dipole operator. As $\vec{\mu}$ is an odd function, the wavefunctions must have opposite parity, if the transition dipole is to be non-zero. This is important, as the parity of a wavefunction with angular momentum, l, is given by $(-1)^l$, so that, unless $\Delta l = \pm 1, 3, 5...$, the transition dipole will be equal to zero. Circularly polarised photons also carry angular momentum [40] which additionally constrains atom-optic interactions. Polarised light can be classified as either linearly or circularly polarised (or elliptically, which describes some combination of the two), and the circular polarisation category is further sub-divided into positively and negatively (also known as right and left handed) circularly polarisations, which denote the direction and the rotation of the optical electric field vector. For linearly polarised light, the direction of polarisation stays fixed, whereas for circular polarisations, the electric field vector rotates around the propagation axis. Positive circular polarisation sees a clockwise rotation of the electric field vector for as seen looking towards the source, and vice-versa for negative.

When interacting with an atom with a fixed quantisation axis (where the projection axis for the m quantum numbers is parallel to the direction of propagation for the photon), the polarisation of light describes the change in angular momentum of the atom, changing the m_l quantum number by ± 1 [41]. σ^+ causes a change of $+1$; σ^-, a change of -1 and a zero change in momentum for π polarised light.

2.4 Atomic selection rules

While a light field is applied to an atom, the bare states of the atom are no longer eigenstates of the system, which causes time evolution of the state populations. In this time, the overlap of the bare ground and excited state wavefunctions gives the probability that the atom will be excited. More intuitively, we can consider the conservation of momentum or parity during these excitation events from which it is clear that some must be forbidden: a consideration which forms the basis of atomic 'selection rules' which define transitions which are not driven by first order photon interactions. Due to the asymmetry of the electric dipole, for example, the change in the angular momentum quantum number must be an odd number if it is to conserve parity.

In addition to the dipole interactions which are described by selection rules, atom-photon interactions also have magnetic, electric quadrupole and other higher order terms, which allowing transitions between states which are not allowed by first order selection rules. These transitions occur with a much weaker coupling strength and are typically 10^5 times weaker for a magnetic quadrupole, and 10^8 times weaker for an electric quadrupole [39]. Table 2.1 is a table of the first order electric selection rules, with a brief description of their origin [41].

2.5 Atomic states nomenclature

Atomic physics has adopted a particular nomenclature for the naming of atomic states, so that each state can be labeled unambiguously. First, the principle quantum number, n, is given, after which the total orbital angular momentum of the state is defined as a capitalised letter: S, P, D, F (G, H and I) represent l= 0, 1, 2, 3, (4, 5 and 6)[2], the total spin and angular momentum, J, is defined as a subscript after the letter. Such a description does not include any information on the F and m_f state quantum numbers. Using Dirac notation, the rubidium lower hyperfine ground state can be written as:

$$|n = 5, l = 0, j = 1/2, F = 2, m_f = 0\rangle$$

or

$$|F = 2, m_f = 0\rangle,$$

for short, assuming prior knowledge of the principal quantum numbers. We will use both of these representations in this thesis.

[2]The letters and numbers in brackets refer to atomic orbitals which are not filled naturally.

Chapter 3

The dissipative effects of light-atom interactions

3.1 The momentum of a photon

Special relativity gives an energy-momentum relation that applies to all particles [42]

$$E^2 = p^2c^2 + m^2c^4, \tag{3.1}$$

where E is the energy of the particle, p its momentum, m its mass and c the speed of light. As the mass of the photon is zero, this can be re-written as:

$$E = \hbar\omega = pc, \tag{3.2}$$

so that

$$p = \frac{\hbar\omega}{c} = \hbar k, \tag{3.3}$$

where ω is the angular frequency of the light, $\hbar$ the reduced Planck constant, and k the wavenumber of the photon. The momentum of the photon is tiny, approximately 10^{-27} Nm, which is around 10^5 smaller than the typical momentum of a rubidium atom at room temperature. It is this that makes the effect of radiation pressure difficult to see.

In addition to the low impulse imparted by the absorption of each photon, most natural light sources posses a low spectral density and therefore only a small fraction of photons are absorbed by each atom. In these cases, unless the light source is very strong, an atom feels very little force. We find that even the famous Crookes radiometer does not turn due to radiation pressure, but as a result of the small difference in temperature

between the black and the white surfaces which causes a difference in the recoil velocities of the gas, and a net pressure on the paddles. Alternatively, by using a laser light source,

FIGURE 3.1: In space comets show up radiation pressure; as a comet passes the sun, gas is boiled from the surface and pushed by strong radiation pressure (as well as solar wind) away from the body into the tail, which point radially away from the sun [43].

for which the linewidth is sufficiently narrow that the probability of absorption is high, the atom can scatter many photons per second; this can cause a significant radiation pressure on the particles.

3.2 Light-atom interactions

In this section, we will give a brief description of the light forces within a one dimensional magneto-optical trap for a two-level atom. The actual system can be seen as an extension to this, and is described elsewhere [44].

3.2.1 The force of light

The force acting on an atom as a result of radiation pressure [45] is a product of the impulse given to the photon during a single interaction and the rate of the interactions:

$$F = \hbar k R_{scat}(\gamma, \delta), \tag{3.4}$$

where $R_{scat}(\gamma, \delta)$ is the photon scattering rate: a function of the detuning, δ and the linewidth, γ.

If we are interested in the steady state force acting on the atom, the rate must include both excitation and decay. For a simplified two-level atom in the limit of high intensities, Einstein showed in his famous paper of 1917 [46] that the population of incoherent[1] atoms in the excited state reaches the limit of a half. Increasing light intensity above this limit causes the rate of absorption to increase at the same rate as stimulated decay. This makes population inversion impossible, putting a limit on the maximum force achievable through incoherent interaction. The impulse given to the atom during stimulated emission is equal and opposite to that of the absorption and thus continuous absorption and stimulated emission causes no change to the atomic momentum. The maximum rate at which photon momentum can be imparted to an atom, therefore, is limited by the rate of spontaneous emission

$$R = \frac{\rho_{ee}}{\tau} = \frac{1}{2\tau}, \tag{3.5}$$

where τ is the decay time of the atom from the excited state, and ρ_{ee} is the population of the excited state, which tends to a half in the limit of intense driving fields. The maximum force acting on an atom is hence

$$F = \frac{\hbar k}{2\tau}. \tag{3.6}$$

In the derivation of equation 3.6, we assumed that the light intensity was sufficiently high that the detuning of the atom from resonance could be neglected. It therefore describes an environment where all atoms are subjected to an equal optical force. Although useful in describing changes in the centre of mass velocity of a cloud [47], it is insufficient to explain the cooling or trapping of gas clouds where the centre-of-mass velocity is already close to zero. For cooling, we require the force to be dependent on the atomic velocity distribution, which can be orchestrated by using the Doppler shift of the stationary energy levels.

[1] We will look at coherence later in chapter 4, but for now it is safe to say that an atomic ensemble that undergoes repeated spontaneous emission will be incoherent.

The detuning dependent absorption rate, $R = R(\gamma, \delta)$, given in equation 3.4, can be obtained from the overlap integral between the initial and final wavefunctions after perturbation by an optical electric field. The full expression is given by [44]:

$$R(\gamma, \delta) = \frac{\gamma s_0}{1 + s_0 + (\delta/\gamma)^2}, \tag{3.7}$$

which results in a force:

$$F = \hbar k \frac{\gamma s_0}{1 + s_0 + (\delta/\gamma)^2}, \tag{3.8}$$

where as before, δ is the detuning of the light source and γ is the natural linewidth. s_0 is the saturation intensity which is given by:

$$s_0 = \frac{2\Omega^2}{\gamma^2}. \tag{3.9}$$

For most laser light sources and thermal atoms, the linewidth is much narrower than the Doppler width and the detuning of each atom can be taken to be the difference between the atomic transition frequency, f_{atom}, and the frequency of the laser, f_{laser}:

$$\frac{\delta}{2\pi} = f_{atom} - f_{laser}. \tag{3.10}$$

The presence of δ in the denominator of equation 3.7 causes a strong dependence of the force on the detuning. For moving atoms (at room temperature, for example), the non-relativistic Doppler effect causes a change in the rest frame transition frequency which can be much greater than a laser linewidth:

$$\Delta f = \frac{v}{c} f = \frac{\sqrt{3 k_B T}}{\sqrt{m} c} \approx 10^{-6} f \approx 1 \text{ GHz FWHM}. \tag{3.11}$$

This shift in frequency causes a significant variation in the detuning of atoms in a thermal cloud, and a corresponding change in the optical force which, as first described by Chu [48], can be used for cooling.

It is also possible to apply a magnetic field gradient which gives a position-dependent shift in the magnetic-sensitive m_f levels and a position-dependent trapping force. The combination of these two effects creates a cooling and trapping force in what is known as a magneto-optical trap.

3.3 The magneto-optical trap

3.3.1 Doppler cooling

The light force acting on atoms due to interactions with a single laser beam is unidirectional, and therefore incapable of cooling atoms with both positive and negative velocities; two identical yet counter-propagating beams, however, enable cooling to occur in both forward and backwards directions.

If the laser is on resonance with the transition frequency of a stationary atom, the photon scattering cross section will be symmetrical around zero velocity. In this case, the atom experiences equal scattering rates from both beams for all velocities, causing no net force on the atom. By detuning the frequency away from resonance the symmetry is broken and the atom sees a difference between the scattering rates for the forward and reverse beams.

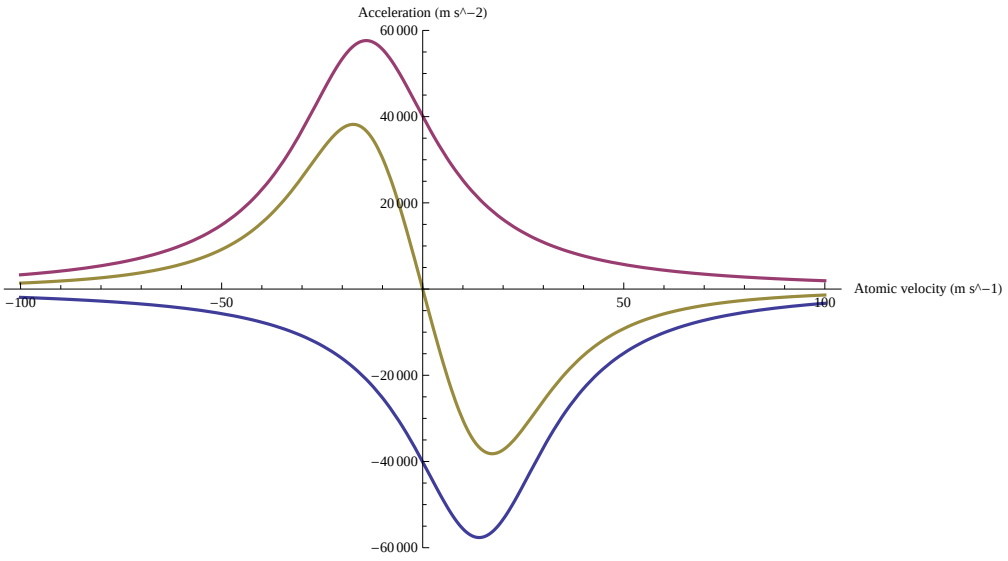

FIGURE 3.2: The acceleration of a rubidium atom as a function of the particle velocity, as a result of interactions with an off resonant beam from the $+v$ direction (blue line), from the $-v$ direction (red line), and the total acceleration on the atom as a result of both beams (yellow line).

For the case of red detuned counter propagating beams, the Doppler shift in the transition frequency causes the atom to scatter more light from the source towards which it is moving. This gives a force which opposes the atom motion in both directions. This scheme can be extended to the cooling of particles in all three directions by using three retro-reflected beams which are orientated along the three orthogonal Cartesian directions (see figure 3.3).

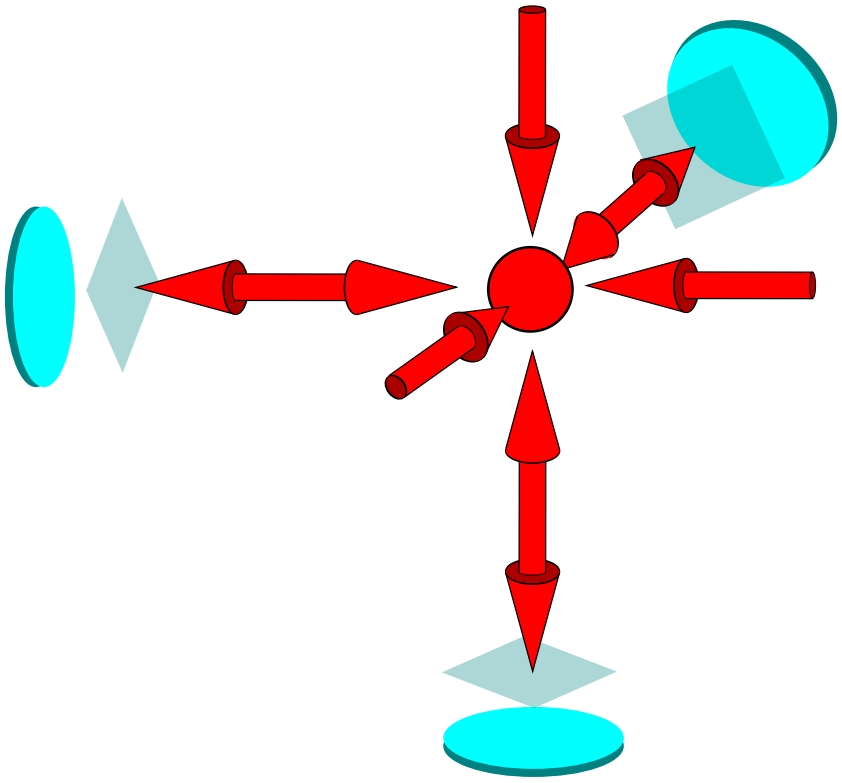

FIGURE 3.3: The standard arrangement for the beams in a MOT uses three beams along orthogonal Cartesian axes, which are retro-reflected after passing through a quarter waveplate.

As it is only necessary for the cooling force to have a component of the optical force along each axis, more complex beam designs can be constructed which allow, for example, better optical access to the MOT cloud [49]. However, as the cooling forces are directional and always orientated along the beam axis, in contrast to the effects of heating, which are isotropic (see section 3.4.3), the minimum temperature achieved in these arrangements is likely to be higher than that for the standard geometries.

3.3.2 The magneto-optical trapping force

In equation 2.7, we saw how the g_F Landè 'g-factor' gives the coupling strength of each of the atomic F states to a homogeneous applied magnetic field. If the ground and excited states of the cooling transition have a different value for g_F, there will be a difference in the magnetic shift of the ground and excited states, and therefore a magnetic sensitive energy shift in all but the $\Delta m_f = 0$ transitions.

Because of this, a magnetic field $B_x = |B|x$ is applied to the region where the beams overlap to cause a position-dependent shift in the transition frequency. As with the cooling forces, this changes the scattering rate of the atom according to the position of

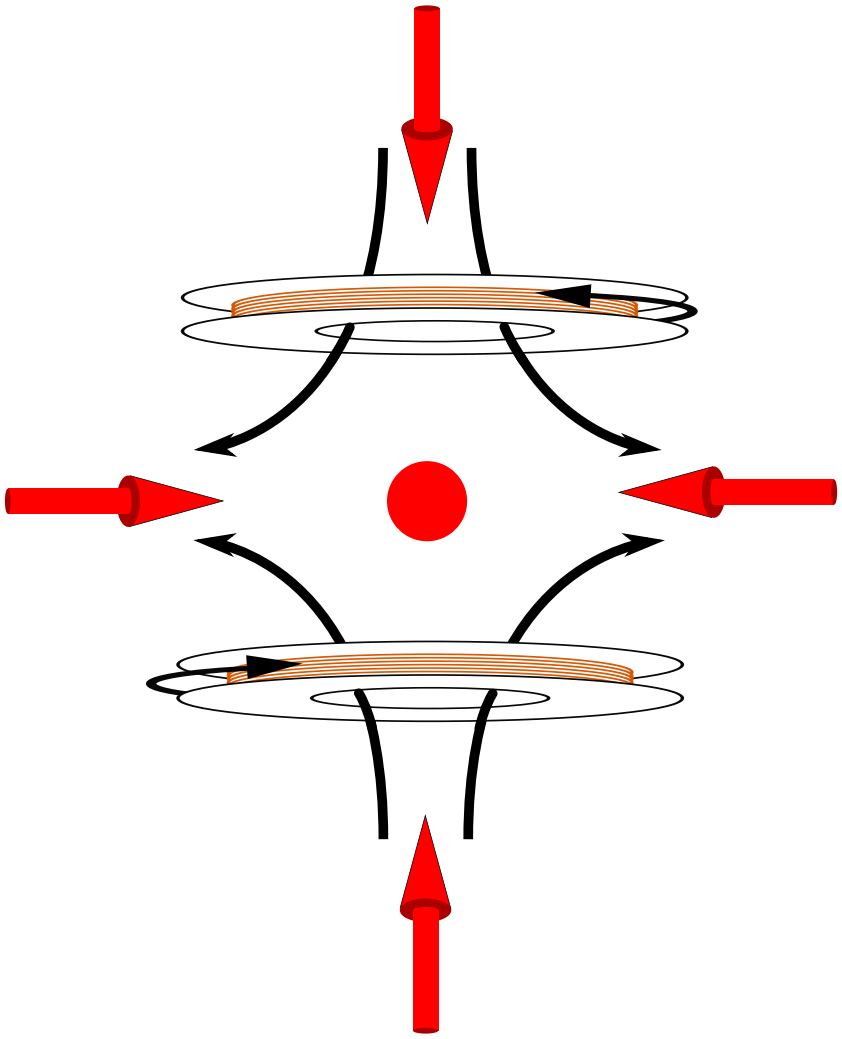

FIGURE 3.4: Two coils, driven in opposite directions, produce a magnetic field (black arrows) with a constant gradient across the trap which, in combination with detuned optical fields, causes the trapping of atoms.

the atom from the centre of the trap, and causes an optical, position dependent trapping force.

In accordance with the atomic selection rules, the direction of the force depends on the polarisations of the beams, which are chosen to have opposite helicity on either side of the trap, bringing about the correct restoring force (figure 3.6). This requires the change in the angular momentum of the atom to be opposite to that of the photon, which is $+\hbar$ for a σ^+ circularly polarised photon, or $-\hbar$ for a σ^- handed circularly polarised photon. This causes the $\mid F = n, m_f = n\rangle \rightarrow \mid F = n+1, m_f = n+1\rangle$ transition, which is resonant in regions of negative magnetic field to be driven by σ^+ polarised light, while the $\mid F = n, m_f = n\rangle \rightarrow \mid F = n-1, m_f = n-1\rangle$, resonant in positive magnetic fields to be driven by σ^- polarised light.

Thus, if we make the beam coming from the negatively circularly polarised, and the

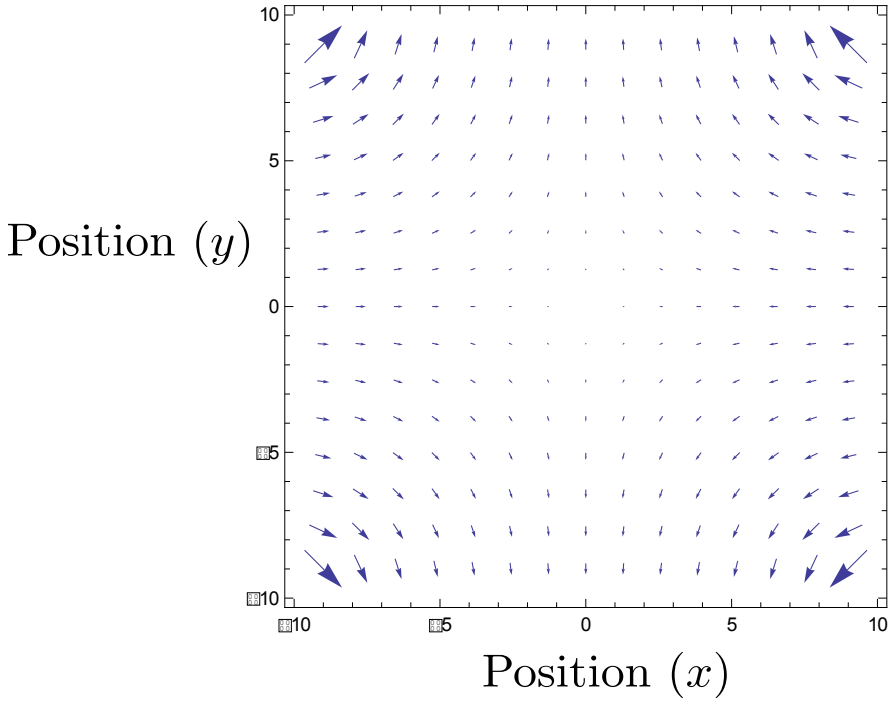

FIGURE 3.5: A 2 dimensional vector plot of the magnetic field produced by an anti-Helmholtz pair of magnetic coils. Note that the field varies linearly around the centre of the pair, where the field is zero, with a stronger gradient along the vertical axis than the horizontal.

Effect	Value	Result
Detuning	18MHz	
Doppler	$\frac{v \times f}{c}$	cooling
Zeeman shift	$\frac{m_f B(x) \mu_0}{h}$	trapping

TABLE 3.1: A summary of the atomic detuning necessary for the Doppler cooling of rubidium 85, and the effect each of these has on the atomic ensemble.

beam from the positively circularly polarised, we make sure that the force always acts towards the centre of the trap.

It is important to note that the force here results from the absorption of photons, and is only indirectly related to the Zeeman shift in energy levels. This is different from the trapping forces which arise as a result of the magnetic dipole interaction of an atom in an externally applied magnetic field [50]. If the laser beams are turned off, the atoms feel some force from this underlying magnetic trap: this is insufficient, however, to support the cloud against gravity.

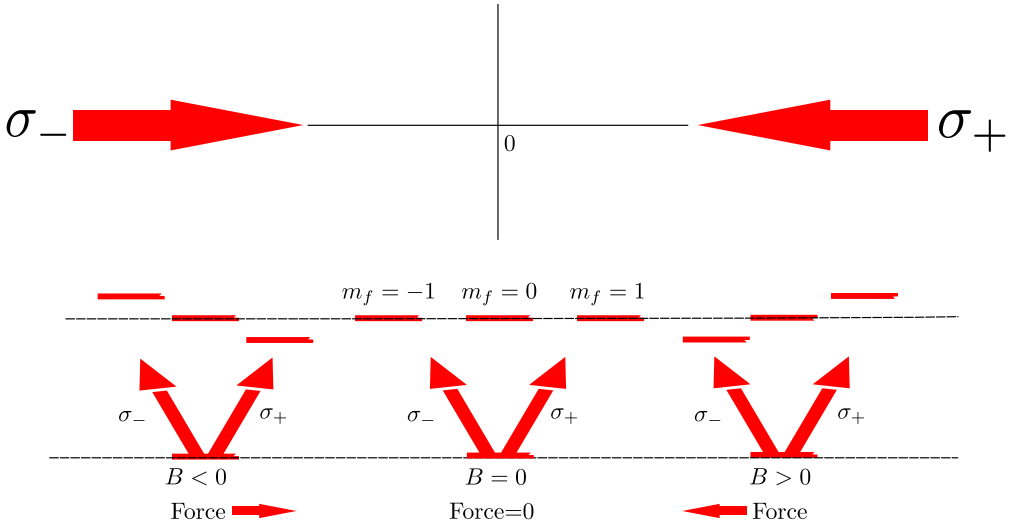

FIGURE 3.6: From applying beams of opposite helicity to the centre of the trap, we cause an opposite optical force on either sides of zero, directed inwards towards the centre, which traps the atoms.

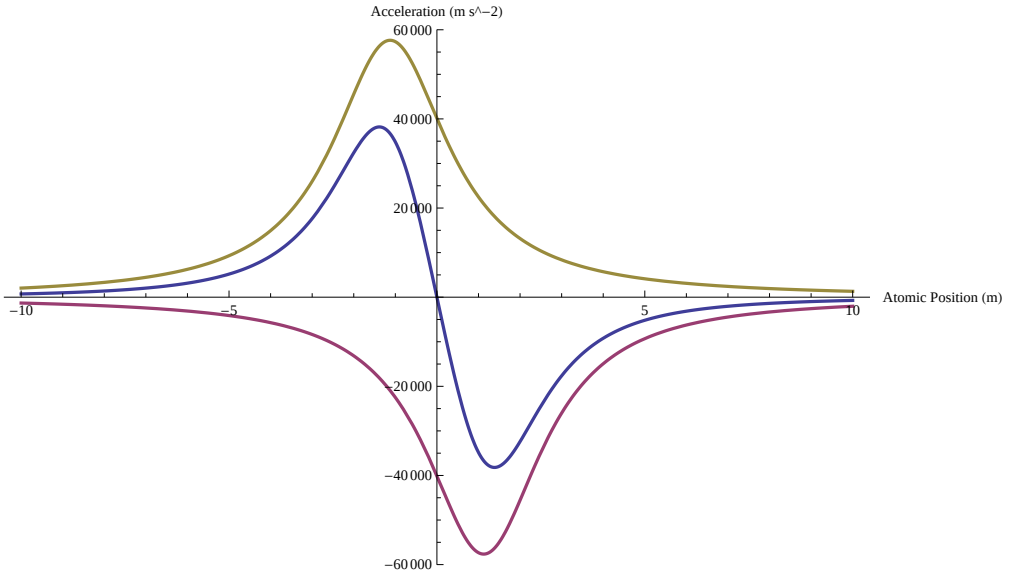

FIGURE 3.7: The acceleration of a rubidium atom due to light interactions detuned from the D2 line while within in a magnetic field gradient of 10 Gauss cm^{-1}.

3.4 Spontaneous emission and the limitations to magneto-optical trapping.

In Doppler cooling, the atom must undergo many absorption-spontaneous emission cycles to reduce its kinetic energy, and dissipate it to the environment. Although it is commonly stated that spontaneous emission is necessary to bring about the positive change in entropy required by the second law of thermodynamics [2], it has been shown

[2]The total entropy of any isolated thermodynamic system always increases over time, approaching a maximum value[51].

[52] that the decrease in entropy of a cooled atom can be made up by the change in the occupation of the light field. Thus, the spontaneous emission is necessary in Doppler cooling, not because of entropy, but irreversibility: without spontaneous emission, the atom would undergo absorption-stimulated emission cycles which cause no change to the atomic momentum [3].

We have considered so far the application of light forces to thermal atoms, with momenta many thousand times that of a photon. In this regime we can solve the equations of motion by considering a continuous optical force, rather then summing the individual photon impulses. In doing this, we neglect the force on the atom due to spontaneous emission on the basis of its random direction which, over many cycles, averages to zero. However, at low temperatures, where the atomic momentum is comparable with the photon momentum, the small effect of spontaneous photons on the atoms in the MOT can no longer be ignored and imposes limits on the minimum temperature, and the maximum density that can be achieved.

3.4.1 Loading rates, atom number and trapping lifetime

The MOT cloud is loaded from a background of low density thermal vapour. However, as the trapping potential in a magneto optical trap is much shallower than the thermal energies of atoms, most collisions between trapped atoms and the background gas supply sufficient energy to the trapped atom to allow it to escape from the trap.

Because of this we can write the loss rate as proportional to the rate of collisions between the trapped and thermal atoms

$$1/\tau = n\sigma v, \tag{3.12}$$

where τ is the decay time, σ the collisional cross section, v the velocity of the atoms and n the density. The loading rate is equivalent to the rate at which atoms within the capture velocity, with an energy less than that of the trap depth, enter the trap region. As the trap depth is a complex parameter, however, which depends on both the atomic position and velocity, the capture velocity cannot be found analytically. It is possible to find numerical solutions of the capture velocity, which consider the position and velocity for an atom at certain discretised points in time, in order to calculate the change in these properties with each time interval. These models, however, depend on the initial starting positions of the atoms, and are difficult to extend to models of the whole of the atomic ensemble.

[3] We go on to look at systems where the light used for absorption and stimulated emission comes from opposite directions, in which case absorption-stimulated emission cycles have been used to impart momentum to a system without spontaneous emission.

The two rates can be combined into a differential equation for the MOT loading rate [53]

$$\frac{dN}{dT} = R_l - \frac{n}{\tau}, \tag{3.13}$$

where N is the total atom number and R_l the loading rate.

It is important to note that the loading rate does not depend on the background pressure of the vapour, as both the loading and decay rates are proportional to the pressure. However, in the limit of low background density, equation 3.12 falls to zero. In this situation, the loss of atoms arises as a result of collisions between trapped, excited atoms in which the electronic energy of the atoms is transfered to the kinetic energy, enabling it to escape. Once in this regime, trapping lifetimes are related to the population in the excited state, and therefore the intensity of the trapping laser. Atoms in this regime have been trapped for up to 100s [54].

3.4.2 The high density limit

During the cooling process, spontaneously emitted photons from purely radiative decay routes are resonant with neighbouring particles, giving a high absorption cross section. This non-zero chance of re-absorption by neighboring atoms causes a repulsive impulse between the emission-absorption partners which in the limit of high atomic densities causes an outwards pressure which opposes the inwards MOT trapping forces, limiting the maximum density.

This allows us to characterise the MOT in two different regimes: high and low density. In the limit of low density, the atoms can be treated individually without interactions with other atoms, in which case they undergo damped harmonic oscillation around the trap centre. In the high density limit, however, the atoms interact heavily, and affect each other on time scales much less than the period of oscillation. Sesko originally gave a description of the magnitude of this re-radiation force in 1990 [55]; a more intuitive explanation can be found in a paper by Steane [56].

If the rate at which one atom scatters photons is [56]

$$R = \frac{\Gamma s}{2(s+1)}, \tag{3.14}$$

then the total force between two atoms can be derived from equation 3.4 to include the interaction cross section, σ_r:

$$F_r = \hbar k \frac{s}{s+1} \frac{\sigma_r}{4\pi r^2}. \tag{3.15}$$

For low densities, the mean distance, r, is large and the re-radiation force is insignificant in comparison with the thermal energies of the particles. In this regime, each atom undergoes damped harmonic oscillation around the trap centre, and the MOT size is related to the temperature [57, 58]:

$$k_B T = m v_{rms}^2 = k z_{rms}^2 \tag{3.16}$$

where z_{rms} and v_{rms} refer to the root mean square position and velocity and k is the spring constant of the trap, given by [57]

$$k = \frac{\mu A}{\hbar k} \beta(\gamma, \delta, \Omega), \tag{3.17}$$

with A being the magnetic field gradient, μ the magnetic moment of the transition, and $\beta(\gamma, \delta, \Omega) = \frac{F_o m}{v}$ describes the spring constant of the trap based on the optical radiation force.

The spring constant of the atoms in the trap is proportional to the m_f state of the atom, which describes the rate of change of the optical force with distance from the trap centre. During the operation of the MOT, the atoms are evenly populated amongst these states and the total density distribution of the cloud must be found from a summation of atoms in each m_f state. Due to the constant optical pumping of the atoms, the population of the m_f levels changes on time scales shorter than the oscillation period. This force acting on each atom is therefore a weighted sum of the forces experienced whilst in each of the m_f states. It is because of this that the spatial extent of the MOT cloud is continuous, without steps which would form as a result of the summation of clouds contained within different trapping potentials.

In the high density limit, the mean distance between particles, r, decreases until the outward repulsive force becomes comparable with the trapping forces. In this regime, the atomic distribution becomes independent of the radius, giving a top-hat distribution [56].

So far in our analysis of the MOT forces, we have assumed that the beams are of a constant intensity, and that the scattering of the beams by the atoms causes a negligible amount of depletion. For atoms in the high density limit, this assumption does not hold, and the intensity depletion of the beam by the scattering from atoms can no longer be ignored. This is especially true if a beam is retro-reflected to form the counter propagating beam, which, for large scattering rates, can no longer be considered to be identical to the incoming beam.

3.4.3 The low temperature limit

The random nature of the spontaneous emission process causes a heating effect which limits the minimum temperature that can be achieved within the magneto-optical trap. Although the directions of the impulses due to spontaneous emission are random, and so average to zero, the root mean square value of the momentum, $\langle P \rangle^2$, does not, and increases as [59]

$$\langle P \rangle^2 = 2D_p t \tag{3.18}$$

where

$$D_p = (\hbar k)^2 R(\gamma, \delta) \tag{3.19}$$

where $R(\gamma, \delta)$ is the scattering rate described by equation (3.7), and t is the time. This causes the atoms in the magneto-optical trap to follow a random walk in both position and momentum, which causes a spreading of the momentum distribution with every absorption-emission cycle. This spread in momentum can be interpreted as a heating term given by

$$\left(\frac{dE}{dt}\right)_{heat} = \frac{(\hbar k)^2 R(\gamma, \delta)}{M}, \tag{3.20}$$

where M is the mass of the particle. This heating term is insignificant at high temperatures, but as the temperature decreases, we reach a point at which it balances with the Doppler cooling rate. If we substitute the definition of $R(\gamma, \delta)$ given in 3.7, we find that the minimum temperature is obtained for a detuning of $\delta = -\gamma/2$, and gives a temperature, known as the Doppler limit,

$$T_{Doppler} = \frac{\hbar \Gamma}{2k_B}. \tag{3.21}$$

The link between the minimum temperature and the linewidth is unintuitive, and so we offer an alternative derivation, originally described by Castin [60]. Because the transition has a natural linewidth, cooling only occurs when the Doppler width is greater than the natural linewidth. This lower limit to the temperature can be written as:

$$E_k = \frac{1}{2} k_B T \approxeq \frac{1}{4} \hbar \Gamma, \tag{3.22}$$

$$T_{Doppler} \approxeq \frac{\hbar \Gamma}{2k_B}. \tag{3.23}$$

3.5 Sub-Doppler cooling

The theories of sub-Doppler cooling were constructed by Dalibard and Cohen-Tonnoudji [61] in order to explain measurements of the temperature of a magneto-optical trap

taken by Bill Phillips and Harold Metcalf [62] which were found to be significantly below the Doppler cooling limit. The mechanisms are found to operate in the regime where the atoms are moving slowly enough so that they experience significant amounts of repumping between light shifted Zeeman levels.

3.5.1 Sisyphus cooling

For well-aligned counter-propagating beams a standing wave is created by each beam and its retro-reflected partner through which the atoms move. For orthogonally polarised, linear-linear beams this also creates a polarisation gradient in the overall electric field, changing from linear to circular right-handed, to linear but orthogonal to the first linear, to circular left-handed, as shown in figure 3.8.

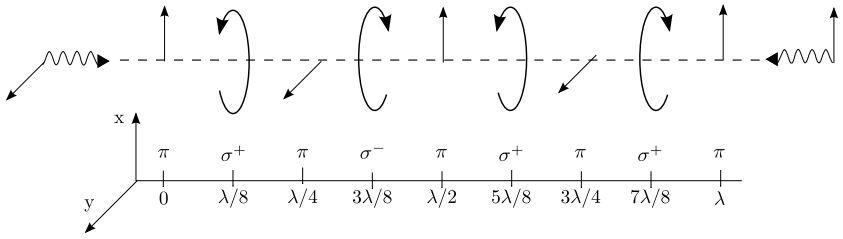

FIGURE 3.8: The counter-propagating lin-lin polarisation beams produce a polarisation gradient that changes from linear to circularly polarised to orthogonally polarised linear to orthogonally circularly polarised and then back to linear in the space of a wavelength. As each of these polarisations has a different coupling strength to the positive and negative m_f levels, this produces the egg-box lattice potential necessary for Sisyphus cooling.

As each of these polarisations is coupled differently to each of the atomic m_f levels, this causes a periodic and spatially-dependent light shift to the energy levels.

The atoms have a kinetic energy which is greater than the depth of the potential and so will continue to move over it; as they do so, the red detuned light preferentially pumps atoms which are least detuned from resonance(figure 3.9). At a later time, spontaneous decay will return the atom to the ground state; in the time between absorption and spontaneous emission, however, the atom will have moved from its original absorption position. The difference in position along with the position dependent light shift causes a difference between the energy of the absorption and emission photons. As the atom is most likely to absorb a photon while at a potential minimum, the decay photon is typically of a higher energy than the absorption photon, a difference which is made up for by the kinetic energy of the atom. Over multiple absorption and emission cycles, the atom loses energy to the light field and is analogous to Sisyphus, the legendary Greek trickster who angered the gods, are cursed to forever move 'uphill'.

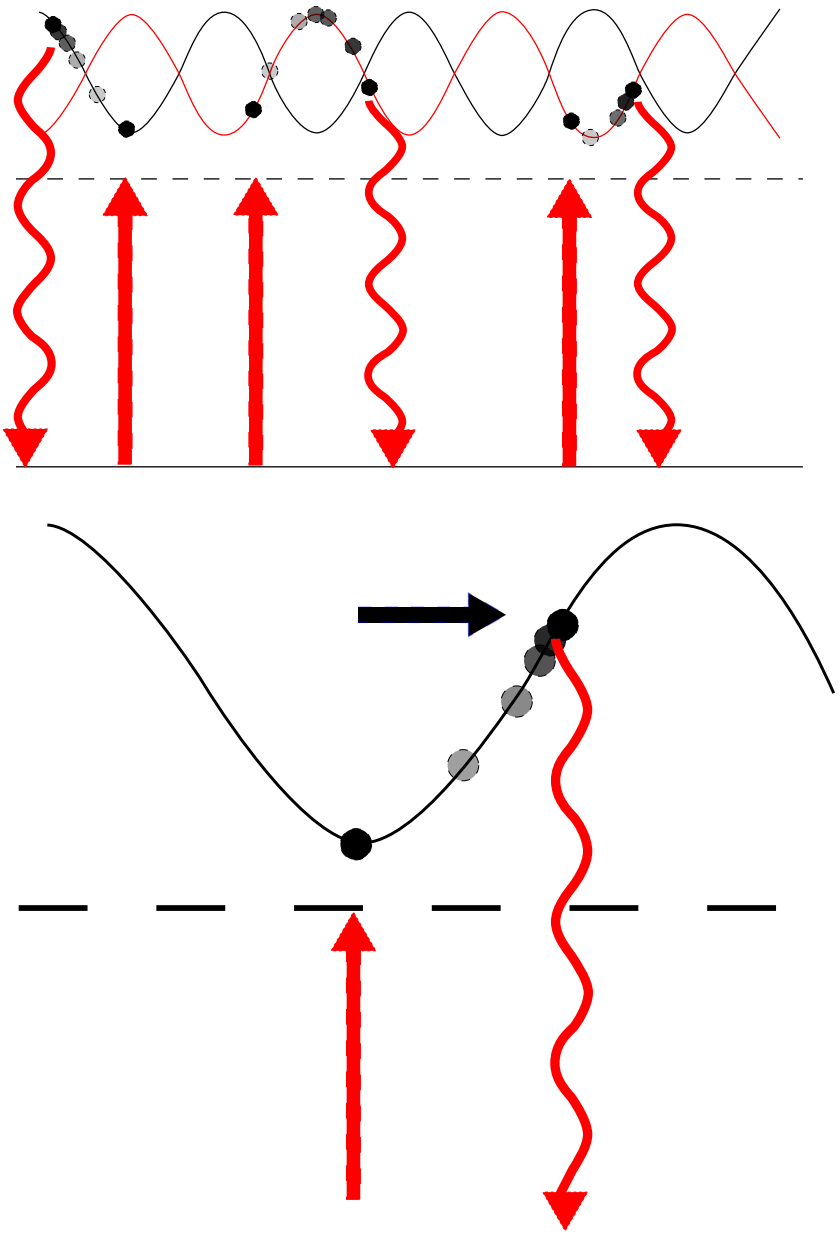

FIGURE 3.9: A diagram showing Sisyphus cooling. The top diagram shows that the atom is excited where the atom is closest to the detuned laser frequency. Because of its residual kinetic energy, the excited atom moves away from the minimum detuning to a location that gives a greater detuning. In this time the atom spontaneously decays, emitting a photon of a higher energy than the exciting photon.

3.5.2 Corkscrew cooling

The case for a standing wave made of circular-circular polarised light is quite different. This time the light maintains a linear polarisation of constant magnitude, where the direction of polarisation corkscrews in space. The origin of the cooling force in this

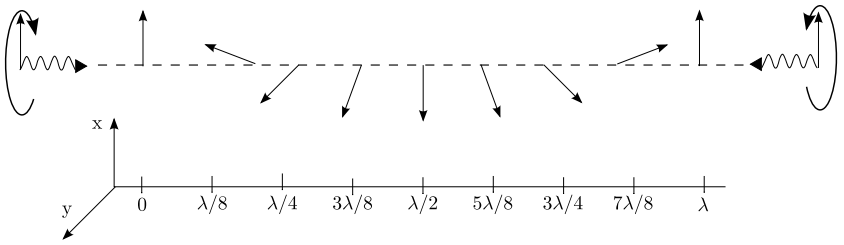

FIGURE 3.10: The counter-propagating, orthogonally circularly polarised light beams create an electric field where the type of polarisation is constant, but where the direction of polarisation changes with space. Unlike Sisyphus cooling, the change in the orientation of the polarisation causes no change in the light shift of the atom.

environment is not intuitive, but was treated at length in the paper by Dalibard and Cohen-Tannoudji [61].

In this regime, the energy eigenvalues of the system can be found by changing to a rotating frame where the direction of polarisation is constant. On changing to this rotating frame, the electric field gains an additional, inertial, term, similar to a static, magnetic field. The inertial term causes a perturbation in the wavefunction of the atom. In their paper, Dalibard and Cohen-Tannoudji found that unlike populations of the stationary atom, which are roughly even, the population of the perturbed m_f states of the moving atoms is uneven and velocity sensitive.

Despite the change in the population distribution, the transition strengths between the energy levels remain constant. The result is a difference in the pumping rates of the two, counter propagating and oppositely circularly polarised beams and an optical force which is far more sensitive to velocity than Doppler cooling.

3.5.3 Sub-Doppler cooling in three dimensions

When we consider the standing wave light field for three counter-propagating beams, the situation is much more complicated than in either of the two situations above. In this three dimensional system, the light produces both a polarisation and an intensity gradient that gives a combination of both corkscrew and Sisyphus cooling. A discussion of this three dimensional system is given by Steane *et al.* [63].

3.5.4 The sub-Doppler cooling temperature limit

As both polarisation gradient cooling and corkscrew cooling provide additional cooling mechanisms which overcome the Doppler limit on the temperature, it is important that we define the minimum temperature which limits these new laser cooling mechanisms. Such a limit can be found by considering the momentum of the spontaneously emitted photon which is always emitted as the last step of any cooling scheme. For atoms below the Doppler limit, this impulse is comparable to the atomic momentum and so cannot be neglected. We can use the momentum of a photon to define a minimum temperature velocity the atom can possess after emission, which can be used to formulate a new temperature limit:

$$T_r = \frac{\hbar^2 k^2}{3mk_B}. \tag{3.24}$$

There are ways to beat this 'recoil' limit, typically used to reach the required phase space density for Bose-Einstein condensation, which do not involve any radiative decay in the last cooling cycle: Raman cooling [64], where a narrow band excitation tuned to zero velocity pumps the atoms to a dark state so that the final cooling stage is not spontaneous emission, and evaporative cooling [65], where a shallow trapping potential is used to eliminate the most energetic particles from an equilibrated cloud.

3.6 The atomic structure necessary for optical cooling

3.6.1 Repumping

Because of the necessity to drive the cooling transition with a narrow linewidth laser, only one transition is coupled to the cooling light. In many systems where the cycle is not perfectly closed however, weak coupling to states outside of the cooling loop causes a loss of population during cycling. In rubidium for example, the detuning of the cooling laser is sufficient to cause a non-zero overlap with the transition to the F=3 state. This, unlike the F=4 excited state which is used for cooling, is free to decay to the F=2 lower hyperfine ground state, which is not coupled to the cooling transition. As the atomic population within the cooling cycle becomes depleted to states not coupled to the cooling laser, the cooling and trapping forces stop. For rubidium 85, the overlap of the driving laser with states outside the cooling cycle is small, and a large fraction of the population stays within the cooling cycle for many thousands of cycles.

It is possible to avoid such a problem by using an additional 'repump' laser at the frequency of the dark state transition, which couples the dark state population to the

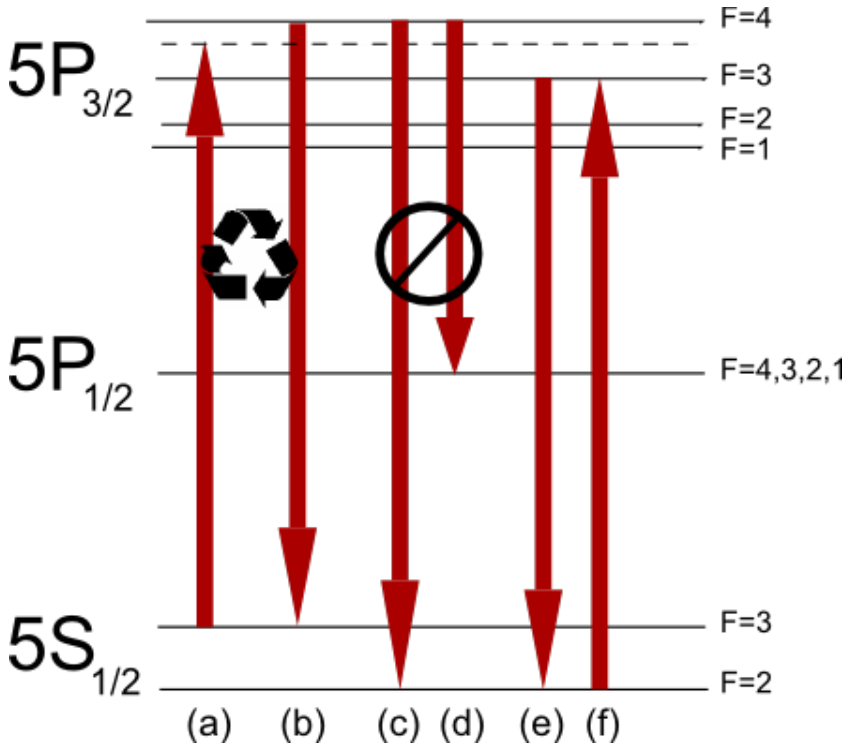

FIGURE 3.11: The repumping scheme for rubidium-85. (a) shows the cooling laser, which is detuned from resonance, which excites the population to the F=4 excited state which then decays (b) to the F=3 upper hyperfine state to form a closed optical loop which cannot decay to either the lower hyperfine state (c) or to the $5P_{1/2}$ state (d). If the atom is excited to the F=3 state it can decay to the lower hyperfine state (e), which is not within the closed cooling loop. The repump laser is used to recouple these atoms back into the cooling cycle (f).

cooling transition, recycling the dark population (figure 3.11). So long as the difference between the two transition frequencies is greater than the thermal Doppler width of the atoms, the Doppler temperature is not affected. In systems such as rubidium, where the decay to states outside the cooling cycle is slow, the power in the repump laser can be significantly smaller than that in the trapping laser, with the repump further from resonance if necessary.

It is conceivable for a large number of repump lasers to be used in a system with a large number of allowed decay paths, such as a molecule. This solution quickly becomes difficult to implement as the number of repump lasers, and hence the cost and experimental complexity, increases. As well as this, absorption from each repump beam causes heating of the atom, so that in the limit of large numbers of repump beams, the time the atom spends in the cooling transition is small, and the minimum achievable temperature is high. Thirdly, repumping does nothing to solve the problems associated with the atom being excited to weakly coupled states, which decay slowly to give small optical forces.

Clearly, although repumping is very useful for many atomic systems, with small leaks in the closed optical cycle, it is not the solution to cooling of atoms without closed

transition loops.

3.6.2 Closed transitions

The huge difference between the atomic and photon momentum makes it necessary for each atom to absorb large numbers of photons to reduce the external energy to the Doppler limit and undergo equal numbers of spontaneous emission events, each of which must return the atom to the ground state after each excitation. With such a large number of cycles, it is necessary that each spontaneous emission event has a high probability of returning the atom to the ground state from which it was excited, so that, on the time scales necessary for cooling, the population stays within the cooling transition:

$$\sum_{routes \nrightarrow ground} \gamma_{routes} < \frac{P_{atom}}{P_{photon}} \gamma_{routes \rightarrow ground}, \tag{3.25}$$

where P_{atom} and P_{photon} are the atomic and photon momentum, γ is the decay time for routes back to the ground state and others which we sum over on the left hand side. It is common for transitions which meet this condition to be described as 'closed'. For ^{85}Rb, the transition is closed for time scales up to:

$$\frac{P_{atom}}{P_{Photon}} \times \gamma = 10^5 \times 26\text{ns} = 2.7\text{ ms}. \tag{3.26}$$

In addition to this, the transition must be strong enough, with a large enough oscillator strength, that spontaneous decay occurs quickly enough for the atoms to be cooled before they leave the apparatus (either by escaping the trapping region for a magneto optical trap, or exiting the end of a linear decelerator in the case of translational cooling).

There are also practical matters to consider when choosing the right species for optical cooling experiments: it is important that the cooling transition is accessible to laser light. Although many such frequencies can by accessed by non-linear processes such as using optical parametric oscillation [66] or harmonic generation [67], these methods are expensive and are limited in their total output power.

As a result of these criteria, the number of species which have been successfully Doppler cooled is severely limited.

3.7 Doppler cooling of Molecules

The majority of molecules are of a greater mass than atoms, which typically gives them a larger inertia. Molecular samples must undergo a comparatively greater number of

absorption-emission cycles than atoms to reduce the temperature by the same amount. This makes the conditions given in equation 3.6, necessary for Doppler cooling, more difficult to achieve.

These problems are exacerbated by the the non-spherical shape of a molecule, which gives additional rotational and vibrational splitting of the electronic energy levels. Although radiative decay between neighboring rotational and vibrational states is slow [68], it is common for electronic decay to be accompanied by changes in the vibrational and rotational state, vastly increasing the number of states to which an excited molecule can decay. Such couping between vibrational modes is described by the Frank–Condon principle [39]. In order for a molecule to be laser cooled, it must possess, to a good approximation, a diagonal Franck–Condon array, such that radiative decay tends to conserve the vibrational state of the molecule.

Despite the difficulty in finding suitable spectra for molecular cooling, a paper by Di Rosa [69] gives a table of ten, mostly hydride, molecules which have an optical cycle that is sufficiently closed to preserve 1-10% of the cooling population after 10^4 optical cycles. Due to the high increase in phase space density achieved by laser cooling in comparison with buffer gas or sympathetic cooling, it has been suggested by Di Rosa that given high loading rates such a technique could be used to form a magneto-optical trap of CaH containing 10^5 molecules trapped over a peroid of $\approx$1 ms, which could then be evaporativly cooled. Di Rosa is not alone in this study of suitable molecular structures, and there have been suggested alternatives in the form of metal oxides, sulphides and carbides [70]. Such work presents interesting and useful alternative methods of cooling.

In comparison to the cooling methods described in this thesis, however, these methods of molecular cooling require significant knowledge of the absorption spectra of the molecular sample, which are not well documented in existing literature. Our work presents methods of cooling which are not as reliant on specific energy level structures, and hence of increased generality.

Chapter 4

Coherent light atom interactions and novel cooling schemes

4.1 Stimulated rather than spontaneous emission

During Doppler cooling, the atom must undergo spontaneous emission to return to the ground state. Although the probability of each decay path is calculable, the path for a given decay is random and unpredictable which makes it necessary for an atom to have a specific energy level structure, which guarantees that the atom will decay back to the state from which it was excited, to form a closed optical loop. Without this closed loop, the atomic population is pumped into states not coupled to the cooling laser, which stops the cooling process.

In addition to the methods of generating samples of cold atoms described in section 1.2, there have been methods suggested for the purely optical cooling of molecules [71–74]. While cooling of atoms has been observed in optical cavities [75], these cooling methods, are still limited in scope and as such have yet to show any advantages over alternative, non-optical methods of cooling. In this chapter we focus on ideas which use coherent particle-light interactions for cooling. 'Coherent amplification of cooling' [31], and 'interferometric cooling' [32] produce cooling forces through tailoring the order and timing of a sequence of coherently interacting laser pulses. For these to succeed, phase stability of both the particle and the laser is crucial.

Let us initially consider a simple model of an atom as a two level system, with a single ground and a single excited state. For most atoms, so long as the detuning and linewidth of the driving laser are smaller than the spacing between the energy levels, only a single transition will be close to resonance at any one time, and an atom with many excited

states can be simplified to a quasi-two level system. We shall look at this assumption later in the chapter.

We start the discussion with the time dependent Schrödinger equation:

$$i\hbar\frac{d\Psi_n}{dt} = (H_0 + H_I)\Psi_n. \tag{4.1}$$

where H_0 is the unperturbed Hamiltonian which acts on eigenvectors of the wavefunction to give the energy:

$$H_0\psi_n = E_n\psi_n, \tag{4.2}$$

and H_I is the interaction Hamiltonian, given by:

$$H_I = er \cdot E = er \cdot E_0\cos(wt). \tag{4.3}$$

This term describes a semi-classical treatment of the system, where the interaction term is formed from a classical electric field–electric dipole interaction; here er represents the atomic dipole, given by the distance of an electric charge, e, from the nucleus, r, and $E_0\cos(\omega t)$ represents the electric field of the incident light of angular frequency ω. The full quantum calculation of the interaction Hamiltonian involves the electron motion and the electromagnetic vector potential, but it can be shown that this becomes approximately equal to equation 4.3 near to resonant frequencies [76]. We represent a general atomic wavefunction as a superposition of the wavefunctions for each state:

$$\Psi = \sum_n c_n(t)\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right), \tag{4.4}$$

which, after substituting 4.4 into 4.1 gives:

$$\sum_n i\hbar\frac{d}{dt}\left[c_n(t)\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right)\right] = \sum_n(H_0 + H_I)\left[c_n(t)\psi_n.\exp\left(\frac{-iE_nt}{\hbar}\right)\right] \tag{4.5}$$

The time dependence of both $c_n(t)$ and the potential means that the chain rule must be applied to the left-hand side of equation 4.4, while on the right, we substitute the definitions for H_0 and H_I:

$$\begin{aligned}\textstyle\sum_n i\hbar\frac{dc_n}{dt}\left[\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right)\right] \quad &+c_n(t)E_n\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right) = \\ &\sum_n\left[E_n + er\cdot E_0\cos(wt)\right]c_n(t)\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right).\end{aligned} \tag{4.6}$$

The second term on the left and the first term on the right are equal, and cancel, leaving:

$$\sum_n i\hbar\frac{dc_n}{dt}\left[\psi_n \exp\left(\frac{-iE_nt}{\hbar}\right)\right] = \sum_n er\cdot E_0\cos(wt)\left[c_n(t)\psi_n\exp\left(\frac{-iE_nt}{\hbar}\right)\right]. \tag{4.7}$$

To find the overlap between state n, and another arbitrary state, m, we multiply both sides by ψ_m^* and integrate over all space. Remembering that the basis wavefunctions are orthogonal we have:

$$\int_0^\infty \psi_m^*\psi_n dr = \delta_{m,n}, \tag{4.8}$$

where $\delta_{m,n}$ is the Kronecker delta function. Thus for each m

$$i\hbar\exp\left(\frac{-iE_mt}{\hbar}\right)\frac{dc_m}{dt} = E_0\sum_n d_{nm}c_n(t)\cos(\omega t)\exp\left(\frac{-iE_nt}{\hbar}\right), \tag{4.9}$$

where

$$d_{nm} = e\int_0^\infty \psi_m^* e\cdot r\psi_n dr, \tag{4.10}$$

which is often not zero, even when $m \neq n$. Thus,

$$i\hbar\frac{dc_m}{dt} = \sum_n d_{mn}E_0\cos(\omega t)\exp\left(\frac{iE_{mn}t}{\hbar}\right)c_n(t), \tag{4.11}$$

where $E_{mn} = \hbar\omega_{mn} = E_m - E_n$. Using the identity

$$2\cos(\omega t) = \exp(i\omega t) + \exp(-i\omega t), \tag{4.12}$$

we can write this as

$$i\hbar\frac{dc_m}{dt} = \frac{E_0}{2}\sum_n d_{mn}c_n(t)\left[\exp(i(\omega_{mn}+\omega))t + \exp(i(\omega_m n - \omega))t\right]. \tag{4.13}$$

At this stage, it is simple to set all bar one of the initial states, $c_n(0)$, to zero in order to determine the time evolution of a system initially in the ground state. In 'Quantum Optics' by Orszag (ref. [77]), this model is used to judge the effectiveness of the two-level system approximation; it is found that the transition probability is sharply peaked at the driving frequency, so that states with transition frequencies very different to the driving field are not excited, supporting our two level assumption.

Let us now assume that the atom has two levels, a and b. Our wavefunction now has the form:

$$\Psi = c_b(t)\psi_b\exp(-i\omega_b t) + c_a(t)\psi_a\exp(-i\omega_a t), \tag{4.14}$$

which we can substitute into equation 4.13 to get:

$$i\hbar\frac{dc_b}{dt} = \frac{E_0 d_{ba}}{2} c_a(t)\left[\exp(-i(\omega_{ab}-\omega)t) + \exp(-i(\omega_{ab}+\omega)t)\right], \tag{4.15}$$

and

$$i\hbar\frac{dc_a}{dt} = \frac{E_0 d_{ba}}{2} c_b(t)\left[\exp(i(\omega_{ab}-\omega)t) + \exp(i(\omega_{ab}+\omega)t)\right]. \tag{4.16}$$

We can find a solution to this analytically by eliminating the high frequency exponential term using what is popularly known as the 'rotating wave approximation'. This assumes that the co-rotating terms are oscillating fast enough, on time scales faster than that of the experiment, that they average to zero. This assumption holds so long as the Rabi frequency is less than the Bohr and field frequencies, which is a suitable approximation for the near infrared, rubidium D2 transition; in various experiments using excitation by microwaves, however, this assumption is not valid. These examples are treated in detail by Shore [78].

We apply a global phase transformation to a rotating frame, in order to remove the time dependence of the Hamiltonian:

$$c'_{a,b} = \exp\left(\pm i\frac{\delta t}{2}\right) c_{a,b}, \tag{4.17}$$

where $\delta = \omega_{ab} - \omega$. This, because of the global phase invariance of the wavefunction, gives this transformation no physical consequences. When applied to equation 4.16 along with the rotating wave approximation, this gives:

$$\frac{d}{dt}\begin{pmatrix} c'_b(t) \\ c'_a(t) \end{pmatrix} = \frac{-i}{2}\begin{pmatrix} -\delta & \frac{E_0 d_{ba}}{\hbar} \\ \frac{E_0 d_{ab}}{\hbar} & \delta \end{pmatrix}\begin{pmatrix} c'_b(t) \\ c'_a(t) \end{pmatrix}. \tag{4.18}$$

There are several interesting ways to proceed with this result. Equation 4.1 tells us that the rate of change of each state wavefunction is related to the Hamiltonian of the system; similarly equation 4.18 can be treated as the same relation, with time evolution absorbed into the coefficients, $c_a(t)$ and $c_b(t)$.

This Hamiltonian has off-diagonal terms not present in the unperturbed atomic Hamiltonian; if the new matrix is diagonalised, we find the new eigenvalues to be:

$$\tilde{\Omega} = \pm\sqrt{\delta^2 + |\Omega|^2}, \tag{4.19}$$

where Ω represents the Rabi frequency,

$$\Omega = \frac{E_0 d_{ba}}{\hbar}, \tag{4.20}$$

which describes the coupling strength between the atom and a resonant laser, and $\tilde{\Omega}$ is the generalised Rabi frequency, which tells us the Rabi flopping frequency in a system with non-zero detuning. The eigenvalues of the Hamiltonian directly relate to the energies of the states, so that the difference between the two eigenvalues tells us the shift in transition energies from the unperturbed atomic states. In the limit of near-resonant interaction, so that $\delta \ll \Omega$, this is equal to

$$\Delta E = \frac{-\hbar\Omega^2}{4\delta}. \tag{4.21}$$

This energy is a manifestation of the electric dipole force, and is also known as the light shift, or the AC Stark shift, which perturb atomic transition frequencies which will be important in section 7.11.

It is also possible to solve equation 4.18 to obtain the solution [77]:

$$\begin{pmatrix} c_b'(t) \\ c_a'(t) \end{pmatrix} = \frac{-i}{2} \begin{pmatrix} \cos\frac{\tilde{\Omega}t}{2} + \frac{i\delta}{\tilde{\Omega}}\sin\frac{\tilde{\Omega}t}{2} & -i\frac{E_0 d_{ba}}{\tilde{\Omega}\hbar}\sin\frac{\tilde{\Omega}t}{2} \\ -i\frac{E_0 d_{ba}}{\tilde{\Omega}\hbar}\sin\frac{\tilde{\Omega}t}{2} & \cos\frac{\tilde{\Omega}t}{2} - \frac{i\delta}{\tilde{\Omega}}\sin\frac{\tilde{\Omega}t}{2} \end{pmatrix} \begin{pmatrix} c_b'(0) \\ c_a'(0) \end{pmatrix}, \tag{4.22}$$

which we can use to find the populations of the states at any given time, given known starting conditions. If we set the initial ground state population $c_b'(0)$ to unity, and the initial excited state population $c_a'(0)$ to zero, we then find that after applying a pulse of radiation for a time, t, the population in the excited state, $|c_a'(t)|^2$ becomes:

$$|c_a'(t)|^2 = \left|\frac{\Omega}{\tilde{\Omega}}\right|^2 \sin^2\left(\frac{\tilde{\Omega}t}{2}\right). \tag{4.23}$$

We see that this oscillates with time between zero and $|\Omega/\tilde{\Omega}|^2$ at a frequency $\tilde{\Omega}$, which for the case of a resonant atom, becomes an oscillation between zero and unity, at the Rabi frequency, Ω.

So long as the pulse area, $\tilde{\Omega}t \ll 1$, equation 4.23 can be rearranged [44] to give a sinc function which is dependent on the detuning of the laser from atomic resonance. This plot of the detuning dependent upper state population (shown in figure 4.1) for coherent interactions is significantly different from the case of an incoherently driven atom, for which we obtain a Lorentzian line-shape.

As it is experimentally feasible to plot the detuning dependence of the upper state population and the upper state populations as a function of pulse area (shown in figure 4.2), this result has huge experimental implications.

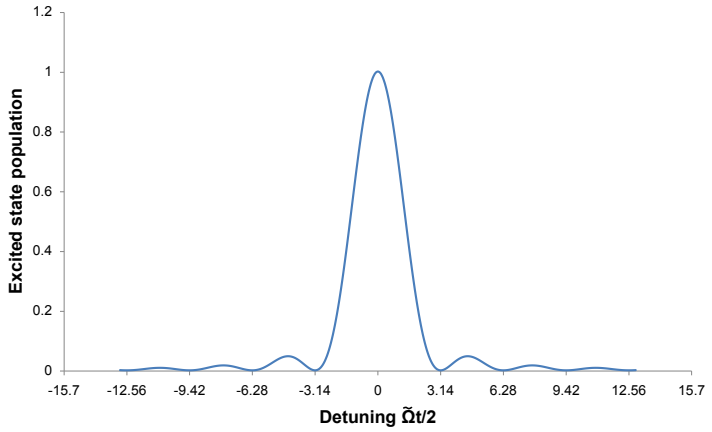

FIGURE 4.1: Excited state population as a function of the detuning of the driving laser from the atomic resonance after a pulse lasting much less than Rabi period.

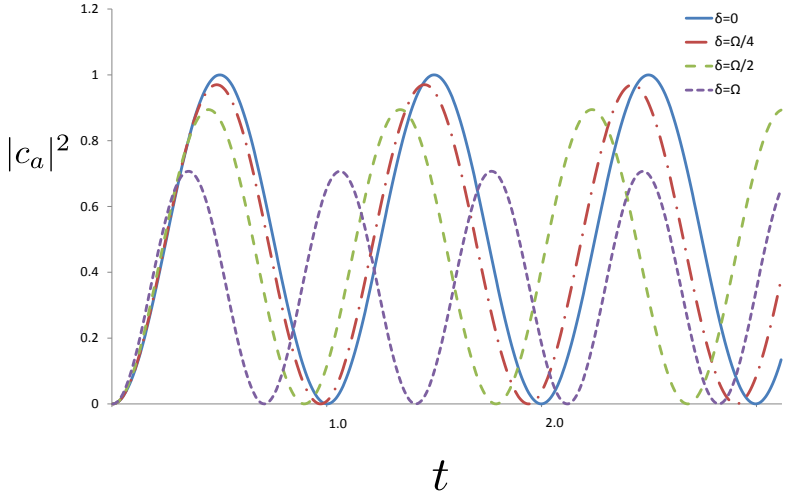

FIGURE 4.2: Coherent population oscillations by a monochromatic light source plotted as against time for the detunings $\delta = 0, \Omega/4, \Omega/2,$ and Ω. Note the difference in frequency and amplitude; this increase in frequency with detuning is useful for measuing the decay time and form in systems where decay occurs much faster than the Rabi period.

The detuning dependence of the upper state population given in equation 4.23 is significant. Off resonant oscillation will occur with a different frequency and amplitude to the resonant case. This is important when attempting to drive Rabi oscillations in ensembles of atoms which are Doppler shifted by varying amounts away from resonance. In these systems, the fidelity of the oscillation will be heavily affected by the distribution of frequencies and amplitudes, as we shall see later.

We see that if we set the detuning, δ, to zero and choose the pulse timing so that it lasts for a specific period of time,

$$t_\pi = \frac{\pi}{\tilde{\Omega}}, \tag{4.24}$$

we cause a population inversion which empties the ground state into the excited state. A pulse of light which meets such a specific condition is known as a 'π pulse'.

We can set the detuning, δ, to zero and substitute t_π into equation 4.22 to obtain a matrix operator which describes the application of a π pulse to an arbitrary initial state:

$$\hat{O}_\pi = \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}, \tag{4.25}$$

which acts on the wavefunction thus:

$$\Psi(t) = \begin{pmatrix} c_1(t) & c_2(t) \end{pmatrix} \begin{pmatrix} |1\rangle \\ |2\rangle \end{pmatrix} = \hat{O}\Psi(0) = \hat{O}. \begin{pmatrix} c_1(0) & c_2(0) \end{pmatrix} \begin{pmatrix} |1\rangle \\ |2\rangle \end{pmatrix}. \tag{4.26}$$

Using this formalism can show that applying a π-pulse to a population initially in the excited state moves the entire population to the ground state, whereas when applied to an atom in an equal superposition it returns in an equal superposition, but with a π phase shift. As with any atom-photon interaction, momentum is imparted to the atom on both excitation and decay which makes it possible to control the motion in order to construct cooling schemes using these pulses.

In order to explain these schemes, we must introduce two additional operators: the $\frac{\pi}{2}$ pulse and free evolution of the atom. A $\frac{\pi}{2}$ pulse is obtained by truncating a π pulse, so that the pulse now satisfies the criterion:

$$t_{\pi/2} = \frac{\pi}{2\tilde{\Omega}}. \tag{4.27}$$

Starting with a ground state atom, the effect of applying such a pulse is to coherently and deterministically move one half of the population to the excited state. This forms an equal superposition of the ground and excited states.

Again, by substituting the pulse time given in equation 4.27 into equation 4.22 we obtain a 2×2 matrix to describe the transformation made by a resonant field on an arbitrary superposition:

$$\hat{O}_{\frac{\pi}{2}} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}. \tag{4.28}$$

In the time between two light pulses, both the phase of the atom evolves with time, as described by the time dependent Shrödinger equation, and the phase of the laser changes with respect to that of the atomic state.

$$\hat{O}_{FE} = \begin{pmatrix} e^{-i\omega_0 t/2} & 0 \\ 0 & e^{i\omega_0 t/2} \end{pmatrix} \begin{pmatrix} e^{i\omega t/2} & 0 \\ 0 & e^{-i\omega t/2} \end{pmatrix} = \begin{pmatrix} e^{-i(\omega_0-\omega)t/2} & 0 \\ 0 & e^{-i(-\omega_0+\omega)t/2} \end{pmatrix}. \tag{4.29}$$

This can be rearranged to give

$$\hat{O}_{FE} = e^{-i(-\omega_0+\omega)t} \begin{pmatrix} 1 & 0 \\ 0 & e^{-i(-\omega_0+\omega)t} \end{pmatrix}, \tag{4.30}$$

where ω_0 is the angular frequency difference between the ground and excited states, ω is the driving frequency of the laser and t is the evolution time. The $e^{-i(-\omega_0+\omega)t}$ term on the right hand side of equation 4.29 describes a global phase change, which can be omitted with no physical significance.

For an atom in an eigenstate $\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ or } \begin{pmatrix} 0 \\ 1 \end{pmatrix}\right)$, the effect of the phase evolution is inconsequential as the relative phase is arbitrary; for an atom in a superposition of states, however, the energy difference between the ground and excited states causes the phases of the two states to evolve at different rates and a time dependent phase difference between the ground and excited states develops. As with a standard interferometer, the outcome of a measurement on the system is heavily dependent on the phase: if we apply a second $\frac{\pi}{2}$-pulse, the relative phase between the laser and the atomic wavefunction determines the result. This sequence can be described by the multiplication of the matrices in order:

$$\hat{O}_{\frac{\pi}{2}}.\hat{O}_{FE}.\hat{O}_{\frac{\pi}{2}} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\phi} \end{pmatrix} \cdot \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} \frac{1}{2} - \frac{1}{2}e^{-i\phi} & -\frac{i}{2} - \frac{i}{2}e^{-i\phi} \\ \frac{i}{2} - \frac{i}{2}e^{-i\phi} & \frac{1}{2} - \frac{1}{2}e^{-i\phi} \end{pmatrix}. \tag{4.31}$$

If $\phi = (\omega_0 - \omega)t$ is equal to zero, which occurs if the pulse separation is equal to zero, for example, the equation 4.31 reduces to that of a π pulse, whereas if $\phi = \pi$, it becomes the identity matrix.

The pulse sequence given in equation 4.31 is the basis for single atom interferometry.

The phase of these atomic interferometers, ϕ, is highly dependent on the local environment of each atom. If the driving laser is detuned from the transition frequency, for example, ϕ becomes dependent on the time the atom has spent in a superposition. It is also possible for the phase of the interferometer to depend on the translational energy of the atom. In this circumstance, given the correct pulse timing, the probability of absorption can be made to be proportional to the momentum, which can result in the cooling of the atomic ensemble. This is described in more detail in section 4.6.1 on coherent cooling.

In describing these processes, we derived the atom-light interaction by employing a sinusoidally varying electric field. In real systems, however, the light field contains phase discontinuities and non-zero line widths, which makes this an unrealistic approximation. In addition to this, the phase of the atom is disrupted by changes in the atomic environment such as through collisions. In order to characterise this, we define a coherence time as a suitable timescale over which disruptions in the phase of both the light and the atom are negligible.

In practical terms, coherence is determined by the time and distance over which two beams maintain a fixed phase relationship. Perfectly coherent systems require perfectly isolated systems and infinitely accurate frequencies: in our experiment we simply ask that the system be coherent for the duration of the experiment.

4.2 The optical Bloch equations, and the Bloch sphere

It is not possible to use the framework we have built up over the last few chapters to describe systems consisting of more than a single atom where we must be able to describe both pure and mixed states. In order to properly describe statistical processes in these systems, such as spontaneous emission or uneven light intensities across a cloud of atoms, which create mixed states, we must use a density matrix. We will then be able to build the equations governing these states into the optical Bloch equations, which we can use to form an intuitive, pictorial method of describing these interactions. We start by defining the density matrix as:

$$\bar{\rho} = \rho_{11}|1\rangle\langle 1| + \rho_{12}|1\rangle\langle 2| + \rho_{22}|2\rangle\langle 2| + \rho_{21}|2\rangle\langle 1|, \tag{4.32}$$

which we then use to derive the optical Bloch equations, and later illustrate using the Bloch sphere: a useful intuitive picture for multi-particle states.

The Bloch vector is a three-dimensional vector, with components along the x,y and z axes given by:

$$r_x = 2\mathrm{Re}(\rho_{12}) \quad , \tag{4.33}$$

$$r_y = 2\mathrm{Im}(\rho_{12}) \quad \text{and} \tag{4.34}$$

$$r_z = \rho_{22} - \rho_{11} \quad . \tag{4.35}$$

The rate of change of the elements of $\bar{\rho}$, as with any state, is given by their commutation with the Hamiltonian:

$$\frac{d\rho}{dt} = \frac{1}{i\hbar}[H_0 + H_I, \bar{\rho}] \tag{4.36}$$

which, as is derived by Mandel and Wolf [79], and first described by Feynman, Vernon and Hellwarth in 1957 [80] can be described by a precession around the field vector, $\bar{Q}$.

$$\frac{dr}{dt} = \bar{Q} \times \bar{r} \tag{4.37}$$

Similarly to the state vector, the field vector, $\bar{Q}$ has components along x and y axis given by the real and imaginary parts of the interaction between the excited and ground states, and a z component, which is given by the detuning of the field from resonance.

$$\begin{aligned} Q_x &= \frac{2}{\hbar}\mathrm{Re}[\langle 1|H_I|2\rangle], \\ Q_y &= \frac{2}{\hbar}\mathrm{Im}[\langle 1|H_I|2\rangle] \text{ and} \\ Q_z &= \omega_0 \end{aligned} \tag{4.38}$$

where as before,

$$\langle 1|H_I|2\rangle = d_{ba} \cdot E = \hbar\Omega \exp[-i(\omega t - \phi)]. \tag{4.39}$$

Before we change to a rotating interaction frame, the state and the field vector precess around the Bloch sphere at the transition and field frequencies respectively. By rotating the frame of reference at the same rate as the field vector, the field vector is fixed to an arbitrary direction, for which we use x, and the state vector is slowed down to make the equatorial precession $\omega - \omega_0$. The state vector precesses now around the field vector at a rate, $\tilde{\Omega}$. π and $\frac{\pi}{2}$ pulses, and free evolution are easily visualised in these representations (figure 4.3).

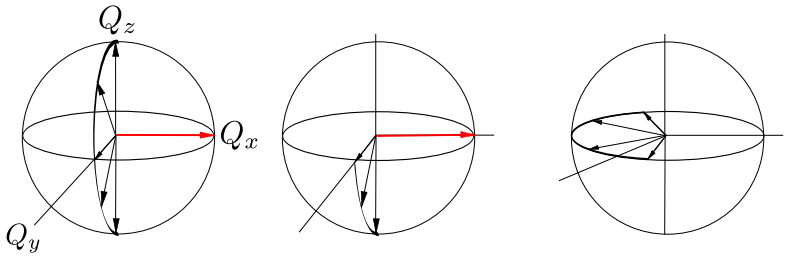

FIGURE 4.3: The Bloch sphere gives us a useful visual representation of π and $\frac{\pi}{2}$ pulses where the state vector (black arrow) precesses around the field vector (red arrow), and free evolution where phase evolution causes a precession of the state vector around the equator.

4.3 Forms of decoherence and dephasing

Like many quantum phenomenon, coherent optical effects are rarely observed in nature. While this observation can be partly explained by the lack of coherent light sources, it is also due to the lack of an atomic environment which gives uninterrupted evolution of the atomic phase for long periods of time. In this section, we look at the many forms of atomic perturbation which disrupt the phase of the atomic evolution, and lower the fidelity of each Rabi oscillation. We look at two different mechanisms: those, which we term 'decoherence' effects, which causes a sudden and randomly timed discontinuity in the atomic phase which distinguishes it from others in the ensemble, and those which we call 'dephasing effects', which cause a slow variation in the atomic phase due to a continuous variation in the atomic evolution across the cloud. Of course, this distinction is rarely so well defined in cases such as a cloud of fast moving atoms where the distinction between the two regimes is not obvious.

4.3.1 Decoherence effects

4.3.1.1 Spontaneous emission

During the Rabi cycle, the population of the excited state is non zero which gives the atom a non-zero chance of decaying via spontaneous emission to the ground state.

As shown in figure 4.4, after decay, the atom starts a new oscillation, which, if the atom is part of an ensemble, will have a random phase with respect to both the driving laser, and other atoms in the cloud. It is only in the absence of any variation in the phase between the atoms that complete population inversion can occur; as the atomic ensemble loses phase coherence, the contrast of the Rabi oscillation decreases. After many spontaneous decays, the phase of the atoms becomes random; in such an incoherent system, the contrast of the Rabi oscillations is zero.

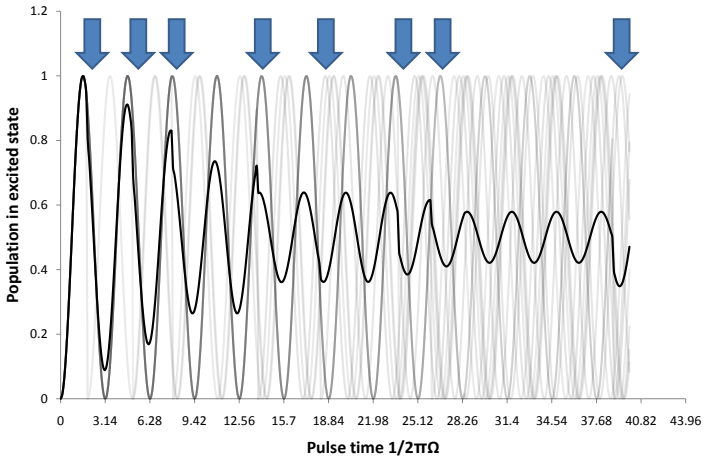

FIGURE 4.4: A simulation showing ten atoms driven coherently. Randomly timed decoherence events cause a loss of phase coherence between atoms which causes the contrast of the oscillation to decrease with time.

4.3.1.2 Collisions

In addition to spontaneous emission, collisions subject atoms to intense Van der Waals electric fields which cause significant perturbations to the energy levels of the atom for the duration of the collision. The outcome of this interaction will depend on the mutual orientation of the collisional partners and their wavefunctions, which are unlikely, due to the undefined coordinate axes of the particles, to be coherent.

4.3.2 Dephasing effects

4.3.2.1 Magnetic field broadening

To maintain a stable transition frequency and avoid magnetic broadening of the Rabi frequency due to external spatially and temporally varying magnetic fields we use the magnetically insensitive $m_f = 0 \rightarrow m_f = 0$ transition for manipulation. As the readout mechanism is based on single photon absorption, with an absorption linewidth much broader than the transition linewidth, it is not possible to construct a readout mechanism sensitive only to changes in the $m_f = 0$ state population. Instead of this, we tune the two-photon Raman frequency to the narrow $m_f = 0 \rightarrow m_f = 0$ transition such that we affect the population of this state only. This can be problematic if the magnetic states are overlapped in frequency space such that we interact with more than the $m_f = 0$ transition. In this case the different coupling strengths and detunings of the transitions

cause the atoms to be driven at different Rabi frequencies. This causes a washing out of the phase and a decrease in the contrast of the Rabi fringes with time.

4.3.2.2 Temporal and spatial variations in beam intensity

In many previous coherent manipulation experiments, the atoms have been highly localised within dipole or magnetic traps [81]. The reasons for trapping are plentiful: highly localised atoms can be subject to a manipulation beam of a smaller waist size, and therefore higher beam intensity, with a smaller intensity variation between atoms. In free-falling clouds of atoms, such as ours, it is necessary for the driving laser radius to be larger to compensate for the lack of localisation, giving a smaller and more varying light intensity across the cloud.

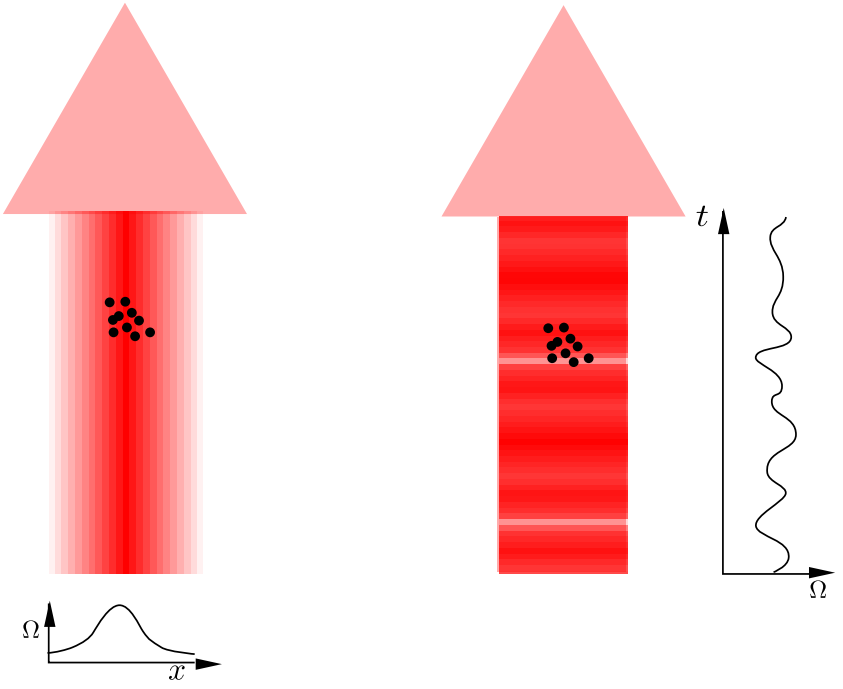

FIGURE 4.5: Temporal and spatial variations in the Raman beam power causes variations in the Rabi frequency during the experiment. Both of these processes cause a dephasing of the atomic population.

The primary reason for a spatial variation across the MOT cloud occurs due to an atomic density distribution which is larger than the laser beam waist. This causes a variation in the power of the Raman beam across the cloud and a corresponding variation in the Rabi frequency. As the readout mechanism is not spatially selective, (unlike [81]) the flopping frequency will be a integral of the different Rabi frequencies across the cloud.

Using a two dimensional projection of a Gaussian cloud given by equation 6.4, the overall signal is therefore:

$$\text{Signal} = \iint \left| \frac{\Omega}{\tilde{\Omega}} \right| \sin^2 \left(\frac{\sqrt{\Omega^2 + \Delta^2}\, t}{2} \right) \, dx \, dy, \tag{4.40}$$

where Ω is given by the position dependent Rabi frequency:

$$\Omega = \sum d_{ij} \sqrt{\frac{2I\rho}{c\epsilon_0 n}} / \hbar. \tag{4.41}$$

Δ is given by the position dependent light shift plus the frequency offset of the driving laser, which we set to zero by tuning to the center of resonance:

$$\Delta = -\frac{\hbar\Omega^2}{4\delta}, \tag{4.42}$$

using the parameters $\sum d_{ij}$ to represent the sum of the coupling strengths between the two states, I the average beam intensity.

In this calculation the atomic density

$$\rho = \sigma_{cz}(t)\sqrt{2\pi} \, \exp\left(-\frac{x^2}{2\sigma_{cx}(t)^2} - \frac{y^2}{2\sigma_{cy}(t)^2} \right), \tag{4.43}$$

and the light intensity

$$I = I_0 \, \exp\left(-\frac{x^2}{2\sigma_{bx}^2} - \frac{y^2}{2\sigma_{by}^2} \right), \tag{4.44}$$

are both described by a spatial Gaussian distribution, where we have used I_0 to refer to the peak intensity, $\sigma_{cx}(t), \sigma_{cy}(t)$ and $\sigma_{cz}(t)$ (the time dependence a result of the free expansion of the cloud) to refer to the cloud radius and σ_{bx} and σ_{by} the beam radius along the x and y (and z for the cloud) directions. The beam is assumed to be perfectly aligned with the MOT cloud in the x, y plane. Equation 4.40 cannot be solved numerically; we must therefore find a solution analytically: plotted in figure 4.7.

Although it is clear from figure 4.7 that the frequency is chirped which makes fitting an exponentially decaying sine function inaccurate for both long and short times simultaneously, it is convenient to be able to characterise the dephasing time using a single parameter: the exponential decay time constant. It is also clear from figure 4.7 that the Rabi frequency is chirped, with a lower frequency at the end than at the beginning. This chirp is a property of the dynamics of the system, where expanding atoms move into regions of lower light intensity, and therefore lower Rabi frequency.

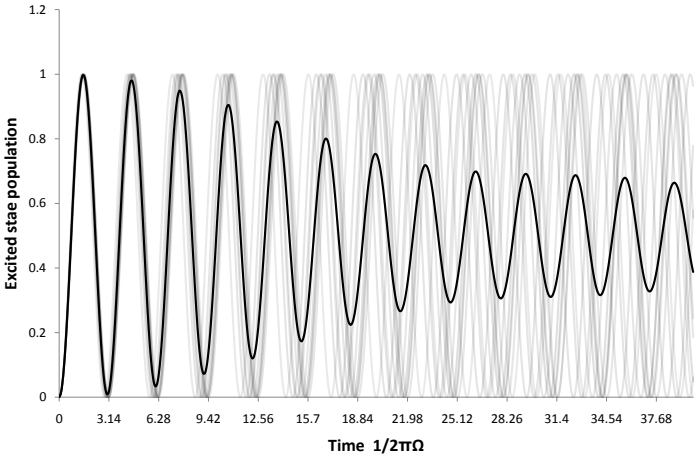

FIGURE 4.6: The effect of a 10% variation in the Rabi frequency on a model consisting of ten atoms

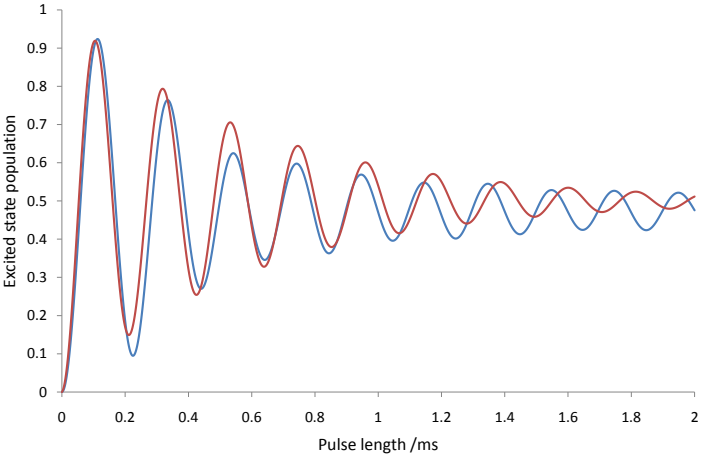

FIGURE 4.7: A comparison between the analytical solution of equation 4.40 (blue) with least squares fitted exponentially decaying sinusoidal oscillation (red) for typical MOT parameters (600 μm cloud radius, 2 mm beam radius, negligible cloud expansion). Note the change in the frequency as well as the difference in the amplitudes for long time scales.

4.3.3 Reversing the effects of dephasing: spin echo experiments

In many situations, where dephased atoms are moving slowly enough for their position to be considered constant, it is possible to reverse the effects of dephasing.

For the case of spatially varying Rabi frequencies, due to intensity variations across the cloud, for example, the Rabi frequency will be a function of the position of each atom in the cloud with respect to the center of the beam. In this situation, a gradient in the phase of the Rabi cycle will develop across the cloud while the driving laser is turned on. While a phase relationship still exists between the atoms in the cloud, the overall phase becomes inhomogeneous, causing a washed out readout. In these situations, it is possible to reverse the effects of the inhomogeneity.

An intuitive method of achieving this would be to use a second beam, with an intensity profile, such as $\sigma_0 - \sigma(x)$ where σ_0 is the maximum intensity of a Gaussian beam and σ the spatially varying beam profile, with higher intensity at the edges than in the center for the same duration of time as the initial pulse. The slower precession of the atoms in the center of the beam with respect to atoms at the edge of the beam causes the spread of state vectors to converge. Such a scheme is experimentally difficult to implement and 'spin-echo-like' experiments [33, 82] remain a more experimentally realisable alternative.

The method for this manipulation is similar to that of spin echo, which is applied to spin states, can be used to cause the state vectors to converge. A driving laser, with a field vector pointing along $+x$, turned on for a period of time t will cause some of the atomic state vectors to precessed further along their Rabi cycle than others. If, after this time, the phase of the driving laser is changed by π, the field vector, Q changes direction to become directed along $-x$. As a result of this switch in the direction, the precession of the Bloch state vector changes direction, but maintains the same rate. This causes a convergence of the phases across the cloud after a time given by the original length of the Rabi pulse.

This process can be visually represented in the Bloch sphere as shown in figure 4.8.

4.4 Methods of population inversion

The 'controlled area' method of coherent population inversion already described, where π and $\frac{\pi}{2}$ pulses are made up of light pulses which truncate the Rabi cycle, gives a conceptually simple introduction to methods of population inversion. Alternative methods of population inversion are available, however, which can, in many situations, offer higher fidelity operations at the expense of the minimum pulse length.

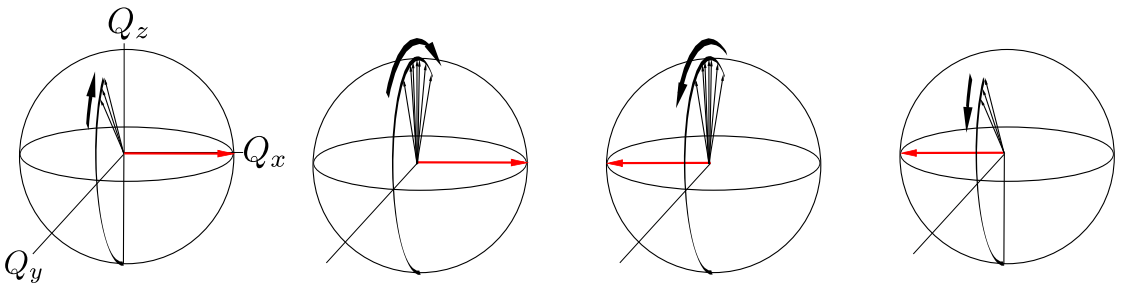

FIGURE 4.8: A Bloch sphere representation of the spin echo experiment: The driving laser is used to drive Rabi oscillations in an atom ensemble with an inhomogeneous Rabi frequency. After a period of time, t, the phase of the driving laser is changed by π, which changes the direction of precession of the state vector so that after another period of time, t, the state vectors realign.

4.4.1 The controlled area Rabi pulse

According to equation 4.24, the length of a π or $\frac{\pi}{2}$ pulse is dependent on the Rabi frequency, and is therefore also related to both the light intensity and the coupling strength between the two states: to achieve high fidelity operations, it is necessary to know the values of these parameters at each point in time and space, in order to apply the correct pulse timing necessary for inversion. When performing these pulses, it is inevitable that variations in the coupling strength, due to changes in beam pointing, stray magnetic fields or timing jitter will preclude the perfect extinction of the pulses. In addition to this, thermal atoms possess a Doppler spread in velocities that shift the resonant frequency of each atom, making it impossible to address each atom in an ensemble resonantly. These variation may occur spatially or temporally, thereby making it impossible to fulfill the criteria (equations 4.24, 4.27) for π and $\pi/2$ pulses with a single frequency. As a result of this, the fidelity of the total operation, which we define as being the fraction of the population that is coherent after the interaction, will suffer.

It is important that the atoms stay coherent throughout the interaction; once the atoms lose coherence they are driven out of phase with the pulse sequence, which makes it difficult to control to momentum transfer to the atoms. For coherent manipulation cooling schemes to be successful, the fidelity of each pulse must be close to one, so that after many pulses, the probability these atoms of being coherent is still much higher than zero. In many situations, for the reasons given above, this can be difficult to achieve using the controlled Rabi area method. An alternative to this is the adiabatic rapid passage method of population inversion, which does not require such judicious control over the atomic environment.

4.4.2 Adiabatic Rapid Passage (A.R.P)

In adiabatic rapid passage the driving laser frequency is swept, or chirped, from far below the atomic resonance (red detuning) to far above (blue detuning). During this process, the un-illuminated atomic energy levels, known as bare states, are perturbed by the electric field and as described by equation 4.18, no longer represent eigenstates of the system. In this environment, we name the new eigenstates 'dressed states'. If we sweep the frequency of the illumination laser, and plot the energies of the dressed states against that of the unperturbed bare states, we see the energies of the states converge to an avoided crossing. If the lasers are then suddenly extinguished, the dressed states revert back to the bare states, and the atom is left in an equal superposition whereas if the laser is swept through resonance to far blue detuning, the bare excited state becomes closer to the populated dressed state. If the lasers are suddenly extinguished at this point, the states revert back to their bare form and the population is left in the excited bare state.

The Bloch sphere picture gives an excellent visualisation of this process (figure 4.9): the field vector starts far off resonance so that the field vector, Q, points towards the south pole of the sphere making a small angle with the state vector. The state vector precesses around the field, tracing out a small cone which, like a spinning top, follows slow variations in Q. If Q is slowly moved from below resonance to far above resonance, the field vector moves from the south to the north pole of the Bloch sphere taking the state vector, which continues to make small revolutions around it, with it.

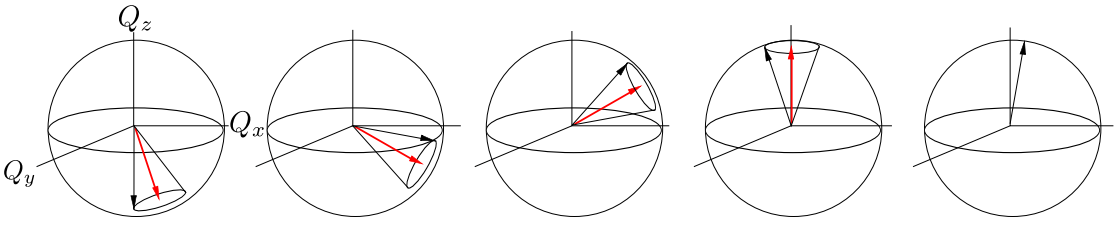

FIGURE 4.9: The Bloch sphere representation of adiabatic rapid passage: The field vector, shown in red, starts far off resonance, pointed at the south pole. The torque acting on the state vector (black arrow) by the field causes it to precess around the field vector, so that it draws out a cone of fixed angle (exaggerated here for clarity). As the field is chirped through resonance, the state vector follows it, dragging the state from the ground to the excited state. The field is quickly extinguished reaching the opposite far detuning, and the population is left in the excited state.

The condition for adiabacity dictates that the rate of precession $\tilde{\Omega}$ must be much greater than the change in angle of the field vector $d\theta/dt$ [83] where θ is the angle the state vector makes with the vertical axis. Unlike in the controlled area method, therefore, the state vector must precesses around the field vector many times during the inversion scheme, and the outcome of the state is determined more by the detuning of the field before and

after the pulse sequence than the timing of the Rabi frequency. Unlike the controlled area method, therefore, perfect fidelity in adiabatic rapid passage is not determined by the timing of the pulse area, but by the detuning of the laser, which can be controlled much more accurately. In addition to this, the condition for adiabacity implies that the final precession angle (the azimuthal angle of the state vector around the field vector) will be indeterminate.

The longer pulses necessary for A.R.P. can be inhibitive in some situations, such as in pulse schemes requiring large number of pulses, or where the state fidelity after a large string of operations is important. A full description of this can be found in reference [84].

Adiabatic passage can also be used to replace $\frac{\pi}{2}$ pulses, by stopping the evolution of the field vector while it is on resonance [85]. However, as with the case with the controlled area method, this is difficult to achieve for atoms with broad linewidths.

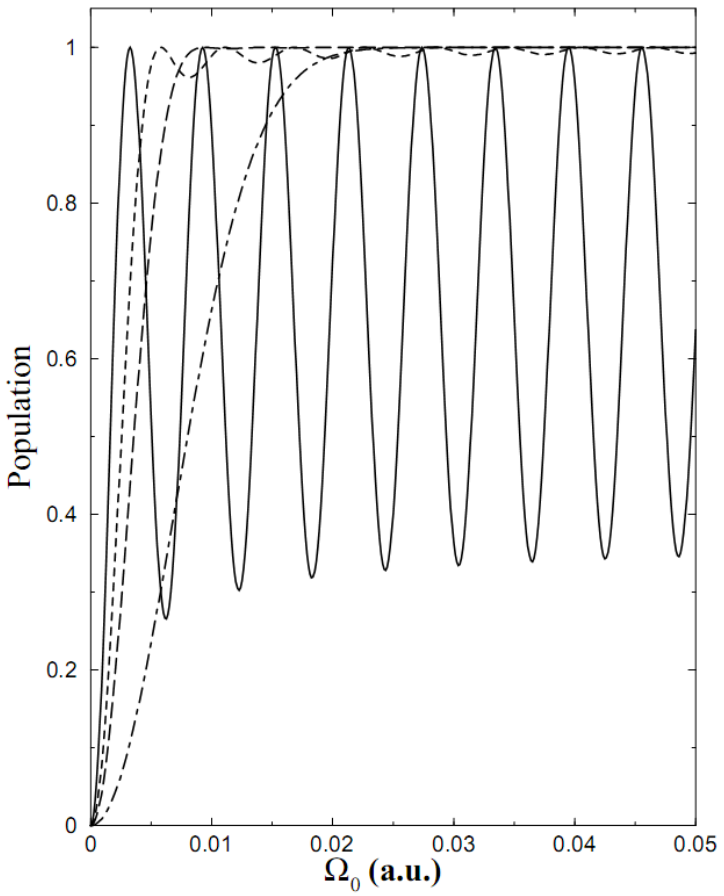

FIGURE 4.10: A comparison of the percentage of the population moved to the excited state for different rates of frequency chip [84].

4.4.3 Raman transitions and stimulated Raman adiabatic passage (STIRAP)

If we are to use light to manipulate our state coherently, it is important that we eliminate as much spontaneous emission from the system as possible; we must therefore choose a system in which the excited state lifetime is much longer that the duration of the experiment. This makes the decision of which atomic states to manipulate an important one. For example, if we choose to manipulate an atom via a dipole interaction, the excited state lifetime is typically of the order of nanoseconds, which makes it necessary for our coherent manipulations to be carried out on time scales much shorter than this. Alternatively, we can manipulate a three-level atom using a Raman transition, which uses two light fields at different frequencies, known as the 'pump' and 'Stokes' beams, to couple two states, which do not possess any direct dipole coupling, via a third state. The advantage in using this technique, as we shall see, is that it is possible to detune the lasers of the two beams sufficiently such that the third, intermediate, state is never populated, making it unable to undergo spontaneous emission. Using this technique, coherent systems can be manufactured with state lifetimes of many hundreds of milliseconds. To see what happens in the Raman lambda configuration, we apply the rate equation 4.13, which describes the rate of change of an n level system, and use it for a three-level system, driven by two frequencies ω_A and ω_B, shown in figure 4.11.

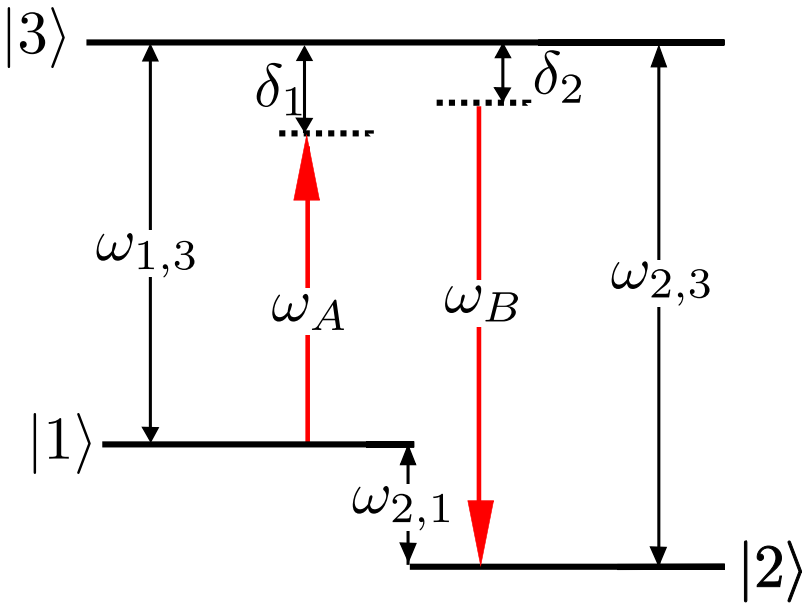

FIGURE 4.11: The two photon Λ configuration. States $|1\rangle$ and $|2\rangle$ are not coupled directly via an electric dipole interaction but indirectly via a third state. The detunings δ_1 and δ_2 are sufficient that the two photon rate is much greater than the single photon rate, which ensures that the state $|3\rangle$ is not populated

Starting with the rate of change of population of a single level, c_1, illuminated by a light field consisting of two frequencies such that $e^{i\omega t} = e^{i(\omega_A+\omega_B)t}$ we get:

$$\begin{aligned} i\hbar\frac{dc_1}{dt} &= \frac{E_0}{2}(d_{12}c_2[\exp(i\omega_{12}-\omega_A)+\exp(i\omega_{12}-\omega_B)] \\ &\quad +d_{13}c_3[\exp(i\omega_{21}-\omega_A)+\exp(i\omega_{21}-\omega_B)]). \end{aligned} \tag{4.45}$$

As the dipole coupling strength between the two hyperfine states is zero, we simplify this by setting the dipole moment to zero:

$$d_{12} = 0, \tag{4.46}$$

and by giving both ω_A and ω_B a common detuning, so that

$$\omega_{13} - \omega_A = \omega_{23} - \omega_B = \delta. \tag{4.47}$$

Let us now set the off-resonant terms to zero; the implications of this will be discussed later, when considering multiple Raman routes. A similar analysis is done for the populations in the remaining two energy levels, so that the behaviour of the system can be written in matrix notation:

$$i\frac{d}{dt}\begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} 0 & 0 & \Omega_{13}\exp(-i\Delta t) \\ 0 & 0 & \Omega_{23}\exp(-i\Delta t) \\ \Omega_{31}\exp(i\Delta t) & \Omega_{32}\exp(i\Delta t) & 0 \end{pmatrix}\begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix}. \tag{4.48}$$

As with the two level case, we apply a global phase transformation as in equation 4.17 to the state populations, which removes any time dependence from the Hamiltonian. We apply an additional rotation to the basis states of the wavefunction. This causes a change in the Hamiltonian and decouples one of the states from the others. A rotation to the eigenvectors:

$$|\psi_B\rangle = \hat{O}\,|\psi_B\rangle\,, \tag{4.49}$$

where

$$\hat{O} = \begin{pmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{4.50}$$

requires us to use a transformed Hamiltonian:

$$H_B = O\left(H\hat{O}^{-1}H - i\frac{d}{dt}\hat{O}^{-1}\right), \tag{4.51}$$

which preserves the original solutions. This corresponds to rotation around a fixed axis by an angle which is time independent, so that the time derivative of the transformation

equation is equal to zero, and the new Hamiltonian can be calculated to be [86]:

$$H = \begin{pmatrix} 0 & 0 & \Omega_{13}\cos\theta + \Omega_{23}\sin\theta \\ 0 & 0 & \Omega_{13}\cos\theta - \Omega_{23}\sin\theta \\ \Omega_{31}\cos\theta + \Omega_{32}\sin\theta & \Omega_{31}\cos\theta - \Omega_{32}\sin\theta & \delta \end{pmatrix}. \tag{4.52}$$

If the rotation angle, θ is set to make $\Omega_{23}\sin\theta = \Omega_{13}\cos\theta$ the equation above becomes:

$$H = \begin{pmatrix} 0 & 0 & \frac{1}{2}\Omega_0 \\ 0 & 0 & 0 \\ \frac{1}{2}\Omega_0 & 0 & 0 \end{pmatrix} \tag{4.53}$$

where $\Omega_0 = \sqrt{\Omega_{13}^2 + \Omega_{23}^2}$. We can then omit the middle state, which has become decoupled from the system, and write the Hamiltonian governing the remaining states as:

$$H = \begin{pmatrix} 0 & \frac{1}{2}\Omega_0 \\ \frac{1}{2}\Omega_0 & 0 \end{pmatrix} \tag{4.54}$$

This system is thus equivalent to a two level system; if we diagonalise the equation 4.54, we can find the eigenvalues to be:

$$\tilde{\Omega} = \sqrt{\Omega_{Raman}^2 + \Delta^2} \tag{4.55}$$

where

$$\Omega_{Raman} = \frac{\Omega_{23}\Omega_{13}}{2\delta} \tag{4.56}$$

and

$$\Delta_{Raman} = \frac{\Omega_{13}^2 - \Omega_{23}^2}{4\delta} \tag{4.57}$$

When we compare equations 4.56 and 4.57 with the two state case, equation 4.19, we see that they represent the transition strength and the light shift, respectively.

In order to avoid unwanted spontaneous emission events we look for long lived states with lifetimes much longer than the duration of the experiment. Particularly useful for this purpose are ground states which are split by the hyperfine interaction. These hyperfine states are ideal for coherent manipulations as they are not coupled via electric dipole transition[1], which gives long excited state lifetimes. The transition frequency between these states typically lies within the microwave frequency range, which makes it possible to couple the states using magnetic interactions inside high power microwave

[1] A transition between these states has a change in angular momentum $\Delta l = 0$ which violates parity conservation

resonators, or via a two-photon Raman interaction driven in a Λ configuration with a third state.

The two-photon transition is a second-order process, and therefore much weaker than the single photon transition. Unless the detuning from single-photon resonance is much greater than that for the two-photon resonance ($\delta \ll \delta_1 - \delta_2 - \omega_{21}$), the effect of the two photon interaction is swamped by the single photon transition. If the detuning from the single photon transition is not sufficient, the upper state population becomes significant. Since state $|3\rangle$ is coupled radiatively to the other two, this induces spontaneous emission which destroys the coherence of the system.

As well as long state lifetimes, there are a number of additional advantages to this two-photon scheme. The momentum imparted to the atom is a function of the angle between the two beams, which makes it possible to map the velocity distribution using velocity selective Raman spectroscopy (described in chapter 6.5). When the beams are counter-propagating, each π pulse delivers $\approx 2\hbar k$ of momentum, where k is the wavenumber of each photon. This is much higher than if we use single microwave photons to drive the transition, in which case each π pulse delivers $1\hbar k$ of momentum, where k is the wavenumber of a microwave photon which is much smaller than the equivalent optical photon momentum.

The process of driving adiabatic rapid passage with Raman beams is known as STIRAP [87]. In order to achieve this, a pulse sequence is applied which first couples states 3 and 2, before then coupling states 3 and 1. This can either be achieved by pulsing the individual Raman beams separately, or by changing the detuning of the two photon beams so that the 2-3 transition is excited before the 1-3 resonance. Because of the long state lifetimes when using these Raman transitions, the adiabatic criterion is much easier to meet.

4.5 Coherent cooling schemes

Now that we have described our tool box of π, $\frac{\pi}{2}$ and free evolution operators, we are able to describe how we can fit these together in different sequences in order to achieve cooling. There are two main schemes we are looking to implement: coherent amplification of cooling, and interferometric cooling. These are similar to each other in their use of coherent interactions, but quite different in many other respects. The final implementation of these schemes may in fact require a mixture both of these techniques to be used, as we will also discuss.

4.6 Amplification of cooling forces

It was suggested by Freegarde *et al.* in 1995 [88], and successfully demonstrated in 1997 by a group in Bonn [89] that a mode-locked Ti:sapphire laser with pulse lengths tuned to give π pulses, can be used to give a spatially selective force to deflect a beam of laser cooled caesium atoms. In the absence of spontaneous emission, the order in which the atom sees the two π pulses determines the direction of the radiation pressure felt by the atom. This order is determined by the time of flight of the photons: atoms on the plane of coincidence between the two pulses see both pulses simultaneously and therefore feel no overall force, atoms away from this plane will see one pulse before the other which will instigate absorption and then stimulated emission (figure 4.12). These two events supply an impulse to the atom which is parallel and then anti-parallel to the direction of the pulse, the impulses from both pulses will, if the atom is initially in the ground state, add to push the atom back towards the plane of coincidence.

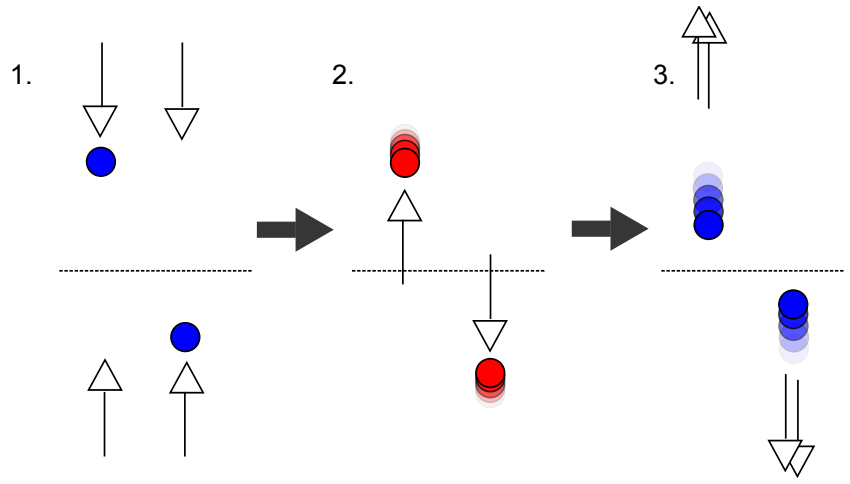

FIGURE 4.12: The interaction order and timing of an atom with two simultaneously released π pulses with opposite wavenumbers depends on the position of the atom relative to the plane of coincidence (dotted line). The first π pulse to reach the atom causes absorption, and an impulse directed towards the plane of coincidence. The second π pulse causes stimulated emission of the atom, and a second impulse again towards the plane of coincidence.

Schemes such as this have been successfully applied to atomic samples in order to impart hundreds of photon momenta, or tens of photon momenta to molecular sodium.

In addition to the ordering of the π pulses, the direction of the force is also dependent on the initial electronic state of the atom. For the focusing experiment above, it is necessary that the atoms start the pulse sequence in the ground state. A similar experiment could be conducted with the atoms initially in the excited state, which would cause the atoms to be repelled from the plane of coincidence.

Freegarde *et al.* [31] describes a situation similar to that above, where instead of the atoms being pumped into a common state at the start, a velocity dependent excitation at the start means that the initial state population is dependent on the momentum of the atom; atoms with positive momentum are excited to the excited state, whereas those with negative momentum are left in the ground. Incident π pulse exert an impulse in the direction of motion of the photon to atoms in the ground state (which undergo absorption, which are then excited to the excited state) and an impulse in the direction opposite to the direction of the photon for atoms in the excited state (as they undergo spontaneous emission and decay to the ground state). As the atomic populations oscillate between ground and excited state with each π pulse, they are also given an impulse which, due to the interleaved nature of the pulses, gives atoms with negative momentum a positive impulse, and atoms with positive momentum a negative impulse, slowly herding the momentum distribution towards zero.

Changes to the momentum distribution are more analogous to 'folding' as opposed to the 'herding' of atoms in Doppler cooling and so the momentum of an atom in the amplified cooling scheme can be reduced using a much smaller number of photons than conventional Doppler cooling. It also means, however that the final momentum distribution is not a Gaussian, as in thermal gases. Some degree of thermalisation must then occur between the atoms before we may accurately characterise the distribution by a temperature. This can be achieved through atomic collisions, which, for dilute gases, can require long periods of time. Feshbach resonances [90] can be used to increase the cross section of collisions, to decrease this time.

The cooling timescale here is determined quite differently from that in Doppler cooling; without the limit imposed by spontaneous decay, the atoms can be cooled on much faster timescales making it possible to apply the technique to narrow transitions, for which the decay time would otherwise be prohibitively long. In addition to this, atoms with a complex structure and many allowed decay routes can be cooled on a closed excitation-stimulated emission loop before they have a chance to decay to non-coupled states.

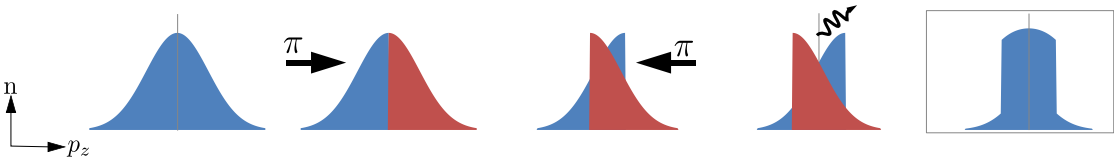

FIGURE 4.13: The amplification of cooling scheme: A velocity selective absorption first splits the velocity distribution, exciting atoms with positive momentum, while leaving those atoms with negative momentum in the ground state. Successive, interleaved π pulses overlap these two states in momentum space, so that when, as the final step, the excited state undergoes spontaneous emission, it leaves behind a narrower momentum distribution.

Without the spontaneous emission necessarily present in conventional Doppler cooling, the Doppler temperature limit does not necessarily apply here. Spontaneous emission is still, however, used as the last step in the cooling process to guarantee irreversibility, and therefore the single photon recoil temperature limit will still apply. However, to reach this low temperature limit, it is important that the right number of pulses are applied to give maximum overlap between the positive and negative velocity states. Once this maximum overlap has been achieved, the ensemble must be de-excited, and undergo a second velocity selective absorption if it is to be cooled further.

In general, instead of dividing the cloud momentum into two states, as per amplification of Doppler forces, we can imagine the atomic momentum distribution as being subdivided into a register of momentum states. In this way it is possible to address the ensemble as a momentum state quantum computer [91], to which divide-by-two algorithms are applied to reduce the overall momentum in fewer steps, achieving a narrower final momentum distribution.

4.6.1 Single atom interferometry

π and $\frac{\pi}{2}$ pulses resemble mirrors and beam splitters in phase space [85], which makes them useful in constructing phase space interferometers in which the atomic wavefunction is placed in a superposition of the ground and excited states. Similar to paths in a light interferometer, these recombine to give an output which depends on the relative phase difference between the two paths. This behaviour is exhibited by equation 4.31, which shows that the application of a sequence of $\frac{\pi}{2}$-free evolution-$\frac{\pi}{2}$ pulses give a state output which is dependent on the phase difference accrued during free evolution.

The beauty of this technique is that, as well as simply being dependent on the time the atom spends in free evolution, the phase on recombination is also dependent on the environmental conditions of each path, which, because of the mass and dipole moment of the particles, makes the interferometer sensitive to changes in gravitational [12] and electric fields[13]. Techniques using atomic interferometry have also been used for frequency referencing and the measurement of fine structure constant [15] and the electron dipole moment [16].

These interferometry techniques can be subdivided into two categories, depending on whether the atom receives an impulse from the light field on absorption, as in single photon interactions. For two photon absorption, the momentum impulse delivered to the atom is dependent on the angular separation of the two beams so that, for a co-propagating/counter-propagating beam orientation, the impulse given to the atom during the interaction is dependent on the difference/sum of the wavenumbers of the

two beams. It is possible to tune the angle between the two beams in order to address a particular velocity class within an atomic ensemble; equivalently, by measuring the absorption of the two beams as a function of their angular separation, one can infer information on the velocity distribution of the atoms. This technique makes use of what are known as recoil induced resonances [92], and shall be covered in more detail in our discussion of MOT characterisation 6.5.

4.6.2 Interferometry with momentum insensitive Raman beams

If the beams are co-propagating, the Doppler shifts for the two beams will be almost identical; we find that a cloud of laser cooled atoms will exhibit two photon Doppler broadening of the order

$$\Delta f = \frac{v}{c}(f_{pump} - f_{Stokes}) = \frac{0.1}{3 \times 10^8} \times 3 \times 10^9 \approx 1 \text{ Hz}, \tag{4.58}$$

where f_{pump} and f_{Stokes} are the frequencies of the pump and Stokes beams, and v is the mean thermal velocity of our laser cooled rubidium atoms. Using this, we find the shift in the linewidth as a result of the moving atoms to be below the resolution of our experiment.

In addition to this, due to the negligible momentum imparted to the atom from a single interaction the momentum of the atom after the interaction can considered to be equal to the momentum before the interaction. Using these velocity insensitive excitations, the atomic ensemble can be manipulated in order to construct single-atom Ramsey interferometers for which the output is dependent on, amongst other things, the temporal spacing of the two $\frac{\pi}{2}$ pulses.

By using counter propagating $\frac{\pi}{2}$ pulses, with a spatially extended cloud of atoms, such as in a vapour cell, the time between the two pulses is dependent on the position of the atom along the cell. Figure 4.14 shows the that the atoms experience the first pulse at a time $\frac{x}{c}$, and the second at a time $\frac{L-x}{c} + t_0$, to give an overall time that each of the atoms spends in superposition as:

$$\Delta t = \frac{L-x}{c} + t_0 - \frac{x}{c}, \tag{4.59}$$

where L is the length of the vapour cell, and t_0 the delay time between the two pulses. This time can be substituted into equation 4.31; we find that the output varies sinusoidally over the length of the cell, with fringe spacing of $c/2\delta$ where δ is the detuning of the light from the atomic transition frequency.

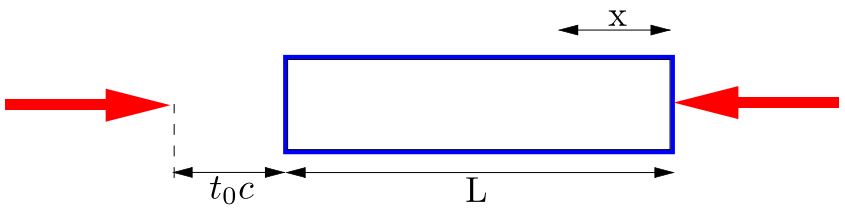

FIGURE 4.14: The time each atom in the vapour cell spends in a superposition, between the application of two $\frac{\pi}{2}$ pulses depends on the position of the atom along the cell, x. The two $\frac{\pi}{2}$ pulses are not released simultaneously, giving the second pulse a time delay, t_0, that guarantees the separation in time of each of the two pulses.

With our typical values for the detuning, $\delta \approx 2$ kHz, we obtain a fringe spacing of $\approx 10^4$ m. In order to make this technique experimentally feasible, therefore, we must find a way of increasing this detuning, such that the fringes fit within the length of a vapour cell.

4.6.3 Interferometric cooling

By changing the angle between the pump and Stokes Raman beams, we vary the momentum imparted during the interaction from $\hbar$ times the difference of the two wavenumbers in the co-propagating to their sum in the counter-propagating regime. We find that the counter-propagating regime behaves significantly differently from the co-propagating regime as momentum is imparted to the excited state path of the interferometer. The consequence of this is an interferometer phase which depends on the energy levels of the two atomic states, as well as the external momentum of the particle and the wavevector of the absorbed photons.

$$\Delta E = E_1 - E_2, \tag{4.60}$$

where E_1 and E_2 are the energies of the two states given by:

$$E_1 = \frac{p^2}{2m} + E_g \text{ and} \tag{4.61}$$

$$E_2 = \frac{(p+\hbar k)^2}{2m} + E_e, \tag{4.62}$$

where p is the atomic momentum. Thus:

$$E = \frac{p^2}{2m} + E_g - \frac{(p+\hbar k)^2}{2m} + E_e = \Delta E - \frac{\hbar k}{2m} - \frac{2p\hbar k}{2m}. \tag{4.63}$$

which we can then use in the time dependent Schrödinger equation to arrive at the energy-dependent phase difference between the ground and excited states, ϕ:

$$\phi = \frac{Et}{\hbar}. \tag{4.64}$$

The third term in equation 4.63 is dependent on the atomic momentum; the outcome of the interferometer is therefore dependent on the momentum of the atom, as well as the energies of the states and the photon wavelength. The outcome of the interferometer determines whether the atom will be given an impulse from the optical field. If the timing of the pulses is chosen carefully, therefore, the velocity dependent output of the interferometer can cool the atomic ensemble. This system forms the basis of interferometric cooling, initially proposed by Weitz and Hänsch in 1999 [32].

This interferometric cooling scheme can be compared to that of Doppler cooling: replacing the Doppler effect, which gives the velocity sensitive absorption, with a velocity sensitive interferometric effect. However, there is a crucial difference between interferometric and Doppler cooling: as the velocity sensitivity in interferometric cooling comes from the free evolution of the atom while the laser is turned off, the probability for a particle to absorb a net photon momentum depends on the particle velocity and on relative phases of the laser pulses, but not on the absolute laser detuning from an optical transition. It is possible, therefore to use broadband $\frac{\pi}{2}$ pulses for interferometry [93], which move the atomic population via adiabatic rapid passage. These chirped pulses have the added advantage that they can be used to couple many transitions at once [94] to simultaneously cool on systems comprised of individually open transitions, but which have small transition amplitudes to states outside the group.

Chapter 5

Cold atom preparation

As a proof of principle, we aim to demonstrate our coherent manipulation schemes on a freely-expanding cloud of ultra-cold Rubidium 85 atoms cooled within a magneto-optical trap. In comparison with room temperature vapours, Doppler and sub-Doppler cooled atoms have much lower densities and velocities and a corresponding lower collision rate. This makes them more suitable candidates for coherent manipulation schemes where it is necessary for the coherence time of the atoms to be equal to the time scales of the experiment.

For our magneto-optical trap, we require frequency stabilised lasers and a vacuum chamber containing a low density Rubidium gas. In addition to this, for coherent manipulation and state readout, we require the capability for the amplitude modulation of the beams, as well as imaging optics sensitive to the cooling transition frequency of the atom.

5.1 Rubidium for coherent control

Rubidium is a group I element with an atomic number of 37. Of its two naturally occurring isotopes, ^{85}Rb makes up 72% with ^{87}Rb making up the remaining 28%. Rubidium has always been a popular candidate for spectroscopy experiments due to its hydrogen-like atomic structure, for which the energy levels are calculable. As a result of this, the spectroscopy of Rubidium has been well documented [95]. The D2 transition, between the ground $|5S_{1/2}\rangle$ state and the excited $|5P_{3/2}\rangle$ is particularly popular for magneto-optical trapping due to the transition wavelength,which at 780.241 nm for ^{85}Rb, is easily accesible by the near-infra-red diode lasers used in compact disc players. In addition to this, the $|5S_{1/2}, F = 3\rangle \rightarrow |5P_{3/2}, F = 4\rangle$ transition forms a closed optical

loop. The short excited state lifetime of 26 ns [95] gives a low Doppler temperature limit (see section 3.6).

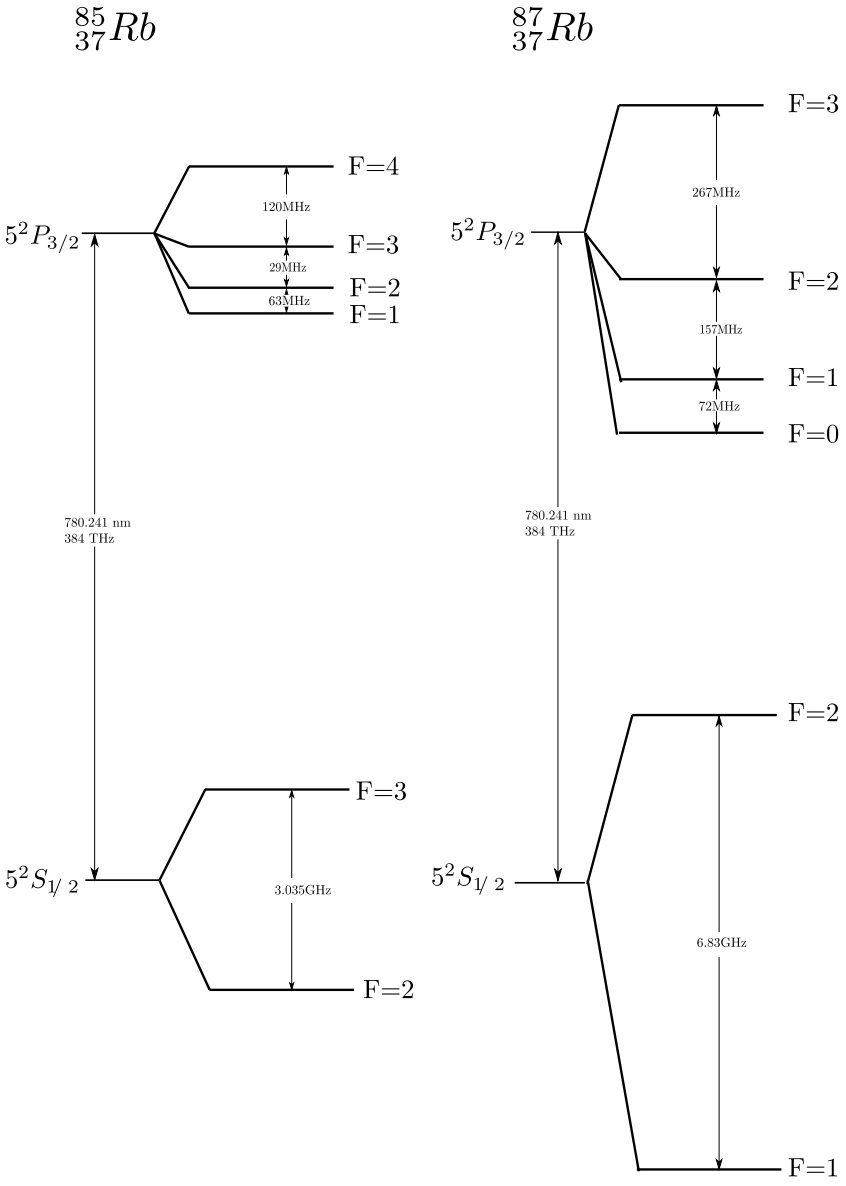

FIGURE 5.1: The D2 line of Rubidium 85, and, for comparison, Rubidium 87.

The atomic structure of ^{85}Rb and ^{87}Rb is similar enough that both can be cooled by a magneto-optical trap using only a single trapping a repump laser. However, the two additional neutrons in the nucleus of Rubidium 87 gives a difference in the nuclear mass and spin for the two isotopes which affects the energy level structure. Thus, the hyperfine interaction, which is a result of the interaction between the electron and the nuclear spin, is more than twice as strong ^{87}Rb as it is for ^{85}Rb. It is for this reason that we choose ^{85}Rb with a ground hyperfine splitting of 3GHz, over ^{87}Rb with one of 6.8GHz; as 3GHz is more easily accessible by equipment designed for the telecoms

industry, the synthesisers are cheaper and more easily available than for the higher ^{87}Rb transition frequency.

5.2 The vacuum chamber

The trapping forces of the magneto-optical trap are shallow enough that collisions with thermal atoms at room temperature supply sufficient energy to cold atoms to enable them to escape; the loading rates, and the total atom number are limited by the collision rate, predominantly between trapped and background atoms. Unless the background pressure is low enough, loss rates will be too high to form a MOT. Although some authors have reported magneto-optical traps in vapour cells [96], this has only been in the presence of an ion pump which reduces the pressure to 10^{-8} mbar, significantly lower than that of a standard vapour cell.

The low collision rates obtained through laser cooling are also useful for experiments which manipulate the properties of atomic coherence where otherwise the high collision rates contribute significantly towards the overall decoherence rate.

For these reasons, our coherent manipulation experiments are performed in a vacuum chamber with a pressure of 10^{-9} millibar at which the density is low enough for the atoms to be treated as ballistic with long trap and coherence lifetimes. Pressures as low as this are only achieved after impurities lining the inner surfaces of the chamber have been removed through cleaning and baking. The pressure is then reduced in two stages, first by a dry scroll pump and then by a turbo pump. Further details of the baking and construction of our vacuum system can be found here in reference [97]. The pressure in the chamber is reduced and maintained using a combination of an ion pump, for which the output current gives an indication of the atomic density within the chamber, and an absorption pump.

While it is necessary that the background pressure of atoms in the chamber is minimised during coherent manipulation, some rubidium background pressure within the chamber is necessary to load atoms into the MOT. However, as with any thermal gas within the cloud, too high a rubidium pressure can cause an increase in the rate of collisions which affects the coherence times of the atoms. The readout of our coherent manipulation is performed with a resonant beam, which excites background gas as well as the coherent atoms (see figure 5.4). If the background gas pressure is too high, the light emitted by these atoms can dominate the signal from the coherent manipulation.

The background pressure of rubidium in the chamber is controlled by getter-based sources which release rubidium 85 into the chamber after heating by a current loop.

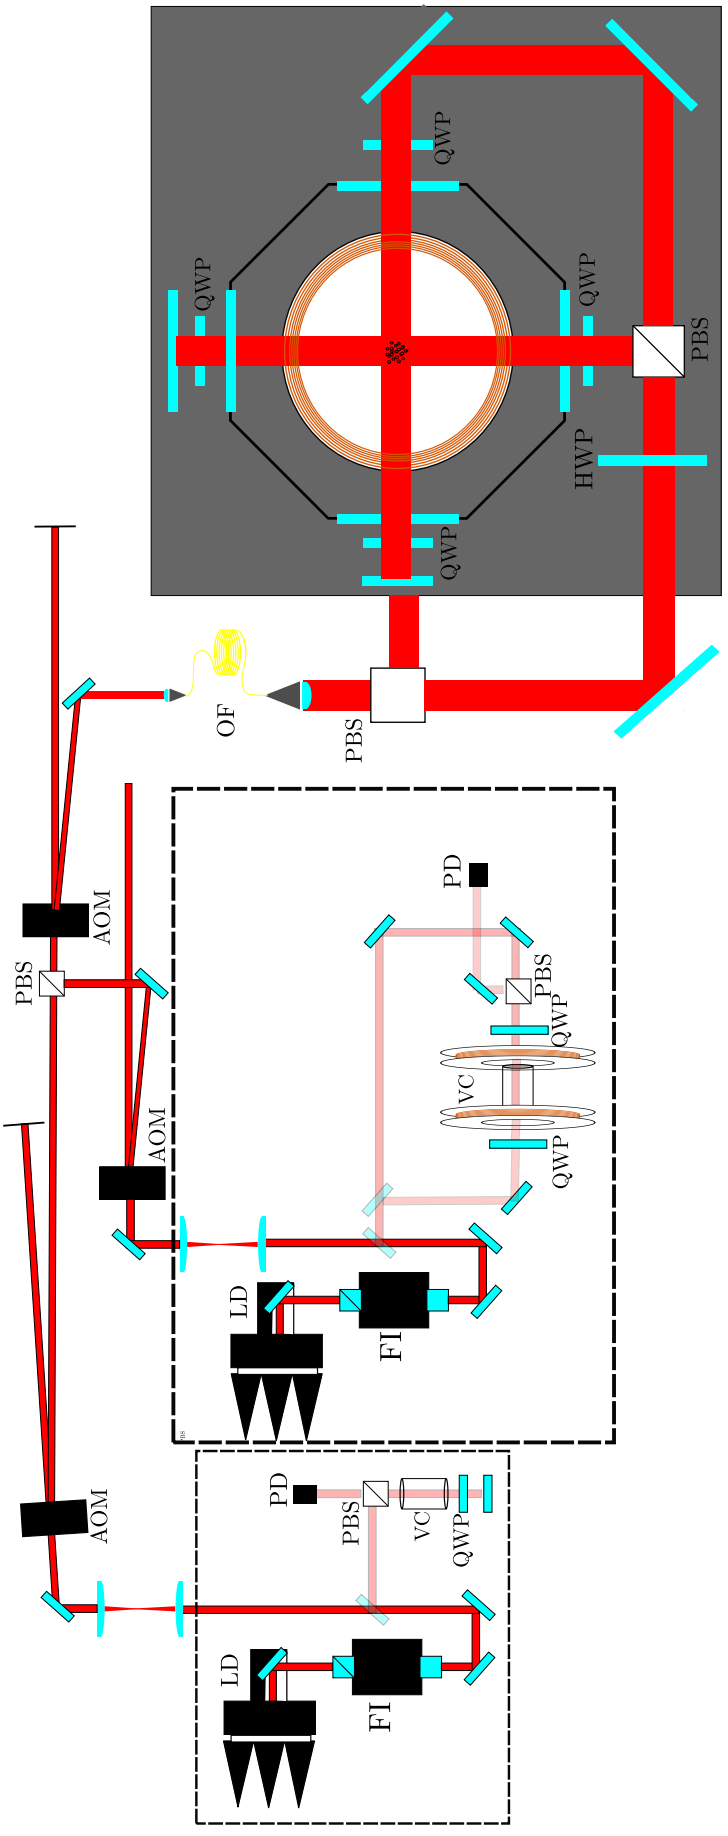

FIGURE 5.2: The trapping laser setup. LD:laser diode, FI:Faraday isolator, PBS:polarising beam splitter, PD:photo-detector, VC:vapour cell QWP: quarter wave plate, AOM-acousto-optic modulator, OF: optical fibre, HWP:half wave plate

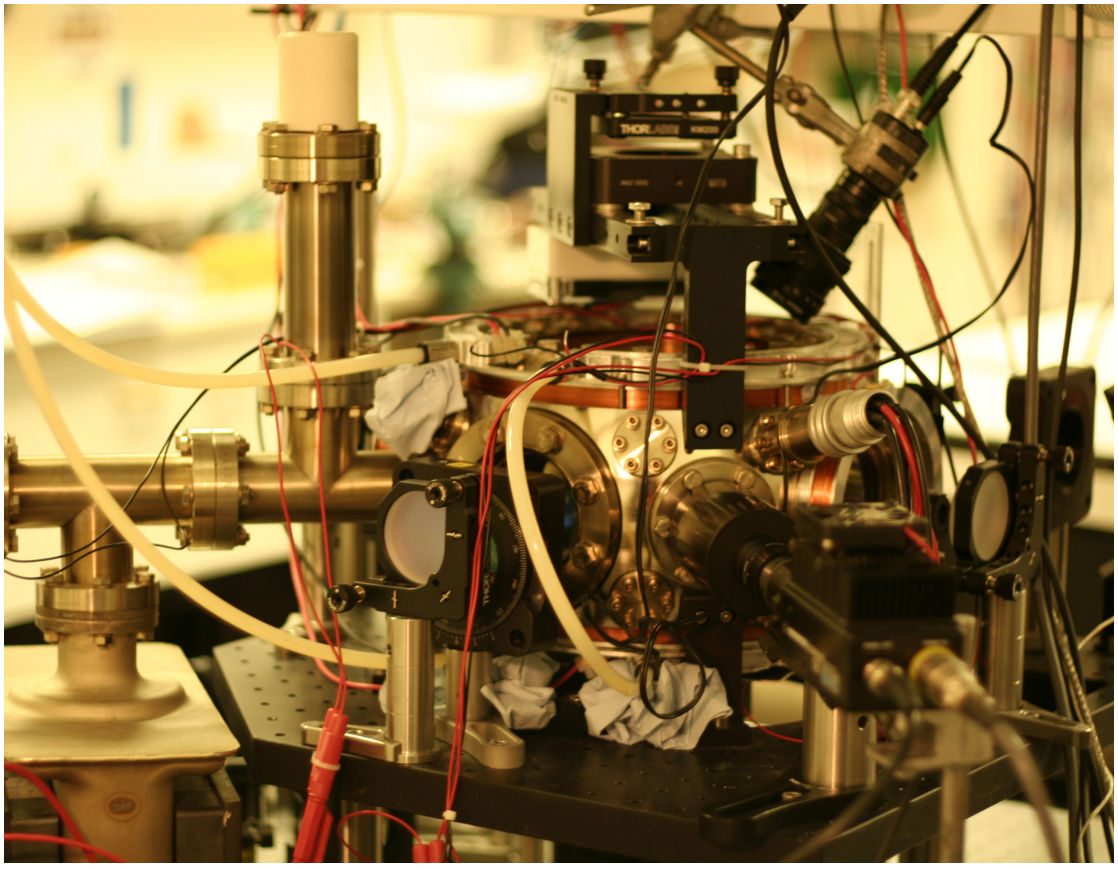

FIGURE 5.3: A photo of our MOT chamber showing three of the six mirrors used to retroreflect the beams, collection optics (foreground), imaging optics (top right) and pumps (left).

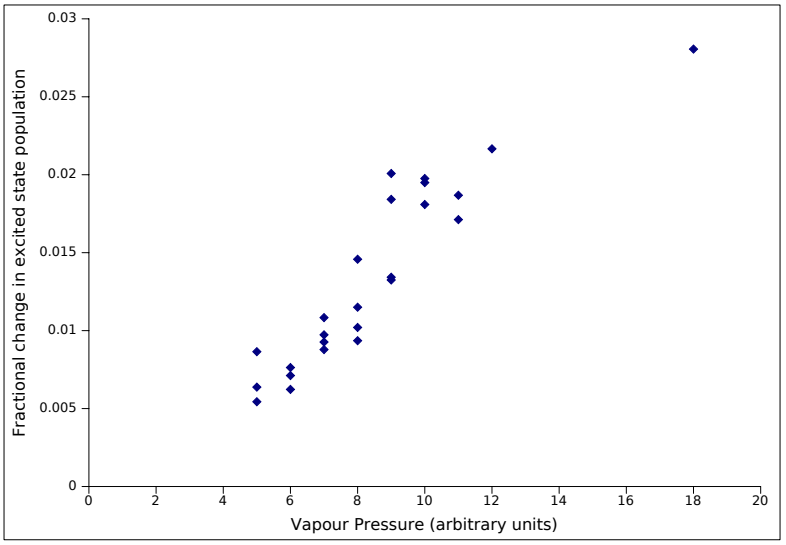

FIGURE 5.4: Applying a current to the getters increases the background pressure of the chamber and increases the number of unpumped atoms which enter the beam region, changing the relative proportion of atoms in the ground state which is important not to confuse with optical pumping.

Rubidium is released from the getters, once a threshold current is reached. In order to minimise the background pressure, after an initial loading of the MOT by the getters on a high current, the current is reduced to just above threshold for the experiments. The chamber is also resistively heated by the magnetic trapping coils, which releases atoms from the side of the chamber, increasing the background pressure, without the need for the getters to be turned on. This release of rubidium can also be achieved by UVA LEDs which cause light induced desorption [98] of rubidium lining the chamber, increasing the rubidium pressure of the cold chamber. The pressure of the chamber responds much faster to changes in the UV light intensity than to changes in the getter current, which makes the UV lights particularly useful in situations when fast changes in the background pressure are necessary.

5.2.1 Magnetic coils

To create the spatial potential for the magneto-optical trap, a pair of anti-Helmholtz coils with 15cm diameter are placed around the circumference of the top and bottom windows of the vacuum chamber. These give a constant magnetic gradient between 10 to 20 G cm^{-1} [97] along the vertical direction, with a zero which is correlated to the geometric centre of the vacuum chamber. A higher magnetic field gradient causes a faster rate of change of the optical force with distance. Wallace *et. al* showed that this causes the MOT radius to be proportional to the square of the magnetic field gradient [58]; this higher degree of localisation can be advantageous in many circumstances, such as for the uniform illumination of the cloud which is discussed later in section 4.3.2.2.

The resistive heating of the coils increases the temperature of the chamber, which can cause stress in the windows and creep in the alignment of the optics if it occurs too quickly. To limit this, we pass cooled water over the coils which limits the rate of heating, and the maximum temperature of the chamber.

The vacuum chamber is not magnetically shielded, and stray fields can move the position of the minimum field from the center of the chamber and cause unwanted structure in the trapping potential. The magnetic field inside the chamber is particularly affected by the strong permanent magnets in the ion pump, which are located close to one side of the chamber. To correct for this, we have a set of four compensation coils, which are placed perpendicularly around the four major horizontal windows in order to add a static offset field which can move the cloud back towards the center of the chamber. Along the vertical direction, we change the current in the trapping coil to achieve the same effect: the top coil operates at a current approximately 5% higher than the bottom

in order to correct for the offset due to gravity, keeping the atoms at the center of the chamber.

5.3 The trapping and repump lasers

Diode lasers are cheap, convenient and reliable sources of coherent light which, combined with an external cavity to form an external cavity diode laser or ECDL, give tunability and a narrower linewidth. These modifications give diode lasers the frequency and intensity stability suitable for cold atom and atom optics experiments.

There are two common arrangements of ECDLs known as the Littrow and the Littman-Metcalf [99] arrangements. The Littrow configuration (shown in figure 5.7) has a grating orientated at an angle such that the first order diffracted beam is directed back into the diode, causing lasing. In comparison, the Littman configuration, reflects the first order beam from a grating onto a mirror which is retro-reflected back onto the grating. The first order diffracted part of the beam is then used to seed the diode. There are a number of differences between the two ECDL arrangements, which must be considered for each application: the Littrow design, suffers a lower output power than the Littman as a result of the unused first order diffracted beam after reflection. In addition to this, the Littman-Metcalf arrangement has a comparatively broader linewidth than the Littrow which has a longer cavity length [99]. The Littrow arrangement also has the added advantage of an output alignment does not change when tuning the cavity.

For the cooling and trapping lasers, we use Sharp 120 mW, 784 nm laser diodes which, in conjunction with an external cavity in a design described by Arnold *et al.* [**?**], can be made to give 40 mW at 780nm. After filtering and amplitude modulation this gives a maximum power in both the trapping and repump beams of up to 14 mW in all three beams, which is then split evenly amongst the three MOT beams with a beam diameter of 7 mm, which gives sufficient optical power to exceed the saturation intensity.

The output facets of these diodes are not anti-reflection coated and reflections allow lasing without an external cavity. Even with an external cavity, feedback of the light from the front surface of the laser into the chip results in a set of laser modes that must be matched to the modes of the external cavity in order to obtain a stable output wavelength.

The wavelength of the laser can be controlled by supplying a voltage to a piezoelectric actuator mounted on the back of the grating such that it changes the angle of the grating, causing a change in the wavelength fed back to the diode. If this feedback is linked to an atomic reference, such that the voltage supplied to the piezoelectric is related to the

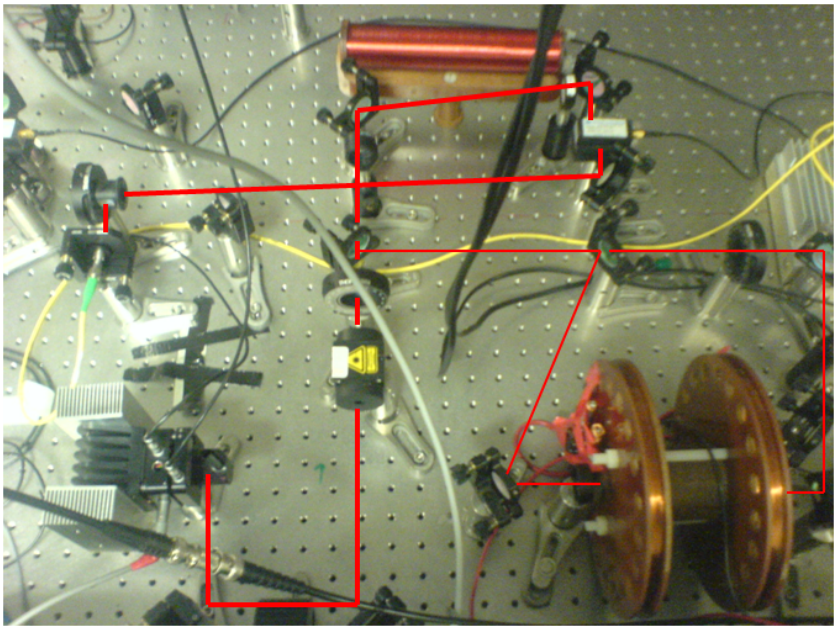

FIGURE 5.5: The trapping beam comes from an ECDL is locked to a saturation absorption spectroscopy reference which is Zeeman shifted by magnetic field coils, intensity modulated by an AOM, combined with the repumper beam by a polarising beamsplitter cube, and then coupled into an optical fibre for spacial filtering

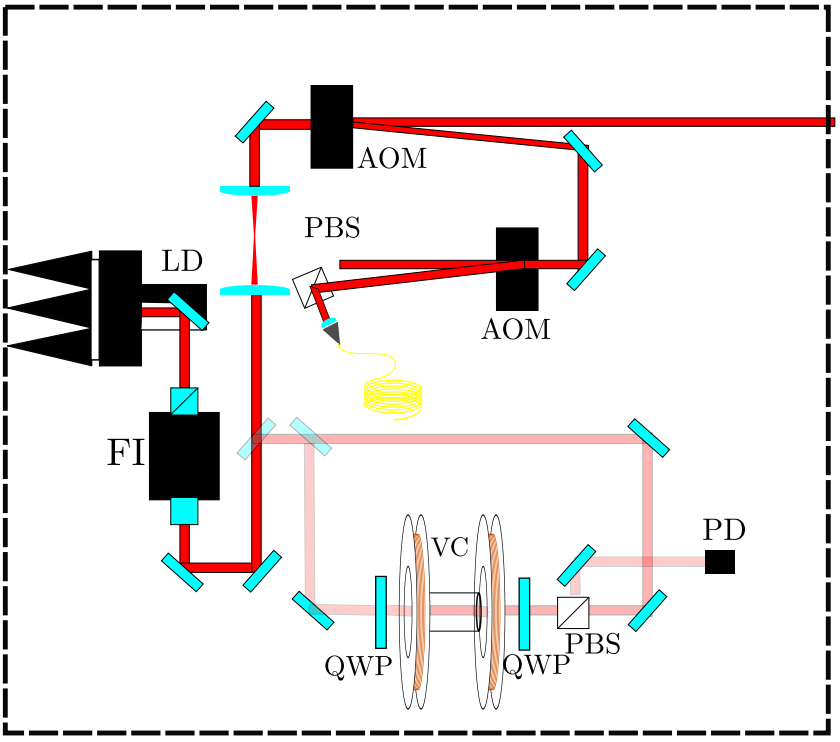

FIGURE 5.6: A schematic of the trapping laser setup. LD: laser diode, FI:Faraday isolator, AOM:acousto-optic modulator PBS:polarising beam-splitter, QWP: quarter waveplate, VC:vapour cell, PD:photodiode.

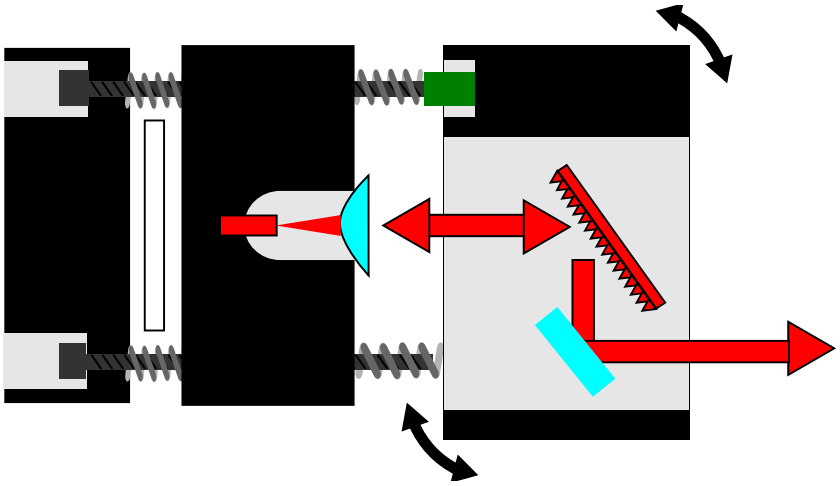

FIGURE 5.7: A grating in the Littrow configuration feeds back the first order diffracted beam to cause lasing in the diode chip. The grating can be rotated to change the cavity length and the grating angle, which affect the lasing output

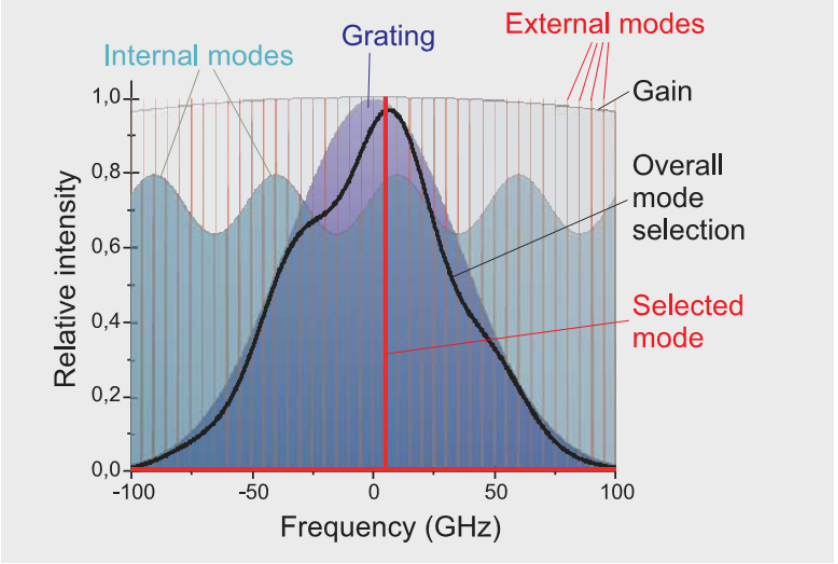

FIGURE 5.8: A figure showing the different mode structures of a typical diode laser in an ECDL setup [100]

detuning of the laser from an atomic transition, the laser can be stabilised to a fixed transition frequency.

In addition to the grating wavelength control, slower variations in the frequency can be corrected for by changing the temperature of the diode itself, which causes an expansion of the chip and a shift in laser modes. To maintain temperature stability of each of the diodes, and enable the temperature control of the wavelength, each laser is maintained at a constant temperature by a Peltier cooler which is mounted between the laser housing and the heat sink, and monitored by a sensor in thermal contact with the diode. Feedback to the Peltier is provided by a Thorlabs laser temperature controller.

For a constant cavity length to be maintained, it is important that the each laser is vibrationally isolated from any sources of mechanical noise. To achieve this, the experiments are performed on a Newport floating table which eliminates most of the vibrations. The lab sits next door to the physics mechanical workshop, however, and vibrations from the workshop machinery are noticeable on the feedback signal to the lasers, causing variations in the output frequency. It is partially due to this reason that most of the measurements presented in this thesis were taken after 5pm. In addition to this, changes in the air density inside the external cavity caused by air currents around the table also change the mode spacing and so we have constructed boxes which fit around the lasers to limit the air flow around the cavity.

As careful as one might be, movement of people in the lab causes an additional air flow around the equipment, and large vibrations occur when Allen keys are dropped; this was a fundamental reason for developing the remote data acquisition software, which allows the experiment to run without any manual presence in the lab.

The laser must also be optically isolated from the rest of the experiment to prevent back-reflections coupling into the diode. This is achieved through Faraday isolators, which are placed after each laser giving 30dB of isolation from back reflected light. Despite this isolation, reflections from the output facets of the optical fibres cause reflections to propagate backwards down the fibre; despite the small amount of power in the reflected beam and the optical isolation, the spatial mode overlap between input and output beams is good enough for a small amount of power to cause interruptions in the mode structure. To avoid this, we use angled-cleaved fibres, where the angled output facets produce reflected beams which are not coupled into the fibre.

5.3.1 Assembling the diode lasers and alignment of the external cavity

The laser construction is designed to maximise the temperature stability and vibrational isolation of the diode. Because of this, the diode is securely enclosed within a housing which is isolated from the optical table by a hollow posts. This maximises the thermal isolation of the diode from the optical bench, while giving the highest stability. A Peltier cooler is used to cool the diode and housing, transferring the heat to an air cooled heat sink. The assembly is held together with rigid springs with high thermal conductivity heat paste placed in between the components to improve thermal contact.

In front of the laser, an aspherical lens collimates the laser output making it necessary for the lens to be situated the correct distance from the diode output to give a parallel output. This is varied by a screw thread. The grating sits on the front of the diode housing, in front of the diode. The angle and distance of the grating to the diode are

roughly set by the geometry of the housing, however fine alignment of both the cavity length and the grating angle is achieved using two fine thread screws which pivot the grating along and perpendicular to the lines of the grating, to achieve changes in the cavity length and grating angle.

The design of the lasers was intended to decouple changes to the cavity length and the grating angle, to allow the cavity length to be set, before the angle of the grating is adjusted. Invariably, however, theses parameters are coupled, and ideal laser operation is achieved only after successive changes to the cavity length and the grating angle are made iteratively.

The laser diode will lase at a lower current in the presence of an aligned cavity, so that the correct cavity length can be found by reducing the current supplied to the diode to just above threshold, and adjusting the cavity length screw. At the correct length, the output of the diode will suddenly increase as a result of the cavity coupling. After the cavity has been adjusted to give the maximum output power, the current is increased to normal laser operating current, and, after a thermal equilibrium has been reached, the cavity angle screw is changed until it reaches the required resonant wavelength; this is inferred from a fluorescing vapour cell, absorption spectra or a wavelength meter. Large changes made to the grating angle change the lasing wavelength more or less linearly, so that it is normally straightforward to tune the grating angle to give a resonant output using the wavelength meter.

After the laser has been tuned to resonance, it is probable, especially if the grating angle screw was changed excessively during tuning, that the length of the cavity will have to be adjusted; this is also necessary after the laser has been knocked, or after many heating and cooling cycles caused by the daily turning off of the lasers. Misalignment of the cavity length can be inferred from discontinuous steps in the absorption spectra as the laser lases on more than a single mode (see figure 5.9).

The position of the grating pivot point means that changes in the angle of the grating also affect the angle of the output beam. This can be a problem for alignment of the lasers, making it necessary to realign the spectroscopy or fibre setup after small changes have been made to the grating angle. This problem is corrected for in a new laser design, which incorporates an additional mirror, which pivots along with the grating to act like a corner cube making, despite small positional changes, the angle of the output beam invariant with changes in the angle of the grating.

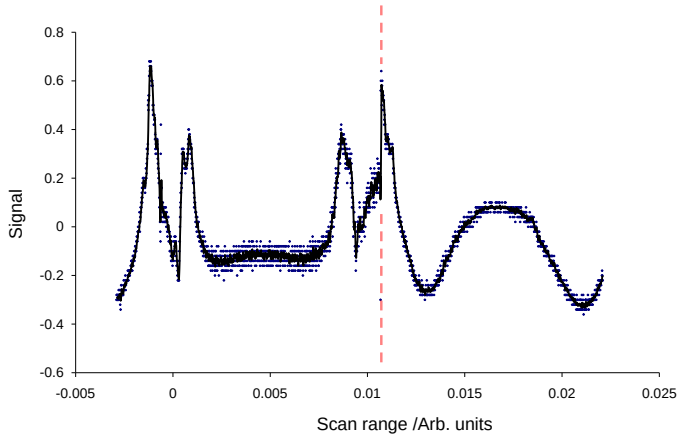

FIGURE 5.9: The lower hyperfine spectra with a discontinuous change in the absorption,shown by the dotted line; this is a good indicator that the external cavity length is poorly set.

5.3.2 Laser stabilisation

As changes in the detuning of the lasers can cause large changes in the MOT loading rate [101], frequency stabilisation of the lasers is vital to avoid any significant changes in the number of captured atoms. To achieve this, we use one of the most convenient, and most accurate frequency references available: the atom. By passing the beam through a vapour cell containing a natural mix of rubidium isotopes, we can obtain a locking signal derived from one of the Rubidium D2 absorption lines, which, when fed back to the laser through the piezo-electric on the external cavity, the temperature, or the current, stabilises the output wavelength.

5.3.2.1 Saturated absorption spectroscopy

Rubidium atoms at room temperature have a Doppler width which is much greater than the hyperfine splitting, which makes it difficult to resolve the individual hyperfine absorption lines. In order to lock to one of these lines, we use a technique known as saturated absorption spectroscopy. This looks at the absorption of a weak 'probe' after the atoms have interacted with a strong, counter-propagating, 'pump' beam. The pump beam causes large changes in the ground state population of resonant atoms, while the probe beam sees a transmission which is heavily dependent on the population of the ground state, although due to its low intensity causes a negligible change in the ground state population itself. For most atoms, with non-zero velocity, the atom sees pump

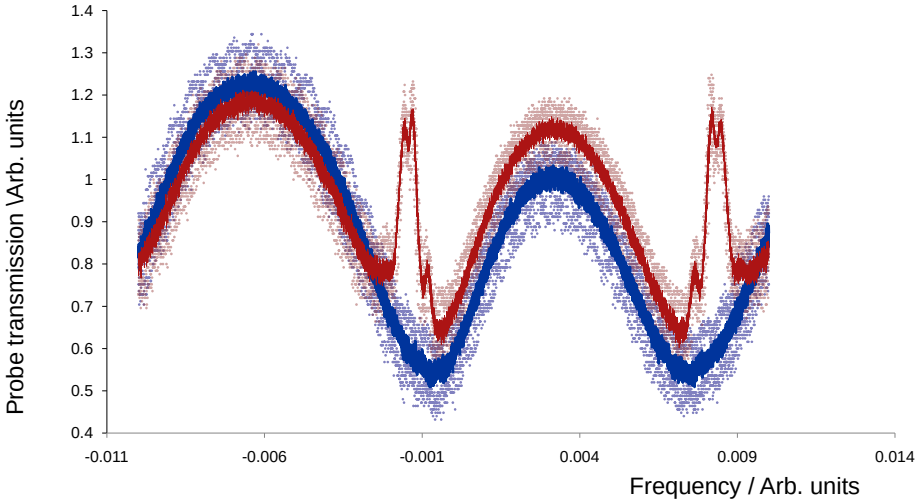

FIGURE 5.10: The lower hyperfine spectra with (red) and without (blue) a pump beam to give the spectra with and without Doppler free saturated absorption features. The red and blue lines show five point moving averages.

and probe beams at different frequencies and the pump beam does not affect the ground state population of atoms resonant with the probe beam. Atoms which are resonant with the pump frequency therefore see a depletion in their ground state populations; if these atoms are also resonant with the probe beam, which occurs if the atom is stationary, or in the case of cross over resonances (see below), a corresponding reduction in the scattering of the probe beam known as a 'Lamb Dip' occurs. This allows states with an energy splitting much less than the Doppler broadening to be resolved.

5.3.2.2 Cross-over resonances

When the Doppler broadened profile contains two or more transitions with a common ground state, one observes more Lamb dips than there are transitions as a result of 'cross-over' resonances. This occurs when the beams interact with atoms which have velocity corresponding to a Doppler shift equal to half the frequency difference between two states. The pump beam, resonant due to the Doppler shift, depletes the ground state via transition one; while the probe beam, interacting with the same atom, interacts due to counter-propagating direction, with the second state. As the ground states are shared

between the two states, the probe beam sees a depletion in the ground state of that atom, and although it is resonant with another transition, sees an decrease in the transmission.

5.3.2.3 Laser Locking Methods

The spectrum measured by saturation spectroscopy has Lamb dips in the Doppler profile which are centered on the transition frequency. These dips have a turning point at the atomic transition frequency to which we are unable to lock. For the repump laser, which does not need to be exactly on resonance, we can choose to offset the locking, and lock to either the side of the Doppler or Doppler-free profile. However, such a scheme is sensitive to variations in beam power, which makes the lock point susceptible to changes in the laser output which can be brought about by changes in the grating angle and air currents.

In order to avoid this, there are schemes which post-process the absorption spectra to create a locking signal which varies linearly around the transition frequency. The feedback from this, can then be fed directly into a servo amplifier, which feeds back the correct voltage to the laser. There are two popular methods of doing this: the Pound-Drever Hall method [102], or the dichroic atomic vapour laser lock (DAVLL) [103] method.

In the Pound-Drever-Hall method, the frequency is modulated, either by a Pockels cell or by modulating the laser driving current, by a small amount to give sidebands, offset by plus and minus the modulation frequency from the the main carrier frequency. These sidebands experience a difference in attenuation around the resonance so that when the light is demodulated, by mixing the light from the photodetector with the original local oscillator modulation signal, the sign of the output is different on each side of the atomic resonance.

In comparison, the DAVLL subtracts the signal after two oppositely circularly polarised light beams are passed through a vapour cell of atomic vapour. If a magnetic field is applied to the atomic vapour to split the m_f levels, the atoms become birefringent causing a difference in the absorption of the two circular polarisations. Subtraction of the two signals from one another in this case gives a linear locking signal centered around the atomic resonance.

Unlike the saturation absorption method, the shape of the DAVLL lock is much broader than the sub-Doppler features of the vapour and depends on the ratio of the intensity of the polarised beams, which can be affected by temperature and other environmental changes. This makes it necessary to monitor the lock point of the DAVLL against a

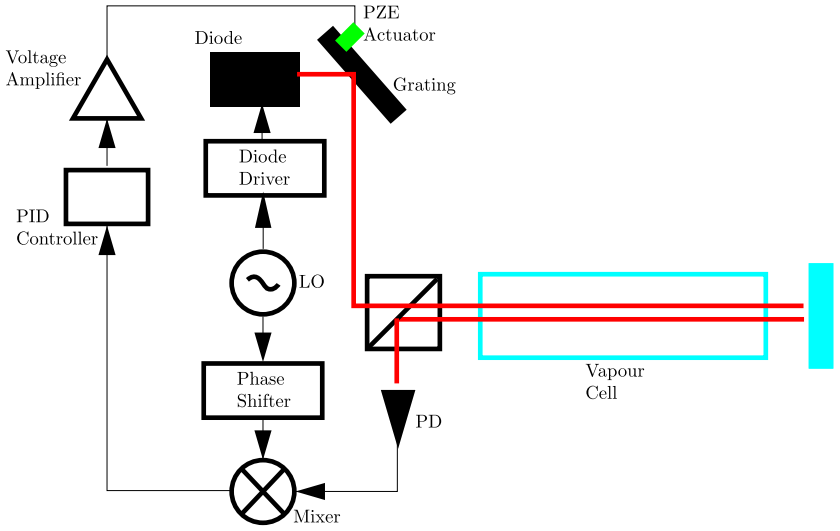

FIGURE 5.11: The setup for the saturated absorption/Pound-Drever-Hall locking technique. PD:photodetector, LO:local oscillator, PZE: piezo electric, PID: proportional, integral and differential.

fixed frequency reference, such as a wavelength meter, or a saturated absorption setup. This broad lockpoint, however gives the DAVLL systems much greater locking times which makes it possible for DAVLL to stay in lock for periods of days, as opposed to hours with an S.A.S. lock.

Due to the added experimental complexity of having to compare the output to an external frequency reference, the experiments performed in this thesis were performed using a saturated absorption lock, with the Pound-Drever-Hall method of locking. Since the initial construction of the S.A.S. setup, comparison of our system has been made with a working DAVLL system, and insufficient differences in the performance were observed to warrant a change-over to the DAVLL system. There are distinct advantages in both the broad error signal offered by DAVLL, and the narrow, atomic based, error signal given by S.A.S.; for this reason our group is interested in combining the error signals of both of these two systems into one: this work is in its early stages and ongoing.

5.3.2.4 Locking with a PID controller

To avoid oscillations of the laser frequency around the lockpoint, we use a feedback based on proportional, integral and differential terms derived from the laser error signal. If these quantities are added to the feedback in the correct amounts, the oscillations are damped, and the laser frequency more closely follows the lockpoint frequency.

The proportional error gives an feedback signal based on the current error, while the integral and differential based on the sum of past error, and prediction of future errors based on current rate of change of the error signal. As a result of this, while the

differential gain useful in systems where signal responds very slowly to changes in the feedback, it is of little use in our locking of laser diodes where the the output wavelength responds quickly to changes in the diode driving current. Changes in the temperature of the diode, however, can also be used to tune the output frequency; in this instance, the temperature feedback would have to be based on the differential of the error signal. Many algorithms are available [104] which find the optimum values of the proportional, integral and differential gain although we find that tuning by hand gives a lock which is sufficient. As a general rule, we repeat iterative steps of finding the maximum value for proportional gain before the output starts to oscillate, then reducing the proportional slightly from this value, and increasing the integral until it stops oscillating.

5.3.2.5 Offset locking, and choosing a suitable detuning for Doppler cooling

For Doppler cooling to be successful, it is necessary to detune the lasers away from resonance. Choosing a suitable detuning is a trade-off between the minimum temperature of the MOT, and number of captured atoms. For atoms at the Doppler cooling limit, the temperature of the atoms is related to the detuning by the relation [62]:

$$k_B T = \frac{\hbar\gamma}{4}\frac{1+(2\Delta/\gamma)^2}{2|\Delta|/\gamma} \tag{5.1}$$

which gives the minimum temperature, $k_B T = \hbar\gamma/2$, when $\Delta = -\gamma/2$. In the sub-Doppler limit, however, different temperature limits apply; empirical evidence in a paper by Lett [62] shows an inverse relationship between the temperature of the atoms in the trap and detuning which suggests that lower temperatures can be achieved if driving laser frequency is set to more than half the linewidth from resonance. As the atom number is also inversely proportional to the detuning minimum, arbitrarily increasing the detuning to obtain lower temperatures is not an option: we must instead find a compromise between temperature and atom number. We find a suitable compromise at 3 times the linewidth, or 18 MHz for the Rubidium D2 cooling transition. In some situations we find that it is advantageous to increase the atom number at the expense of a higher temperature; if we reduce the diameter of the trapping beams, for example, it is also necessary to change the detuning to increase the atom number in order to maintain a constant signal to the photo-detectors.

The linewidth of the locking signal produced by the Pound-Drever-Hall method is much narrower than 18MHz, and as such it is not possible to simply offset the lockpoint to stabilise the laser at this detuning. While the frequency shift of the trapping and repump AOMs can be used to achieve a detuning closer to 18MHz, although the range of offset frequencies we can produce by this method is limited. As it is already necessary

to modulate the amplitude (section 5.3.4) of the beams by passing them through two AOMs before the chamber, we use these to generate a 30MHz shift in the trapping and repump frequencies. If we then lock to an atomic transition which has been Zeeman shifted by 12MHz, we can achieve the necessary detuning. We apply this magnetic shift using a using a pair of coils in a Helmholtz configuration around the vapour cell. The current necessary to produce this shift can be calculated, from the geometry of the coils and the magnetic dependence of the state. It is much simpler, however, to calibrate the energy shift by comparison with the transition frequencies in order to apply the correct offset.

In order for us to be able to resolve the individual resonances, it is important that the magnetic field causes an energy shift which is homogeneous along the length of the cell to within the line spacing of the atomic resonance. Despite this, a slightly inhomogeneous field can be advantageous, since this inhomogeneity broadens the spectral line and thereby gives a wider locking signal, therefore increasing the lifetime of the lock.

The separation of the hyperfine states is small, so that it is possible that the magnetic shift is sufficient to overlap the opposite extreme m_f states for neighbouring F states, which can limit the resolution of each individual spectral line. In order to avoid this, the vapour cell is pumped with circularly polarised light so that the atomic population is pumped into one of the extreme m_f states, eliminating the interference from the overlap of the other transitions.

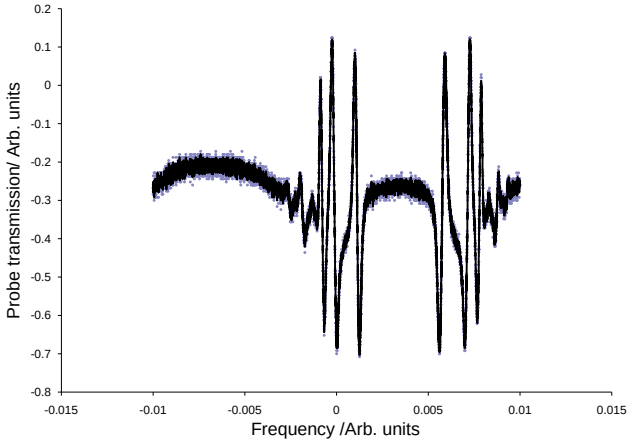

The magnetic coils used for offset locking are also useful for heating the vapour cell, which increases the vapour pressure within the cell to give a bigger signal to the photodetector. For this reason, the stability of the laser lock noticeably increases after the magnetic coils have had time to increase the temperature of the vapour cell.

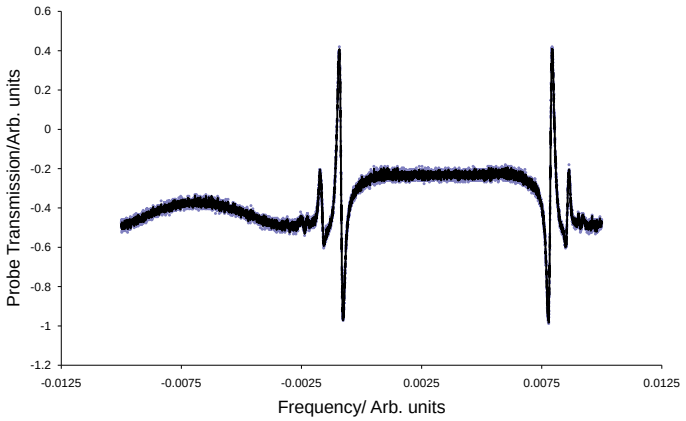

FIGURE 5.12: The Doppler free saturated absorption spectra after demodulation by the Pound Drever Hall method, with and without a magnetic field applied to shift the energy levels. Note that with an applied magnetic field, the population is pumped into the most positive two m_f levels by the circularly polarised light, and therefore only two spectral lines are observed.

5.3.3 Measurement of the laser linewidth of the trap and repump lasers

The tera-Hertz frequencies of both the trapping and repump lasers makes it experimentally unfeasible to measure the change in electric field with time in order to directly measure the linewidth of the trap and repump laser. Instead of this, we couple the repump and the trapping beams into an optical fibre and set the polarisations so that they are equal. With the two beams spatially overlapping, we lock the lasers to separate atomic lines and see a beat note which modulates the intensity of the light at the difference in the transition frequencies. As the lasers and locking electronics are identical, we can assume that the linewidth of the two lasers is identical, so that a measurement of the variation in the beat note between the two frequencies can be used to infer the linewidth of the individual lasers (more experimental details on this can be found here [105]). In addition to this, by looking at the spacings between the zero point crossings, we can measure the Allan variance [106], which describes the deviation of an oscillator over a given time interval. The stability between the lasers was found to be 220 $\pm$ 130 kHz on time scales from 1μs $\rightarrow$ 10 ms.

5.3.4 Intensity Modulation: Acousto-optic modulators

While the MOT beams are turned on, the $F = 4$ and to a lesser extent the $F = 3$ excited states are populated, and spontaneous emission occurs to either of the two ground

hyperfine levels. In order to maintain the coherence of our atoms during manipulation, we must ensure that the excited states remain unpopulated which we achieve by extinguishing the MOT beams during the coherent manipulation. Independent control over the trapping and readout beams is also required for processes such as state preparation, readout of the coherent states (section 8.1) and release-and-recapture temperature measurements (section 6.5).

We perform the switching with acousto-optic modulators (AOMs): non-linear crystals with traveling compression waves perpendicular to the light wavevector. The light, with a speed many times that of the compression wave, sees a slowly moving grating, which causes diffraction. By switching the amplitude of the compression wave, we can turn these diffracted beams on and off at the rate at which it takes the sound wave to pass over the beam front (typically tens of nanoseconds). Although small when compared with the speed of the light, the speed of the compression wave causes a frequency shift of the diffracted light by an amount given by the product of the compression frequency and the diffraction order, which makes it possible to induce both positive and negative frequency shifts.

Each AOM offers diffraction efficiencies up to 90% for beams with flat wavefronts and a Gaussian profile while at an angle that obeys the Bragg condition; although our diode lasers give out a less than perfect beam profile, we still achieve efficiencies upwards of 70%. A single AOM gives an extinction ratio better than than 40 dB which is not sufficient to be able to neglect single photon interactions with the MOT cloud while the AOM is turned off. To compensate for this, we use a second AOM which doubles of the extinction of the trapping and repump beams during coherent manipulation.

5.3.5 Optical fibres for spatial filtering

Although close to a Gaussian profile, the transverse beam quality of the diode lasers is not perfect. This is exacerbated by the AOMs which diffract only parts of the beam profile which meet the Bragg condition. In order to correct for this, we sacrifice beam power for a good beam shape and filter the light by coupling into an optical fibre. Once the light is filtered, the MOT cloud takes on a more uniform shape. In addition to this, using a fibre between the laser and the chamber allows the spatial alignment of the beam through the switching optics, which is affected by changes in the grating, to be decoupled from the alignment of the beam onto the MOT cloud; in doing this we know that, so long as the fibre output coupler does not move, once the beams are coupled into the fibre, they will be aligned in the correct positions for a MOT.

5.3.6 Alignment of the MOT beams

Fine and coarse alignment of the MOT beams is crucial for the success of our coherent manipulation techniques; the correct balance of optical forces obtained from a well aligned MOT gives a higher number of colder atoms. In order to achieve sub-Doppler temperatures, the atoms must sit in a standing wave between the incoming and outgoing beams, which causes a light shift greater than that of the magnetically induced Zeeman shift; if the MOT beams are not counter propagating, or if the MOT is offset from centre to a region of non-zero magnetic field, sub-Doppler cooling will not occur.

Our MOT design uses one beam, which is a mixture of the trapping and repump beams, linearly polarised and in a good Gaussian mode, which is then split into three beams by two polarising beam splitters. The ratio of the splitting is determined by the polarisation of the light, which we set to give equal powers in all three beams. Three quarter-waveplates are then used to give the correct circular polarisations necessary for trapping. Each of these beams is sent though the vacuum chamber so that they coincide with the geometric center of the chamber, which also overlaps with the zero of the magnetic field. These beams are then retro-reflected through a quarter-waveplate by a mirror, which reverses the direction and polarisation of the light. The alignment is non-trivial and we find the only practical method of aligning the beams with the zero of magnetic field to be through geometrical methods.

The symmetrical nature of the chamber is such that if the beams pass through the centre of each window they will also coincide with the zero of the magnetic field. In order to achieve this, we use circular plastic cutouts, which are a tight fit over the windows, with small holes cut into the centre. If the MOT is aligned correctly, the small amount of light will pass through the hole on the input window, through the small hole on the opposite window, and then be retro-reflected along the same direction through the holes on the way back. For the larger windows, which are too constrained by equipment to use the circular cutouts, we use a fine thread which is wrapped around the symmetrical nuts of the chamber to indicate the centre. For sub-Doppler cooling, it is essential that the the beams are perfectly counter propagating, and to achieve this, we place a iris at the output of the fibre which reduces the beam diameter. If the beams are perfectly counter propagating, both the outwards and retro-reflected beams will pass through the iris. By meticulous alignment of the beams, we achieve a cloud with a higher atom number, at a significantly colder temperature than the cloud for loosely aligned beams.

The lowest temperatures of a magneto-optical trap are achieved when the beams are perpendicular to one another. For ease of alignment, we achieve this using vertical and horizontal beams through the six biggest windows of the vacuum chamber. Such an

arrangement uses up valuable optical access that could be put to better use in increasing the optical access for the imaging optics. Our chamber design makes it possible to 'rotate' the beams so that they enter the chamber through some of the smaller windows, which would allow the large vertical windows to be used for readout giving greater optical access to the cloud.

Chapter 6

Characterisation of the atomic cloud

6.1 Why a MOT for coherent manipulation?

Although well documented that the gas clouds produced by magneto-optical trapping give low collision rates, before we neglect these processes ourselves, we must calculate the collision rate in our atom cloud which we can achieve from measurements of the cloud size, atom number and temperature. If we find this rate to be much greater than the time scales of the experiment, we can justifiably neglect them from our interaction picture.

In addition to helping to estimate the collision rate, we will go on to discover how many of the additional decoherence processes described in section 4.3 are affected by the properties of the MOT cloud. Regular measurements of the atom cloud are needed to characterise the change in coherence with changes in these parameters. So that regular characterisation does not take time away from the main running of the coherent manipulation experiments, it is preferable for these measurements to be quick and automated. Finally, as the objective of this work is to implement changes in the momentum distribution, it is vital that we are able to measure the shape of the velocity distribution both before and after manipulation.

The main characterisation parameters are the position, size of the MOT cloud, shape (density distribution) and temperature (velocity or momentum distribution), many of the which have been covered in previous theses by other group members [97, 105, 107].

Although stable on short time scales for the duration of individual experiments, slow drifts in the experimental setup can cause daily variations in the MOT parameters.

The lock points of the lasers, for example, are derived from the attenuation of a probe beam inside a rubidium vapour cell which is dependent on the temperature of the cell. Although a constant temperature of the lab is maintained to within a degree by an air conditioning unit, the many amplifiers, lasers and magnetic coils located around the optical table result in local variations in the temperature. These local temperature variations have also been found to cause small changes in the alignment of the optics, which affect the MOT beam parameters.

6.2 Measuring the atom cloud shape

In normal operation, the MOT acts as a source of spontaneously emitted, near infra-red (NIR) photons. Because of this, if loaded to a high atom number, it is possible to see the cloud with the naked eye. Under normal loading conditions, however, the wavelength of the emitted light, which is at the far end of the eye's response, and the small size of the MOT make it necessary to use imaging optics which are more sensitive in the NIR.

We image the MOT cloud onto a Watec CCD camera, which we then display on a CRT screen. We use this image to give a qualitative indication of the alignment, the size and atom number of the cloud and the background pressure within the vacuum chamber.

If the pressure of Rubidium in the chamber is sufficiently high and the lasers on resonance, Rubidium within the beam path fluoresces making it visible to the cameras; this can be used to fine tune the alignment of the cloud to coincide with the centre of the beams. The shape of the MOT also gives a good indication of this alignment, with many exotic toroidal and flared shapes indicating misalignment of the beams (figure 6.1). Despite independent laser and camera triggering, the image of the MOT beams while the trapping beams are pulsed on and off can be used to quickly infer some information on the velocity distribution of the cloud, particularly if it is non-spherical.

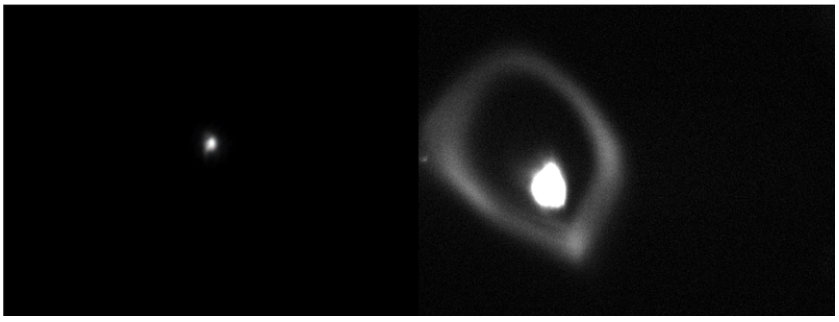

FIGURE 6.1: If the MOT is over loaded and misaligned, it is possible to form toroidal shaped clouds.

From a single camera, it is only possible to infer information from a 2D projection of the MOT cloud, which makes it necessary to use two imaging cameras orientated orthogonally to ensure that the cloud is spherical from all directions. During alignment we use a combination of both the lens system described above, and a separate CCD camera which we attach to our readout optics.

6.3 Measuring the atom cloud size

Using a CCD camera with a known pixel size it is possible to measure the MOT size using the image on the CCD. In this case, the magnification of the image is given by the ratio of the distance from the cloud to the lens, u, and the lens to the CCD, v:

$$M = \frac{u}{v}. \tag{6.1}$$

Alternatively, if we take an image of the MOT cloud and then use the same collection optics to image a ruler, we find that the two images will have the same magnification, enabling a direct measurement of the field of view which can be used to calculate the MOT size. Such a system relies on a shallow depth of field, to be certain of the exact distance at which the camera is in focus; this can be achieved with a large enough distance, v.

For the limit of low intensity trapping beams and non interacting particles, the density distribution of the MOT cloud centered on the coordinates (x_0, y_0, z_0) takes the form of a 3 dimensional Gaussian given by:

$$\rho(x,y,z) = \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2} - \frac{(z-z_0)^2}{2\sigma_z^2}\right), \tag{6.2}$$

which can be re-written as a product of a two dimensional Gaussian, $\rho(x,y)$:

$$\rho(x,y,z) = \rho(x,y)\ \exp\left(-\frac{(z-z_0)^2}{2\sigma_z^2}\right), \tag{6.3}$$

where σ_x, σ_y and σ_z are the widths of the Gaussian cloud along the x, y and z directions respectively. For imaging optics with a depth of field significantly larger than the cloud width, the direction along the z axis is integrated to give:

$$\int_{-\infty}^{+\infty} \rho(x,y,z)dz = \int_{-\infty}^{+\infty} \rho(x,y)exp\left(-\frac{(z-z_0)^2}{2\sigma_z^2}\right) = \sigma_z\sqrt{2\pi}\rho(x,y). \tag{6.4}$$

Which tells us that the width of the cloud can be inferred by measuring the width of the two dimensional projection of the three dimensional cloud.

Figure 6.2 gives the fluorescence cross section for a typical MOT cloud. From this, we can measure the $1/e^2$ value of the MOT size to be 328 ± 3 μm and 357 ± 3 μm along the two axis of the imaging system.

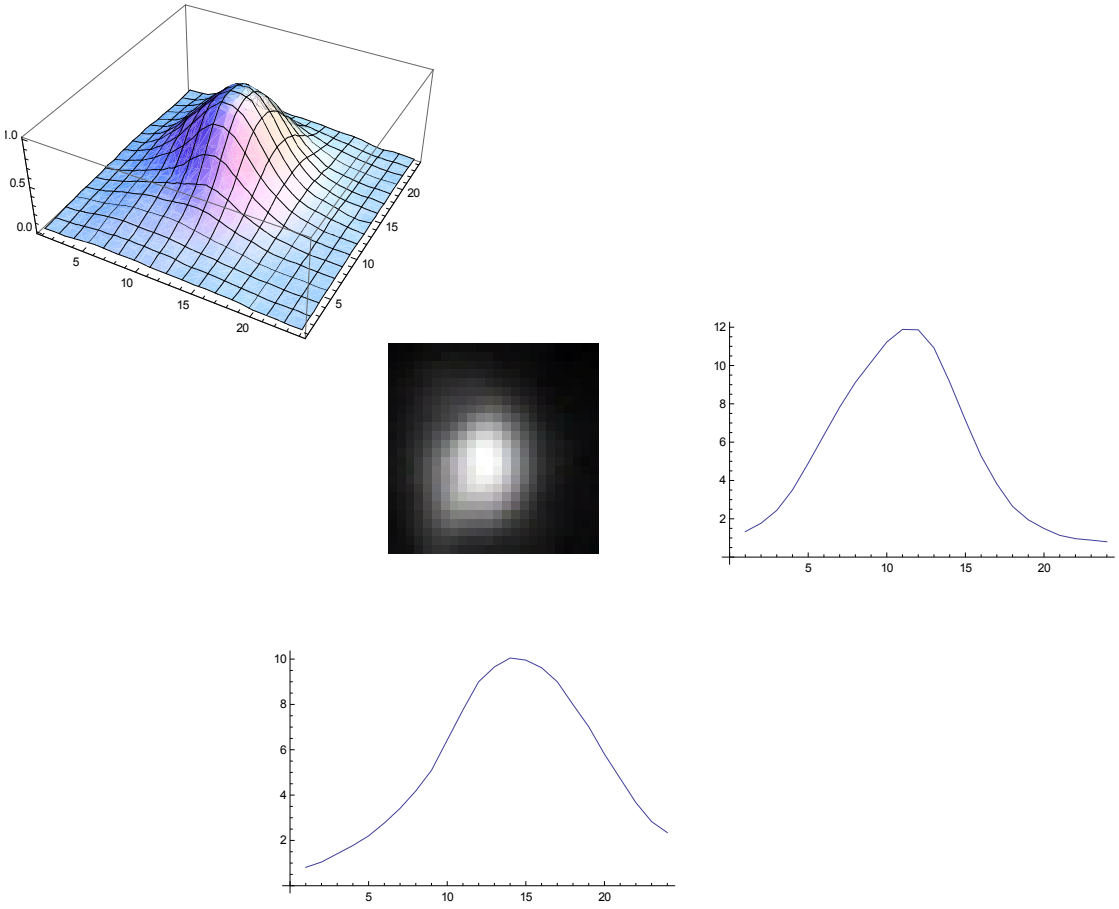

FIGURE 6.2: A CCD camera can be used to infer the shape and a symmetry of a 2-dimensional projection of the MOT cloud.

6.4 Measuring the atom number

Measurement of the number of atoms in the cloud cannot be made directly, and instead we must rely on an indirect measurement of the voltage from a photomultiplier tube which is proportional to the number of atoms. Relative measurements of the atom number, such as are necessary for our coherent manipulation schemes, are a straightforward experimental procedure requiring no changes to the magneto-optical trap setup. On the other hand, absolute calculations of the number of atoms contained within a cloud are far more uncertain, requiring knowledge of the intensity of the scattering beam, and the fraction of photons reaching, and then measured by, the detector. These calculations use a number of parameters which are difficult to measure, such as the collection efficiency of photons which have been emitted by off axis atoms, the amount of scatter from windows

and lenses and the the quantum efficiency of the detector, which added together give a large uncertainty.

A typical signal from the MOT cloud, after collection and amplification is 50 mV. After applying the conversion factors which can be found in reference [107] in conjunction with the settings on the PMT and current amplifier, we find the atom number for the cloud measured in section 6.3to be $(1.6 \pm 0.1) \times 10^7$.

6.4.1 Density of the MOT

The atomic density of the MOT can be calculated from the measurements of the atom number and size described in the previous sections. As both the MOT size, the atom number of the vary from day to day, it is important that the density is calculated from measurements of the size and atom number taken at the same time.

Using these numbers, we arrive at a value for the density of 7.1×10^{17} m^{-3}: significantly less than the $\approx 10^{25}$ m^{-3} number density of air at room temperature and pressure. This (at a typical MOT temperature of 80μK) gives a phase space density of 6.7×10^{-6}, which is significantly lower than that of ≈ 2.6 [101] required to achieve Bose-Einstein condensation.

As the intensity in the trapping beams is increased, the density profile has been predicted [56] to evolve into a top hat as the force from re-radiated photons becomes comparable to the trapping forces. We have failed to observe such a transition, and alongside other authors continue to observe a Gaussian density distribution at high powers.

6.5 Measuring the velocity distribution of the MOT

If we are to perform schemes which affect the velocity distribution of the atoms, it is important that we have a method of measuring the velocity distribution of the atoms before and after the pulse schemes. The velocity distribution can also be used to infer the temperature of the atom cloud.

The simplest method of measuring this rate of expansion is to take several snapshots of the MOT during freefall, and measure the change in the size of the atom cloud during each of these intervals. This can be achieved with a fast camera which is triggered to coincide with a flash of the trapping beams after a period of expansion. By comparing the initial and final distribution of atoms one can deduce the velocity distribution for both thermalised and non-thermalised distributions.

This method, however, is limited by the sensitivity and linearity of the camera which is likely to saturate at high densities, and become dominated by noise at low densities. We can avoid this by increasing the exposure time for each measurement, this however also extends the time taken for each measurement as well as discards any information regarding the shot to shot variation in the temperature. The flash pulse required for the atoms to fluoresce to produce each image, also causes changes to the atomic velocity distribution, which may make it difficult to resolve any small cooling effects.

The 12:8.5 ratio of the distance of the lens from the atom cloud and the lens to the imaging camera gives our imaging setup a magnification of 0.7. The CCD detector array has 480 x 640 pixels across an area of 5 mm^2 so that, in conjunction with our imaging optics, we obtain a field of view for each pixel of 7.3 μm. For Rubidium 85, an impulse of one $\hbar k$ gives the atom an acceleration which causes the velocity of the atom to increase by 10^{-2} ms^{-1}. It is therefore possible to measure a change in the velocity distribution of the atoms by $1\hbar k$ after 0.73 ms of freefall. Clearly, for this to work to this accuracy, it is necessary for us to know the velocity distributions of the atoms perfectly before manipulation, which is experimentally unfeasible. Despite this, it is interesting to see that with remarkably simple equipment, we can potentially measure such small changes to the velocity distribution. In reality, we would hope to impart many units of momentum to the atomic species before readout, and it is feasible to upgrade our CCD to one with a much higher resolution, which would make this technique more experimentally viable.

Several different orientations of the freefall atom cloud and trap and repump beams can be used to measure the temperature. As the freefall atoms traverse a readout region, illuminated by the repump beam, the change in fluorescence with time can be used to infer the temperature of the cloud. Two particularly note worthy orientations, presented in Paul Lett's excellent paper [62] are known as 'release and recapture' and 'time of flight' (TOF).

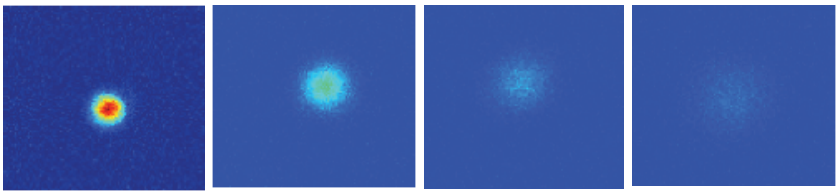

FIGURE 6.3: Four false colour shapshots [97] of the magneto-optically cooled cloud under free fall and expansion 0(left) to 7 (right) ms from the extinction of the trapping beams.

For release and recapture, difference in the cloud populations are measured before and after the the trapping beams are extinguished for a time, t. If the extinction time is shorter than the time of flight for atoms to leave the trapping region, the MOT loads from

residual laser cooled atoms, and the loading rates are significantly shorter than for longer extinction times, or hotter atoms, when the MOT is primarily loaded from background gas. Comparison of the recorded data with Monte–Carlo models of the MOT can be used to measure the temperature to be between 30-100μK for a well aligned MOT, and significantly above this otherwise (see figure 6.4).

For the time of flight scheme, the readout beam is stretched into a pancake shape and displaced below the MOT. When the trapping beams are turned off, the atoms fall, due to gravity, through the readout beam. Again, comparison of the fluorescence with Monte-Carlo simulations with different temperature atoms can be used to infer a temperature of the atom cloud.

There are errors in measuring the velocity distribution using these methods: during release and recapture method, the system is symmetrical around the trap centre, and the temperature recorded is the average temperature along all axes. In contrast, the time of flight measurement is symmetrical around the vertical axis, such that only the velocity distribution along this axis can be inferred. This unidirectional measurement can be problematic for more exotic cloud shapes, for which one must rely on techniques which are more sensitive to asymmetry [108].

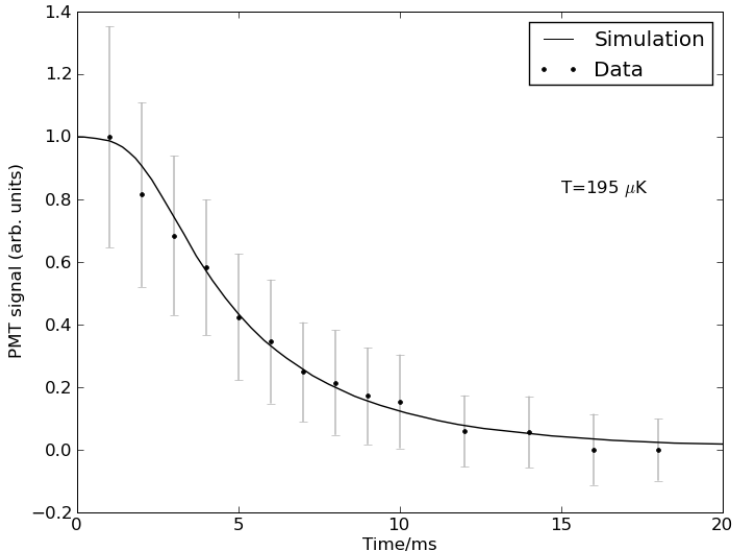

FIGURE 6.4: The signal from release and recapture, showing the decay in atom number with time, as atoms move out of the readout volume. The data is fitted to the closest of a set of Monte-Carlo simulations with different initial cloud temperatures [107] which gives a higher than average temperature measurement of 195 μ K.

Both the release and recapture and the time of flight methods rely on the fitting of data to Monte-Carlo simulations of the fluorescence; of the three methods suggested,

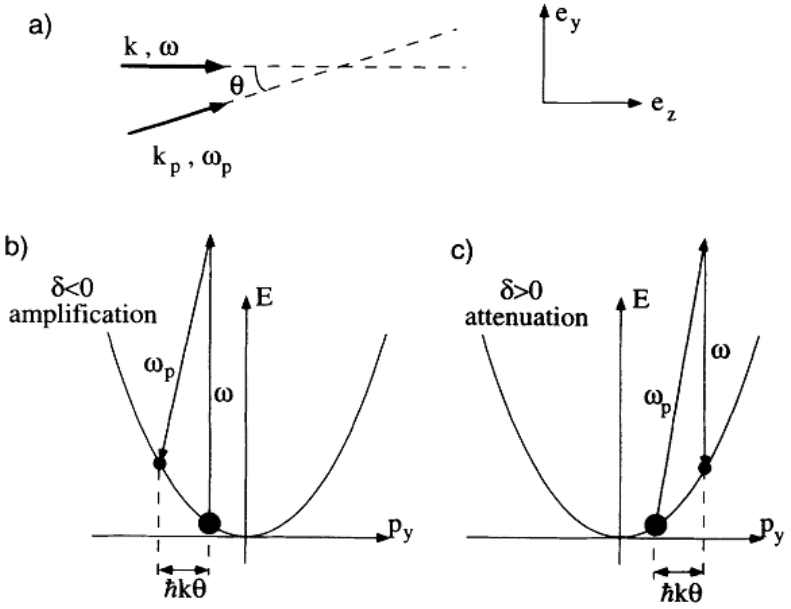

FIGURE 6.5: Changing the angular separation, θ between pump and Stokes beams changes the momentum imparted to the atom during absorption For particular values for the angular separation and frequency, the beams couple two separate points on the dispersion curve; measuring the attenuation of the beam for a given angle tells us the difference in the population at these two points. [109]

therefore, the method of imaging the cloud using CCD cameras is alone in being able to infer information without a priori assumption of the form of the velocity distribution. As the atomic momentum distribution during coherent cooling can be significantly different from that of thermal cooling, we should not assume that the outcome will take the form of a Gaussian function; this makes it difficult to use either the release and recapture or the time of flight methods to decide on the success of the coherent cooling schemes.

A final method, known as recoil induced resonance spectroscopy [109], can be used to overcome the limitations of the techniques mentioned earlier. By varying the angular separation between two Raman beams, we change the momentum imparted during absorption from the difference to the sum of the pump and Stokes momenta. The transition energy, however, which is a scalar quantity, is invariant with changes in the angle. By varying the angular separation of the two beams, we map the available density of states of the atoms within the trap allowing the population density for a particular velocity class to be inferred from the absorption of the beams at a particular angular separation. Due to the improvement in measuring the velocity distribution of non-spherical distributions, our group plan to implement this method of temperature measurement. Before we can achieve this, however, we must first develop the capability of applying the pump and Stoke's beams separately, either through separate amplification, or splitting by Mach-Zehnder interferometry (see section 7.7).

Until we start manipulating the momentum state of the cloud, a combination of the CCD and the release and recapture methods has proven to be sufficient: while the MOT is properly aligned, the shape of the cloud will be, to a good approximation, spherical which gives consistent temperature measurements. Unlike the time of flight and the recoil induced resonance methods, the release and recapture does not require any additional beams or changes to the optical setup, and so is preferred due for it's simple implementation.

6.6 Calculating the mean free path of atoms in the MOT

Using measured values for the velocity, which can be used to calculate the volume each atom draws out in a given time interval, and the the density, it is possible to calculate the mean free path:

$$\lambda = \frac{1}{2\pi r^2 n}, \tag{6.5}$$

where n is the number density of for particles in the cloud (before expansion, and so therefore representing an upper limit on the mean free path) and r represents the atomic collision radius given by the range of van der Waals forces. For an atom cloud containing 10^6 atoms in a Gaussian distribution of width 350 μm, we find the maximum mean free path, for atoms at the centre of the cloud, to be 5.3 m, which for a large enough cloud of Rubidium atoms at 80 μK, gives a mean free time of 0.66 s. From this calculation, it is clear that for the 1 ms duration of our experiments, it is acceptable to neglect collisions.

Chapter 7

Cold atom manipulation

To demonstrate coherent manipulation, we drive a two-photon Raman transition between the two $5S_{1/2}$ hyperfine states of rubidium 85. To achieve population inversion it is necessary for the population to be driven resonantly which requires a light field consisting of two frequencies which are separated in frequency by the hyperfine splitting. It is crucial for the coherence of the manipulation that these two beams maintain a fixed phase relationship such that the beat note is without discontinuities. To drive the atoms at a constant Rabi frequency, the two beams must also possess a stable intensity and detuning over the whole of the atomic cloud for the duration of the experiment. To implement the adiabatic rapid passage method of population inversion, described in section 4.4, it is necessary for us to have control over the frequency detuning of the beam so that we can create chirped pulses. Our manipulation setup is designed to meet these criteria; the Raman beams are derived by electro and acousto-optical modulation (EOM and AOM) of a free running Sacher Tiger 1W laser system. The RF signal supplied to the electric-optical modulators is also used as a reference for the RF supply to the acousto-optical modulator, which maintains the stability and phase coherence of the Raman lasers. By controlling the phase and waveform of the driving RF signal, we can change the modulation parameters to achieve sub-microsecond control over the optical waveform.

Intensity modulation of the Raman beam is achieved, as with the MOT beams, using AOMs, which allow high extinction of the beams with low timing jitter. This fast and reliable switching allows us to run the experiment continuously, and therefore average the results over many experimental runs. To achieve the powers necessary for two-photon manipulation, we amplify the beam via an injection locked diode and a commercial tapered amplifier before manipulation.

Low-noise readout of the state after coherent manipulation is performed using a large aperture lens to project an image of the cloud onto a photomultiplier tube which is amplified by a low noise current amplifier, before averaging and readout by a fast oscilloscope

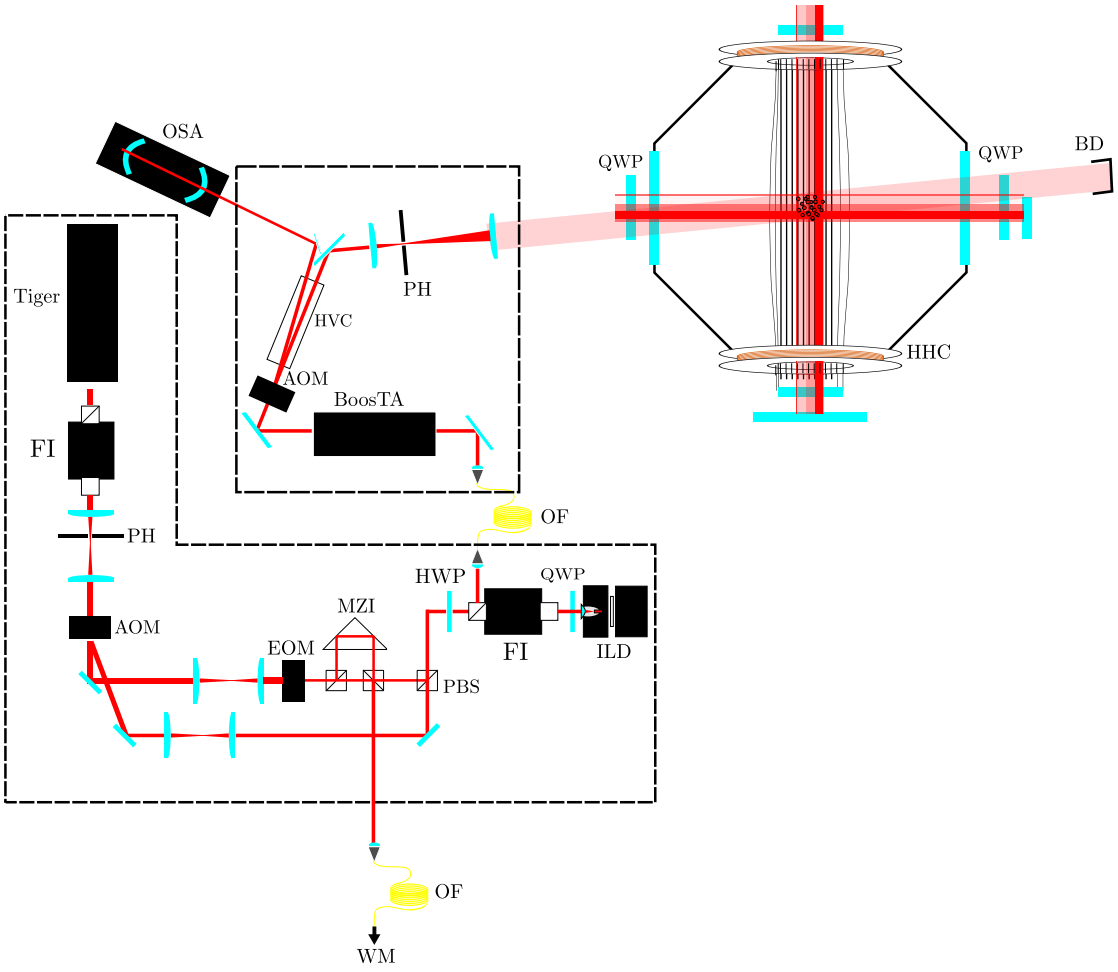

FIGURE 7.1: The manipulation laser setup. FI: Faraday isolator, PH: pin hole, AOM: acousto-optic modulator, EOM: electric-optic modulator, MZI: Mach–Zehnder interferometer, PBS: polarising beam splitter, ILD: injection locked diode, HVC: heated vapour cell QWP: quarter wave plate, OF: optical fibre, HWP: half wave plate, WM: wavelength meter, OSA: optical spectrum analyser, BD: beam dump, HHC: Helmholtz coils.

7.1 Methods of hyperfine manipulation

There are two common techniques of hyperfine manipulation of rubidium: via a magnetic interaction with the fields of a microwave cavity, and through a two-photon interaction.

Magnetic interactions have low coupling rates, typically 10^5 times smaller than that of electronic interactions. For this reason it is necessary to use high powers, inside optical

cavities for example, which the atomic beam traverses [110]. The temporally varying field seen by the atom as it traverses the cavity is tailored to apply π or $\pi/2$ pulses. This cavity method is constrained by the geometry of the cavity, which makes it impractical for atom manipulation. These high finesse cavities are also slow to respond to changes in the incident electric field strength, which makes them unsuitable for applying fast pulse schemes.

Alternatively, the hyperfine states can be manipulated using a two-photon interaction, resonantly enhanced by proximity with the single photon D2 transition, known as 'resonant Raman spectroscopy' [111]. The simplest method of constructing beams of light for Raman manipulation is by phase locking two diode lasers by using a fast photodetector [112], although this can make it difficult to perform frequency chirps necessary for adiabatic rapid passage. Electro-optic modulation of a carrier beam as described by Lee *et. al* [113] also gives a two frequency light field with additional control over the phase and waveform. As summarised by Lee, this technique may encounter problems with the number of unused frequencies, as we will discuss.

7.2 Electro and acousto optic modulation for the generation of Raman beam

An electric-optic modulator (EOM) consists of a non-linear crystal, such that it exhibits a change in refractive index with electric field, inside a RF or microwave cavity. Light crossing the cavity, aligned with the optical axis of the crystal, will experience phase modulation as a result of the microwave field, which adds sidebands to the carrier frequency at plus and minus the modulation frequency at an amplitude given by the Bessel function of the modulation depth[76].

Acousto-optical modulation causes a spatial separation between the unmodulated and modulated (+310MHz) beam, so that it is possible to electrically modulate the undiffracted beam at 2.7GHz separately from the diffracted beam. This gives a sum frequency difference between the AOM modulated beam, and the first order EOM sideband of 3GHz, equal to the hyperfine splitting of Rubidium. Due to the higher frequency of the 310MHz AOM, in comparison to the lower frequency 110MHz and 80MHz modulators used in the MOT setup, the diffraction efficiency is lower, at around 65%. In addition to this, the apeture window is smaller, emphasising the need for good allignment.

We use an AOM, driving a transition between the AOM-EOM beams, to break the symmetry of the two sidebands around the carrier: as the two sidebands are symmetrical around the central carrier, setting the carrier-EOM beams to resonance would cause a

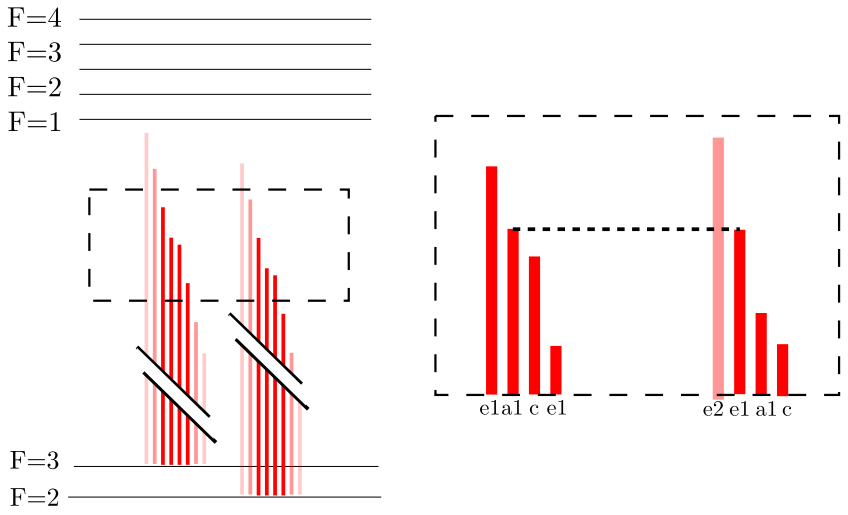

FIGURE 7.2: The detuning of each of the frequency components, carrier (c), AOM (a1) and EOM 1st and 2nd and 3rd order sidebands (e1, e2, e3), from resonance with the ground and excited states states, highlighting the two-photon resonance between the EOM first order (e1) sideband with the AOM first order diffracted beam (a1).

resonance between both the plus and minus sidebands simultaneously. The detuning of these carrier-EOM resonances from the single photon transition, however, would differ by 3 GHz, affecting the Rabi frequency of each oscillation. Each atom would therefore be driven by the sum of the two oscillations, making it difficult to achieve population inversion. More importantly, as the plus and minus sidebands are 180° out of phase with respect to each other [76], the effects of each of these oscillations would cancel.

7.3 Setting the power of the Raman beams

If the sidebands are to be used for Raman transitions, as in our case, it is important that the power in the modulated sidebands is maximised while the power contained in higher order sidebands, which are not used for the Raman transition, is minimisd.

To choose the modulation depth which satisfies these criteria, we refer to figure 7.3. We see that at a modulation depth of 1.8 radians the first order sidebands can be created with an amplitude three times that of the carrier after modulation or 30% of the initial carrier power. To generate 4 W (36 dBm) of RF power necessary to achieve this $1.8 - 2.4$ radian modulation depth, we we use a Miteq low phase noise $2 - 4$ GHz RF amplifier to amplify the output of the Agilent signal generator from 63 mW (18 dBm). After applying the required power to the EOMs, however, the conversion efficiencies are still far less than we expect: the maximum output power which has been observed in the first order sidebands has been limited to 20 % of the height of the carrier after modulation. The reasons for this poor conversion efficiency are unclear; the RF power

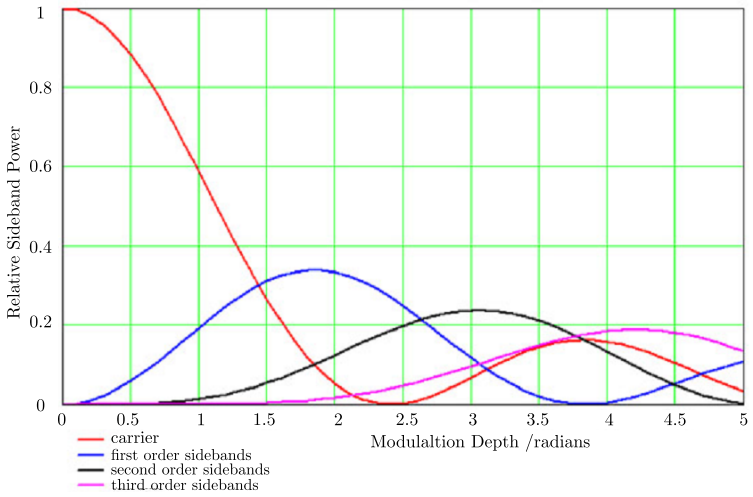

FIGURE 7.3: Relative strengths of the first to third order sidebands, and the original carrier beam, as a function of the modulation depth [114]. It can be seen from this that the power in the first order sideband is maximsed for a modulation depth of 1.8 radians, but at this depth, large amounts of the carrier power and second order sideband power are present. Conversely, at 2.4 radians, all of the carrier power is converted from the carrier.

supplied to the EOM has been checked, and is sufficient to give the modulation depth necessary for maximum conversion; the polarisation, beam quality and divergence have been checked to be sufficient for maximum conversion. The remaining option is that the crystal has undergone thermally induced degradation [115] where past heating of the crystal due to the optical or RF field causes local changes in the crystal structure which limit the electrical response. Such crystal defects can be corrected by annealing the crystal although the time scales involved in this process are long and at present would unnecessarily disrupt the experiment.

7.3.1 Frequency generation and stability

The EOM RF supply is supplied by an Agilent arbitrary frequency generator which can generate extremely low phase noise frequencies between 9 kHz and 4 GHz with an accuracy of 3 ppm and a frequency drift of up to ± 2 ppm/year or 6kHz/year. The Agilent is also used to supply a 10MHz reference to the Marconi frequency generator used to drive the Raman AOM, which makes it particularly important that we are certain that it is calibrated. For this, we use a high speed counter in conjunction with a very precise 1 Hz pulse received from a GPS system. The counter is reset and read-out after each timing pulse from the GPS which corresponds to a frequency. We consistently measure 10^7 pulses per second from the 10MHz reference output to the waveform generator over time scales of one second to many hours, with no measurable deviation [107].

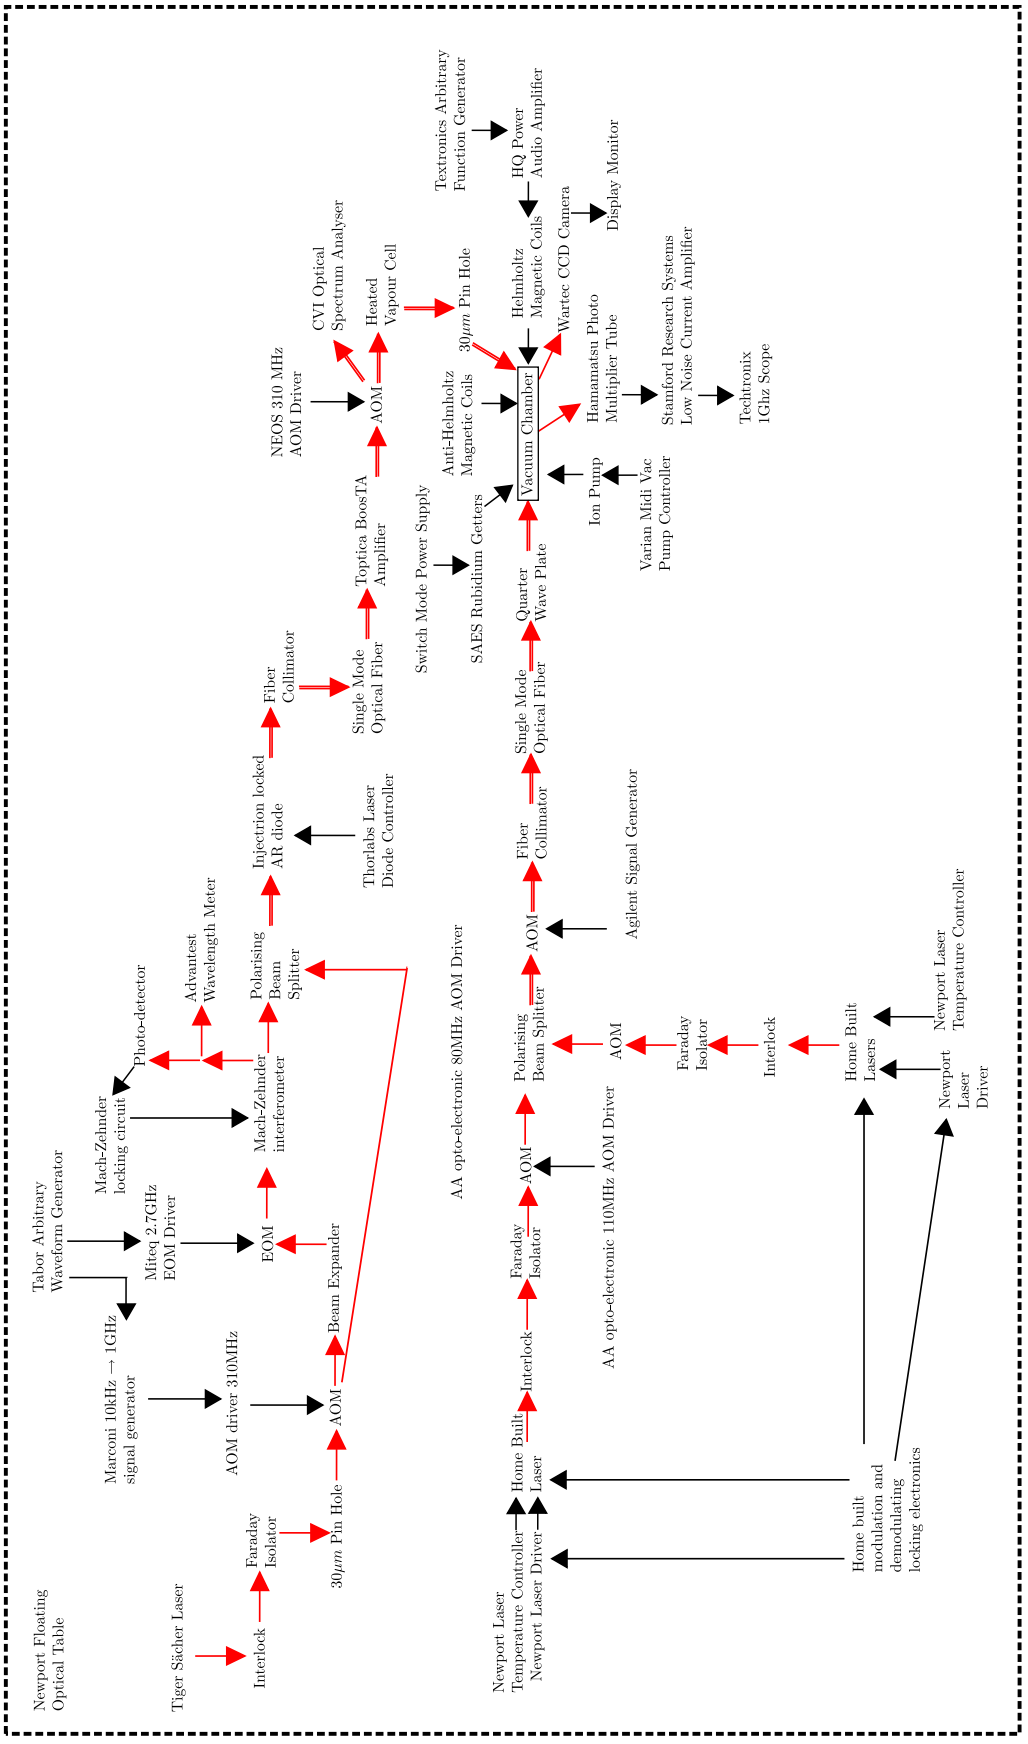

FIGURE 7.4: A schematic of the setup in the lab including the drivers and electronics

7.3.2 I & Q modulation for phase shifting the Raman beams

To produce the frequency sweeps necessary for adiabatic rapid passage and spin echo experiments, we use an I & Q modulator, which adds an arbitrary phase offset to the frequency supplied to the EOM. This modulator takes the input waveform, splits it and adds a 90 ° phase shift to one path. Each path is then mixed with the signal supplied to the Q(for the 'quadrature' path) and I (for the 'in phase' path) inputs, before recombination. By choosing suitable choices for the inputs, arbitrary changes to the input waveform can be made; a phase shift of ϕ can be achieved by applying $\sin(\phi)$ to the Q input and $\cos(\phi)$ to the I input.

With no voltage applied to the I and Q ports of the modulator, no output is obtained. However, by supplying the modulator with the output from an arbitrary signal generator, a 5V pulse can be sent to the I & Q inputs which causes a 180 degree change in phase of the EOM RF phase.

7.4 The Sacher Tiger laser

The carrier light is supplied from a commercial Sacher Tiger laser system which uses a tapered gain medium to generate up to one Watt of optical power at 780 nm. The device has transparent front and rear facets with the external cavity grating positioned in a Littrow configuration (see section 5.3) behind the diode, causing the majority of the light to be amplified in a single forward pass through the diode. To increase the maximum amplification of the chip, the diode is tapered to maintain a constant power per unit volume throughout the length of the chip. In theory, this geometry also gives a single mode output beam found in non tapered diodes due to the narrow input facet of the chip, which spatially filters the input mode before amplification. The grating orientation, on the opposite side of the diode from the output, gives a stable output which, unlike our home-made laser systems, does not depend on the angle of the grating. The laser controller, also purchased from Sacher allows the manual and external control of the current supplied to the diode, as well as voltage feedback to the grating and temperature.

7.4.1 Problems with the Tiger

The design of the Tiger makes the device passively stable, and easily tunable over large wavelengths. Despite this, however, it is let down by the manufacturing of the chip, which does not meet the power or beam quality given by the specification sheet. The

output power whilst specified for 1 Watt, gives a maximum of 350mW of which 100mW is contained within a good, TEM_{00} spatial mode. The single pass of the laser also makes it sensitive to feedback, and it is necessary to supplement the 30 dB of optical isolation built into Tiger with an additional, external isolator to give a total of 60 dB of isolation.

7.4.2 Tiger Stabilisation

Due to the modulation method we use to generate the Raman beams, the Raman Rabi frequency is determined primarily by the difference between the EOM and AOM frequencies as opposed to their absolute frequencies. This should ensure that, so long as the Tiger is detuned sufficiently far enough away from single photon resonance, such that the ratio of the variation in the frequency to the detuning is small, we can neglect variations in the Rabi frequency. Large changes in the frequency of the Tiger remain problematic, however, due to changes in the transmission of light through the Mach–Zender interferometer (section 7.7), which is heavily dependent on frequency of the laser.

The large detuning of the Tiger from single photon resonance and the absence of spectral absorption features at large detunings makes locking of the Tiger laser more complex than for the MOT lasers. As a result of this, we use the output of our Avantest wavelength meter to measure the Tiger wavelength to an accuracy of 50MHz; this is read by a computer, which then corrects the Tiger grating angle. This has been found to successful in correcting for the slow drifts in the output frequency. More complicated methods of locking the laser have also been considered, such as to an external cavity. Methods such as this require a large commitment of time which, given the small effect on the Rabi oscillations by our wavelength-meter locked laser, for the moment is considered unnecessary.

7.5 The optical spectrum analyser, beam spectra and powers

To measure the spectrum of the light after modulation, we use an optical spectrum analyser (OSA) inside of which is a Fabry-Pérot etalon to give a strongly frequency and length dependent transmission. By sweeping the cavity length, with a piezo-electric actuator, the transmission maps out the frequency spectrum, which is measured by the photodetector after the cavity. The free-spectral range is an important characteristic of each OSA as it defines the difference in frequency between cavity modes. Our OSA has a free spectral range of 2 GHz, which is smaller than the separation of the sidebands

causing the spectrum of the Raman beam to be folded back on itself making it difficult to identify specific frequencies on the spectrum or to use the OSA to measure the spacings between the different frequency components.

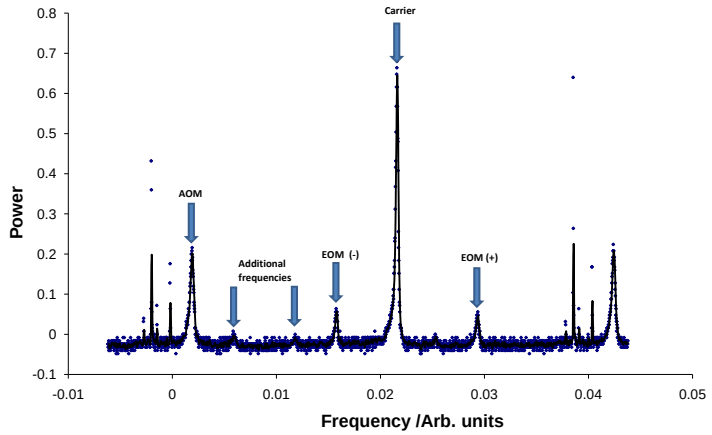

FIGURE 7.5: The spectrum, as measured by the OSA, of the Raman beams after amplification. Notice the central carrier frequency, the plus and minus first order EOM sidebands symmetrical around the central carrier, and the AOM beam to the left of those. Notice also, the presence of additional frequency components, especially noticeable between the minus sideband, and the AOM beam. The (unlabeled) peaks on the right of the plot are not significant and are formed from the OSA scanning back over the same spectral range.

By comparing the heights between the different frequency components, and selectively extinguishing the beams, identification of the main frequency components is possible, despite the small free spectral range of the OSA. The linearity of the OSA photo-detector can then be used to measure and set the relative powers of the Raman beams to be equal. It can also be used to check for any slow temporal variations in the intensities, and to study the spectrum of the Raman beams, which changes during modulation and amplification (section 7.9.3.2).

During an experimental run, the OSA trace can be recorded alongside the data allowing the spectrum to be analysed using software that measures the values of the properties described above. This is possible by marking the positions of the four fundamental frequency components; the separations of the peaks on the spectrum, along with the fixed-parameters which give the separations of the components in frequency space can then be used to find the value of the free spectral range, and identify the positions of the four laser frequency components on the OSA spectrum, even when the spacing of the frequencies is greater than the free-spectral range. This can be used to automatically measure the relative powers in the fundamental frequency components, or if the span

of frequencies is much greater than the free spectral range, measure the changes in the free spectral range.

A Newport power meter is used to measure the total power in the beam. This is then divided according to the relative magnitudes of the individual frequency components measured by the OSA. There are in inaccuracies with this method, however, due to the AC coupling of the OSA which makes it is difficult to judge the power contained in light with a linewidth much greater than the free-spectral range, such as that of amplified spontaneous emission (ASE). The power contained within the many second and higher order sidebands and their harmonic derivatives is also omitted from this calculation. These frequency components, although less powerful than the four fundamental frequency components, may be significantly more numerous, and therefore may contain a significant fraction of the total beam power. This detail may be important when it comes to explaining the discrepancy which arises between the measured and calculated Rabi frequency in section 9.3.2.

As the tapered amplifier can be highly non-linear, adding additional frequency components and changing the relative heights of existing frequencies, it is important that the analysis of the spectrum and calculation of the power in the beams is performed after amplification.

7.6 Measuring the coherence of the beam

Although the electric field of each frequency component of the Raman light oscillates at a frequency which is not realistically measurable by laboratory equipment, the frequencies contained within the Raman beam beat at a much lower frequency, which is more directly measurable by fast photo-detectors. The frequency components necessary for Raman manipulation, for example, beat at 3GHz, which is within the 30GHz bandwidth of ultrafast GaAs photo-detectors from Hamamatsu. The output from the photodetector, after amplification and band pass filtering, is mixed with the 2.7 GHz RF signal supplied to the EOM. The 310MHz output is sent to an oscilloscope, where we are able to take direct measurements of the relative change in amplitude, phase and frequency of the components of the Raman beam, which we can use to infer a coherence time for the light source. Analysis of the waveform shows good phase coherence, lasting many μs, but a slow drift in the intensity of the beams which we have been able to correlate to uneven air flow from the air conditioning unit and eliminate with draught exclusion.

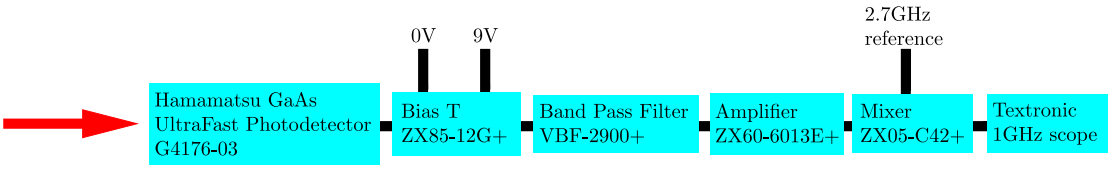

FIGURE 7.6: The setup for our fast photo-detector: the bias T biases the photo-detector, the signal is then sent through a 2.9-3GHz band-pass filter, amplified and then mixed with a 2.7GHz reference signal. The resulting beam is analysed on a Tecktronic 1GHz scope.

7.7 The Mach–Zehnder interferometer

Phase modulation adds additional frequencies to the injected carrier frequency. Of the four largest contributors to the beam power, only two frequencies are necessary to drive the two-photon Raman transitions while the remaining two are superfluous. While these frequency components are not resonant with the single or two-photon resonances, and therefore do not affect the coherence of the system, they give a number of unwanted side effects. Firstly, the power in these beams causes a light shift which shifts the transition frequency. These shifts can make it difficult to find the transition frequency as narrow spectral scans often miss the frequencies of interest, while broader scans lack the point density necessary to resolve the resonance peaks. Secondly, as the Raman beam is amplified by a injection-locked amplifier operating near saturation, the unused frequency component causes unnecessary depletion in the gain medium, which reduces the overall gain seen by the two Raman frequencies In addition to this, non-linear effects within the gain medium have been shown [116] to cause frequency mixing and slow oscillations in the output intensity.

To avoid these effects, we use a Mach–Zehnder interferometer in a similar configuration to that used by Dotsenko *et al.* [117], to cancel the unused EOM sideband and reduce the height of the unused carrier components present after modulation.

Figure 7.7 shows the construction of our Mach–Zehnder interferometer: the input beam is split into two paths by a non-polarising beam splitter. One of these paths is then reflected through 180° by a corner cube, so that it can be recombined with the other by a second non-polarising beam splitter. The outcome of recombination is determined by the path length difference between the two beams, which, for the beam transmitted by the first prism of the Mach–Zehnder is:

$$\phi = 2a + 2nd + nc - b \tag{7.1}$$

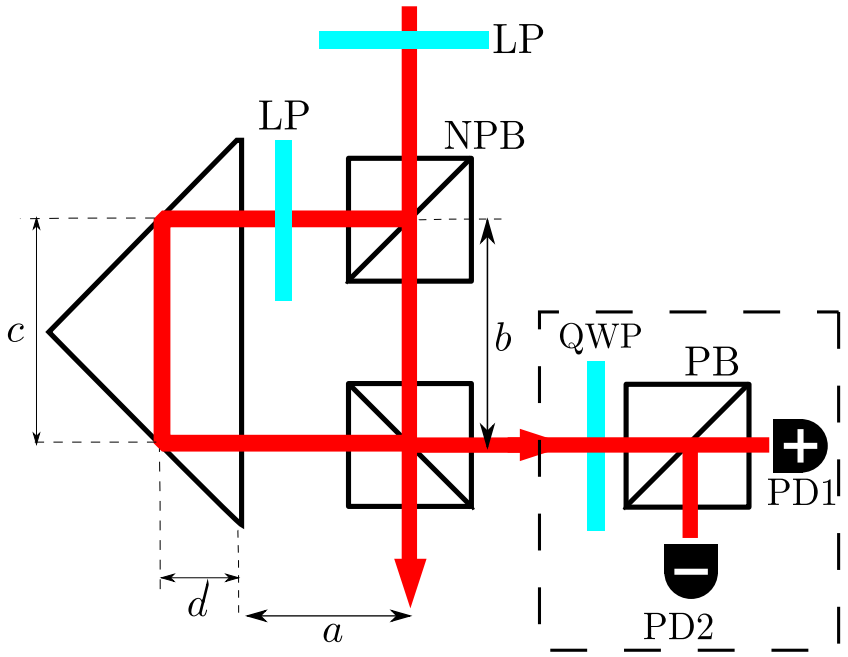

FIGURE 7.7: A schematic diagram of the Mach–Zehnder interferometer, showing the distances, a, b, c and d necessary for path length calculations. The dotted box in the bottom right corner shows the setup for the polarisation locking, where the plus and minus represent the positive and negative inputs of a balanced photodetector. LP: linear polariser, NPB: non-polarising beamsplitter, QWP: quarter wave plate PB: polarising beamsplitter.

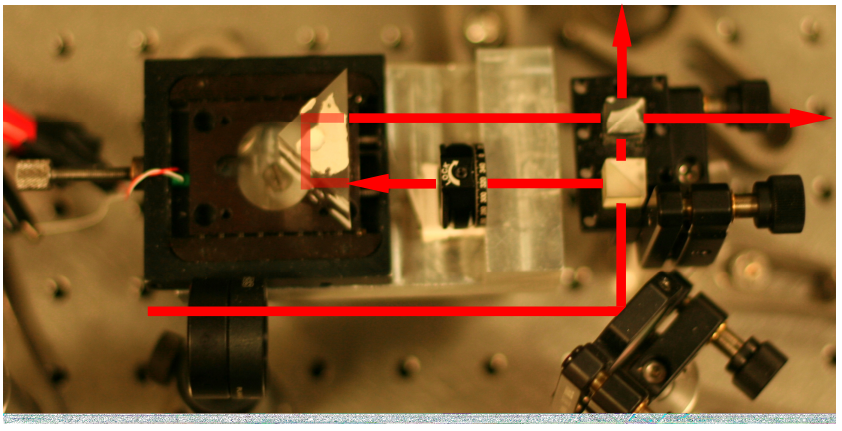

FIGURE 7.8: A photo of our Mach–Zehnder interferometer showing the two non-polarising beamsplitters, corner cube reflector, translation stage, and the piezo-electric actuator used for stabilisation

and for the reflected beam:

$$\phi = 2a + 2nd + nc - b + \frac{\lambda}{2} \tag{7.2}$$

where the values of a, b, c and d are shown in figure 7.7, λ is the wavelength and n the refractive index of the corner cube. The $\lambda/2$ difference between the beams exiting from the two output ports means that while the beam exiting one output interferes constructively, the other will interfere deconstructively. It is an interesting point, not often mentioned, that it is only after the rotational asymmetry of each beam-splitter, a result of the reflection boundary with non-zero thickness [118], is included in the description of the Mach-Zehnder that the paths possess this necessary phase difference.

Equations 7.1 and 7.2 show that the output of the beams at each port is determined by the phase difference gained by the two beams at recombination, which gives control of the output intensity varying the distance a, between the beam splitters and the corner cube. As the optical path length is also dependent on the wavelength, it is possible to apply the correct phase delay such that a beam containing two wavelengths can be separated into the two wavelengths.

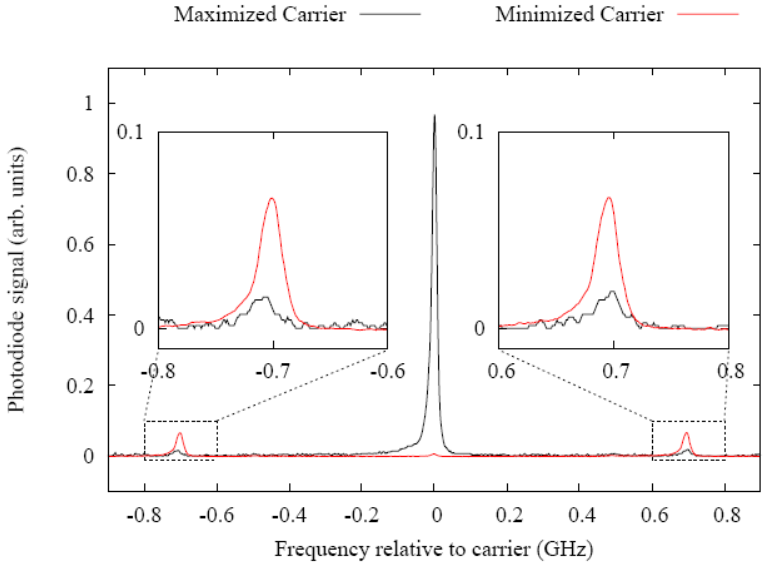

FIGURE 7.9: A plot of the OSA spectrum of the electrically modulated beam after transmission through the Mach-Zehnder interferometer. The separation of the carrier frequency (centre) and the EOM sidebands can be seen, while the Mach Zehnder is set to transmit the carrier (black) or the sidebands (red) [119].

The corner cube is mounted on a translation stage with both manual and piezo adjustment which allows fine or coarse changes to the relative path difference. To avoid further adjustments having to be made to the Mach–Zehnder after it has been made, the components are glued onto the translation stage, rather than mounted. While some of the parameters necessary for proper alignment of the Mach–Zehnder can be controlled

externally, by changing the position and angle of the incoming light after the glue has set, others, such as any change in the angle of the beam after the first beamsplitter, cannot. It is therefore essential that the Mach–Zehnder is aligned with good fringe contrast, before the glue is set. To make this easier, we use a low expansion UV curing glue which allows the glue to be applied and the Mach–Zehnder aligned before the glue is set.

Alignment of the Mach–Zehnder is achieved with phase modulation turned off, so that the spatial alignment can be determined through the output power. With only a single frequency, the fringes are visible when the beams are nearly overlapping, which can be used for coarse adjustment; fine adjustment can be achieved by scanning the piezo voltage and optimising the fringe contrast using a photo-detector. If the Mach-Zehnder is to be used for mixing or splitting beams of different frequencies, the necessary path length can be calculated to give a rough indication of the distance required, and then fine tuned using an optical spectrum analyser to see the change in output power of different spectral components.

If the direction of the light incident on the corner cube is exactly parallel to the axis along which the corner cube is translated, changes to the path length become decoupled from the spatial alignment. While this decoupling may be achieved for small translations, large changes will inevitably cause mis-alignment. It is important, therefore that the path length is set to roughly the right value before gluing the prisms. As any angular change in the beams inside the Mach–Zehnder cannot be compensated for externally, it is also important for alignment that the two beam-splitters are parallel to each other and parallel to the corner cube. All three must have a common angle to the horizontal.

7.7.1 Methods of Mach–Zehnder locking

Although the Mach–Zehnder is passively stable on time scales less than a second, variations in air currents and external vibrations cause a slow drift in the path length which must be actively compensated for. The stability of the Mach–Zehnder is important, as any variation in the output causes changes in the amplitudes of the first order EOM sidebands, changing the relative powers of the Raman beams and hence large changes in the Rabi frequency.

7.7.1.1 Optical spectrum analyser

The output of the optical spectrum analyser (OSA) can be used to provide a good visual indicator of the success of the Mach–Zehnder locking. However, locking the

Mach–Zehnder using this technique alone would require stabilisation of the OSA cavity as well as the Tiger frequency. Using software to identify the heights of the peaks as described in section 7.5 is too slow for locking.

7.7.1.2 Voltage Modulation

Similar to the locking of the lasers via the Pound-Drever-Hall method, the voltage on the piezo can be modulated around the lock point. This causes a modulation in the output power which, once demodulated, gives a locking signal that can be fed via servo controller, back to the piezo. In modulating the Mach–Zehnder, however, the powers in the Raman components are also modulated, which creates an unwanted time dependence in the Rabi frequency.

7.7.1.3 Computer control

Due to the passive stability ($\approx$ 1s) of the Mach–Zehnder, it is possible to stabilise the output of the Mach-Zehnder using a data acquisition card (DAQ), which can read the power of the output, and perform an algorithm to minimise it. Although similar to the voltage modulation method mentioned above, the computer can perform the modulation and locking in the time between experiments, while data is downloaded from the oscilloscope. Once the algorithm has found the minimum lock point, the modulation stops and the experiment proceeds. As locking does not occur during the running of the experiments, however, the output of the Mach–Zehnder is particularly vulnerable to air currents and small changes in the frequency of the carrier laser at this time. The 8-bit digitisation of the DAQ means that the modulation applied by the DAQ is greater than that for voltage modulation; this digitisation also limits the accuracy with which the lock point can be both found and set. Finally, the nonlinear response of the piezo with voltage causes unwanted effects such as creep and hysteresis after the modulation has been turned off, which causes the interferometer to drift during data acquisition.

7.7.1.4 Polarisation (Hänsch-Couillaud) locking

The birefringence of the corner cube in the Mach–Zehnder causes a change in the polarisation axis of the light if initially offset from the axis of birefringence of the beam splitter. As a result of this, the light reflected from the first non-polarising beamsplitter (which passes through the corner cube) will reach the second, recombining beam-splitter with a polarisation direction rotated with respect to that of the non-reflected light. This

change in the polarisation, as originally applied by Hänsch *et al.* [120] to cavities can be used for locking.

If the Mach–Zender has the correct path length necessary for constructive interference, the reflected and the non reflected light have the same phase at the second recombining beamsplitter, and will add together to give linearly polarised light. Alternatively, if there is some difference in the path length for the two paths which is not an integer number of wavelengths long, the reflected and non-reflected beam paths recombine at the second beamsplitter out of phase. Due to the rotation of the reflected light relative to the non-reflected, the phase difference between the reflected and non-reflected paths will cause a slight elliptical polarisation. This change in polarisation gives a direct indication of the Mach–Zehnder offset from the lockpoint, which can be measured with a balanced photodetector (plus and minus photodetectors on figure 7.7). The sign of the output depends upon the direction of the Mach–Zehnder offset, which gives an excellent locking signal without modulation. Further description of the locking scheme, which also describes how the locking can be made to be invariant to the addition of carrier sidebands, can be found here [119].

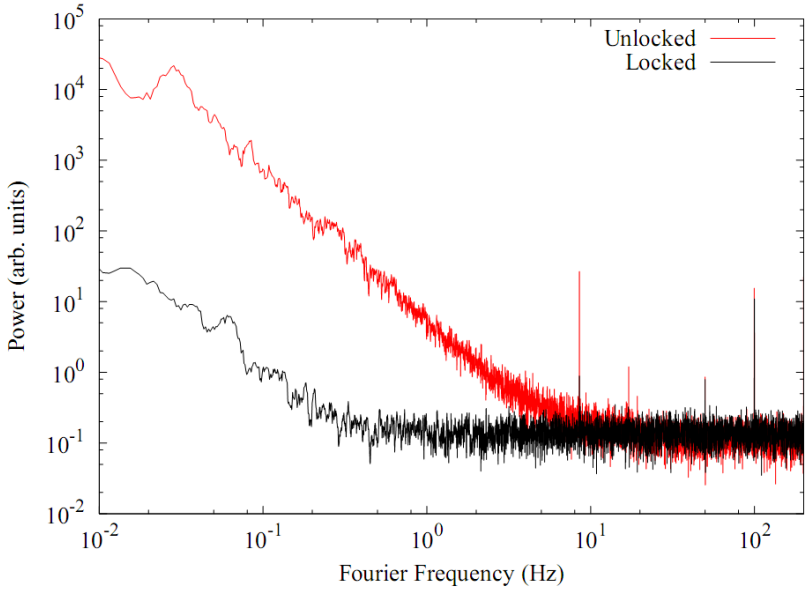

FIGURE 7.10: A Fourier transform of the Mach–Zehnder output intensity both with (black line) and without (red line) Hänsch-Couillaud locking. Locking significantly suppresses the noise for 5 Hz or less, although frequencies less than this, there is some electronic noise. The graph also shows spikes at 10, 60 and 310 Hz from our air conditioning unit, mains pickup, and pickup from our AOM amplifier respectively, which are also reduced by locking.

7.7.2 Mach–Zehnder for beam mixing

By setting the path length of the Mach–Zehnder so that it coincides with the constructive interference of the AOM beam, while still giving the same deconstructive interference for the EOM beams as described above, the Mach–Zehnder can be used to mix the two beams together, while avoiding the losses present when using beam splitters. However, as this coincidence only occurs at pathlengths of 310MHz/c, the dimension a needs to be of order ≈ 30 cm, at which length stability is much harder to achieve than for shorter path lengths.

7.8 Equipment necessary for future adaptations to the experiment

To implement the coherent cooling schemes, it will be necessary to be able to apply π and $\frac{\pi}{2}$ pulses to the atomic cloud from two opposing directions. This could be achieved using two identical sets of manipulation equipment at additional effort and experimental complexity. Alternatively, a Pockels cell, with an electrically controlled refractive index, similar in principle to that of the electro-optic modulators, allows fast switching of the polarisation of beam can be used in conjunction with a polarising beam splitter to switch the beam direction. For large changes in the polarisation, voltages on the order of 200V are necessary, and the design of high speed switching drive electronics is non-trivial.

7.9 Amplification and ASE

The poor output power and beam quality of the Tiger laser and the losses during modulation and Mach–Zehnder filtering significantly decrease the optical power of the Raman beams so that it is common to achieve only a single milliWatt of output power in the Raman beams. Amplification is necessary therefore, if we are to achieve a Rabi frequency large enough for cooling.

Amplification is performed after the acousto and electro-optically modulated beams have been recombined by a polarising beam splitter cube to create a combined 'Raman' beam. This is then used to seed an injection locked, anti-reflection coated diode; secondary amplification is supplied with a Toptica BoosTA tapered amplifier. For both of these amplification mechanisms, the gain bandwidth is much greater than the modulation frequency which allows the frequency components to be amplified simultaneously.

7.9.1 The injection locked diode

An injection-locked diode is used to boost the 1-2 mW of power after modulation and filtering to 30 mW, which is then sufficient to saturate the gain of the BoosTA. For this, we use an anti-reflection coated diode from Thorlabs which we seed with the modulated beam from the Tiger laser. Using this technique, the phase correlation of the master and the slave lasers has been shown [121] to be negligible over many seconds. We separate the ingoing and outgoing amplified beams using a polarising beam splitting cube, which, due to a quarter wave plate before the diode are of orthogonal polarisations. An anti-reflection diode is used to give amplification without interference from cavity modes within the diode. Despite the anti-reflection coating, however, a small amount of reflection still occurs, causing some additional mode structure to be present. As with the standard ECDLs, the current supplied to the diode can be used to control the frequency and power in these additional modes. A Faraday isolator optically isolates the diode from any back reflections and spontaneously emitted photons released from the input of the BoosTA amplifier.

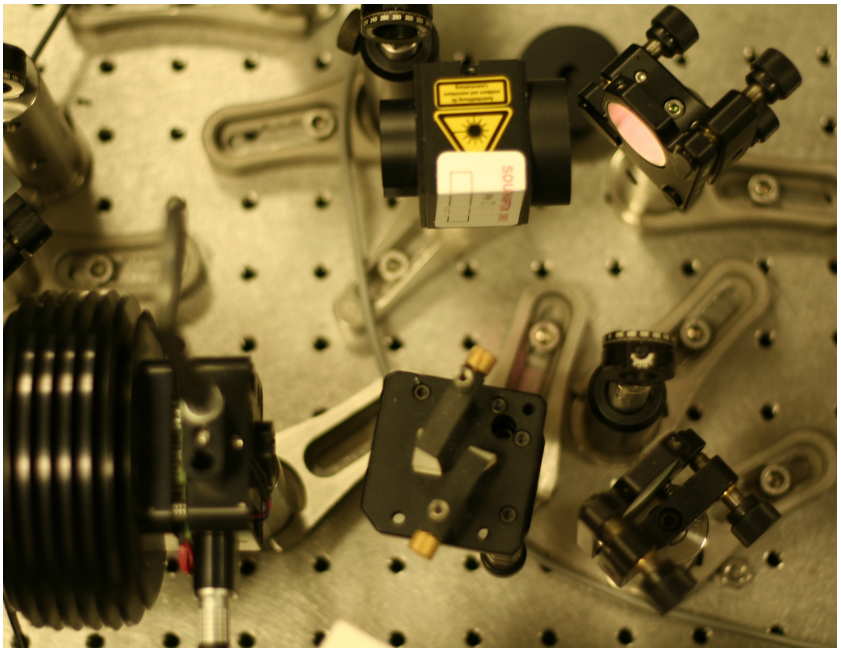

FIGURE 7.11: A photo of our injection locking setup, an anti reflection coated diode is seeded with power, isolated from back reflections by a Faraday isolator. The beam quality is made more circular by an anamorphic prism pair.

7.9.2 The Toptica BoosTA

We use a Toptica BoosTA tapered semiconductor amplifier to amplify the optical power in a single pass from 30 mW to 750 mW. To get as close as possible to the 40mW

damage threshold, without exceeding it, we use a fibre to couple the beam power into the amplifier: once the fiber mode is matched to that of the BoosTA, we can accurately measure with a power meter, the output power of the fiber, which is coupled into the BoosTA.

Other than this, the control electronics for the BoosTA are simple and effective, consisting of a single 'on' switch, and a gain control which is set to maximum. We must also ensure that no optical power is reflected back into the output of the BoosTA, which would cause damage.

7.9.3 Problems with amplification

The nature of the single pass method we use for amplification gives a number of unwanted side effects which, unless rectified, cause a destruction of our coherent state. This occurs as a result of amplified spontaneous emission (ASE) and distortion of the beam profile, which will will cover in the following section.

7.9.3.1 Amplified spontaneous emission

Amplified spontaneous emission is particularly problematic in lasers with low finesse cavities, as a result of the lower seed power. This therefore includes the Tiger, the injection locked laser, and the BoosTA, where the amplification is mostly performed in a single pass.

As ASE arises from spontaneous emission, the frequency spectrum is broadband, mirroring that of the gain profile, which can be several nanometers wide. Although the power in ASE is much lower, and the collimation much poorer than that of the laser beam, ASE has still been found to cause significant single photon absorption within the atom cloud (see figure 7.12) which cannot be neglected. We can reduce the relative amount of ASE by increasing the seed power, or turning down the gain on the amplifier. It is impossible, however, to eliminate it completely. Instead, we use several methods to reduce the intensity of the transmitted ASE and in particular, that of the resonant light which causes incoherent absorption.

As the emission direction of the ASE is determined by the dimensions of the chip, it is far less collimated than the laser beam. As a result of this, by increasing the distance between the equipment in the setup to reduce the solid angle we find that the ASE disperses much faster than the laser, which allows us to decrease the ratio of ASE to laser power: this is important on either side of the BoosTA.

By positioning a rubidium vapour cell in the beam path after the BoosTA, we can use the atoms as a filter for light resonant with single photon transitions. This filtering occurs due to the isotropic scattering of the resonant light, for which the solid angle reaching the atom cloud after re-radiation is extremely small. In order to maximise the attenuation of the resonant light, while simultaneously minimising the attenuation of the Raman frequencies, we use an anti-reflection coated vapour cell; as the optical quality of the windows used for these cells is also much higher, this gives the added advantage of decreasing the distortions to the beam shape. By heating the cell to $\approx 200°C$, we increase the vapour pressure and the Doppler absorption width of atoms in the cell which increases the attenuation of the resonant light. It is important, during this heating that the windows are at a higher temperature than the body of the cell so that condensation does not occur on the windows, which would reduce the throughput of light. Thermal air currents surrounding the cell are also seen to cause phase variations in the beam, and changes in the beam pointing. The effects of this air flow can be limited by restricting the air flow around the cell. Alternatively, if thermal air currents are unavoidably disrupting the phase sufficiently to be affecting the coherence of our manipulations it would be possible, though unpractical due to the long distances involved in separating such small frequency differences, to replace the cell with an alternative method of filtering, such as a grating.

After the heated vapour cell, we pass the beam through a 30 μm pinhole situated at the focal point of a telescope, through which it is unlikely that spontaneous emission will pass. This has the added advantage of spatially filtering the beam profile, which is distorted by both amplification and the imperfect windows of the vapour cell. We use aspheric lenses to give a diffraction limited focus, which allows a higher throughput of light through the pinhole.

7.9.3.2 Gain modulation and frequency mixing

As we have already discussed, amplification of the frequencies within the beam varies non-linearly with input power, giving proportionally higher amplification for frequencies with a low input power. The more unwanted frequency components there are, therefore, the more these will be amplified at the expense of amplifying the Raman frequencies. This problem is exacerbated by modulation sidebands produced as a result of gain modulation if the difference between the amplified frequencies is small. Ferrari [116] found that significant gain modulation will occur if two frequencies separated by less than 2 GHz are amplified simultaneously. For our system, unless the carrier in our system is extinguished, the 310 MHz frequency difference between the carrier and AOM beam will cause this gain oscillation. These sidebands extend for several GHz from

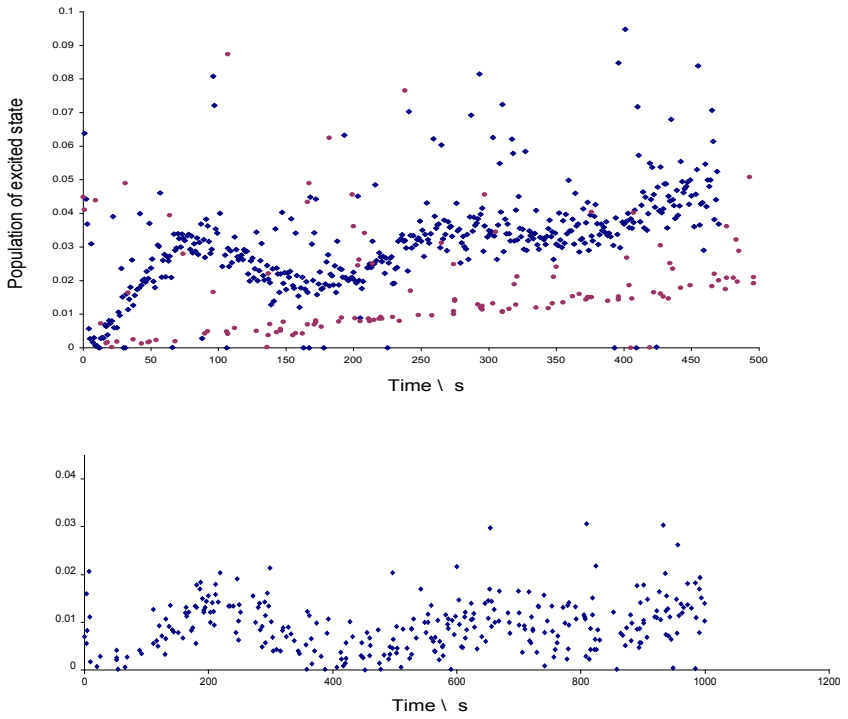

FIGURE 7.12: Unfiltered ASE from amplification causes non coherent repumpng the the cloud. On the top: Rabi oscillations driven by a Raman beam straight from amplification stages (blue) contain both coherent light as well as non-coherent, resonant light. If the amplification stages are unseeded, the illumination light becomes purely incoherent ASE, and we observe incoherent repumping (red dots), whereas if a heated vapour cell is placed between the amplification stages and the atom cloud, we eliminate resonant, incoherent light through scattering, and achieve much higher fidelity Rabi oscillations (bottom graph), albeit with a longer period as a result of a lower intensity.

resonance, which limits the minimum detuning which we can set our laser to avoid single photon absorption.

Independent amplification of each of the two beams by two separate diode-BoosTA setups would give the total output power of 750mW per frequency, and so a combined power in the Raman beam of 1.5W, without any non-linear mixing. Our group is currently working on designing an in house tapered amplifier system to achieve this.

7.10 Quantisation axis

It is important that, during manipulation, the atoms see an applied magnetic field which lifts the degeneracy of the Zeeman states. With degenerate m_f levels, there are many paths an atom can take between the upper and lower hyperfine states which, due to the difference in coupling strengths, have different Rabi frequencies. Without the splitting of the Zeeman levels by an applied field, all of these routes are resonant with the laser frequency. As the readout is sensitive only to the hyperfine state, rather than the m_f

states, these all contribute to the signal, causing a mixing of the Rabi frequencies, and a washing out of the Ramsey fringes. If the magnetic field is applied parallel to the beam direction, it also acts as a quantisation axis, defining the projection axis for the m_f states.

Applying a field is difficult within a magneto-optical trap, as the trap centre, where cooled atoms accumulate, is at zero magnetic field. Although it is possible to shift the position of the trap centre using unequal beam powers, this was found to disrupt the shape of the atom cloud as a result of the non-spherical optical forces. Applying an AC magnetic field, with a period of oscillation much higher than the motional frequency of the atoms in the trap can be used to apply a field, while maintaining a spherical MOT shape. By observing the behaviour of the MOT cloud after applying magnetic fields with different frequencies, we deduce that the oscillations are negligible for changes in the magnetic field on time scales shorter than ≈ 10 ms.

The magnetic field necessary to split the degeneracy is applied using a pair of Helmholtz coils, which we fit, perpendicular to the axis of the anti-Helmholtz trapping coils, parallel to the axis of the Raman beam, around one pair of windows. An arbitrary waveform generator, in combination with an audio amplifier, supplies the current used to drive the coils. The audio amplifier is an inexpensive method of amplification which can be used so long as the output impedance is high enough to avoid damage.

Rather than applying a continuous oscillating magnetic field to the atoms, we pulse the waveform generator to give a square wave output, $\approx$ 1ms long, which is triggered to coincide with the extinction of the MOT beams. Although we manipulate the magnetically insensitive $m_f = 0 \rightarrow m_f = 0$ transition, which negates the need for a homogeneous field, it is better to avoid any coupling between the different m_f states which may arise if we start sweeping the magnetic field. In practice, the low bandwidth and non-linearities of the audio amplifier make a square wave difficult to achieve, and the waveform before amplification must be manually tuned to find the best square-wave fit after amplification. We measure the induced current in a neighbouring coil of wire to determine the current in the magnetic coils, and use this to determine the homogeneity of the applied magnetic field; this we find to be within 5% of the total magnetic field over the duration of the entire pulse.

7.11 Choosing the right detuning for the Raman beams

By changing the current or the grating angle of the Tiger, it is possible to change the detuning of the Raman beams from single photon resonance. It is important that this

parameter is chosen carefully: a frequency too close to resonance will cause single photon effects (which, for a power broadened Lorenzian linewidth, will decrease with the square of the detuning) to dominate, whereas too far from resonance and the generalised Rabi frequency (which falls off with reciprocal of the detuning) will become too slow to perform a meaningful number of operations before the atoms fall from the manipulation volume. We run several temporal scans of the Rabi oscillation at different detunings and from this are able to determine the minimum detuning we can achieve in our system without observing decoherence effects from single photon absorption: we find this minimum detuning to be around 125 GHz. This is a much higher detuning than one would expect to need and indicated that there are perhaps additional frequency components contained within the Raman beam which are causing absorption even when detuned many GHz away from the desired two photon resonance (see section 7.9.3).

In addition to this, a light shift will affect the transition frequencies of the atom for all detunings, which makes it vital for coherent interactions that we minimise variations of the light shift across the cloud. This is non trivial to achieve as a result of the variations in the intensity of the Gaussian manipulation beam across the cloud. This variation, therefore, causes a shift in the Rabi frequency for two reasons: as a result of a change in the electric field strength, and also in the generalised Rabi frequency, as a result of the variations in the detuning from two-photon resonance across the cloud. As a result of the narrow linewidth of the Raman transition, the variation in the detuning of the atom cloud can cause a more significant change in the generalised Rabi frequency than intensity variations.

Figure 7.13 shows the dependence of the light shift on single photon detuning, which is obtained by summing the light shifts, over all paths, between the $F = 2,\ m_f = 0$ and the $F = 3,\ m_f = 0$ states. At our operating detuning of ≈ 125GHz, the light shift varies linearly with detuning, with a value of 0.543 ± 0.001 Hz mW^{-1} cm^2. For our beam of power 500 mW, this gives a frequency shift of 10 MHz.

7.12 Polarisations necessary for Raman transitions

In addition to the magnetic field we must apply to the atoms, we must also ensure that the Raman beams used to manipulate the atoms have the correct circular polarisations. As shown in figure 7.14, there are two routes between the $m_f = 0$ states for π polarised light, either via the $F = 2$ or $F = 3$ states. However for each of the two routes, one of the transitions has zero coupling strength, and as the total Raman transition strength is given by the product of the single photon transition strengths, both routes have a

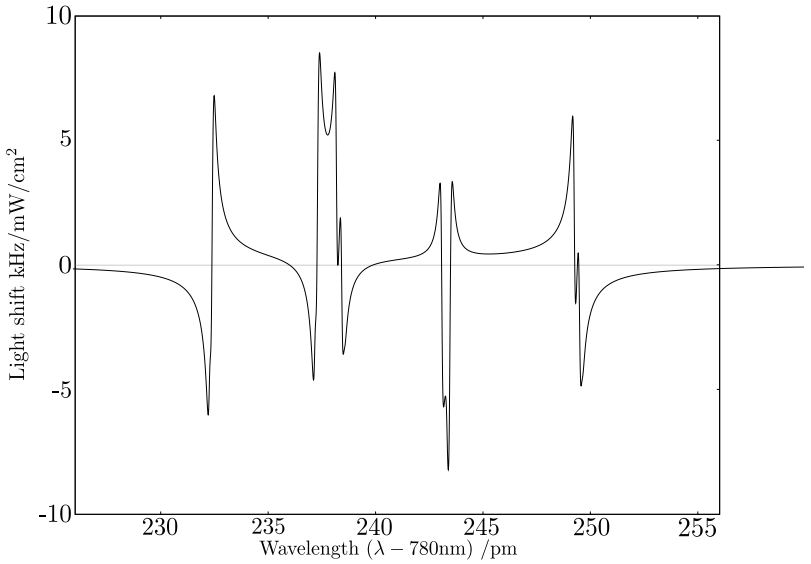

FIGURE 7.13: A graph showing the light shift in the Raman transition frequency as a function of the single photon wavelength [107]. The detunings showing in this plot are close to resonance, in regions where single photon absorption is dominant. We typically operate at detunings of 125GHz (λ = 780.5nm) at which point the light shift varies linearly with a value of 0.543 $\pm$ 0.001 Hz mW^{-1} cm^2.

zero transition amplitude. Circularly polarised light, however, has a non zero transition amplitude for both transitions and can therefore be used to couple the states.

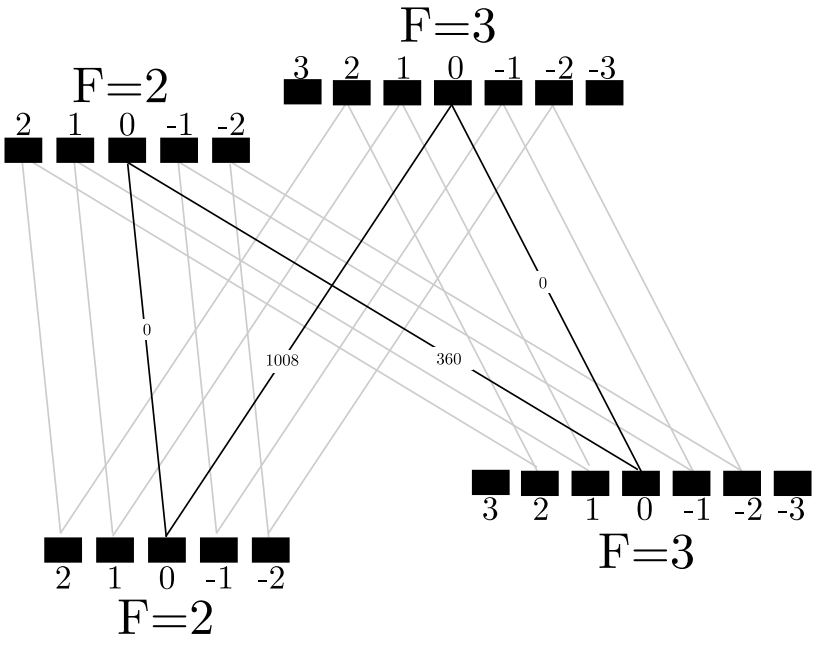

FIGURE 7.14: Two routes are available for the atom between the lower hyperfine $m_f = 0$ states, either via resonant enhancement with the F=2 state, or the F=3. While driving these transitions with linearly polarised light, however, the transition amplitudes for both paths (proportional to the products of the one photon transition strengths) are zero. This makes it necessary that we use circularly polarised light for which the transition amplitudes are non-zero (not shown).

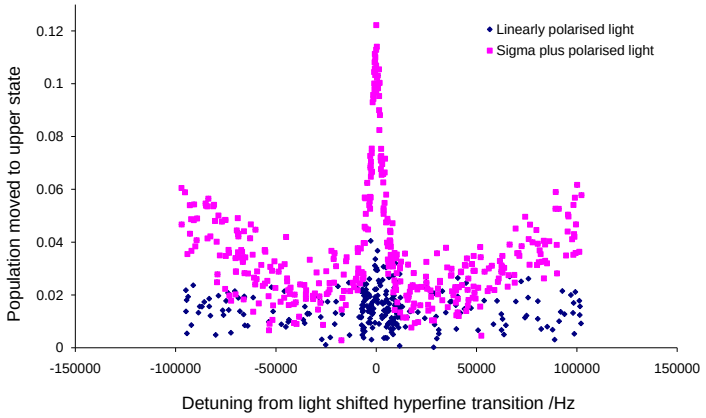

FIGURE 7.15: Experimental data showing the difference in absorption between Raman beams composed of $\pi - \pi$ polarisations (blue dots) and $\sigma - \sigma$ (red dots). The circularly polarised light shows a large resonance at the hyperfine splitting, whereas for the linear polarised light, all of the possible routes between the upper and lower hyperfine states with $m_f = 0 \rightarrow m_f = 0$ have a zero transition strength.

Chapter 8

Readout and analysis

To characterise the Raman manipulation, we sandwich a Raman pulse between two readout pulses and measure the repumping rate as a function of the frequency or pulse length of the beams. We refer to these scans as 'spectral' and 'temporal' scans, respectively. This process is automated by controlling the intensity of the trapping, repump and Raman manipulation beams, as we will discuss in the following section.

Since our readout is purely based on the state of the atom after manipulation, it difficult to distinguish between atoms repumped through single and two-photon mechanisms. Modulation of one of the Raman beams, while observing changes in the repumping, can be used to distinguish between two photon and single photon repumping mechanisms, although as the two photon rates are typically smaller than that of the single photon they can be difficult to measure accurately. In addition to this, the method we use to generate our Raman beams means that it is not possible to control the intensity of each of the Raman components independently. For these reasons, we change the relative frequencies of the Raman beams, and measure the effect of the Raman manipulation through the change in the excited state population. The large detuning of the Raman beams from single photon resonance gives a slower change in the absorption rate than for two photon absorption, which enables the narrower two photon resonances to be resolved.

8.1 The collection optics

To maximise the solid angle subtended by our collection optics, we use a large aperture lens, with a long focal length to give a large solid angle and high magnification, to project an image of the atom cloud onto an aperture, which we close around the MOT

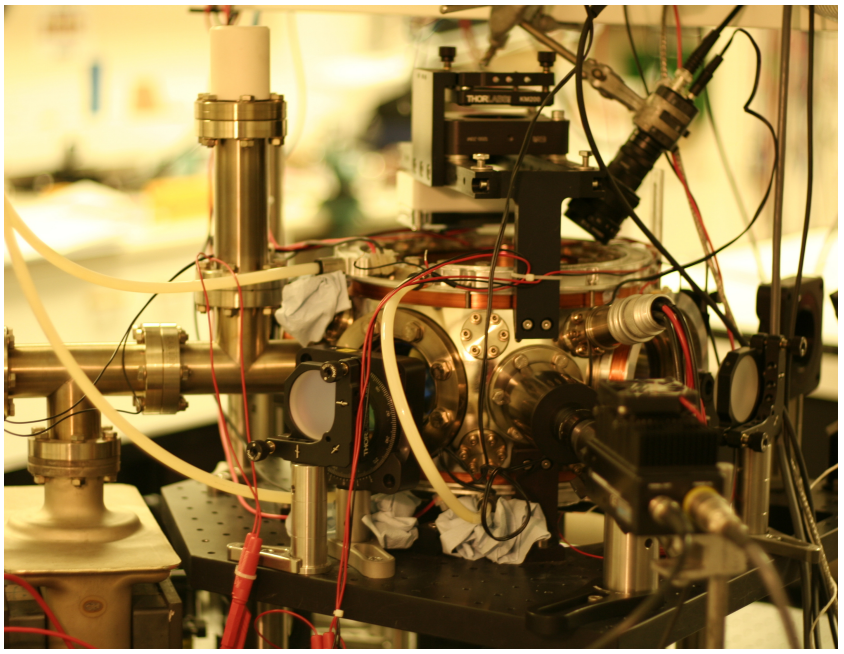

FIGURE 8.1: A photo showing the position of the imaging and readout optics surrounding the chamber. Both of these systems use lens extension tubes between the lens and the detector to increase the distance from the atom cloud to the lens such that it is comparable to the distance from the lens to the detector, giving a large magnification.

to spatially inhibit scattered light from hitting the detector. To achieve alignment, the detector is replaced with a CCD camera to look at the position of the cloud relative to the centre of the aperture. The correct aperture size is a trade off between noise and MOT signal, which is also dependent on the temperature of the cloud.

A spherical lens is used as the primary optical element in the readout system primarily due to its availability at large sizes. Due to the failure of the paraxial approximation in these lenses, we have calculated an increase in the signal with aspherical alternatives which we plan to implement in the future.

The detector consists of a fast Hamamatsu photo-multiplier tube (PMT), which we amplify using a Stanford Research Systems current amplifier. The signal from the current amplifier is sampled by a GHz oscilliscope, which averages and then saves the data over a local network.

8.1.1 The readout pulse sequence

The automated sequence starts by extinguishing the trapping beam, pumping the population into the $F = 2$ hyperfine state. During pumping, both the fluorescence of the MOT cloud and the number of atoms left in the transition loop decreases exponentially. After the population left in the transition loop has become negligible, the trapping laser

is extinguished and both the Raman laser and the quantisation-axis magnetic field are turned on. To readout the proportion of the population which has been affected by the manipulation beam, the trapping laser is turned on, and the height of the exponentially decaying fluorescence is compared to the height of the first exponential fit. Finally, the repump beam is turned on again, and the MOT is left to reform.

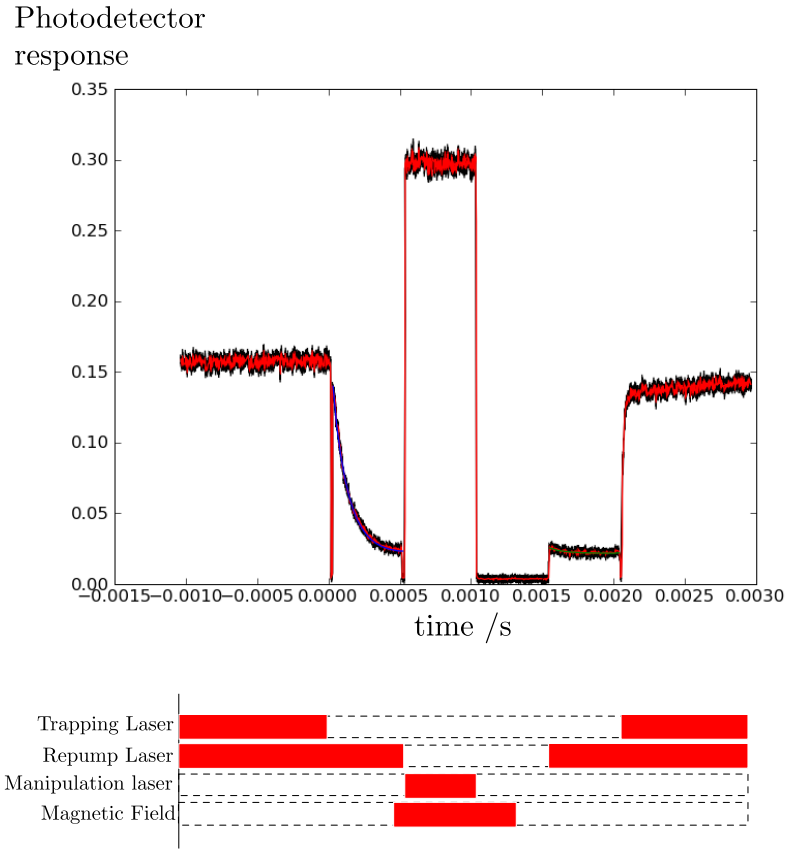

FIGURE 8.2: The pulse sequence (bottom) and signal(top) from a typical spectral or temporal scan. The atoms are optically pumped into the $F = 2$ state, after which a pulse from the manipulation laser is used to illuminate the atom. Finally the atoms are optically pumped back into the $F = 2$ state, before reloading the MOT. The ratio of the amplitudes from the two exponential fits tells us the fraction of the population moved during manipulation. The MOT magnetic field is left on for the entire duration of the pulsing.

8.2 Noise sources

It is vital for the successful measurement of coherent oscillations that we are able to measure small signals with sufficient accuracy that we can measure small changes in the excited state population. One simple method of achieving this is to increase the solid angle subtended by the collection optics which increases the number of measured photons up to a limit given by the maximum available optical access to the MOT. After

this, we rely on automation, which allows the averaging of many experimental runs, and other subtle techniques to measure these population changes accurately.

8.2.1 Stray light

The main source of noise is from light from both scattering of the Raman beam and background sources can be seen on the photo-detector while the experiment is running causing noise which must be eliminated.

The relatively high intensity of scattered light from the Raman beams causes large spikes in the output of the photo-detector. While passing the Raman beam through an optical fibre before the MOT cloud, and using a proper beam dump to extinguish the Raman beam helps to reduce this scatter, it is still found to be significantly greater than the signal from the preparation and readout pulses. As a result of this, the PMT exhibits a recovery time of up to 100 μs after large changes in the intensity, during which time the detector output is highly non-linear. Although possible to compensate for the instrument response when measuring the fluorescence of the repump beam, we instead omit signal recorded within 100 μs of the end of the Raman pulse. Despite this fix, the scatter is still undesirable, making it difficult to infer information on the fluorescence of the atoms during the pulse, information which could be used to determine the rate of spontaneous decay from atoms in the cloud.

8.2.2 Temperature changes and varying cloud atom densities

As well as from the scatter from the windows of the chamber, room temperature rubidium atoms which traverse the detection region also add noise to the signal. As these atoms have not been confined within the MOT region during pumping, they are equally populated across the two hyperfine states, causing an offset to the signal which is proportional to the density of the background gas. The density of this background gas changes with the temperature of the MOT chamber, a property which is difficult to hold costant as a result of heating by the magnetic coils, which causes anomalous results. In addition to this, the changing density within the chamber causes the atom cloud number to fluctuate.

As well as temperature changes affecting the density of atoms in the chamber, temperature changes cause a variation in the trapping and repump laser intensity and detuning as a result of changes in the atomic vapour cell used for locking and beam pointing. This causes a variation in the output of the cloud atom number, which, if variations are significant during the experiment, can affect the result. The use of air conditioning to

stabilise the temperature of the lab can cause significant air currents which are detrimental to the stability of the beam pointing. To avoid this, we enclose the beam path away from air currents and rely on the thermalisation of the equipment inside of these enclosed areas to reach a stable temperature before staring experimentation.

8.2.3 Electrical pickup

Finally, between the photon detection and readout, the electrical signal from the photodetector is also affected by pickup in the cables and amplifiers. This is particularly problematic as a result of the large number of RF drivers and amplifiers used across the laboratory. Fourier analysis of the Raman signal, shows a significant spike at 310MHz, due to pickup from the amplifier used for the Raman AOM. Shielding the detection optics from this pickup was achieved through electrical isolation of the driver housing from the optical table, as well as encasing the driver in metal. Shortening of the cable length between the PMT and the amplifier also resulted in significant reduction of pickup. The current amplifier which we use to amplify the PMT adds a noticeable modulation to the signal at 50 Hz due to mains pickup; we find a significant reduction in this pickup if the amplifier is run off an internal battery.

8.3 Methods of noise reduction

In order to reduce the significance of noise on our data, we use two general methods: increasing the manipulated atom number and averaging. While overloading the MOT can be used to increase the total atom number, we discuss here methods of increasing the atom number with which we interact, which can be used to give much more significant improvements, without increasing the size or temperature of the atom cloud. Averaging is only made possible as a result of the automated pulse sequencing, and is cruicial in order to allow us to visualise Rabi fringes.

8.3.1 Automation and randomisation of data taking

The modulation method of generating the manipulation beams, using arbitrary waveform generators, makes it possible to automate the experiment so that it runs autonomously, selecting the frequency for each individual run randomly from within the scan range. In using this random method of data acquisition, many of the slowly varying noise sources discussed in the previous chapter are randomised and made uncorrelated in time so that they show-up as stochastic noise on the data set. Automation of the

experiment also allows the experiment to be repeated many times within a short time scale for each single parameter. This data is then averaged on an oscilloscope to give a single data point.

For each run, we record 10^6 points, which are each averaged in parallel over 3 s; each of the exponentially decaying readout fluorescent segments of the overall readout is therefore sampled by 10^5 data points.

Despite the efforts taken to reduce the noise, without averaging it would be difficult to distinguish our signal from background, and it would certainly be impossible to discern subtle information such as any possible chirping of the Rabi frequency. This averaging, however, is not without its faults: we find that we lose information which could otherwise be used to characterise the decoherence mechanisms within our cloud.

When the dephasing mechanisms are occurring on time scales faster than measurement time for each of the points, we obtain a continuous yet heavily damped Rabi oscillation (illustrated in the top figure 8.3). Alternatively, dephasing mechanisms occurring slower than the measurement time for each point show multiple Rabi frequencies which are randomly sampled (top of figure 8.3). From observations of our data, we conclude that the decoherence mechanisms are occurring faster than the rate at which we take data: 3 seconds.

8.3.2 State preparation

During the operation of the magneto-optical trap, the atoms are mostly contained within the $5S_{1/2}\ F = 3 \rightarrow 5P_{3/2}\ F = 4$ cooling cycle and the $5S_{1/2}\ F = 2 \rightarrow 5P_{3/2}\ F = 3$ repumping cycle, equally divided amongst the m_f levels. Because of this, before commencing the coherent manipulation schemes, we must pump the atoms into a single hyperfine ground state. As we only interact with a single m_f transition, we also discuss a process which would allow pumping of atoms into a single m_f state, which would increase the signal by a factor of 5.

8.3.2.1 Optically pumping the hyperfine levels

For similar reasons to those which make neutral Rubidium suitable for magneto-optical trapping, the simple energy level structure makes it straightforward to perform high-fidelity initial state preparation. Extinguishing one of the two MOT beams causes the population to be incoherently cycled while optical pumping occurs to the state not coupled to that transition. This causes the absorption and fluorescence to decrease with time, and this can be easily observed with the detection optics.

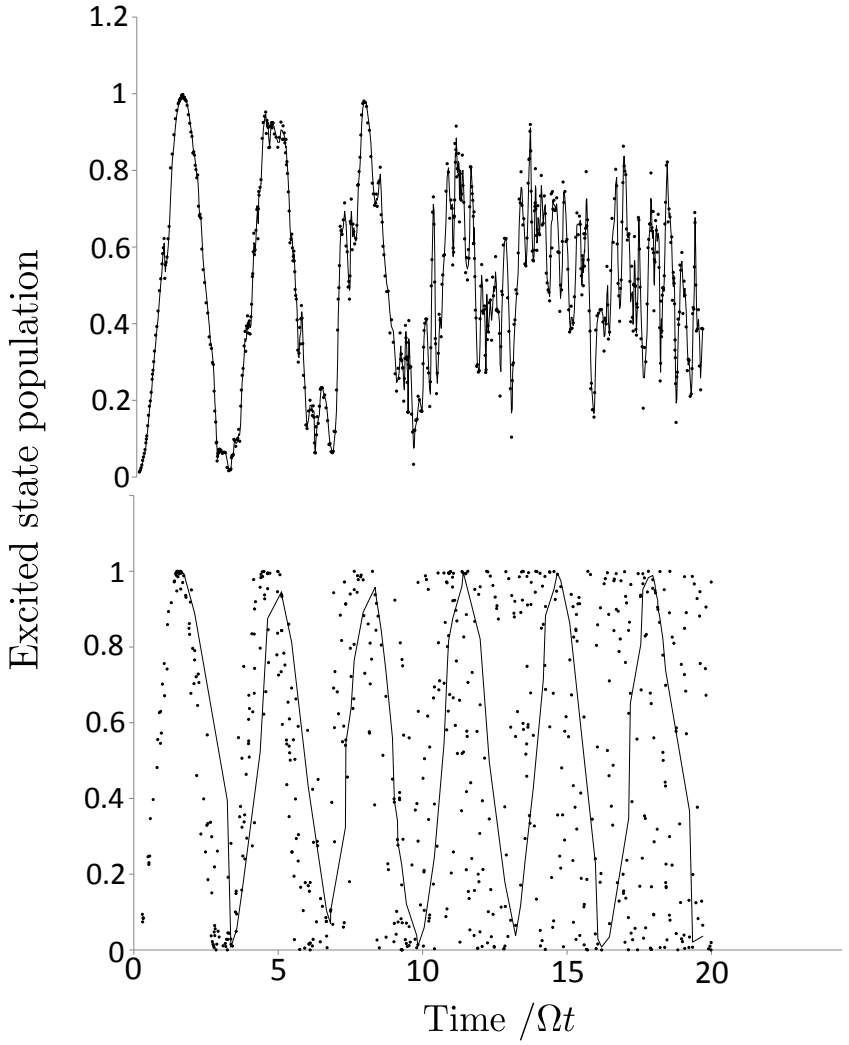

FIGURE 8.3: A comparison of two systems with a temporal variation in the Rabi frequency used to illuminate the atom cloud. The top graph shows the outcome of a simulated system with a variation on timescales slower than the averaging time, whereas the bottom graph shows the effect of a variation in the Rabi frequency which is faster than the averaging time. The solid line in the top graph shows a two point moving average of the data while in the bottom, a thirty point moving average.

For the trapping transition, which very nearly forms a closed optical loop, the atom decays to the $F = 2$ state after several microseconds of being cycled on the $F = 3 \rightarrow F = 4$ transition. Decay from the repumping cycle is much slower, and occurs on time scales of a nanosecond. We therefore use the decay of the fluorescence during pumping from the $F = 3$ state into the $F = 2$ state to determine the state populations, in which decay occurs much slower, allowing more accurate fitting to the exponential decay.

8.3.2.2 Optically pumping m_f levels

We start our manipulation experiments with the atoms equally divided amongst the F=2 Zeeman states. As we interact with only one of the five Zeeman states, we find

it desirable to concentrate the atoms in the $m_f = 0$ state, such that the available population in that level is closer to the total population of the cloud. While it is simple to pump the atoms into an extreme m_f state, ($m_f = \pm 3$) using circularly polarised light, pumping to the $m_f = 0$ state is far less trivial.

One possibility, illustrated in figure 8.4, shows the coupling rates between the $5S_{1/2}F = 2$ and $5P_{3/2}F = 1, 2, 3$ states for π polarised light. We see that the coupling rate is zero for the $F = 2 \rightarrow F = 2, m_f = 0$ transition. Because of this π polarised light, tuned to the $F = 2 \rightarrow F = 2$ transition will pump the population into the $m_f = 0$ state. This would require a separate beam, which unless retro-reflected would subject the atoms to additional optical forces, which would displace the cloud from the trap centre.

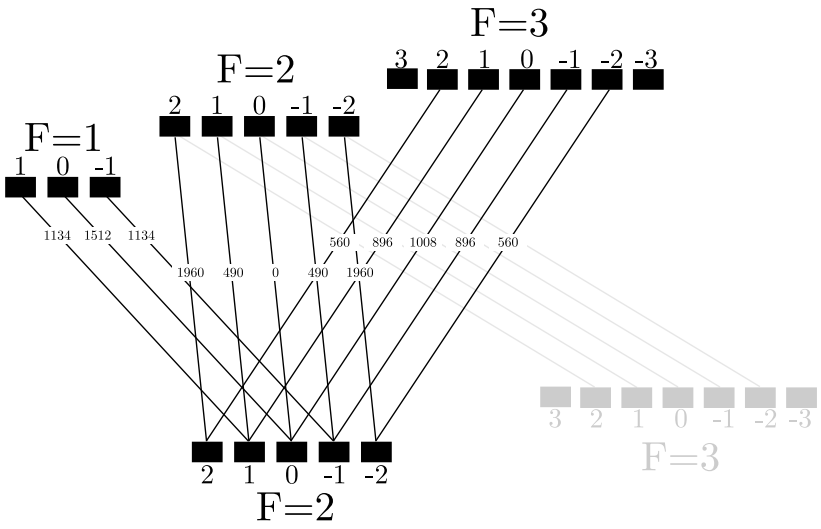

FIGURE 8.4: A graph showing the transition strengths from the $5S_{1/2}F = 2$ state in Rubidium 85, for π polarised light. The zero coupling strength between the F=2 $m_f = 0$ and F=2 $m_f = 0$ gives optical pumping to the m_f=0 state.

In our setup, it is difficult to find optical access to the MOT cloud which does not pass through a quarter waveplate, and however well the retro-reflected pump beams are aligned, the standing wave pump field causes damage to the structure of the cloud. Since our experiment is currently more sensitive to inhomogeneities in the cloud shape than to the signal to noise ratio, we have decided not to implement this m_f pumping scheme until we have achieved better coherence.

Chapter 9

Results

In order to drive two photon transitions coherently, it is necessary for us to supply a constant light intensity over the whole of our atomic ensemble, with a constant frequency detuning relative to the state we are driving. In addition to this, population inversion can only be achieved with driving light fields which are resonant with the atomic transition. The main content of this section is dedicated to the work done in finding the conditions required to meet these criteria for the entire duration of the experiment.

9.1 Non-coherent manipulation of Rubidium

To begin with, we find the position of the two photon resonance by measuring the repumping of Rubidium atoms as a function of the change in the two photon frequency difference. Depending on the coherence of the system, the shape of this resonance can either be a power-broadened Lorentzian, or a sinc squared profile, centered around the transition frequency (see figure 9.1 below).

9.1.1 m_f states

In a broad, 1MHz wide scan of the resonance, it is possible to resolve the individual m_f sublevels, which are split by their interaction with both the applied, pulsed magnetic field and the trapping magnetic field gradient. So long as the two Raman beams are of the same polarisation, and the single photon detuning is sufficient to ensure that the upper $5P_{3/2}$ state remains unoccupied, co-propagating Raman beams impart no change to the angular momentum of the atom. It is because of this that Raman transitions with a well defined polarisation by a correct quantisation axis, parallel to the propagation direction of the beam, will only occur between transitions with $\Delta m_f = 0$. By changing

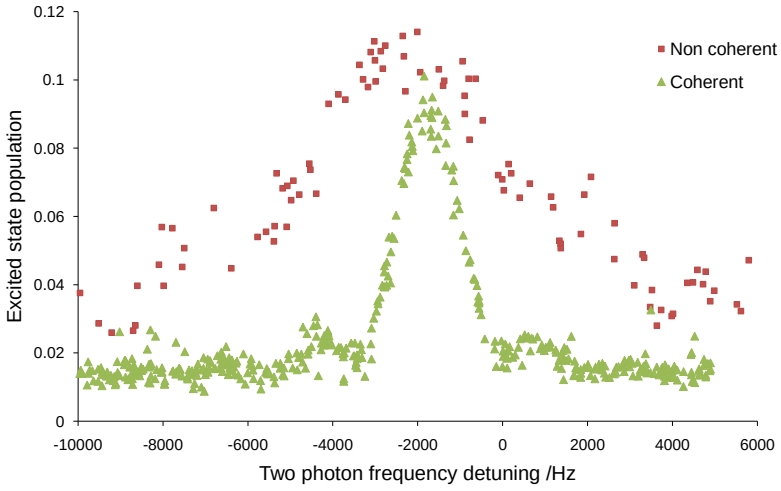

FIGURE 9.1: A comparison of the central $m_f = 0 \rightarrow m_f = 0$ transition for long (green) and short (red) coherence times. The different beam powers explain the different line widths, and different, light shifted, resonant frequency. Notice the side lobes on the coherent interaction, not present for the incoherent interaction.

the value of the applied magnetic field, and noting the changes in the positions of the resonances (as shown in figure 9.2 and 9.3), we are able to verify that the resonances on the spectral scans are a result of interactions with the different m_f transitions.

The gradient in the magnetic trapping field across the cloud causes inhomogeneous broadening of the $m_f \neq 0$ resonances. It is for this reason that we choose the $m_f = 0 \rightarrow m_f = 0$ transition for our coherent manipulation experiments which does not exhibit variations in the Rabi frequency as a result of the variation in the Zeeman shift across the cloud.

9.1.2 Non-coherent interactions for MOT size measurements

As described in the paper by Savard *et al.* [122], the frequency and broadening of the magnetically sensitive $m_f \neq 0$ transitions can be used in conjunction with the strength of the magnetic field gradient and the Zeeman shift in the atomic energy levels with magnetic field strength to measure the spatial distribution of atoms along the beam axis.

Using the fit shown in figure 9.4, we find the cloud width for this scan to be 550 μm which is comparable to that found in section (6.3).

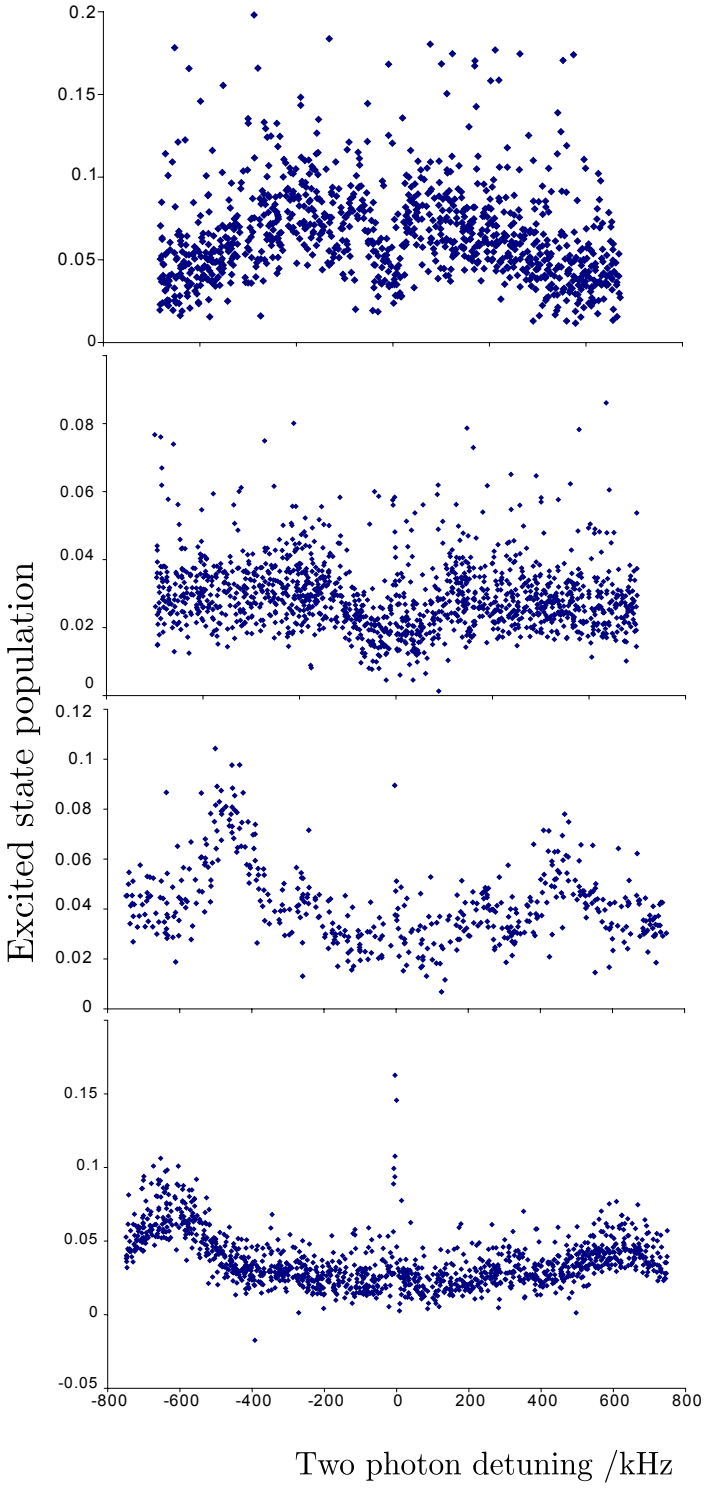

FIGURE 9.2: Four Raman spectral scans with a 0.05(top), 0.133, 0.266 and 0.399 Gauss (bottom) magnetic field applied to the atoms through the quantisation axes coils. An increase in the spacing between the resonances with magnetic field strength is clearly evident, which supports our conclusion that they are a result of the different m_f transitions.

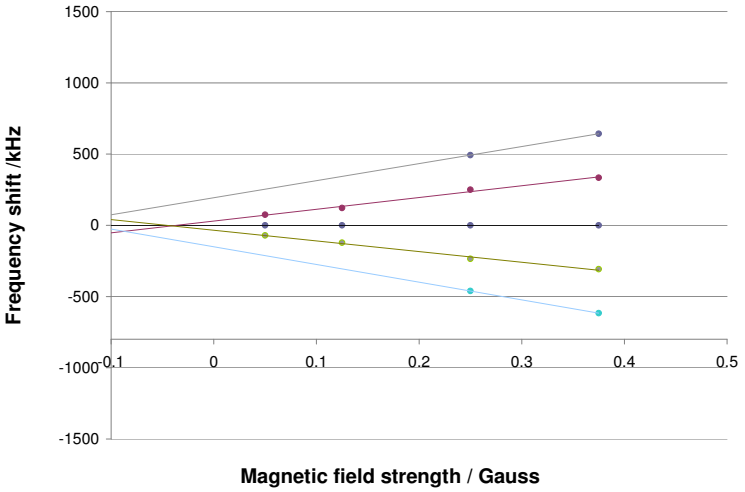

FIGURE 9.3: The detuning of the m_f levels from two photon resonance plotted against the voltage supplied to the quantisation axes coils.

9.2 Choosing the right m_f state for manipulation

In order to guarantee interactions with a single m_f state at a time, it is necessary for the resonances to be separated by more than the linewidth of the transition: hundreds of kHz for the magnetically broadened non-$m_f = 0$ states . If more than one transition is resonant at any one time, the difference in the two photon detuning causes the atoms to be driven at different Rabi frequencies, reducing the contrast of the Rabi oscillation. The non-degeneracy of the m_f levels can be easily verified from the separation of the resonances in figure 9.4; if the resonances are found to be overlapping, as in the case of figure 9.4, the amplitude of the applied magnetic field can be increased to separate the states further.

9.3 Coherent Raman interactions

9.3.1 Calculation of the expected two photon Rabi frequency

Using the numbers given in the previous chapter, we can calculate the Rabi frequency expected from our two photon interactions.

Property	Value
Total Beam Power	750 mW
Individual Beam Power	250 mW
One Photon Detuning, $\Delta/2\pi$	0.25nm $\approx$ 120GHz
Beam Radius	5 mm
Average Electric field strength	1.5×10^3V m^{-1}
Peak Electric field strength	3×10^3V m^{-1}
Dipole Moment-pump beam	2.64×10^{-29} C m
Dipole Moment-stokes beam	1.48×10^{-29} C m
$\Omega_P/2\pi$	4.38×10^8 Hz
$\Omega_B/2\pi$	7.811×10^8 Hz
$\tilde{\Omega}/2\pi$	2.2×10^5 Hz

TABLE 9.1: The numbers used in the calculation of the two photon Rabi frequency for our system, and the resulting value of the frequency, $\tilde{\Omega}$. To give the maximum two photon Rabi frequency, the powers of the pump and Stokes beam are made equal.

In this calculation, the value for the two photon dipole interaction strength is found from a summation of the dipole matrix elements over all of the individual routes between the two states. These routes are summarised by figure 9.5.

9.3.2 Temporal scans and Rabi oscillations

After finding the resonance frequencies of the Raman interactions from the spectra, we set the difference frequency of the Raman beams to coincide with the centre of the $m_f = 0$ transition; the beam intensity is then modulated using AOMs to measure the change in the atomic population with pulse length. As we predicted in chapter 4, if the decoherence decay time is comparable to the Rabi period, this shows some oscillatory behaviour in the output state of the atom as a function of time.

Figure 9.6 shows the result of one of these temporal scans which exhibits a large damping factor. Due to the finite Rabi frequency of the atoms, this decay inhibits the number of coherent manipulation operations which can be applied to the atomic system; it is therefore crucial for the coherent manipulation schemes that we reduce the dephasing mechanisms as much as possible in order to manipulate a large fraction of the population through at least two full Rabi oscillations: the minimum requirement for any coherent cooling scheme.

As described in section 4.3.2.2, the rate of decay to the Rabi oscillation due to this distribution is not an exponential, which makes it difficult to measure the decay time constant. This is inconvenient: making it difficult to compare the decay rates for the different decay mechanisms. For short times, of the order of one to two times the Rabi period, the decay rate can approximated by an exponential for which a time constant

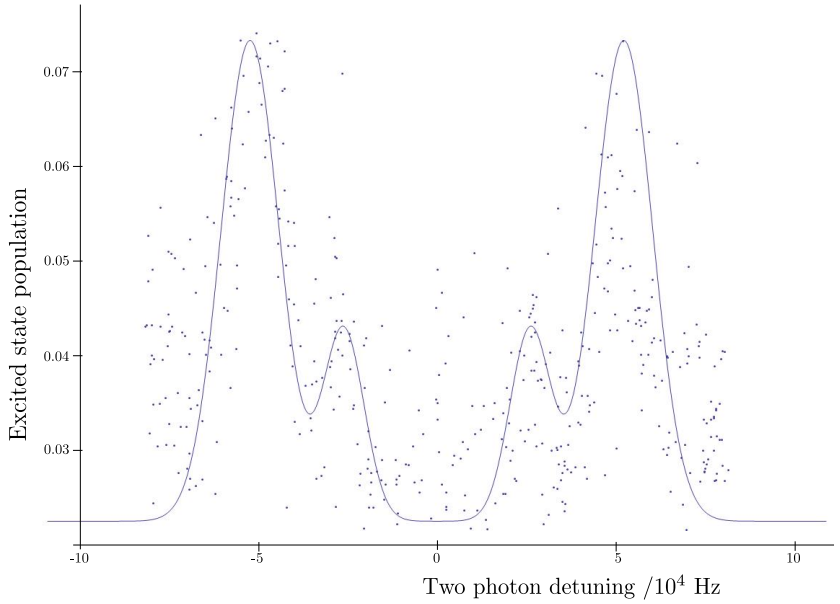

FIGURE 9.4: Four Gaussian functions, offset from zero detuning by the MOT's offset from the position of zero magnetic field and broadened by the magnetic gradient, are summed to give the total repumping spectrum as a function of the detuning. The central resonance is not magnetically broadened which makes it too narrow to be resolved in this scan.

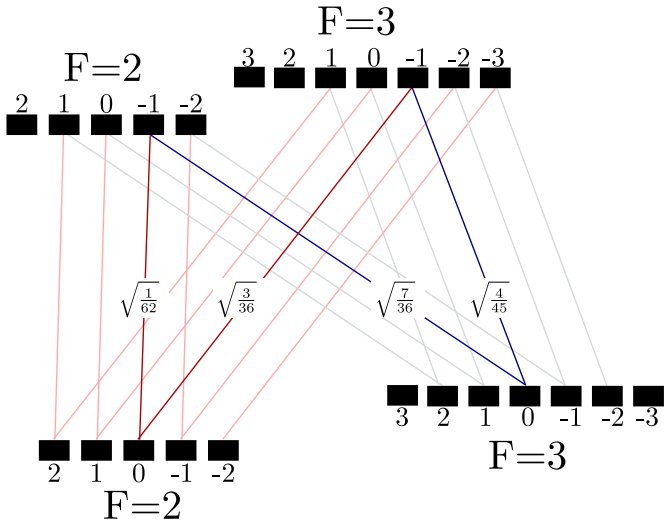

FIGURE 9.5: The paths available for two photon transitions between the F=2 and F=3 lower hyperfine states by σ^+ polarised light. The paths between the $m_f = 0$ states are shown in bold, with the value of the dipole moment, divided by 3.6×10^{-29} C m, printed alongside [95].

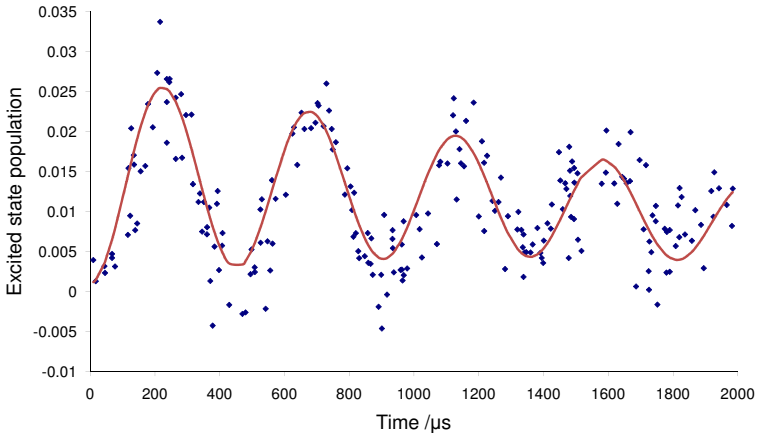

FIGURE 9.6: A figure of a Rabi scan date (blue dots) with fitted exonentially decaying sine (red line). Least squares fitting gives a period of 453±9 μs, which equates to a Rabi frequency of 2.2±0.035 KHz, and decay constant of 2.1±0.2 ms.

can be measured, and used for comparison. We find a suitable decay constant via least squares fitting, which fits the decay in coherence of the system during the first few oscillations.

We fit an exponentially decaying sine function to the data shown in figure 9.6, and find the Rabi frequency to be 2.19 kHz with a decay constant of 2.1±0.2 ms. The oscillation frequency is significantly less than that predicted in table 9.1 and can be explained by an over estimation of the power in the Raman beams. As discussed in section 7.5, it is likely that as a result of amplification, additional frequency components are generated in the Raman beam, which, despite containing a significant amount of power are not used in driving the Raman transition. As the method we use to measure the power in the Raman beams is sensitive to the power contained within these superfluous frequencies, the expected calculation of the Rabi frequency may be over-estimated.

9.3.2.1 Inconsistencies between the expected and the measured Rabi frequencies and light shifts

The 1MHz expected Rabi frequency calculated in table 9.1 is three orders of magnitude larger than the 2.2kHz frequency we measure above. While this difference between the measured and calculated Rabi frequency can be partially attributed to the inaccuracies

in measuring the total power in the Raman beams, described in chapter 7.5, as a result of the coupling of broadband ASE into the BoosTA, non-linear amplification, power loss during spatial filtering and poor transmission through windows. Together, these processes are unlikely to result in more than a 50% reduction in the power contained within resonant spectral components, which is not enough to explain the apparent three orders of magnitude reduction in resonant beam power: a few milliWatts.

Similarly, a discrepancy arises between the predicted and the expected two photon light shift, which again differs by a factor of $\frac{\Delta_{calculated}}{\Delta_{observed}} = \frac{1\text{MHz}}{4\text{kHz}} \approx 10^3$.

To explain this, it is important for us to remeasure the Raman beam spectrum, with properly calibrated measurements of the spectral density, allowing us to determine the amount of power contained within broadband ASE. Due to the narrow linewidth of the transition frequency, it is plausible that the low Rabi frequencies are a result of driving the atoms off-resonance with the driving laser, which causes large changes in the Rabi frequency with the detuning of the driving frequency. We must therefore re-check the calibration of our electronic frequency references. It may also be necessary for us to check the spectrum before and after each of the amplification stages in order to see that no significant changes to the spectrum of either beams is occurring, and if necessary, use separate amplification stages for the pump and stokes beams.

Overall, the reduction in power contained in the Raman frequency components is catastrophic if we want to maximise the number of coherent operations we can make to the atomic cloud. High beam power is essential for two reasons: firstly it enables us to create a large diameter manipulation beam, such that the intensity varies slowly with distance around the centre, causing a small variation in the Rabi frequency across the cloud. Secondly high Rabi frequencies allow a greater number of operations to be performed before the atoms move out of the trapping region. If the optical power is limited, we are forced to trade off these two conditions to achieve the maximum number of oscillations with the best fidelity.

By increasing the effective Rabi frequency through detuning the two photon frequency further from resonance, we can increase the effective Rabi frequency (equation 4.2). While increasing the frequency of the oscillations, detuning also causes a reduction in the maximum inverted population. This can be useful when characterising the form of the decoherence mechanisms, although the large reduction in amplitude for off resonant Rabi oscillations causes a reduction in the percentage of the total population moved between the ground and the excited states making it unsuitable for the implementation of coherent cooling schemes.

9.3.2.2 Decay rate

The shape of the decay of these Rabi oscillations gives a good indication of the decay mechanism. By changing the radius of the Raman beams used to drive the transition, for example in figure 9.7, we see that the rate of oscillation changes with beam radius, as expected due to the change in the optical intensity, but also that the overall shape of the curve is constant, albeit stretched. This implies that the rate of decoherence is proportional to the driving field intensity or inversely proportional to the square of the Raman beam radius.

In order to further ascertain the cause of the decoherence, we run Ramsey-Bordé [**?** **?**] interferometry and spin echo experiments, which can be use to measure the relative decay constants for non-optical and non-stochastic forms of decoherence respectively.

9.3.3 Ramsey interferometry

Once we have achieved high fidelity manipulation over a quarter of a Rabi oscillation, it is possible to use two $\frac{\pi}{2}$ pulses to construct a Ramsey-Bordé [**?** **?**] interferometer. As the interferometer relies on measuring the evolution of the atom while the beams are turned off, the success or failure of this technique gives an insight into the non-optical forms of decoherence, such as collisions, which occur regardless of the presence of manipulation laser light. In addition to this, it is necessary for the $\frac{\pi}{2}$ pulse area to be constant in order for the superposition to be completely collapsed on application of the second pulse; any movement of the atoms between the two pulses, however, will cause a change in the light intensity seen by the atom, which will prevent this. Figure 9.9 shows a plot of the signal from a Ramsey interferometer, where the time between two $\frac{\pi}{2}$ pulses is plotted against the excited state population.

Using a least squares fit, the decay rate in the Ramsey fringe visibility is calculated to be 4.9±0.1 ms. As the decay in this system occurs while the driving laser is turned off, this rate of decoherence is indicative of non-radiative mechanisms such as through collisions or spontaneous emission. The decay rate of 1/660 ms^{-1} measured on these graphs is comparable to that calculated in section 6.6. In addition to this, figure 9.9 shows a linear increase in upper state population with time, with a time constant of 35±3 ms; this indicates pumping, which could possibly be due to single photon absorption followed by spontaneous emission.

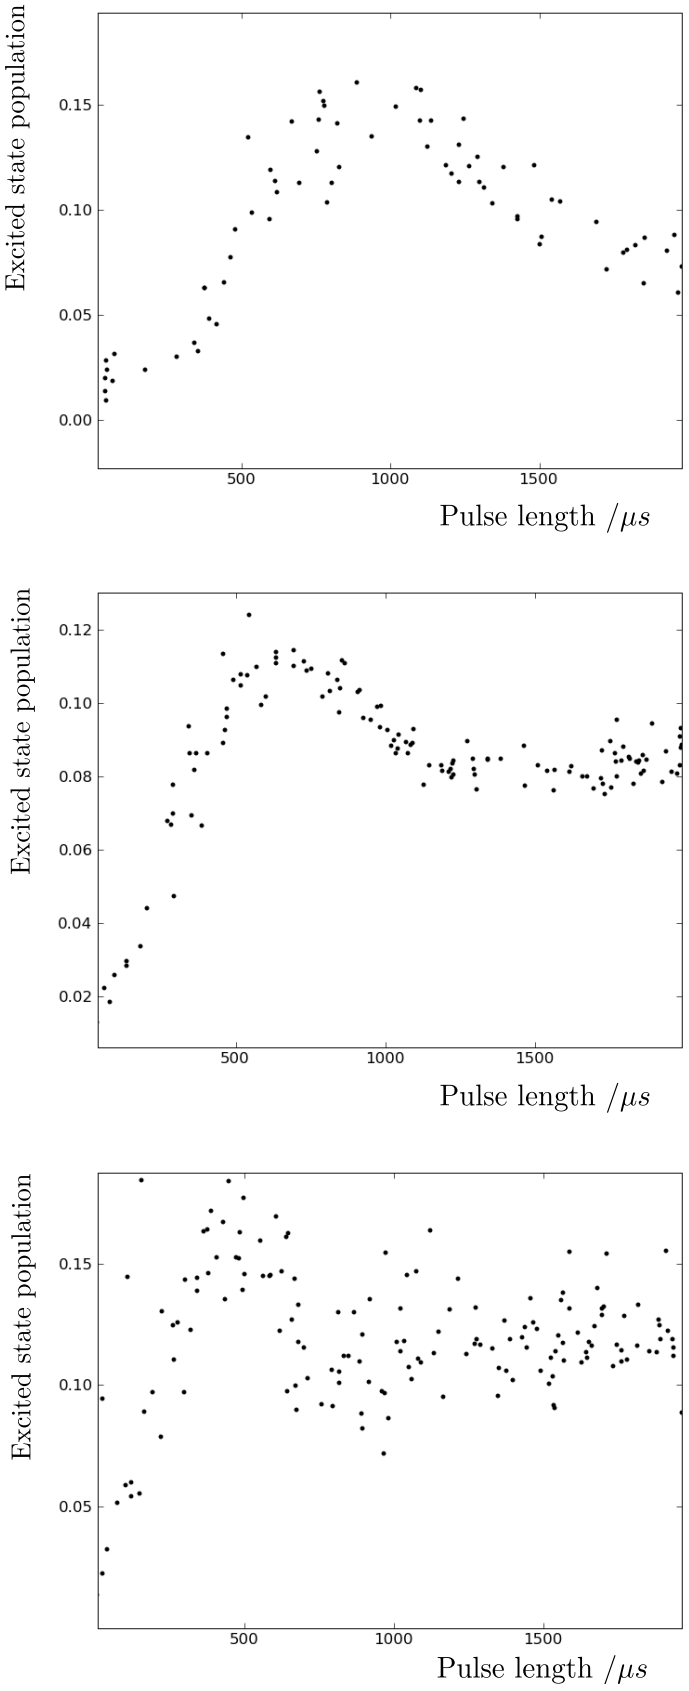

FIGURE 9.7: A comparison of three Rabi frequency scans with Raman beam waists of 7mm (top), 5 and 3mm (bottom), keeping the total power and detuning of the beam beam constant. The Rabi frequency decreases with optical intensity, but so does the decay rate, causing the curves to take the same shape, but with different frequencies.

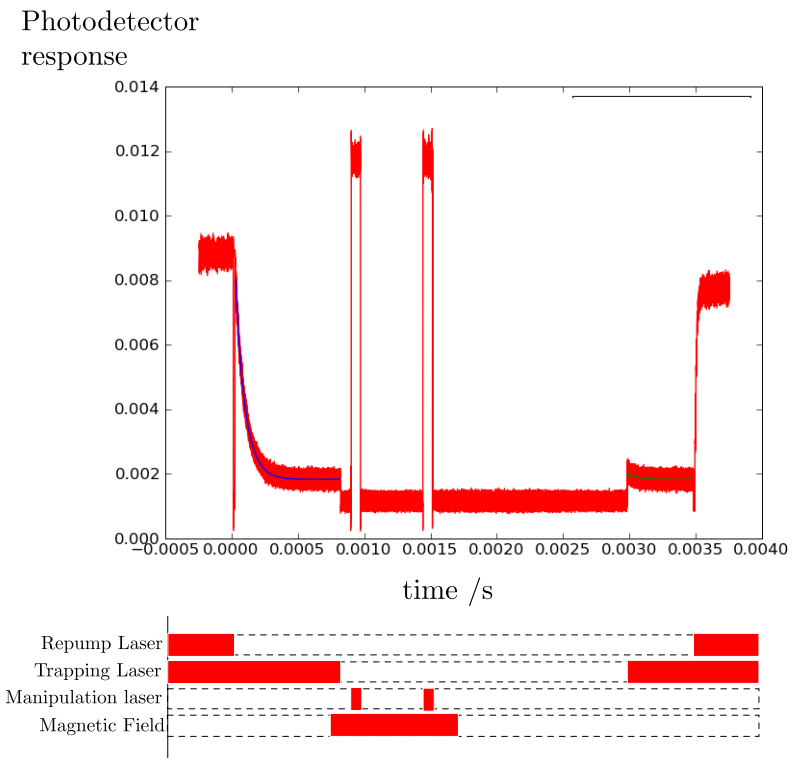

FIGURE 9.8: A trace and corresponding pulsing scheme for a typical Ramsey interferometer run. As with the Rabi scans, the change in the atomic population due to two $\frac{\pi}{2}$ pulses is plotted against the separation between the pulses.

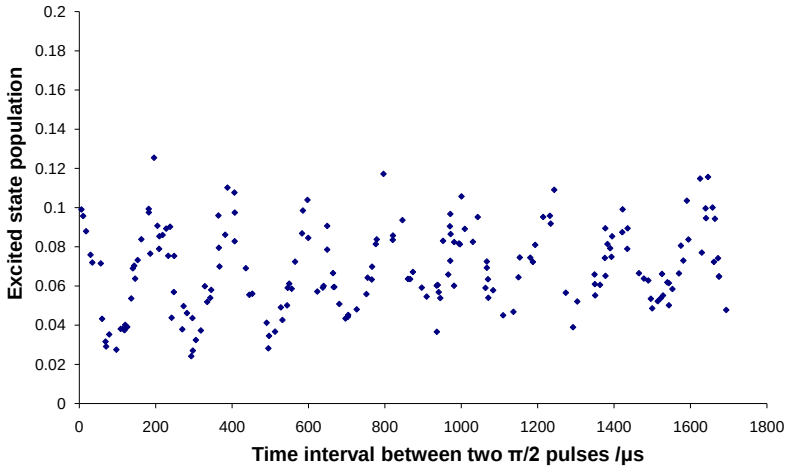

FIGURE 9.9: Oscillations in the excited state population, 'Ramsey fringes', as a function of the spacing between two $\pi/2$ pulses. The small decay in the fringe contrast, compared with the large decay in the Rabi oscillations, indicates that the dephasing is primarily an optically induced phenomenon.

9.4 Spin echo

After the success of the Ramsey fringe experiments in demonstrating the effect of non-radiative forms of decoherence, such as collisions and thermal broadening, we investigated the effects of radiative, stochastic processes on decoherence with a pseudo spin-echo technique. As described in section 7.3.2, spin echo experiments are performed with an I& Q modulator in series with the RF supply for the EOM to causes a 180 degree change in phase of the EOM RF phase, and a corresponding 180 degree change in the phase of the Raman light.

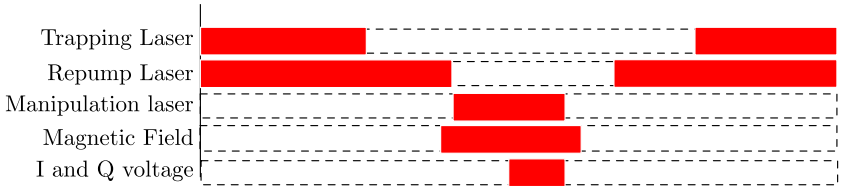

FIGURE 9.10: The spin-echo pulse sequence

We choose the timing of the I & Q pulse, $T/2$, so that it arrives after the system has seen significant decrease in fringe contrast. Thus for pulses up to $T/2$ in length, the graph resembles a standard Rabi scan. After $T/2$ however, the driving laser changes phase by 180 degrees which reverses the direction of precession of the state vector on the Bloch sphere as described in section 4.3.3. Although the direction changes, the rate of precession for each state vector around the Bloch sphere is not affected, causing the spread in state vectors to converge as the time approaches T.

For atoms off resonance with the driving laser, this reversal of the phase is not so straightforward. The zenith angle of the of the field vector of the Bloch sphere depends on the ratio of the Rabi frequency to the detuning, and not at all on the phase of the driving laser. As a result of this, reversing the phase for atoms which are off resonant with the driving field will not cause the convergence of the state vectors. As light shifts can occur as a result of a light shift in the atomic energy levels, due to a spatially varying Raman beam, we must make sure that the light shift is much less than the Rabi frequency in our system if the state vectors are to converge.

We can quickly test the ratio of the light shift to the Rabi frequency from the shape of the central $m_f = 0$ resonance on the spectral scans. The Lorentzian, rather than Gaussian shape of the central resonances in our system implies that the rate of power broadening is greater than the light shift-broadening. For the rest of this chapter, we shall therefore assume that the light shift is negligible in comparison to the Rabi frequency.

Figure 9.11 shows the result of a successful spin echo sequence where the phase of the driving laser is reversed after 100 μs. The dip around 200 μs shows that decoherence

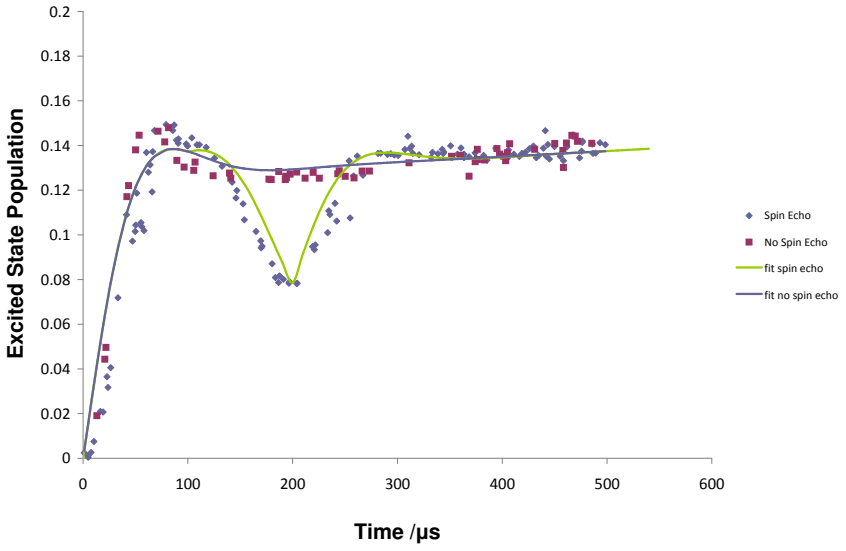

FIGURE 9.11: A spin echo experiment with (blue dots) and without (red dots) a phase reversal of the Raman beam after 100 μs along with fitted data. This phase change causes a the Raman scan to be reflected around 100 μs, which causes a dip in the population at 200 μs. The presence of non-reversible stochastic forms of decoherence causes a time reversal asymmetry around 100 μs, which stops the population returning to zero at 200 μs.

is partially a result of non-stochastic forms of decoherence, for which the most likely candidate is intensity variations of the Raman laser across the cloud. The spin echo graph does not return to zero at 200 μs, demonstrating that some stochastic, non-reversible processes must also contribute to coherence loss in the system. These could be a result of many other processes, such as intensity variations of the driving laser field, brought about by changes in beam pointing or coupling of light through optical elements due to air currents; more likely, however, is the effect of thermal motion of the atoms on timescales similar to that of the experiment.

We use a least squares algorithm to fit an exponentially decaying sine curve to the spin echo data. After the phase of the driving laser is reversed, we use a product of a sine curve with two exponentials, with different decay times, one positive while the other negative, to describe the data: representing the reversible, non-stochastic decay time and a non-reversible, stochastic decay time respectively. These we find to be 45 μs and 140 μ seconds respectively.

The total decay time for atoms in this spin echo experiment was found to be 34 μs, which is significantly shorter than 2.1 ms for the decay time of atoms in the Rabi scan experiment described earlier in section 9.3.2. This difference in the overall Rabi

Experiment	Decay Type	Timescale
Ramsey interferometry	Non-radiative	4.9±0.1 ms
Rabi oscillations	Total	2.1±0.2 ms
Spin echo	Non-Stochastic	140 μs
Spin echo	Stochastic	45 μs

TABLE 9.2: A summary of the key dephasing mechanisms, and the experiment used to identify them. It is difficult to draw conclusions from the spin echo results as a result of the higher than average level of dephasing across the cloud.

frequency is significant as it shows us that while a combination of stochastic and non-stochastic decay mechanisms contribute to dephasing in fast decaying clouds such as that shown above, the information cannot be used to infer the relative decay constants in clouds with a longer decay time such as figure 9.6. The reason for the difference in the two decay rates is likely to be trivial; small changes in the alignment can increase the temperature by a factor of ten, while changes in the orientation of the beams can significantly increase the drift in beam pointing, both of which can significantly affect the coherence times of the cloud. Increasing the time spent on properly aligning the beams can give significant improvements in the coherence time of the atoms, and a longer decay constants.

9.5 Summary of decoherence mechanisms

We have been able to infer, through Ramsey interferometry and spin echo experiments, that the decoherence of the system is a result of radiative, non-stochastic processes. As a result of this it is likely that the decoherence in our system is a result of the non zero initial size of the cloud, which is driven non-homogeneously by a Raman beam with a non-infinite beam diameter. This can cause the generalised Rabi frequency to vary either as a result of a change in $\Omega = d_{ij}E^2/\hbar$, due to the change in the electric field strength of the light, or as a result of a change the light-shift which changes the two-photon detuning. As these processes are both results of the electric field of the light acting on the atom, there is no way to separate them in order to measure their effects separately.

9.5.1 Eliminating the decay in the Rabi cycle

There are a number of solutions to address the problems of non uniform manipulation across the atomic distribution. Increasing the diameter of the driving laser beam or decreasing that of the MOT density profile both cause an increase in the uniformity of the driving field across the cloud, increasing the homogeneity of the interaction.

Alternatively, we can apply a spatially selective readout method which can be used to only read out atoms which have been manipulated homogeneously.

9.5.1.1 Illuminating the atoms uniformly

Increasing the diameter of the driving laser beam for a given total intensity causes a smaller gradient in the optical electric field. This homogenises the atom-light interactions, minimising variations in the Rabi frequency across the cloud. In expanding the beam, however, the power in the beam is spread over a greater area, greatly diminishing the power left in the part of the beam used to manipulate the atoms. This causes a slower average Rabi frequency, reducing the number of operations that can be made to the atoms before they fall from the manipulation region.

Alternatively, diffractive optics can be used to change the beam intensity profile into a top hat, giving a constant intensity which can be used to illuminate the MOT cloud, with a higher power than is obtained through increasing the beam diameter. The diameter of the top hat is an important variable as any atoms moving over the boundary see an average of the total beam intensity, which smooths the step function causing a variation in the Rabi frequency in similar manner to the effects of a Gaussian intensity distribution. Is therefore important that the diameter of the beam is larger that the time of flight of the atoms during the manipulation stage. Using these values to calculate the rms time of flight for a 100μK cloud gives 0.2ms^{-1}, or 200 μm over 1 ms. As a result of this, a 600μm beam radius, three times the rms value, would give constant illumination for 99.6% of atoms within the velocity distribution.

In order to maximise the power in the top-hat, phase masks [123] are used which apply a phase pattern to the beam causing a far field top hat beam profile without sacrificing any power of the beam. Devices such as spatial light modulators (SLMs) can also be used to perform this transformation, which also enable changes to be made to the beam profile at refresh frequencies of 60Hz. The damage threshold of these devices is not sufficient to avoid damage from the Raman beams, and the high diffraction efficiency outweighs any advantage given by the more uniform beam intensity.

9.5.1.2 Spatially selective readout

Our readout is performed using the MOT trapping laser, which gives a large trapping region which encompasses all atoms in the cloud. If we engineer the readout beam to be narrow in comparison to the diameter of the manipulation laser, and aligned with the propagation axis of the manipulation beam, it is possible to selectively read-out

atoms which are at the centre of the Raman beam during manipulation. There are some potential difficulties with this technique such as the need to use a separate readout beam from the multi-directioned trapping beams. Due to the smaller volume of the trapping beam, it is likely to cause a decrease in the overall signal.

Reetz-Lamour *et al.* have described a two-photon readout scheme which can be used to give a 3 dimensional readout region [124]. Such a scheme could be easily applied to the rubidium $5D_{5/2}$ state [49].

Chapter 10

Conclusions

This thesis reports the results of investigations into the coherent manipulation of Rubidium 85 using a freely expanding cloud of rubidium atoms after laser cooling and trapping by a magneto-optical trap. Successful characterisation of this phenomenon is important to facilitate the generation of π and $\frac{\pi}{2}$ pulses for use in coherent, optical cooling schemes in clouds of untrapped atoms.

The Rabi frequency was measured to be 10^2 times smaller than was calculated from the measured light intensity. While this can be partially attributed to the electro-optical modulation method of generating the two photon transition frequencies, specifically in the generation of additional unused frequency components, it is unlikely to account for it entirely. A discrepancy of around three orders of magnitude was measured between the light shift of the $F = 2$ and $F = 3$ hyperfine states and that calculated from the measured light intensity. There are two possible explainations for this: either we have underestimated the loss in power contained within the Raman manipulation beam, through processes such as scattering of light from the windows, or conversion efficiency of our modulators. Alternatively, the beam is off resonant with the two photon transition frequency, which may occur due to a light shift in the atomic transition frequency during manipulation. These causes are currently under investigation: we plan to use large bandwidth measurements of the spectral density of the Raman beams, before and after amplification to directly measure the resonant power in the Raman beams.

The reduction in the Rabi frequency inhibits the application of complex pulse sequences, applying a limit to the number of operations which can be applied to atoms in free fall. Before more complex manipulation schemes are attempted, it is imperative that the Rabi frequency is increased.

An increase in the Rabi frequency has been achieved through the elimination of closely separated frequency components which caused gain oscillations in the amplification stages. This was achieved by Mach-Zehnder interferometry, locked by a Hänsch-Couillard style polarisation technique. In addition to this, the coupling of ASE into downstream amplifiers was reduced, further increasing the proportion of resonant light in the Raman beam.

Future separation of the EOM and AOM beams before amplification will eliminate frequency mixing through gain oscillation which will allow increases in the Rabi frequency through smaller detuning without single photon absorption, and a greater ratio of power contained within the Raman frequency components to total beam power.

In addition to the slow rate of the Rabi oscillation, a decay in the visibility of the Rabi fringes was observed, with a decay constant that appears related to the Rabi period. This limits the fidelity of sequences containing multiple π and $\frac{\pi}{2}$ pulses.

A Ramsey interferometry pulse scheme was applied to our atom cloud, and an oscillation in the excited state population as a function of the time interval between two $\frac{\pi}{2}$ pulses was observed. The long decay constant for the decay in the contrast of these fringes was used to infer that dephasing must be primarily a result of the illuminating light, rather than non-radiative collisional or thermal based effects. Similarly, spin echo experiments, in which the phase of the driving Raman light is reversed mid-way through a Rabi scan experiment to restore any non-stochastic forms of decoherence showed signs of the revival of Rabi oscillations, demonstrated a stochastic component to the decoherence. However, as the spin echo experiment was performed on Rabi graphs which were displaying a higher than average amount of decoherence, it was not possible to use this data to comment on the exact proportion of stochastic to reversible dephasing decay rates in slower decaying clouds. As a result of this, it is imperative that we re-run the spin echo experiment with an optimised system exhibiting a low dephasing rate. If the dephasing is a result of non-stochastic mechanisms, which cause reversible decay in the spin echo experiment, we will know that it will be possible to achieve population inversion with adiabatic rapid passage. On the other hand, stochastic mechanisms, represent very different problems which will make it necessary for us to look more closely at the dynamic aspects of our atom cloud; thermal motion of the atoms and beam pointing issues from air currents.

The most likely candidate for non-stochastic dephasing is due to inhomogeneous variation of the Rabi frequency across the cloud. This occurs due to an intensity variations across the cloud, which causes changes to the generalised Rabi frequency as a result of both the light shift from resonance, and the change in the Rabi frequency. By reducing the size and temperature of the atom cloud, increasing the size of the Raman beam

waist and better alignment between the beam and cloud, to achieve lower cloud temperatures, significant improvements in the fidelity of the Rabi oscillation have been, and will continue to be seen. Despite these improvements, the minimum decay constant seen during Rabi oscillation is still of the order of 1-2 times the Rabi period, which inhibits the application of complex pulse schemes.

Significant changes in the momentum of a room temperature atom of molecule require large numbers of coherent π pulse interactions. At present, our π and $\frac{\pi}{2}$ pulses are too long, and with insufficient fidelity to noticably change the momentum distribution of a cloud, which makes them ineffective in cooling using the schemes that were mentioned earlier in this thesis. For the cooling of molecules, where the decay time from the excited state is faster than in our proof-of-principle rubidium system, short pulse lengths are of even greater improtance.

Previous work in this field has been concerned with the manipulation of atoms which, as a result of dipole trapping, have a narrow spatial distribution. For this reason, we have considered adapting our system so that the atoms are loaded into a dipole trap before manipulation. This allows a smaller cloud diameter to be used for manipulation, which would reduce the variation in the driving Raman light during the experiment, and longer times in which to carry out the operations. Increasing the power in the beams would alternatively allow us to use a bigger Raman beam waist, which would have a more gentle spatial variation. Using separate amplifier stages for the pump and stokes beams would bring about a change in the optical power by a factor of two, although more significant increases in the power are more difficult to achieve. This large increase in power could be achieved through a pulsed rather than continuous laser system although adapting our method of electro-optically generating the Raman frequency components for a pulsed carrier beam, would present challenges due to the difficulty in achieving temporal overlap and compensating for dispersive effects.

Bibliography

[1] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. Cold bosonic atoms in optical lattices. *Phys. Rev. Lett.*, 81(15):3108–3111, Oct 1998.

[2] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hänsch, and Immanuel Bloch. Quantum phase transitions from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.

[3] M.W. Zwierlein, J. R Abo-Shaeer, A. Schirotzek, C.H Schunck, and W. Ketterle. Thermodynamic observation of first-order vortex-latice melting transition. *Nature*, 375(6530):373–376.

[4] Immanuel Bloch. Quantum gases. *Science*, 319(5867):1202–1203, 2008.

[5] T.V. Tscherbul and R.V. Krems. Manipulating spin-dependent interactions in rotationally excited cold molecules with electric fields. *The Jourrnal of Chemical Physics*, 125(194311), 2006.

[6] Myung Hwa Kim, Brian D. Leskiw, Lei Shen, and Arthur G. Suits. Velocity map imaging mass spectrometry. *International Journal of Mass Spectrometry*, 252(1): 73 – 78, 2006.

[7] S. Knoop, F. Ferlaino, M. Berninger, M. Mark, H.-C. Nägerl, R. Grimm, J. P. D'Incao, and B. D. Esry. Magnetically controlled exchange process in an ultracold atom-dimer mixture. *Phys. Rev. Lett.*, 104(5):053201, Feb 2010.

[8] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller. Entanglement of atoms via cold controlled collisions. *Phys. Rev. Lett.*, 82(9):1975–1978, Mar 1999.

[9] N. P. Robins, C. Figl, S. A. Haine, A. K. Morrison, M. Jeppesen, J. J. Hope, and J. D. Close. Achieving peak brightness in an atom laser. *Phys. Rev. Lett.*, 96(14): 140403, Apr 2006.

[10] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. Observation of interference between two Bose condensates. *Science*, 275(5300):637–641, 1997.

[11] M. Fattori, C. D'Errico, G. Roati, M. Zaccanti, M. Jona-Lasinio, M. Modugno, M. Inguscio, and G. Modugno. Atom interferometry with a weakly interacting Bose-Einstein condensate. *Phys. Rev. Lett.*, 100(8):080405, Feb 2008.

[12] A Peters, K Y Chung, and S Chu. High-precision gravity measurements using atom interferometry. *Metrologia*, 38(1):25, 2001.

[13] G. Nogues, A. Rauschenbeutel, S. Osnaghi, M. Brune, J.M. Raimond, and S. Haroche. Seeing a single photon without destroying it. *Nature*, 400(6741): 239–242, 1999.

[14] Alan Lenef, Troy D. Hammond, Edward T. Smith, Michael S. Chapman, Richard A. Rubenstein, and David E. Pritchard. Rotation sensing with an atom interferometer. *Phys. Rev. Lett.*, 78(5):760–763, Feb 1997.

[15] Andreas Wicht, Joel M Hensley, Edina Sarajlic, and Steven Chu. A preliminary measurement of the fine structure constant based on atom interferometry. *Physica Scripta*, 2002(T102):82, 2002.

[16] B. C. Regan, Eugene D. Commins, Christian J. Schmidt, and David DeMille. New limit on the electron electric dipole moment. *Phys. Rev. Lett.*, 88(7):071805, Feb 2002.

[17] Rainer Blatt and David Wineland. Entangled states of trapped atomic ions. *Nature*, 453:1008–1015.

[18] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Neutral atom quantum register. *Phys. Rev. Lett.*, 93 (15):150501, Oct 2004.

[19] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85(10):2208–2211, Sep 2000.

[20] Judah Levine. Introduction to time and frequency metrology. *Review of Scientific Instruments*, 70(6):2567–2596, 1999.

[21] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Measurement of the electron electric dipole moment using YbF molecules. *Phys. Rev. Lett.*, 89(2): 023003, Jun 2002.

[22] M. G. Kozlov and D. DeMille. Enhancement of the electric dipole moment of the electron in PbO. *Phys. Rev. Lett.*, 89(13):133001, Sep 2002.

[23] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, S. Riedl, C. Chin, J. Hecker Denschlag, and R. Grimm. Bose-Einstein condensation of molecules. *Science*, 302(5653):2101–2103, 2003.

[24] Johann G. Danzl, Manfred J. Mark, Elmar Haller, Mattias Gustavsson abd Russel Hart, Jesus Aldegunde, Jeremy M. Hidson, and Hanns-Christoph Nägerl. An ultra-cold high density sample of ro-vibronic ground state moleules in an optical lattice. *Nature Physics*, 6:265–270, 2010.

[25] Kevin E. Strecker, Guthrie B. Partridge, and Randall G. Hulet. Conversion of an atomic fermi gas to a long-lived molecular Bose gas. *Phys. Rev. Lett.*, 91(8): 080406, Aug 2003.

[26] A. Fioretti, D. Comparat, A. Crubellier, O. Dulieu, F. Masnou-Seeuws, and P. Pillet. Formation of cold cs_2 molecules through photoassociation. *Phys. Rev. Lett.*, 80(20):4402–4405, May 1998.

[27] Thorsten Köhler, Krzysztof Góral, and Paul S. Julienne. Production of cold molecules via magnetically tunable feshbach resonances. *Rev. Mod. Phys.*, 78 (4):1311–1361, Dec 2006.

[28] M. Mudrich, S. Kraft, K. Singer, R. Grimm, A. Mosk, and M. Weidemüller. Sympathetic cooling with two atomic species in an optical trap. *Phys. Rev. Lett.*, 88 (25):253001, Jun 2002.

[29] John M. Doyle, Bretislav Friedrich, Jinha Kim, and David Patterson. Buffer-gas loading of atoms and molecules into a magnetic trap. *Phys. Rev. A*, 52(4): R2515–R2518, Oct 1995.

[30] E. Vliegen and F. Merkt. Stark decleration of hydrogen atoms. *Journal of Physics B*, (39), 2006.

[31] Tim Freegarde, Geoff Daniell, and Danny Segal. Coherent amplification in laser cooling and trapping. *Phys. Rev. A*, 73(3):033409, Mar 2006.

[32] M. Weitz and T. W. Hänsch. Frequency-independent laser cooling based on interferometry. *EPL (Europhysics Letters)*, 49(3):302, 2000.

[33] A. Lengwenus, J. Kruse, M. Volk, W. Ertmer, and G. Birkl. Coherent manipulation of atomic qubits in optical micropotentials. *Applied Physics B: Lasers and Optics*, 86:377–383, 2007.

[34] M.A. Nielsen and I.L. Chuang. *Quantum computation and quantum information.* Cambridge University Press, 2000.

[35] E A Hinds and I G Hughes. Magnetic atom optics: mirrors, guides, traps, and chips for atoms. *Journal of Physics D: Applied Physics*, 32(18):R119, 1999.

[36] Ian Fleming. *Goldfinger.* Jonathan Cape, 1959.

[37] T.A. Littlefield and N. Thorley. *Atomic and Nuclear Physics: 3rd edition.* Van Nostrad Reinhold Co. Ltd, 1979.

[38] Ayan Banerjee, Dipankar Das, and Vasant Natarajan. Precise fine-structure and hyperfine-structure measurements in Rb. arXiv:physics/0209019v1, 2002.

[39] Peter Atkins and Ronald Freidman. *Molecular Quantum Mechanics.* Oxford University Press, 2007.

[40] C.V. Raman and S. Bhagavantam. Experimental proof of the spin of the photon. *Indian J. Phys.*, 6:353–366, 2010.

[41] Mark Fox. *Quantum Optics.* Oxford University Press, 2006.

[42] David Halliday, Robert Resnick, and Jearl Walker. *Fundamentals of physics (Sixth edition: extended).* John Wiley and Sons, Inc., 2001.

[43] Philipp Salzgeber. Comet hale bopp. March 1997.

[44] Christopher. J. Foot. *Atomic Physics.* Oxford University Press, 2005.

[45] A. Ashkin. Acceleration and trapping of particles by radiation pressure. *Phys. Rev. Lett.*, 24(4):156–159, Jan 1970.

[46] A. Einstein. On the quantum theory of radiation. *Classics of science*, 1917.

[47] T.V. Zueva and V.G Minogin. *Sov. Tech. Phys. Lett.*, 7:411, 1981.

[48] Steven Chu. Laser cooling of neutral atoms. *Opt. Photon. News*, 1(12):40–43, Dec 1990.

[49] Hamid Ohadi, Matthew Himsworth, André Xuereb, and Tim Freegarde. Magneto-optical trapping and background-free imaging for atoms near nanostructured surfaces. *Opt. Express*, 17(25):23003–23009, Dec 2009.

[50] Alan L. Migdall, John V. Prodan, William D. Phillips, Thomas H. Bergeman, and Harold J. Metcalf. First observation of magnetically trapped neutral atoms. *Phys. Rev. Lett.*, 54(24):2596–2599, Jun 1985.

[51] Yunnus A. Cengel and Michael A. Boles. *Thermodynamics.* McGraw-Hill, 1998.

[52] Harold Metcalf. Entropy exchange in laser cooling. *Phys. Rev. A*, 77(6):061401, Jun 2008.

[53] José W. R. Tabosa, Sandra S. Vianna, and Clayton A. Benevides. Loading mechanism in a two-beam magneto-optical trap. *Optics Communications*, 116(1-3):77 – 80, 1995.

[54] D. Sesko, T. Walker, C. Monroe, A. Gallagher, and C. Wieman. Collisional losses from a light-force atom trap. *Phys. Rev. Lett.*, 63(9):961–964, Aug 1989.

[55] D. W. Sesko, T. G. Walker, and C. E. Wieman. Behavior of neutral atoms in a spontaneous force trap. *J. Opt. Soc. Am. B*, 8(5):946–958, May 1991.

[56] A. M. Steane, M. Chowdhury, and C. J. Foot. Radiation force in the magneto-optical trap. *J. Opt. Soc. Am. B*, 9(12):2142–2158, Dec 1992.

[57] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*, volume 62. Springer, January 1999.

[58] C. D. Wallace, T. P. Dinneen, K. Y. N. Tan, A. Kumarakrishnan, P. L. Gould, and J. Javanainen. Measurements of temperature and spring constant in a magneto-optical trap. *J. Opt. Soc. Am. B*, 11(5):703–711, May 1994.

[59] Colin Webb and Julian Jones. *Laser technology and Applications Volume III:Applications*. Institute of Physics, 2004.

[60] Y. Castin, H. Wallis, and J. Dalibard. Limit of doppler cooling. *J. Opt. Soc. Am. B*, 6(11):2046–2057, Nov 1989.

[61] D. J. Wineland, J. Dalibard, and C. Cohen-Tannoudji. Sisyphus cooling of a bound atom. *J. Opt. Soc. Am. B*, 9(1):32–42, Jan 1992.

[62] Paul D. Lett, Richard N. Watts, Christoph I. Westbrook, William D. Phillips, Phillip L. Gould, and Harold J. Metcalf. Observation of atoms laser cooled below the doppler limit. *Phys. Rev. Lett.*, 61(2):169–172, Jul 1988.

[63] A. M. Steane and C. J. Foot. Laser cooling below the doppler limit in a magneto-optical trap. *EPL (Europhysics Letters)*, 14(3):231, 1991.

[64] J. Reichel, O. Morice, G. M. Tino, and C. Salomon. Subrecoil raman cooling of cesium atoms. *EPL (Europhysics Letters)*, 28(7):477, 1994.

[65] D. Guery-Odelin, J. Soeding, P. Desbiolles, and Jean Dalibard. Strong evaporative cooling of a trapped cesium gas. *Opt. Express*, 2(8):323–329, Apr 1998.

[66] Angus Henderson, Ryan Stafford, and Paul Hoffman. Tunable Sources: High Power CW OPOs span the spectrum. *Laser Focus World*, 2009.

[67] F Y Loo, A Brusch, S Sauge, M Allegrini, E Arimondo, N Andersen, and J W Thomsen. Investigations of a two-level atom in a magneto-optical trap using magnesium. *Journal of Optics B: Quantum and Semiclassical Optics*, 6(1):81, 2004.

[68] Benjamin K. Stuhl, Brian C. Sawyer, Dajun Wang, and Jun Ye. Magneto-optical trap for polar molecules. *Phys. Rev. Lett.*, 101(24):243002, Dec 2008.

[69] M. D. Di Rosa. Laser-cooling molecules. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 31:395–402, 2004.

[70] Jonathan D. Weinstein, Robert deCarvalho, Thierry Guillet, Bretislav Friedrich, and John M. Doyle. Magnetic trapping of calcium monohydride molecules at millikelvin temperatures. *Nature*, (6698):148–150, 1998.

[71] Benjamin L. Lev, András Vukics, Eric R. Hudson, Brian C. Sawyer, Peter Domokos, Helmut Ritsch, and Jun Ye. Prospects for the cavity-assisted laser cooling of molecules. *Phys. Rev. A*, 77(2):023402, Feb 2008.

[72] P. Maunz, T. Puppe, I.Schuster, N.Syassen, P.W.H. Pinkse, and G. Rempe. Cavity cooling of a single atom. *Nature*, 428(6978):50–52, 2004.

[73] Peter Horak, André Xuereb, and Tim Freegarde. Mirror-mediated cooling of a particle by coupling to its own reflection. *Submitted for publication*, 2009.

[74] D. Kielpinski. Laser cooling of atoms and molecules with ultrafast pulses. *Phys. Rev. A*, 73(6):063407, Jun 2006.

[75] Hilton W. Chan, Adam T. Black, and Vladan Vuletić. Observation of collective-emission-induced cooling of atoms in an optical cavity. *Phys. Rev. Lett.*, 90(6): 063003, Feb 2003.

[76] Dieter Meschede. *Optics, light and lasers.* Wiley-VCH, 2004.

[77] M. Orszag. *Quantum Optics.* Springer, 2010.

[78] Bruce W. Shore. *The theory of coherent atomic excitation: Volume 1.* Wiley, 1990.

[79] Leonard Mandel and Emil Wolf. *Optical Coherence and Quantum Optics.* Cambridge University Press, 1995.

[80] Richard P. Feynman, Jr. Frank L. Vernon, and Robert W. Hellwarth. Geometrical representation of the schrödinger equation for solving maser problems. *Journal of Applied Physics*, 28(1):49–52, 1957.

[81] M Reetz-Lamour, J Deiglmayr, T Amthor, and M Weidemüller. Rabi oscillations between ground and rydberg states and van der waals blockade in a mesoscopic frozen rydberg gas. *New Journal of Physics*, 10(4):045026, 2008.

[82] M. P. A. Jones, J. Beugnon, A. Gaëtan, J. Zhang, G. Messin, A. Browaeys, and P. Grangier. Fast quantum state control of a single trapped neutral atom. *Phys. Rev. A*, 75(4):040301, Apr 2007.

[83] C. N. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics*. Wiley, 1999.

[84] V.S. Malinovsky and J.L. Krause. General theory of population transfer by adiabatic rapid passage with intense, chirped laser pulses. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 14:147–155, 2001. ISSN 1434-6060. 10.1007/s100530170212.

[85] James Bateman and Tim Freegarde. Fractional adiabatic passage in two-level systems: Mirrors and beam splitters for atomic interferometry. *Phys. Rev. A*, 76 (1):013416, Jul 2007.

[86] James Bateman, André Xuereb, and Tim Freegarde. Stimulated raman transitions via multiple atomic levels. *Phys. Rev. A*, 81(4):043808, Apr 2010.

[87] R. Unanyan, M. Fleischhauer, B. W. Shore, and K. Bergmann. Robust creation and phase-sensitive probing of superposition states via stimulated Raman adiabatic passage (STIRAP) with degenerate dark states. *Optics Communications*, 155(1-3):144 – 154, 1998.

[88] T. G. M. Freegarde, J. Walz, and T. W. Hänsch. Confinement and manipulation of atoms using short laser pulses. *Optics Communications*, 117(3-4):262 – 267, 1995.

[89] A. Goepfert, I. Bloch, D. Haubrich, F. Lison, R. Schütze, R. Wynands, and D. Meschede. Stimulated focusing and deflection of an atomic beam using picosecond laser pulses. *Phys. Rev. A*, 56(5):R3354–R3357, Nov 1997.

[90] Vladan Vuletić, Andrew J. Kerman, Cheng Chin, and Steven Chu. Observation of low-field feshbach resonances in collisions of cesium atoms. *Phys. Rev. Lett.*, 82 (7):1406–1409, Feb 1999.

[91] Tim Freegarde and Danny Segal. Algorithmic cooling in a momentum state quantum computer. *Phys. Rev. Lett.*, 91(3):037904, Jul 2003.

[92] M C Fischer, A M Dudarev, B Gutiérrez-Medina, and M G Raizen. Fm spectroscopy in recoil-induced resonances. *Journal of Optics B: Quantum and Semiclassical Optics*, 3(4):279, 2001.

[93] Martin Weitz, Brenton C. Young, and Steven Chu. Atomic interferometer based on adiabatic population transfer. *Phys. Rev. Lett.*, 73(19):2563–2566, Nov 1994.

[94] Jianshu Cao, Christopher J. Bardeen, and Kent R. Wilson. Molecular " π pulse" for total inversion of electronic state population. *Phys. Rev. Lett.*, 80(7):1406–1409, Feb 1998.

[95] Steck Daniel A. Rubidium 85 D line data. available online at http://steck.us/alkalidata (revision 0.1.1, 2nd May 2008), 2009.

[96] D.M.B.P. Milori, M.T. de Araujo, I.Guedes, S.C. Zilio, and V.S. Bagnato. Analysis of a vapour-cell magneto-optical trap for sodium atoms based on a two-level doppler theory. *Brazilian Journal of Physics*, 27(2), 1997.

[97] Sunil Patel. *A Chirped, Pulsed Laser System and Magneto-Optical Trap for Rubidium.* PhD thesis, School of Physics and Astronomy, 2009.

[98] C. Klempt, T. van Zoest, T. Henninger, O. Topic, E. Rasel, W. Ertmer, and J. Arlt. Ultraviolet light-induced atom desorption for large rubidium and potassium magneto-optical traps. *Phys. Rev. A*, 73(1):013410, Jan 2006.

[99] Michael G. Littman and Harold J. Metcalf. Spectrally narrow pulsed dye laser without beam expander. *Appl. Opt.*, 17(14):2224–2227, Jul 1978.

[100] Toptica. Tunable diode lasers. http://www.toptica.com/products/itemlayer/32/BR-101-029-C-ProductCatalog-2009-05-TDL.pdf, 2010.

[101] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard. Phase-space density in the magneto-optical trap. *Phys. Rev. A*, 52(2):1423–1440, Aug 1995.

[102] Eric D. Black. An introduction to pound–drever–hall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, 2001.

[103] Kristan L. Corwin, Zheng-Tian Lu, Carter F. Hand, Ryan J. Epstein, and Carl E. Wieman. Frequency-stabilizes diode laser with zeeman shift in an atomic vapour. *Applied Optics*, 37(15), 1998.

[104] Paul Horowitz and Winfield Hill. *The Art of Electronics.* Cambridge University Press, 1989.

[105] Matthew Himsworth. *Coherent Manipulation of Ultracold Rubidium.* PhD thesis, School of Physics and Astronomy, 2009.

[106] Jr. William J. Riley. *Handbook of frequency stability analysis.* Hamilton Technical Services, 2007.

[107] James Bateman. *Novel schemes for the optical manipulation of atoms and molecules*. PhD thesis, School of Physics and Astronomy, 2008.

[108] M. T. de Araujo, R. J. Horowicz, D. Milori, A. Tuboy, R. Kaiser, S. C. Zilio, and V. S. Bagnato. Time of flight technique to determine temperature of atoms in a ring-shaped trap. *Optics Communications*, 119(1-2):85 – 89, 1995.

[109] J.Y. Courtois, G. Grynberg, B. Lounis, and P. Verkerk. Recoil-induced resonances in cesium: An atomic analog to the free-electron laser. *Phys. Rev. Lett.*, 72(19): 3017–3020, May 1994.

[110] A. Clairon, C. Salomon, S. Guellati, and W. D. Phillips. Ramsey resonance in a zacharias fountain. *EPL (Europhysics Letters)*, 16(2):165, 1991.

[111] Roland K Appel, Chris D Dyer, Anita W Jones, and John N Lockwood. Enhanced raman spectroscopy using collection optics designed for continuously tunable excitation. *Measurement Science and Technology*, 13(3):411, 2002.

[112] Jürgen Appel, Andrew MacRae, and A I Lvovsky. A versatile digital ghz phase lock for external cavity diode lasers. *Measurement Science and Technology*, 20(5): 055302, 2009.

[113] P. J. Lee, B. B. Blinov, K. Brickman, L. Deslauriers, M. J. Madsen, R. Miller, D. L. Moehring, D. Stick, and C. Monroe. Atomic qubit manipulations with an electro-optic modulator. *Opt. Lett.*, 28(17):1582–1584, Sep 2003.

[114] Unknown Author. *Thorlabs Electro-optic Phase Modulators: Operating Manual*, 2010.

[115] I. P. Khristov, I. V. Tomov, and S. M. Saltiel. Self-heating effects in electro-optic light modulators. *Optical and Quantum Electronics*, 15:289–295, 1983. ISSN 0306-8919.

[116] Gabriele Ferrari, Marc-Oliver Mewes, Florian Schreck, and Christophe Salomon. High-power multiple-frequency narrow-linewidth laser source based on a semiconductor tapered amplifier. *Opt. Lett.*, 24(3):151–153, Feb 1999.

[117] I. Dotsenko, W. Alt, S. Kuhr, D. Schrader, M. Müller, Y. Miroshnychenko, V. Gomer, A. Rauschenbeutel, and D. Meschede. Application of electro-optically generated light fields for raman spectroscopy of trapped cesium atoms. *Applied Physics B: Lasers and Optics*, 78:711–717, 2004. ISSN 0946-2171.

[118] K.P. Zetie, S. F Adams, and R. M. Tocknell. How does a mach-zehnder interferometer work? *Teaching Physics*, 35(1), 2000.

[119] J. E. Bateman, R. L. D. Murray, M. Himsworth, H. Ohadi, A. Xuereb, and T. Freegarde. Hänsch–couillaud locking of mach–zehnder interferometer for carrier removal from a phase-modulated optical spectrum. *J. Opt. Soc. Am. B*, 27(8): 1530–1533, Aug 2010.

[120] T.W. Hänsch and B. Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. *Optics Communications*, 35(3):441 – 444, 1980.

[121] M. Gertsvolf and M. Rosenbluh". Injection locking of a diode laser locked to a zeeman frequency stabilized laser oscillator. *Optics Communications*, 170(4-6):269 – 274, 1999. ISSN 0030-4018.

[122] T. A. Savard, S. R. Granade, K. M. O'Hara, M. E. Gehm, and J. E. Thomas. Raman-induced magnetic resonance imaging of atoms in a magneto-optical trap. *Phys. Rev. A*, 60(6):4788–4795, Dec 1999.

[123] Gordon D. Love and David F. Buscher. Focal spot generator without sidelobes and its application in coronagraphy. *Journal of Modern Optics*, 51(14):2159–2165, 2007.

[124] M. Reetz-Lamour, T. Amthor, J. Deiglmayr, and M. Weidemüller. Rabi oscillations and excitation trapping in the coherent excitation of a mesoscopic frozen rydberg gas. *Phys. Rev. Lett.*, 100(25):253001, Jun 2008.