UNIVERSITY OF SOUTHAMPTON

Novel schemes for the optical manipulation of atoms and molecules

by

James Bateman

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

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

March 2009

UNIVERSITY OF SOUTHAMPTON

Abstract

Faculty of Engineering, Science and Mathematics

School of Physics and Astronomy

Doctor of Philosophy

Novel schemes for the optical manipulation of atoms and molecules

by James Bateman

The range of atoms which can be cooled by lasers is limited to those which have a closed two level structure. Several schemes have been proposed which aim to extend this range by using coherent control of the particle momenta, but none have yet been demonstrated. We hope to implement these and other coherent manipulation schemes, and we begin with a system which is well understood and over which we can exert precise control.

This thesis covers the design and construction of an experiment to demonstrate coherent manipulation of cold rubidium atoms collected in a magneto-optical trap. The lower hyperfine levels of these cold atoms very closely mimic the ideal two-level atom, and we use carefully crafted laser pulses to prepare, manipulate, and read their quantum state.

The hyperfine levels are coupled using two fields whose frequency difference is equal to the hyperfine splitting. The way in which these Raman coupled levels can be used to emulate a two-level atom is explored, and the experimental apparatus used to create and control the driving fields is described in detail. The amplitude, frequency and phase of these fields is programmable, and complex manipulation schemes can be implemented merely by programming a computer.

We have observed Raman transitions in the cold rubidium atoms, and the experimental methods used to detect these features amidst large experimental noise are discussed. Although we have not yet seen Rabi oscillations, we are confident that we can now have sufficient control to begin to implement simple interferometric sequences. However, there remain significant challenges if we are to coherently manipulate the momentum, and the prospects for such manipulation are discussed.

Contents

List of Figures

Declaration of Authorship

I, **James Bateman**, declare that this thesis titled **Novel schemes for the optical manipulation of atoms and molecules** and the work presented in it are my own. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;

- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

- where I have consulted the published work of others, this is always clearly attributed;

- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

- I have acknowledged all main sources of help;

- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

- parts of this work have been published as:

    – J. E. Bateman and T. G. M. Freegarde, Phys. Rev. A 76(1):013416 (2007)

Signed:
_____

Date:
_____

Acknowledgements

If you really want to understand something, the best way is to try and explain it to someone else. That forces you to sort it out in your mind. And the more slow and dim-witted your pupil, the more you have to break things down into more and more simple ideas. And that's really the essence of programming. By the time you've sorted out a complicated idea into little steps that even a stupid machine can deal with, you've learned something about it yourself.

Dirk Gently's Holistic Detective Agency (1987), Douglas Adams

Chapter 1

Introduction

Cooling [1–6] and trapping [7–10] of atoms with light provides us with extremely cold samples far denser than possible by any other means and enables us to study and control these systems in unprecedented detail [11, 12]. We can not only learn about these systems, but can devise experiments of such precision that we test the underlying physics. These techniques have revolutionised spectroscopy [13], metrology [14, 15], and have allowed us to explore states of matter that have never previously existed [16, 17]. It is unfortunate therefore that the most effective techniques—Doppler cooling and the Magneto Optical Trap—rely on a simple energy level structure found only in a handful of elements.

1.1   Conventional Cooling

To cool a distribution requires dissipation, for without this irreversible decay the entropy of a system cannot decrease. For Doppler cooling, this dissipation is provided by spontaneous emission of a photon by an atom in an excited state. The atom is placed in this excited state by absorption of laser light, and the frequency of the light is carefully tuned so that absorbed photons are likely to carry momentum which reduce the atom's speed. The cycle of absorption and re-emission must repeat many times to significantly slow the atom, so we require that after spontaneous emission the internal state of the

atom returns to the state it was in before the initial absorption. The cycle must be closed, and it is this requirement that most elements and all molecules fail to meet.

In fact, real atoms can only ever approximately meet this criterion. Additional levels into which the atom can decay do exist, but if this occurs rarely the atom can be returned to the 'closed' transition using a repumping laser, and cooling can still take place. Indeed, laser cooled atoms and ions may require many repumping fields, and we may wonder, by extension, whether the many possible routes in other atoms and molecules could, in principle, be repumped. Unfortunately, as repumping unavoidably introduces heating the prospects for molecules are bleak [18]. If we are to study cold molecules with anything like the precision which has become commonplace for atoms, we must look beyond techniques which require a closed optical transition.

It is possible, using photoassociation [19–21] or Feshbach resonances [22, 23], to construct cold molecules from cold atoms. However, considerable effort is needed to explore new molecules and the range of molecules we can investigate is restricted to those which we can construct from conventionally cooled atoms. A cold molecular sample can also be obtained by filtering a much hotter distribution. For example, a quadrupole electric field guide with a bend will guide only those molecules with less than a certain speed; faster moving molecules will not be guided around the corner and will be discarded [24]. Unfortunately, this simple technique cannot approach the temperatures and densities routinely found in a MOT.

1.2   Coherent Control

To extend optical cooling techniques we must do more than apply well tuned and stabilised lasers to a sample. With modern lasers and electronics we can envisage applying engineered sequences of pulses to influence the particle dynamics in otherwise impossible ways. These light pulses carry momentum, and if their duration is less than the decay time of the excited state, we can coherently manipulate the momentum of the particles. It is this idea that the experiments described in this thesis were designed to explore.

One scheme which exploits the coherence of particles is amplified cooling [25]. During the time between absorption and spontaneous emission (in, for example, Doppler cooling) a particle is subjected to a sequence of π pulses from alternating directions. Each pulse imparts a photon momentum, but the direction depends on whether the particle is in the ground or the excited state. In the ground state, the particle absorbs the photon and its momentum, but in the excited state, the pulse stimulates emission of a photon, and the overall recoil is opposite to the momentum of the pulse. The internal state of the atom changes after each pulse, and if we send consecutive pulses from opposite directions, each pulse imparts momentum in the same direction, and this direction depends on whether the particle was *initially* in the ground or the excited state – the sequence can be seen as amplifying the effect of an initial absorption. This greatly increases the momentum imparted per spontaneous emission event, increasing the cooling effect relative to other heating effects, and so, in principle, makes it possible to cool a wider range of species.

A second scheme which relies even less on a closed level scheme is interferometric cooling [26]. Here, the velocity dependent absorption, which is due to red-detuned excitation in Doppler cooling, is replaced by a Ramsey interferometer sensitive to kinetic energy. A $\pi/2$ 'beam-splitter' pulse couples a lower state to an upper state and, after a delay, a second pulse recombines the two paths. The overall absorption of a photon now depends on the accrued phase difference and, since in one path the particle has one extra photon momentum, this phase depends on the *initial* particle momentum. Time delays and phases can be chosen to ensure absorption is such that the particle is slowed and, crucially, if we use broadband beam-splitter pulses, it does not matter between which states in an upper and lower manifold that the transition is driven.

Interferometric cooling is undoubtedly very insightful, but if the pulse scheme it uses is seen as an example of a more general sequence of quantum computation-like operations, it is clearly a very simple sequence and is not necessarily the best. With modern technology, and specifically computer controlled laser modulation, we are not restricted to such simple schemes; by using the language of quantum computation we can envisage and analyse more complicated schemes [27] which, since there can be a clear correspondence between computational gates and laser pulses, can be implemented simply by programming signal generators.

1.3   Thesis Overview

The experiments described in this thesis were designed to demonstrate coherent manipulation of atomic momentum, including these coherent cooling schemes. We begin with cold atoms and use Raman transitions between the long-lived hyperfine groundstates to emulate the canonical two-level system. We intend to use counter-propagating Raman beams so that the atom absorbs the difference in photon energy but the (scalar) sum of the two photon momenta. Many of the steps towards this goal have been achieved, but we are yet to demonstrate manipulation of the momentum. The necessary improvements to the experiment are outlined.

The thesis is composed of four main parts. In Chapters 2 and 3 we review the basic theory behind a magneto-optical trap, and then show how the two hyperfine groundstates of rubidium 85 can be used to emulate a two level system; in Chapter 4 we explore how we can coherently manipulate these states with imperfect experimental apparatus.

In Chapter 5 we begin discussion of the experiment with the lasers, spectroscopy, magnetic fields, and vacuum systems necessary to create our magneto-optical trap. Next, in Chapter 6, we describe the lasers and modulators used to create the fields which drive the Raman transitions.

The measured characteristics of the MOT cloud are reported in Chapter 7, and experiments which manipulate the hyperfine levels of these cold atoms are presented in Chapter 8.

Finally, in Chapter 9, we summarise what has been achieved and outline the improvements we must make to the apparatus if the experiment is to progress. We end with a reassessment of the reasoning which led us to this new coherent manipulation approach to cooling, and I mention an intriguing alternative technique.

Chapter 2

Magneto Optical Trap Theory

The magneto-optical trap, which has become the staple of cold atom physics, is neither a pure radiation trap (such as a dipole trap) nor an electric or magnetic trap (such as a Paul or Penning ion trap), but a hybrid; it uses the scattering force due to near-resonant red-detuned light for Doppler cooling and imposes a position dependence via the Zeeman shift due to a spatially-dependent magnetic field. A MOT was first demonstrated in 1987 [8], and the workings have been described in detail elsewhere [7, 11, 28], particularly by Metcalf and van der Straten [29]. An understanding of the underlying mechanisms, and specifically an understanding of some of the subtleties which give rise to non-ideal MOT behaviour, is so crucial to the experiments which follow that it is important to present a description of it here.

We first show how we can describe a quantum mechanical system which includes decay, and we then describe the statistical properties of light emitted by such a system. Using these tools, we derive the Optical Bloch Equations which provide some general results which we can use to describe the behaviour of atoms in a magneto-optical trap.

2.1 The Density Matrix

Quantum mechanics is a remarkably powerful description of nature, but the simplistic picture of systems existing in pure eigenstates and superpositions thereof is not sufficient to describe all the properties that we can observe. For example, a large ensemble of quantum mechanical particles need not all exist in the same state, but using wave functions we are unable to calculate an average, and hence predict the outcome of a measurement, without performing a calculation for each particle. A similarly computationally intractable problem is to account for all the mechanisms by which a system in an excited state may decay. In the case of an atom, this includes spontaneous emission of a photon even in the absence of any applied field.

The density matrix allows us to treat an ensemble or statistical average and quantum mechanical probability amplitudes in the same mathematical object. If we begin with the conventional expansion of a quantum mechanical state in terms of some basis set, but include an additional parameter ε which labels each of the possible wave functions

$$|\psi(\varepsilon)\rangle = \sum_n c_n(\varepsilon) |n\rangle \tag{2.1}$$

we can modify the equation for the expectation value of an operator $\hat{A}$ to include an average over the ensemble using $p(\varepsilon)$, the classical probability the system is in state ε:

$$\begin{aligned}
\left\langle \hat{A} \right\rangle &= \int p(\varepsilon) \langle\psi(\varepsilon)| \hat{A} |\psi(\varepsilon)\rangle \, \mathrm{d}\varepsilon \\
&= \int p(\varepsilon) \sum_n c_n^*(\varepsilon) \langle n| \hat{A} \sum_m c_m(\varepsilon) |m\rangle \, \mathrm{d}\varepsilon \\
&= \int p(\varepsilon) \sum_{n,m} c_n^*(\varepsilon) c_m(\varepsilon) A_{n,m} \mathrm{d}\varepsilon \,.
\end{aligned} \tag{2.2}$$

The values $A_{n,m}$ are independent of ε, and we can re-write the expectation value of the operator $\hat{A}$ as the trace of a matrix $\boldsymbol{A}$ composed of elements $A_{n,m}$ multiplied by the density matrix $\boldsymbol{\rho}$ composed of elements $\rho_{n,m}$:

$$\left\langle \hat{A} \right\rangle = \mathrm{Tr}\left(\boldsymbol{A}\boldsymbol{\rho}\right) \text{ where } \rho_{n,m} = \int p(\varepsilon) c_n^*(\varepsilon) c_m(\varepsilon) \mathrm{d}\varepsilon. \tag{2.3}$$

This density matrix can represent the quantum mechanical and statistical properties of an ensemble, or of a system which includes some irreversible decay. Clear and thorough treatments can be found in Shore [30] and Cohen-Tannoudji [31].

2.1.1 Time Evolution

The time evolution of the density matrix can be deduced from the Schrödinger equation, and the result is known as the quantum Liouville equation:

$$i\frac{\partial}{\partial t}\boldsymbol{\rho}_{n,m} = [\boldsymbol{H}, \boldsymbol{\rho}]_{n,m} \ . \tag{2.4}$$

The overall Hamiltonian $\boldsymbol{H}$ is typically composed of a free atom part $\boldsymbol{H}_0$ and an interaction part $\boldsymbol{H}_I$. Throughout this thesis we use the semi-classical approach and ignore energy in the field. Also, for simplicity, we use the Hamiltonian to represent the *frequency*, and not the energy, of a state or coupling. Hence, the free Hamiltonian of the atom is

$$\boldsymbol{H}_0 = \sum_n \omega_n \left| n \right\rangle \left\langle n \right| \ . \tag{2.5}$$

where we have summed over each the states $\left| n \right\rangle$ which have Bohr frequency ω_n. We use matrix representations and Dirac's butterfly operator notation interchangeably, with the operator $\left| n \right\rangle \left\langle m \right|$ equivalent to the identity matrix with columns n and m interchanged.

For this Hamiltonian, each element $\rho_{n,m}$ of the density matrix oscillates at the frequency $\omega_{n,m} = \omega_n - \omega_m$. (For the real-valued diagonal elements, which represent populations, this frequency is zero.) To remove this time-dependence, we make the substitution

$$\tilde{\rho}_{n,m} = \rho_{n,m} e^{i\omega_{n,m}t} \ . \tag{2.6}$$

We can also represent this element-wise multiplication as a matrix transformation:

$$\tilde{\boldsymbol{\rho}} = \boldsymbol{T}\boldsymbol{\rho}\boldsymbol{T}^{-1} \tag{2.7}$$

where $T_{n,n} = e^{i\omega_n t}$ and $T_{n,m} = 0$ for $n \neq m$. We now differentiate the elements of $\tilde{\boldsymbol{\rho}}$ and use the Liouville equation to find the time evolution of this transformed density matrix:

$$
\begin{aligned}
\dot{\tilde{\rho}}_{n,m} &= i\ \omega_{n,m}\tilde{\rho}_{n,m} + \dot{\rho}_{n,m}e^{i\omega_{n,m}t} \\
&= i\left(\omega_{n,m}\tilde{\rho}_{n,m} - [\boldsymbol{H},\boldsymbol{\rho}]_{n,m}e^{i\omega_{n,m}t}\right) \\
&= i\left(\omega_{n,m}\tilde{\rho}_{n,m} - \left[\boldsymbol{T}[\boldsymbol{H},\boldsymbol{T}^{-1}\tilde{\rho}\boldsymbol{T}]\boldsymbol{T}^{-1}\right]_{n,m}\right) \\
&= i\left(\omega_{n,m}\tilde{\rho}_{n,m} - [\tilde{\boldsymbol{H}},\tilde{\boldsymbol{\rho}}]_{n,m}\right) \qquad (2.8)
\end{aligned}
$$

where $\tilde{\boldsymbol{H}} = \boldsymbol{THT}^{-1}$ is the similarly transformed Hamiltonian. This is a transformation to the interaction picture [30].

This transformation removes the inherent fast phase evolution of the off-diagonal density matrix elements. Also, if two states are coupled by an oscillating external field, the frequencies which remain after this transformation are the sum of and difference between the excitation and the transition frequencies. For near resonant excitation, we can use the rotating-wave approximation, which assumes the system cannot respond to the rapidly oscillating sum frequency, and keep only the slow difference frequency. This can greatly simplify the task of finding the density matrix time evolution.

In general, the interaction Hamiltonian $\boldsymbol{H}_I$ due to a number of externally applied fields can be described by

$$
\begin{aligned}
\boldsymbol{H}_I &= \sum_{n,m}\sum_{j}\Omega_{j,n,m}\cos\left(\omega_j t + \phi_j\right)|n\rangle\langle m| \qquad (2.9) \\
&= \frac{1}{2}\sum_{n,m}\sum_{j}\Omega_{j,n,m}\left[e^{+i(\omega_j t + \phi_j)} + e^{-i(\omega_j t + \phi_j)}\right]|n\rangle\langle m|
\end{aligned}
$$

where, for each field, labelled by j, we have defined an oscillation frequency ω_j, a phase ϕ_j, and a coupling strength for each pair of levels n, m. This permits us to absorb all selection rules, Clebsch-Gordan coefficients (see Section 3.3), and other considerations into the coupling strength $\Omega_{j,n,m}$. The decomposition of the cosine to complex exponentials is merely to emphasise that one of these components can be dropped via the rotating-wave approximation. The butterfly operator $|n\rangle\langle m|$ converts state $|m\rangle$ to state $|n\rangle$ and, as before, could also be written as a matrix.

The transformation as presented uses Bohr frequencies, but the problem can often be simplified further if instead we use the frequency of the applied field, as described in Shore [32, §14.2]. This transformation is used in Sections 3.1 and 3.5 to reduce the problem of a three or more level atom to an effective two-level system with a time independent Hamiltonian.

2.1.2 Decoherence and Spontaneous Emission

An enormous advantage of the density matrix prescription over the simpler wave functions is the ability to describe irreversible, incoherent processes. We introduce decoherence $\mathbf{D}$ and spontaneous emission $\mathbf{S}$ by adding two new terms to the rate of change of the density matrix, as in Demtröder [33]:

$$\dot{\tilde{\rho}}_{n,m} = i(\omega_{n,m}\tilde{\rho}_{n,m} - [\tilde{\boldsymbol{H}}, \tilde{\boldsymbol{\rho}}]_{n,m}) + D_{n,m} + S_{n,m} . \tag{2.10}$$

Dephasing—the decay of coherence between two states—is modelled as a decay of the off-diagonal elements:

$$D_{n,m} = -A_{n,m}\tilde{\rho}_{n,m} \text{ with } A_{n,n} = 0 \text{ and } A_{n,m} = A_{m,n}. \tag{2.11}$$

The description of spontaneous emission is slightly more involved as it must couple two elements of the density matrix, rather than acting on elements in isolation. If we let $\Gamma_{n,m} > 0$ be the rate at which population decays into $|n\rangle$ from $|m\rangle$, and $\Gamma_{m,m} < 0$ be the rate at which level $|m\rangle$ loses population, then to conserve population $\sum_n \Gamma_{n,m} = 0$ for each m, and

$$S_{n,m} = \delta_{n,m} \times \left(\sum_a \Gamma_{n,a}\tilde{\rho}_{a,a} \right) \tag{2.12}$$

where $\delta_{n,m}$ is the Kronecker delta. These two processes are related, and we can describe dephasing in terms of spontaneous emission rates:

$$A_{n,m} = \begin{cases} \frac{1}{2}\Gamma_{n,m} + \gamma_{n,m} & n > m \\ 0 & n = m \\ \frac{1}{2}\Gamma_{m,n} + \gamma_{m,n} & n < m \end{cases} \tag{2.13}$$

Note that $\boldsymbol{A}$ is symmetrical, but $\boldsymbol{\Gamma}$ is not. We parameterise any decay of coherences due to mechanisms not associated with spontaneous emission by the additional term $\boldsymbol{\gamma}$. This includes, for example, collisions between particles which interrupt the phase of an oscillation, but do not change the population.

This general description is used in Section 2.3 to derive the Optical Bloch Equations.

2.2   Coherence and Linewidth

We now consider the coherence properties of light which these spontaneously decaying atoms emit. This is of particular use in understanding the atomic absorption spectra, described in Section 5.1.2, to which we reference our lasers.

Light emitted or absorbed by atoms does not consist purely of a single frequency, but rather has a finite spectral width. This is also the case for light emitted by a (continuous wave, narrow band) laser, and the following analysis is useful for both cases.

Typically, we cannot use physical arguments to make statements about the electric field itself, but we may be able to make statements about the coherence properties of the field and use this to infer other properties, such as the spectral width. This is described further by Meschede [34].

To understand a broadening about a central frequency, we begin by defining the auto-correlation function:

$$g(t) := \int_{-\infty}^{+\infty} E^*(t' + t)E(t')dt' \qquad (2.14)$$

where $E(t)$ is the time varying electric field. This quantifies for how long a field maintains a deterministic phase relationship, and this time is known as the coherence time. The cross-correlation theorem—or, in this instance, the Wiener-Khintchine theorem [33]—provides a direct link from this correlation function to the intensity profile:

$$\tilde{g}(\omega) = \tilde{E}^*(\omega)\tilde{E}(\omega) \Rightarrow \Re\left[\tilde{g}(\omega)\right] \propto \tilde{I}(\omega) \qquad (2.15)$$

where $\tilde{I}(\omega)$ is the intensity at some frequency ω, measured relative to a central frequency, and the over-tildes indicate that we have Fourier transformed to the frequency domain.

2.2.1 Natural linewidth

In the absence of any other broadening mechanism, a transition will have a minimum linewidth dictated by the finite lifetime of the transition. To correctly understand spontaneous emission and this finite lifetime, one must appeal to full Quantum Electro-Dynamics where the modes of the field as well as the states of the atom are quantised. However, an adequate semi-classical picture can be found by treating the atom as an oscillator whose phase is interrupted by interactions with the ever-present vacuum fluctuations. Hence, the electric field from a spontaneously decaying atom can be modelled as a pure frequency whose phase is interrupted in a stationary[1] stochastic way, with a mean interval between interruptions given by the lifetime of the state $1/\gamma$. This model is described in Shore [30, §1.8].

The correlation function for the electric field from a spontaneously emitting atom is $g(t) = e^{-i\omega_0 t}e^{-\gamma t}$, the Fourier transform of which (found via integration by parts) gives a Lorentzian profile, and this Lorentzian is typically parameterised by the experimentally observable decay of intensity, rather than electric field. Intensity is proportional to the square of the electric field, and hence decays as twice the rate: $\Gamma = 2\gamma$ and $\tau = 1/\Gamma$ is the *radiative* lifetime:

$$\tilde{I}(\omega) \propto \frac{1}{\pi}\Re\left[\frac{1}{\gamma - i(\omega_0 - \omega)}\right] = \frac{1}{\pi}\frac{\Gamma/2}{(\Gamma/2)^2 + (\omega_0 - \omega)^2} \ . \tag{2.16}$$

The full-width at half-maximum is, for a Lorentzian, the linewidth: FWHM $= \Gamma$.

2.2.2 Doppler Broadening

The same rationale can be applied to derive the spectral line broadening due to the different Doppler shifts of individual atoms in a sample which has a Maxwell-Boltzmann

---

[1]This is meant in the statistical sense that the mean and standard deviation (and other moments) are invariant to shifts in time.

velocity distribution. The correlation function is the oscillating field scaled by the average frequency shift: $g(t) = e^{-i\omega_0 t}\langle e^{i\boldsymbol{k}\cdot\boldsymbol{v}}\rangle$. As before the spectrum is found via the Fourier transform:

$$\tilde{I}(\omega) = \frac{1}{\Delta_D\sqrt{\pi}}e^{-\left(\frac{\omega-\omega_0}{\Delta_D}\right)^2}. \tag{2.17}$$

The broadening is characterised by the Doppler width $\Delta_D = \omega_0\sqrt{\frac{2k_BT}{mc^2}}$ and, for this Gaussian profile, the full-width at half-maximum is FWHM $= 2\sqrt{\ln(2)}\Delta_D \simeq 1.67\Delta_D$.

A measured linewidth typically has significant contributions from natural linewidth (Lorentzian) and Doppler broadening (Gaussian). The combined effect of the mechanisms is found by multiplying the two auto-correlations, and so the combined linewidth is found by convolving and in general we have a Voigt profile. Pressure broadening, as is sometimes significant, arises from collisions which interrupt the phase, and hence one can follow the same derivation as for natural linewidth.

2.3 The Optical Bloch Equations

We can now apply the general density matrix described in Section 2.1 to the useful and manageable case of a two-level atom illuminated by a single monochromatic field. This provides results which we can use to understand the MOT. Treatments of this can be found elsewhere, notably in Allen and Eberly [35]. We begin with the 2×2 density matrix and the 2×2 Hamiltonian and apply the transformation to the rotating frame as prescribed by Equation 2.7:

$$\boldsymbol{\rho} = \begin{pmatrix} \rho_{0,0} & \rho_{0,1} \\ \rho_{1,0} & \rho_{1,1} \end{pmatrix} \Rightarrow \tilde{\boldsymbol{\rho}} = \begin{pmatrix} \rho_{0,0} & \rho_{0,1}e^{-i\omega_{1,0}t} \\ \rho_{1,0}e^{+i\omega_{1,0}t} & \rho_{1,1} \end{pmatrix} \tag{2.18}$$

$$\boldsymbol{H} = \begin{pmatrix} \omega_0 & \Omega\cos(\omega t) \\ \Omega^*\cos(\omega t) & \omega_1 \end{pmatrix} \Rightarrow \tilde{\boldsymbol{H}} = \begin{pmatrix} \omega_0 & \frac{1}{2}\Omega e^{+i\Delta t} \\ \frac{1}{2}\Omega^* e^{-i\Delta t} & \omega_1 \end{pmatrix} \tag{2.19}$$

where $\omega_{1,0} = \omega_1 - \omega_0$, $\Delta = \omega - \omega_{1,0}$ and, in the last stage, we have used the rotating wave approximation. For this system, decoherence and spontaneous emission are specified by

$$\boldsymbol{\Gamma} = \Gamma \begin{pmatrix} 0 & +1 \\ 0 & -1 \end{pmatrix} \text{ and } \boldsymbol{\gamma} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \tag{2.20}$$

and hence

$$\boldsymbol{A} = -\frac{\Gamma}{2} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{2.21}$$

We can now expand the commutator of Hamiltonian and density matrix

$$[\tilde{\boldsymbol{H}}, \tilde{\boldsymbol{\rho}}] = \frac{1}{2} \begin{pmatrix} \Omega\tilde{\rho}_{0,1}e^{+i\Delta t} - \Omega^*\tilde{\rho}_{1,0}e^{-i\Delta t} & \Omega e^{-i\Delta t}\left(\tilde{\rho}_{0,0} - \tilde{\rho}_{1,1}\right) - 2\tilde{\rho}_{0,1}\omega_{1,0} \\ \Omega^* e^{+i\Delta t}\left(\tilde{\rho}_{0,0} - \tilde{\rho}_{1,1}\right) + 2\tilde{\rho}_{1,0}\omega_{1,0} & \Omega\tilde{\rho}_{1,0}e^{-i\Delta t} - \Omega^*\tilde{\rho}_{0,1}e^{+i\Delta t} \end{pmatrix} \tag{2.22}$$

and use Equation 2.10 to find the time evolution of each element:

$$
\begin{aligned}
\dot{\tilde{\rho}}_{0,0} &= -\tfrac{i}{2}\left(\Omega\tilde{\rho}_{1,0}e^{-i\Delta t} - \Omega^*\tilde{\rho}_{0,1}e^{+i\Delta t}\right) &&+ \Gamma\tilde{\rho}_{1,1}, \\
\dot{\tilde{\rho}}_{1,1} &= +\tfrac{i}{2}\left(\Omega\tilde{\rho}_{1,0}e^{-i\Delta t} - \Omega^*\tilde{\rho}_{0,1}e^{+i\Delta t}\right) &&- \Gamma\tilde{\rho}_{1,1}, \\
\dot{\tilde{\rho}}_{0,1} &= +\tfrac{i}{2}\Omega\ e^{-i\Delta t}\left(\tilde{\rho}_{1,1} - \tilde{\rho}_{0,0}\right) &&- \tfrac{\Gamma}{2}\tilde{\rho}_{0,1}, \\
\dot{\tilde{\rho}}_{1,0} &= -\tfrac{i}{2}\Omega^* e^{+i\Delta t}\left(\tilde{\rho}_{1,1} - \tilde{\rho}_{0,0}\right) &&- \tfrac{\Gamma}{2}\tilde{\rho}_{1,0}.
\end{aligned}
$$

Finally, we can make the replacement $\tilde{\tilde{\rho}}_{0,1} = \tilde{\rho}_{0,1}e^{+i\Delta t}$ (and the corresponding complex conjugate for $\tilde{\rho}_{1,0}$) and hence remove the oscillating exponentials. To simplify notation, we revert to using the symbol $\boldsymbol{\rho}$ for this transformed density matrix, and we find the Optical Bloch Equations for a two level atom illuminated by a monochromatic field:

$$
\begin{aligned}
\dot{\rho}_{0,0} &= &&-\tfrac{i}{2}\left(\Omega\rho_{1,0} - \Omega^*\rho_{0,1}\right) &&+ \Gamma\rho_{1,1}, \\
\dot{\rho}_{1,1} &= &&+\tfrac{i}{2}\left(\Omega\rho_{1,0} - \Omega^*\rho_{0,1}\right) &&- \Gamma\rho_{1,1}, \\
\dot{\rho}_{0,1} &= +i\Delta\rho_{0,1} &&+\tfrac{i}{2}\Omega\ \left(\rho_{1,1} - \rho_{0,0}\right) &&- \tfrac{\Gamma}{2}\rho_{0,1}, \\
\dot{\rho}_{1,0} &= -i\Delta\rho_{1,0} &&-\tfrac{i}{2}\Omega^* \left(\rho_{1,1} - \rho_{0,0}\right) &&- \tfrac{\Gamma}{2}\rho_{1,0}.
\end{aligned}
$$

We can use these equations to find the steady state populations and coherences under constant illumination. In the presence of decay, the rate of change of the populations and

the rate of change of coherence terms will tend to zero. We set the left hand terms in the above equations to zero and solve to find the populations and coherences. Additionally, we introduce the parameter w which describes the population inversion $w = \rho_{0,0} - \rho_{1,1}$ (hence $\rho_{1,1} = \frac{1}{2}(1 - w)$ and $\rho_{0,0} = \frac{1}{2}(1 + w)$ using $\rho_{0,0} + \rho_{1,1} = 1$), and we find

$$w = \frac{1}{1 + s} \text{ where } s = \frac{|\Omega|^2}{2\left|\frac{\Gamma}{2} - i\Delta\right|^2} \tag{2.23}$$

and s is known as the saturation parameter. For the special case of on-resonance excitation, $s = s_0$ and we can define a saturation intensity I_{sat}, for which $\frac{1}{4}$ of the population is in the excited state: $s_0 = I/I_{\text{sat}} = 2|\Omega|^2/|\Gamma|^2$. The general off-resonance saturation parameter is easily expressed in terms of this on-resonance parameter:

$$s = \frac{s_0}{1 + (2\Delta/\Gamma)^2} \tag{2.24}$$

and, although they are not used in this work, the coherences are

$$\rho_{0,1} = -i\frac{\Omega}{2}\frac{1}{(\Gamma/2 - i\Delta)}\frac{1}{1 + s} . \tag{2.25}$$

These general results for a two-level system are used extensively throughout this thesis, beginning with the next section where we derive the force on an atom.

2.4 Force on a Two-level atom

For near resonant excitation, where the scattering force and not the dipole force dominates, an atom experiences a force due to the photons it absorbs and re-emits. If we assume that the re-emission process is due entirely to spontaneous emission, which will dominate for low saturation $s \ll 1$, emission from the atom is isotropic and will, over many events, average to zero. In this regime, therefore, the force on an atom is given by the momentum carried by a photon $p = h/\lambda = \hbar k$ (where λ is the wavelength of the light and $k = 2\pi/\lambda$ is the wave-vector) multiplied by the rate of absorption. This rate is the steady state fraction of population in the excited state, multiplied by the rate at which this population decays via spontaneous emission – in steady state, the rate of

spontaneous emission is equal to that of absorption. A similar derivation can be found in Metcalf and van der Straten [29].

$$F = \hbar k \times \rho_{1,1} \Gamma = \hbar k \frac{s_0 \Gamma / 2}{1 + s_0 + (2\Delta/\Gamma)^2} = \frac{F_0}{1 + s_0 + (2\Delta/\Gamma)^2} \tag{2.26}$$

where we have introduced the characteristic force $F_0 = \frac{1}{2}\hbar k s_0 \Gamma$.

The approach here implicitly treats the atom as a classically moving point-like particle. For this approximation to hold, several conditions must be met. Firstly, the position of the particle must be well defined relative to other length scales in the problem: namely, the wavelength of the light: $\Delta x \ll \lambda$.

Secondly, the velocity of the particle must be well defined, and the appropriate quantities to compare are the Doppler shift $\Delta\omega$ due to the recoil when a photon is emitted, and the natural width of the transition Γ. The Doppler shift due to a recoil is $\Delta\omega = \omega_0 \Delta v/c$ with $M\Delta v = \hbar k$, where $\omega_0 = ck$ is the transition frequency and M is the atomic mass. We find so-called Broad-line condition, described by Dalibard [36]:

$$\frac{\hbar^2 k^2}{M} \ll \hbar\Gamma. \tag{2.27}$$

The Heisenberg Uncertainty relation, $\Delta x \Delta p \geq \frac{\hbar}{2}$, shows this is compatible with the constraint on the uncertainty in position. For rubidium, this condition is satisfied by a factor of $\simeq 800$ and hence we can safely continue with this classical point-like picture of the external degrees of freedom of the atom.

2.4.1 Perturbed force on a Two-level atom

If we allow the detuning Δ to be composed of some offset in the laser frequency $\alpha\Gamma$ and some other (typically smaller) contribution $\delta\Gamma$ (due, for example, to a Doppler or Zeeman shift), we can derive some useful results. We express both components of the detuning in terms of the state linewidth Γ, which is the natural frequency scale for this problem. This approach is semi-classical, and so has a limited range of applicability.

This approximation is useful here, but the situation can be described more generally and fully quantum-mechanically, and a treatment can be found in Gordon and Ashkin [37].

From the Taylor series expansion of the force in δ,

$$\frac{F(\Delta)}{F_0} = \frac{F(\alpha\Gamma + \delta\Gamma)}{F_0} = \frac{1}{1 + s_0 + (2\alpha)^2} - \frac{8\alpha\delta}{[1 + s_0 + (2\alpha)^2]^2} + \dots , \qquad (2.28)$$

we see that if we two counter-propagating beams, each well below saturation $s_0 \ll 1$ [2] and whose absorption has opposite dependence on δ, the net force on the atom is, to first order, linear in δ and either positive or negative depending on the sign of α (i.e. whether the laser is red or blue detuned):

$$F_{\text{total}} = F(\alpha\Gamma + \delta\Gamma) - F(\alpha\Gamma - \delta\Gamma) = -F_0 \frac{16\alpha}{\left[1 + (2\alpha)^2\right]^2} \delta + \dots \qquad (2.29)$$

This addition implicitly assumes that we may add intensities, not electric fields, and this ignoring of coherence is permissible only for derivations of effects on length scales much larger than the transition wavelength. Specifically, this cannot account for sub-Doppler cooling mechanisms, which are described in Section 2.5.

The dependence of this first order term is shown in Figure 2.1, and is extremised to $\partial F_{\text{total}}/\partial \delta = \mp F_0 3\sqrt{3}/2$ at $\alpha^{\pm}_{\text{max}} = \pm 1/(2\sqrt{3})$; for this detuning, the force is within 10% of linear for $|\delta| \lesssim 0.2$. The form of the force on a two level atom, for this maximum detuning $\alpha^{-}_{\text{max}}\Gamma$, is shown in Figure 2.2.

If we make δ depend on velocity (via the Doppler shift) we can engineer a cooling force; if we can contrive to make δ depend on position, we can make a trap. In the Magneto-Optical trap, we have both.

---

[2]For non-small saturation, one beam will affect the population seen by the other in a way not accounted for in this treatment. Additionally, we have not accounted for the possibility of the atom absorbing a left moving photon and being stimulated to emit by a right moving photon; however, this is a two-photon event and, as such, is equally likely for both directions regardless of the balance of power in the beams, and so averages to zero.
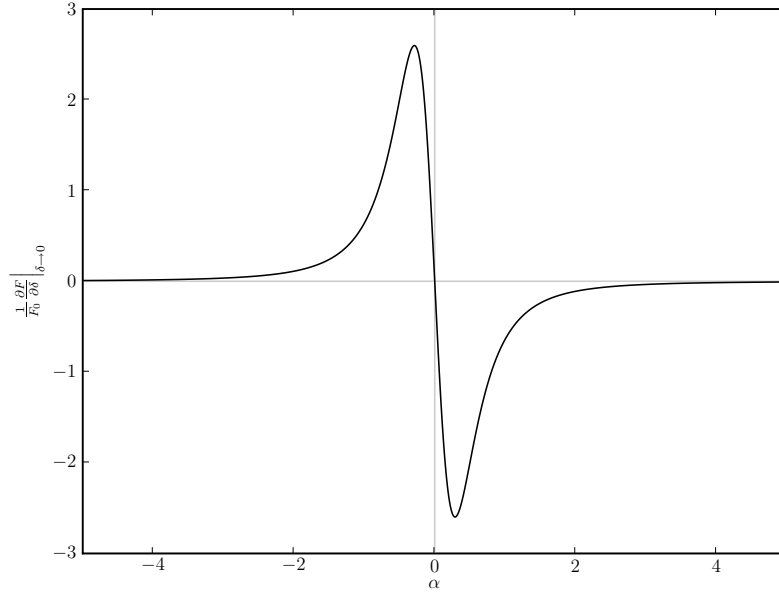
FIGURE 2.1: The gradient of the force due to a single beam for small atom detuning $\delta\Gamma = 0$ (from e.g. a Doppler shift or a Zeeman shift) for a range of laser detunings $\alpha\Gamma$. For negative laser detuning, the gradient is positive and if we can make δ depend negatively on the parameter which the force affects, the force will act to reduce the detuning (by e.g. slowing or displacing).
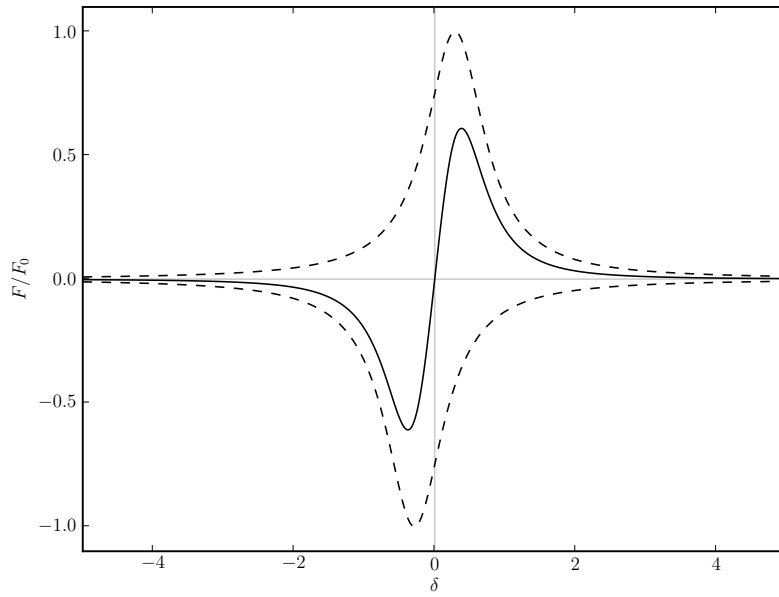
FIGURE 2.2: The force on an atom due to two counter-propagating beams with opposite dependence on atom detuning $\delta\Gamma$, for laser detuning $\alpha\Gamma = \alpha_{\max}^{-}\Gamma$. The larger (dashed) curves are the contributions from individual beams and the smaller (solid) curve is the combined force. The 1$^{\text{st}}$ order Taylor series expansion of Equation 2.28 reproduces the slope around $\delta = 0$ (i.e. atoms which move slowly or are near the origin see a linear cooling or restoring force).

2.4.2 Optical Molasses

With these tools we can understand the cooling mechanism known as optical molasses. A two-level atom is illuminated from all sides by red-detuned light ($\alpha < 0$) such that, as it moves, it sees the light source it moves towards Doppler shifted towards its resonance, and so preferentially scatters photons whose momentum oppose it.

For a given light wave, motion in the negative direction (opposite to propagation of the light) means the atom moves across the wavefronts of the beam faster than if it were stationary and hence gives a positive Doppler shift: $\delta\Gamma = -kv$ for $v \ll c$ and

$$F = F_0 \frac{16\alpha}{[1 + (2\alpha)^2]^2} \frac{kv}{\Gamma} \text{ for } \delta \ll \alpha .$$ (2.30)

For a negative detuning, this force damps the motion and, while not always linear, retains the correct sign for all δ; all atom velocities are damped, but the force is only significant for atoms below a certain speed.

2.4.3 Magneto-Optical Trapping

Optical Molasses cools atoms towards zero velocity, but it does not define a preferred position; atoms are cooled, but they are not trapped. In order to define a point in space towards which atoms can be driven, we impose a spatially dependent Zeeman shift using the inhomogeneous magnetic field generated by a pair of coils in the anti-Helmholtz configuration [38] – two co-axial coils with current flow in opposite directions create a magnetic field zero at the mid-way point, and we choose their separation so that the field increases linearly with displacement away from this zero. This combination forms a magneto-optical trap and the same tools as used to understand Optical Molasses can be used to analyse this situation: δ is now spatially, not velocity, dependent.

The magnetic field used in our experiment is described in Section 5.3.1 and, near the centre of the trap, the magnitude increases from zero with distance away from the centre. For displacement along axis x only ($y = z = 0$), $|B| \propto x$ and, for the transition in rubidium $\left|5^2\mathrm{S}_{1/2}, F = 3, m_F\right\rangle$ to $\left|5^2\mathrm{P}_{3/2}, F' = 4, m'_F = m_F + 1\right\rangle$ (described in more

detail in Section 3.2), $\partial\omega/\partial B = 2\pi \times 230\text{kHz/G}$ and we see a spatially dependent Zeeman shift:

$$\Delta_z = \Delta + \frac{\partial\omega}{\partial B}\frac{\partial B}{\partial x}x. \tag{2.31}$$

where $\partial\omega/\partial B$ is the Zeeman shift, and $\partial B/\partial x$ is the field gradient in our experiment.

The final trick required to convert this spatial inhomogeneity to a stable trap is to use a selection rule (here, a statement of conservation of angular momentum) to ensure that only photons with positive momentum interact with atoms displaced to negative position. We use left-circularly polarised light (labelled $q = +1$ or σ_+ and carrying positive angular momentum) propagating from left to right (carrying positive linear momentum), which will drive only the transition $m_F \rightarrow m'_F = m_F + 1$. By choice of the sign of the magnetic field gradient, we ensure that atoms displaced to the left see a Zeeman shift which brings their transition closer to resonance with the red-detuned laser light, and hence atoms displaced to negative position interact more strongly with light carrying positive momentum, and, overall, experience a restoring force.

With the addition of this selection rule, we have, as in Section 2.4.1, forces from opposite directions with opposite dependences on δ:

$$\delta\Gamma = -\zeta x \text{ where } \zeta = \frac{\partial\omega}{\partial B}\frac{\partial B}{\partial x} \tag{2.32}$$

$$F = F_0\frac{16\alpha}{[1 + (2\alpha)^2]^2}\frac{\zeta x}{\Gamma} \text{ for } \delta \ll \alpha . \tag{2.33}$$

Thus, for a red-detuned laser $\alpha < 0$, we have a force which acts to push atoms towards $x = 0$. For the typical value $\zeta = 2\pi \times 5\text{MHz/cm}$ (and $\Gamma = 2\pi \times 6\text{MHz}$), the trapping force is linear over $\sim 1\text{cm}$.

We can picture an atom near the centre of the trap $(B = 0)$ displaced slightly radially (in the cylindrical geometry defined by the orientation of the magnetic field coils). This displacement causes the atom to see a slightly increased magnetic field, hence the degeneracy of its m_F levels is lifted and $m_F > 0$ is brought closer to resonance with σ_+ light travelling towards the centre of the trap. For vertical displacement the picture is

the same, except the field now decreases to negative B away from $z = 0$, and the atom sees its $m_F < 0$ levels brought to resonance with the inbound σ_- light.

For the combination of Doppler and Zeeman detuning shifts found in a Magneto-Optical trap for *slow* atoms near the *origin*, we add these contributions: $\delta\Gamma = -\zeta x - kv$ and

$$F = F_0 \frac{16\alpha}{[1 + (2\alpha)^2]^2} \frac{(\zeta x + kv)}{\Gamma} \text{ for } \delta \ll \alpha \tag{2.34}$$

or $M\ddot{x} = \beta\dot{x} + \kappa x$, with the obvious substitutions. This is the equation of a damped harmonic oscillator with damping rate β/M and natural frequency $\sqrt{\kappa/M}$. Additionally, since in equilibrium the energy of the system will be distributed across the degrees of freedom of the system, this allows us to connect the temperature (via the velocity) with the size of the cloud [39, p. 158]: $k_{\mathrm{B}}T = mv_{\mathrm{RMS}}^2 = \kappa x_{\mathrm{RMS}}^2$.

2.4.4 Doppler Temperature

From Equation 2.30 we see the velocity dependent force on an atom is linear for small velocities, and for red detuning acts as a friction. Additionally, we know the atom spontaneously emits photons, each of which imparts a random momentum to the atom. These dynamics very closely mimic Brownian motion of a particle and so, while we can expect significant cooling, we note that there will be some residual random motion which can be interpreted as a finite temperature.

In addition to the cooling force of Equation 2.30, we include a force to model spontaneous emission: $F_s(v, t)$. The time- or ensemble- average of this force is zero, but the variance is not. The correct tool for understanding the behaviour of the velocity distribution of an ensemble subject to these dynamics is the Fokker-Planck equation, an excellent treatment of which can be found in Risken [40]. The evolution of a time-evolving velocity distribution $W(v, t)$ is described by[3]

$$\frac{\partial}{\partial t}W(v, t) = \left[-\frac{\partial}{\partial v}D^{(1)}(v) + \frac{\partial^2}{\partial v^2}D^{(2)}(v) \right] W(v, t) \tag{2.35}$$

[3]The partial derivatives *operate* on W rather than simply multiplying.

where $D^{(1),(2)}$ are the drift and diffusion coefficients, here modelling friction and spontaneous emission respectively. If these coefficients are independent of time, the stationary solution is:

$$\langle W(v) \rangle \propto \frac{1}{D^{(2)}(v)} \exp \left[\int_0^v \frac{D^{(1)}(v')}{D^{(2)}(v')} dv' \right] \ . \tag{2.36}$$

For Doppler cooling, the drift is linear $D^{(1)}(v) = F_0 \frac{16\alpha}{[1+(2\alpha)^2]^2} \frac{kv}{\Gamma}$ and the diffusion is constant $D^{(2)}(v) = \Gamma s \left(\hbar k\right)^2 / m$ [39, §5.4 & §7.2] and Equation 2.36 becomes

$$\langle W(v) \rangle \propto \exp \left[\left(\frac{1}{2} m v^2 \right) / \left(\frac{\hbar \Gamma}{4} \frac{1+(2\alpha)^2}{2\alpha} \right) \right] \ . \tag{2.37}$$

For $\alpha < 0$ this fits a Maxwell-Boltzmann distribution, and for $\alpha = -\frac{1}{2}$ gives a minimum temperature, known as the Doppler limit, described by Wineland and Itano [41]:

$$T_{\text{Doppler}} = \frac{\hbar \Gamma}{2k_{\text{B}}} \ . \tag{2.38}$$

The actual temperature depends on detuning, and we have assumed the saturation parameter is small. Additionally, if the coefficients are not linear and constant, the momentum distribution will not fit a Maxwell-Boltzmann distribution and the notion of temperature is not so clearly defined. Notably, this approach is not valid for sub-Doppler cooling because in this regime the one-photon recoil is not a small fraction of the drift; a more general approach is presented by Castin, Wallis, and Dalibard [42].

2.4.5 Spatial Distribution

All of the analysis so far has treated atoms in isolation; we have assumed we can describe the dynamics without considering the interaction of atoms with one another. This is approximately true for the diffuse background gas from which the cloud is formed, but is certainly not true within the cloud where the sample becomes optically dense; here, the probability that a photon spontaneously emitted by one atom is absorbed by another is significant, and hence an atom feels a pressure exerted by its neighbours and mediated by real photons, analogous to an electrostatic repulsion mediated by virtual photons. The attenuation of the fields as they penetrate into the optically dense cloud clearly indicates

that the scattering rate, and hence repulsive pressure, will be less from atoms deeper within the cloud, but as a first approximation, and especially for large laser detunings, we can treat the intensity of the trapping fields as constant throughout the cloud.

The absorption cross-section for an atom is $\sigma = \frac{3\lambda^2}{2\pi}$ and the fraction of a sphere of radius r which this cross-section covers is $\frac{\sigma}{4\pi r^2}$. For a spherically symmetrical distribution, the divergence theorem [43] tells us that particle inside the sphere will feel the pressure due only to other particles deeper within the sphere:

$$F_p(r) = F_0 \frac{16\alpha}{[1 + (2\alpha)^2]^2} \frac{\sigma}{4\pi r^2} \int_0^r \rho(r')4\pi r'^2 \mathrm{d}r' . \qquad (2.39)$$

This is balanced by the spatially dependent trapping force of Equation 2.33, and hence the number density $\rho(r)$ is constant: $\rho = 3\zeta/(\sigma\Gamma)$.

As discussed in Section 5.3, the trapping force is not spherically symmetrical, but this treatment nevertheless provides insight and estimates of the scales involve. The distribution can, of course, be far more detailed than this uniform density sphere, and many interesting shapes have been described [44–46]. The spatial distribution of our cloud is characterised in Section 7.1, and we show an interesting shape which is seen when the retro-reflected beam is slightly misaligned.

2.4.6   Capture from Background Gas

Equation 2.34 suggests the situation is simple and that an atom at any position and with any velocity will experience trapping and cooling forces. This equation is, however, only an approximation, and is invalid far from the origin. For instance, an atom displaced so far from the origin that its positive Zeeman shift ζx is sufficient to overcome the laser negative red-detuning $\alpha\Gamma$ will be see light *blue*-detuned from the resonance, and so will experience a negative Doppler cooling force: a heating. Hence we define the MOT region as the region in which a slow moving atom experiences both trapping and cooling forces.

The full form of the force on an atom (Equation 2.29) has a maximum at $\alpha \simeq -\delta$ and hence a stationary atom displaced further from the origin than $x_{\max} = |\alpha\Gamma/\zeta|$ would find

itself in a region of blue-detuning [4]. Similarly, an atom at the origin will no longer feel a trapping force if the speed exceeds $v_{\max} = |\alpha\Gamma/k|$. For our laser detuning $\alpha\Gamma = -3\Gamma$ and field gradient 5G/cm (see Section 5.3.1), $x_{\max} = 160$mm.

With our magnetic field gradient, the MOT region is an ellipsoid with vertical minor axis 320mm and horizontal major axis 640mm, but at these scales the field is no longer linear with position. The capture volume is restricted to the much smaller volume defined by the intersection of the three orthogonal laser beams. As described in Section 5.5.1, we use Gaussian beams with 7mm waist, and the overlap volume is ~ 190mm^3. Using room temperature $T = 300$K and typical background pressure $P = 10^{-4}$Pa, the ideal gas law gives a number density $\sim 10^7$mm^{-3}, and hence the capture region contains $\sim 10^9$ atoms. Of these, approximately $\sim 1\%$ are moving sufficiently slowly ($< v_{\max}$) to be trapped, and, very crudely, we can expect $\sim 10^7$ atoms in our cloud. (An atom moving with the maximum velocity must absorb ~ 800 photon impulses to bring it to rest, and this can easily occur before it has traversed the trapping region.) This estimate is reasonably close to the 13×10^6 reported in Section 7.2.

We can do little better than these estimates using only the approximations to the force for slow moving atoms near the origin. To explore the dynamics of capture far from the origin and for fast moving atoms, where these parameters cannot be treated separately, we must appeal to numerical simulation [29]. The dependence on various MOT parameters has also been studied experimentally [48].

2.5 Sub-Doppler Cooling

To conclude this chapter, we briefly mention the mechanisms by which atoms can be cooled to temperatures below the Doppler limit.

The analysis thus far is sufficient to describe the process of atoms being captured in the trap, and the important process of cooling from room temperature to the μK regime,

---

[4]If the steady state number of atoms trapped is assumed to be proportional to this length scale, this dependence agrees with that found experimentally by Hope *et al* [47, Figure 2].

but is insufficient to describe the additional 'sub-Doppler' mechanisms which come to dominate near zero velocity and the centre of the trap.

Doppler cooling uses the very strong near-resonant scattering force to quickly cool from room temperature to the μK regime, but even in the low intensity limit, the lowest attainable temperature is [33]:

$$T_{\text{Doppler}} = \frac{\hbar\Gamma}{2k_{\text{B}}} \tag{2.40}$$

To reach lower temperatures, we must appeal to more subtle mechanisms which, fortuitously, are compatible with the conditions necessary for a MOT. In 1988, Lett *et al* [9] reported the unexpected observation of temperatures below the Doppler limit and there was a subsequent flurry of theoretical work to explain this observation; the two landmark papers are Dalibard and Cohen-Tannoudji [10] and Ungar, Weiss, Riis, & Chu [49].

At the Doppler temperature, atoms move sufficiently slowly that they begin to resolve the wavelength structure of the red-detuned standing waves in which they are immersed. We consider the case of two equal intensity counter-propagating beams with orthogonal polarisation. If these beams are linearly polarised, the ellipticity of the resulting polarisation changes with position along the optical axis: from linear, to right circular, to orthogonal linear, to left circular. The various Zeeman sub-levels in our atom (for these sub-Doppler mechanisms require our model to include more than just two levels) are hence coupled more strongly for different positions, resulting in a spatially dependent lightshift (discussed in Section 8.3.2). The coupling strengths also ensure that atoms are preferentially pumped into, since the light is red-detuned, the lightshifted Zeeman sub-level with the lowest energy. If we now allow the atoms to move, we see that they are pumped to the lowest energy, and that they must climb out of this potential well, where they are once again pumped to the lowest energy, after which the cycle repeats: this is Sisyphus cooling [33, p. 790].

However, in the MOT, we have not orthogonal linear, but orthogonal circularly polarised beams. In this case, the resulting polarisation is linear, with a direction which depends on position with azimuthal angle $\phi = kz$. The intensity and the coupling does not

depend on position, and Sisyphus cooling does not occur. Rather, Dalibard and Cohen-Tannoudji describe a different mechanism in which atoms moving in this field see, in their frame, a rotating linear polarisation. This rotating polarisation, as can be found by rotating to the frame where this polarisation is fixed, gives rise to spatially dependent coupling, resulting in spatially dependent pumping and hence a higher chance that moving atoms will absorb counter-propagating photons [50].

The situation in the MOT cannot be classed as either of these two theoretically neat examples. The beams are indeed orthogonally polarised and counter-propagating, but three such pairs, one for each dimension, converge on the centre; the resulting polarisation changes in ellipticity and in direction on a wavelength scale [39, §11.4.2] and atoms experience a mixture of the cooling effects.

We should note there are many other sub-Doppler techniques, and surprisingly one can even beat the recoil limit with optical methods (in two [51] and three dimensions [52]). Methods include Velocity-Selective Coherent Population Trapping [39, 53, 54], Raman sideband cooling [55–59], and, of course, evaporative cooling [60], famously used to create the first Bose-Einstein condensate [16, 17].

Chapter 3

Coherent Manipulation Theory

The accumulation of cold atoms in the MOT provides us with a slow moving collection of atoms which are isolated from the rest of the world. For the few milliseconds of an experiment, these particles will not interact and any coherence we create in them will not be disturbed. This chapter describes how this coherence is embodied by the levels in the atom, and how we can manipulate it.

3.1   Effective Two Level System

The rubidium atom has two lower 'ground' states and a number of upper states, split by the hyperfine interaction. The details of these levels is discussed in Section 3.2, but for now it is sufficient to picture the system as having two lower levels, between which there can be no direct dipole coupling, and a single upper level, via which the two lower levels can be indirectly coupled; direct electric dipole transitions between these ground-levels are forbidden, with the fortunate consequence that the states are extremely long-lived. These long-lived states are coupled by two fields, both of which are nearly resonant with the upper level, and the transition is pictured as being driven *via* this upper level, while never actually populating it. We intend to use this system to emulate a two-level atom, and this section describes how this three-level two-field problem can be recast as an effective two-level system coupled by an effective photon.

The problem is relatively straight forward when the difference between the field frequencies equals the energy difference between the two indirectly coupled states [61, 62] and the problem is treated semi-classically. We follow this procedure, and extend the treatment in Section 3.5. It is possible to approach the problem fully quantum mechanically [63, 64], and to solve for the case where the system is not on two-photon resonance, but the semi-classical theory and the on-resonance case are sufficient here.

This so-called Raman transition inherently requires treatment of more than two levels, but we can emulate a two-level system by ensuring all but two of the levels remain unpopulated. As the simplest system in which a Raman transition can occur, we take a three-level system with two stable states ($|0\rangle$ and $|1\rangle$) and one radiative state ($|2\rangle$) which remains unpopulated and via which the Raman transition is driven. It is often useful to write these states as column vectors (and operators as matrices) and we use the correspondence:

$$|0\rangle \equiv \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \ ; \ |1\rangle \equiv \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \ ; \ |2\rangle \equiv \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \tag{3.1}$$

As in Section 2.3, we can write down a matrix describing spontaneous emission $\mathbf{\Gamma}$ and from it obtain the decay of coherences $\boldsymbol{A}$:

$$\mathbf{\Gamma} = \begin{bmatrix} 0 & 0 & +\Gamma_{0,2} \\ 0 & 0 & +\Gamma_{1,2} \\ 0 & 0 & -\Gamma_{0,2} - \Gamma_{1,2} \end{bmatrix} \ ; \ \boldsymbol{A} = \begin{bmatrix} 0 & 0 & \Gamma_{0,2} \\ 0 & 0 & \Gamma_{1,2} \\ \Gamma_{0,2} & \Gamma_{1,2} & 0 \end{bmatrix}. \tag{3.2}$$

We see that coherence between $|0\rangle$ and $|1\rangle$ does not decay except indirectly, via $|2\rangle$. Hence, if we ensure the population of $|2\rangle$ remains negligible, we can neglect decay and treat this problem with wave functions rather than a full density-matrix.

Before we begin, a general note: given the Schrödinger equation

$$i\frac{\partial}{\partial t}|\psi\rangle = \hat{H}|\psi\rangle \tag{3.3}$$

we are at liberty to make the substitution $|\psi\rangle = \hat{O} |\psi'\rangle$, which, after application of the chain rule to the right hand side, migration of the time-derivative from right to left, and finally left-multiplication by the inverse of $\hat{O}$, gives

$$i\frac{\partial}{\partial t} |\psi'\rangle = \hat{H}' |\psi'\rangle \;\; \text{where} \; \hat{H}' := \hat{O}^{-1} \left(\hat{H}\hat{O} - i\frac{\partial}{\partial t}\hat{O} \right). \tag{3.4}$$

The task of recasting the atom as an effective two-level system is tackled with these kind of transformations, and Equation 3.4 is used repeatedly in the following derivation.

The bare Hamiltonian for this system is found from the energies of the levels (as described in Equation 2.5) and the interaction Hamiltonian is found from the couplings of these levels (as described in Equation 2.9):

$$\boldsymbol{H} = \boldsymbol{H}_0 + \boldsymbol{H}_I = \begin{pmatrix} \omega_0 & 0 & \Omega_P \cos\omega_P t \\ 0 & \omega_1 & \Omega_S \cos\omega_S t \\ \Omega_P \cos\omega_P t & \Omega_S \cos\omega_S t & \omega_2 \end{pmatrix}. \tag{3.5}$$

We have assumed fields P and S couple only the states with which they are nearly resonant; for two-photon Rabi frequencies much less than $\omega_1 - \omega_0$, the rotating wave approximation removes any coupling even if it is present. Additionally, since we are emulating a two-level atom, we insist that the coupling strengths maintain a constant ratio: that is, the Ωs may have different amplitudes but if they vary in time, they must share a common time-dependence.

Although it is possible to treat the general case [65], we now restrict ourselves to the case of two-photon resonance where $(\omega_P - \omega_S) = (\omega_1 - \omega_0)$. We use Equation 3.4 and, similarly to Equation 2.6, use a diagonal matrix whose elements are complex exponentials oscillating at the frequencies ω_0, ω_1, and $\omega_v = \omega_2 + \Delta$. This last term is significant; we have chosen the frequency with which the fields are resonant, not the frequency of the energy level. This is illustrated in Figure 3.1 and, after the transformation and the
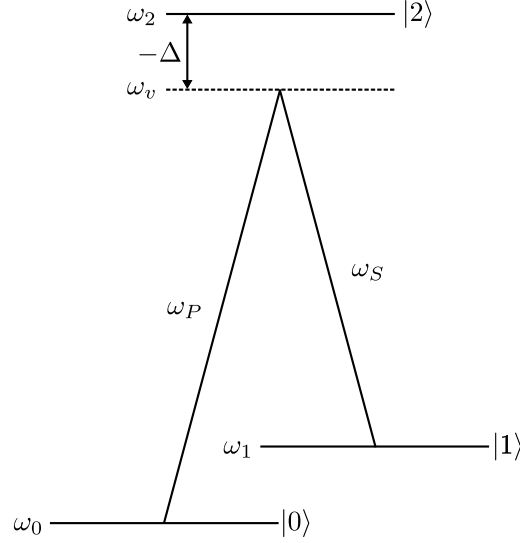
FIGURE 3.1: Three level atom with states $|0\rangle$ and $|1\rangle$ coupled via the virtual state (in the sense that it does not contain any population [66]) $|2\rangle$. The frequency difference of the two lasers is equal to the separation between $|0\rangle$ and $|1\rangle$, and their detuning from the upper state is Δ. Note that, as drawn, the detuning is negative.

rotating-wave approximation, we find a very conveniently formed Hamiltonian:

$$\tilde{\boldsymbol{H}} = \begin{pmatrix} 0 & 0 & \frac{1}{2}\Omega_P \\ 0 & 0 & \frac{1}{2}\Omega_S \\ \frac{1}{2}\Omega_P & \frac{1}{2}\Omega_S & \Delta \end{pmatrix} . \tag{3.6}$$

It is now possible to transform a basis where one coefficient does not change with time. The so-called bright and dark states are found via the following rotation:

$$\begin{pmatrix} |0\rangle \\ |1\rangle \\ |2\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} |d\rangle \\ |b\rangle \\ |2\rangle \end{pmatrix} \tag{3.7}$$

where $\tan\theta = \frac{\Omega_S}{\Omega_P} = \frac{\Omega_0\cos\theta}{\Omega_0\sin\theta}$, with $\Omega_0 = \sqrt{\Omega_P^2 + \Omega_S^2}$. We can use this rotation operator in Equation 3.4 to find the transformed Hamiltonian which, after simplification using

trigonometric identities, yields a Hamiltonian which does not couple the dark state $|d\rangle$:

$$\boldsymbol{H}_R = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2}\Omega_0 \\ 0 & \frac{1}{2}\Omega_0 & \Delta \end{pmatrix}. \tag{3.8}$$

Hence, we can drop this state from our description and we are left with the effective two-level system we sought:

$$i\frac{\partial}{\partial t}\begin{pmatrix} c_b \\ c_2 \end{pmatrix} = \begin{pmatrix} 0 & \frac{1}{2}\Omega_0 \\ \frac{1}{2}\Omega_0 & \Delta \end{pmatrix}\begin{pmatrix} c_b \\ c_2 \end{pmatrix}. \tag{3.9}$$

We can now use familiar tools to solve this two-level problem. For example, we can move to the dressed-state basis by rotating by $\phi = \frac{1}{2}\arctan\left(\Omega/\Delta\right)$, using

$$R(\phi) = \begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix} \tag{3.10}$$

in Equation 3.4, as with the previous rotation. This dressed-state Hamiltonian is

$$\boldsymbol{H}_D = \begin{pmatrix} \frac{1}{2}\Omega_- & 0 \\ 0 & \frac{1}{2}\Omega_+ \end{pmatrix} \tag{3.11}$$

where $\Omega_\pm = \Delta \pm \sqrt{\Omega_0^2 + \Delta^2}$. The solution in this basis is simple, and we can now concatenate all the transformations to discover the time evolution in the original bare state basis.

In summary, the transformations

$$
\begin{pmatrix} c_0 \\ c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} e^{-i\omega_0 t} & 0 & 0 \\ 0 & e^{-i\omega_1 t} & 0 \\ 0 & 0 & e^{-i(\omega_2 - \Delta)t} \end{pmatrix} \begin{pmatrix} \tilde{c}_0 \\ \tilde{c}_1 \\ \tilde{c}_2 \end{pmatrix}
\tag{3.12}
$$

$$
\begin{pmatrix} \tilde{c}_0 \\ \tilde{c}_1 \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} c_d \\ c_b \end{pmatrix} \; ; \; \tan\theta = \frac{\Omega_P}{\Omega_S} = \frac{\Omega_0 \sin\theta}{\Omega_0 \cos\theta}
\tag{3.13}
$$

$$
\begin{pmatrix} c_b \\ \tilde{c}_2 \end{pmatrix} = \begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix} \begin{pmatrix} c_- \\ c_+ \end{pmatrix} \; ; \; \tan 2\phi = \frac{\Omega_0}{\Delta}
\tag{3.14}
$$

convert the three-level two-photon system to a basis where the time evolution is

$$
i\frac{\partial}{\partial t}\begin{pmatrix} c_d \\ c_- \\ c_+ \end{pmatrix} = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta - \sqrt{\Omega_0^2 + \Delta^2} & 0 \\ 0 & 0 & \Delta + \sqrt{\Omega_0^2 + \Delta^2} \end{pmatrix} \begin{pmatrix} c_d \\ c_- \\ c_+ \end{pmatrix}.
\tag{3.15}
$$

The overall solution in the limit of far detuning $(\Delta \gg \Omega_0)$ is

$$
\tilde{c}_0(t) = \cos\theta c_d(0) + \sin\theta c_b(0) e^{\frac{i\Omega_- t}{2}}
\tag{3.16}
$$

$$
\tilde{c}_1(t) = -\sin\theta c_d(0) + \cos\theta c_b(0) e^{\frac{i\Omega_- t}{2}}
\tag{3.17}
$$

$$
\tilde{c}_2(t) = \tilde{c}_2(0) = 0
\tag{3.18}
$$

with the initial values $c_b(0)$ and $c_d(0)$ found by setting $t = 0$ and solving using the given values for $\tilde{c}_0(0)$ and $\tilde{c}_1(0)$. This is, as expected, an oscillation and the frequency Ω_- can be simplified in this limit. The effective Rabi frequency (the frequency at which the population oscillates) becomes

$$
\begin{aligned}
\Omega_{\text{Raman}} = \quad \Omega_- &= \frac{\Delta - \sqrt{\Omega_0^2 + \Delta^2}}{2} = \frac{\Delta - \Delta\sqrt{1 + (\Omega_0^2/\Delta^2)}}{2} \\
&\simeq \frac{\Delta - \Delta\left[1 - \frac{1}{2}\left(\Omega_0^2/\Delta^2\right)\right]}{2} = \frac{\Omega_0^2}{4\Delta}
\end{aligned}
\tag{3.19}
$$

or, expressed in terms of the individual Rabi frequencies,

$$\Omega_{\text{Raman}} = \frac{\Omega_0^2}{4\Delta} = \frac{\Omega_P^2 + \Omega_S^2}{4\Delta} \tag{3.20}$$

This, quite unexpectedly, suggests that there can be some coupling between states $|0\rangle$ and $|1\rangle$ when one of the beams is not present. This is clearly unphysical, and the problem is resolved if we consider the *amplitude* of these oscillations.

If we choose $c_d(0)$ and $c_b(0)$ to be real, which we can do without loss of generality by including an additional phase term in the complex exponential, the peak-to-peak amplitude of oscillations in the probability $|\tilde{c}_0(t)|^2$ is $4\sin\theta\cos\theta c_b(0)c_d(0)$. Given all population is accounted for by states $|d\rangle$ and $|b\rangle$, the maximum value of $c_b(0) \times c_d(0)$ is $\frac{1}{2}$, corresponding to an equal superposition of bright and dark states. The maximum peak-to-peak modulation depth is therefore $\sin 2\theta$. This is largest for equal coupling strengths Ω_P and Ω_S, and falls to zero as either becomes small.

In the case of equal coupling strengths $\Omega_P = \Omega_S = \Omega$ (and $\Omega_P^2 + \Omega_S^2 = 2\Omega^2$), the Rabi frequency of the effective two-level atom is [32, Eqn 13.4-4]:

$$\Omega_{\text{Raman}} = \frac{\Omega^2}{2\Delta} = \frac{\Omega_P \Omega_S}{2\Delta}. \tag{3.21}$$

The effective Rabi frequency in general depends on the individual coupling strengths and is not in simply their product divided by twice the detuning. Many treatments (notably Shore [32, §13.4, Eqn 13.4-3]) assume equal coupling strengths, but this is not necessarily true experimentally. Similar derivations can be found in Atomic Physics by Foot [67, Appendix E] and Linskens *et al* [68, § II.A].

In Section 3.5 we extend this treatment to account for more than one intermediate state, as is the case for most Raman transitions in the alkali atoms.

3.2 Atomic Structure

Atomic states can be calculated as solutions to the time-independent Schrödinger equation in the potential of the charged atomic nucleus, and can be parameterised by a succession of quantum numbers. The general many-electron solution is extremely involved and here we restrict our discussion to hydrogen-like atoms of which the alkali metal atoms, with closed shells and one outer electron, are a good approximation.

The first such quantum numbers are the principal quantum number n, corresponding roughly to the radial extent of the orbital, and the magnitude of the orbital angular momentum L, on which the angular shape of the orbital depends.

Classically, the angular momentum is $\boldsymbol{L} = \boldsymbol{r} \times \boldsymbol{p}$ which, in a quantum mechanical treatment, becomes $\hat{\boldsymbol{L}} = -i\hbar \boldsymbol{r} \times \boldsymbol{\nabla}$. The potential of the nucleus is spherically symmetrical so we need only consider one component of the vector $\boldsymbol{L}$, and we choose the component along the z axis: $\hat{L}_z = -i\hbar(\partial/\partial\phi)$. This tells us the direction of the angular momentum, and we can use a similar approach to find $|\boldsymbol{L}|^2$, which tells us the magnitude. Neither of these acts on the radial part of the wave-function, and the angular functions which are simultaneously eigenfunctions of both are the spherical harmonics: $\hat{L}_z Y_L^{m_L} = \hbar m_L Y_L^{m_L}$ and $\hat{L}^2 Y_L^{m_L} = \hbar^2 L(L+1) Y_L^{m_L}$. The value of L is restricted to integers or half-integers, and the projection m_L to integer-spaced values between $\pm L$. The shapes of these orbitals have tangible shapes, and their overlap integrals account for the coarse features of the dipole transition strengths.

The electron does not only have orbital angular momentum, but also an intrinsic spin $|\boldsymbol{S}| = \frac{1}{2}\hbar$. We must appeal to relativistic quantum mechanics to account for this property, but we can accept its existence and include it ad-hoc in the non-relativistic treatment. This spin magnetic dipole interacts with the orbital magnetic dipole (due to the 'current-loop' of electron orbiting around the nucleus), and the direction of each dipole changes in time. The system is closed, however, and the *total* angular momentum $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$ remains constant. We hence parameterise the state by this total angular momentum and we call it a 'good' quantum number. Its magnitude and projection are subject to

the same rules as for L and m_L. Strictly, one should speak of angular momentum operators and eigenvectors of the Hamiltonian, but this 'vector-model' provides an accessible classical analogy to the underlying physics; this is presented in Woodgate [69, §4].

The coupling between $\boldsymbol{L}$ and $\boldsymbol{S}$ splits the energy levels to states where $\boldsymbol{S}$ is either parallel or anti-parallel to $\boldsymbol{L}$, and these are typically many THz apart. We can easily address particular transitions by tuning our laser. For our experiments, a far more relevant interaction is the hyperfine interaction. Here, we add the angular momentum of the nucleus $\boldsymbol{I}$ and the total angular momentum becomes $\boldsymbol{F} = \boldsymbol{J} + \boldsymbol{I}$. This splitting is typically GHz, and as described in Chapter 6 we use frequencies derived from a single laser to address both levels.

Now, finally, we consider the energy shift when the atom is subjected to an external magnetic field $\boldsymbol{B}$. The total angular momentum $\boldsymbol{F}$ interacts with this external field, and precesses about it at a rate depending on the projection m_F. This is the Zeeman interaction; each level is split into $2F + 1$ equally spaced levels, one for each of the possible m_F states between $-F$ and $+F$. By applying an external magnetic field we create a sufficiently large frequency shift that we can address these levels by tuning the frequency difference of our Raman beams.

Spectroscopic notation is a convenient way to represent quantum numbers up to LS coupling; they are written are written as $n^{2s+1}L_J$ with the replacement $L = (0, 1, 2, 3 \ldots) \mapsto (\text{S}, \text{P}, \text{D}, \text{F} \ldots)$. We use this and the numbers F, m_F to label our states. We are typically interested in states $\left|5^2\text{S}_{1/2}, F = (2, 3), m_F\right\rangle$ and $\left|5^2\text{P}_{3/2}, F' = (1, 2, 3, 4), m'_F\right\rangle$.

3.3   Dipole Operator

The absorption of light by an atom depends on the overlap of the field and the absorption spectra, but the overall strength depends on the dipole moment. This property can be found using the dipole operator $\hat{\mu}$ via the overlap integral $\mu = \langle\psi_{\text{final}}|\, \hat{\mu}\, |\psi_{\text{initial}}\rangle$; $\hat{\mu} = q\hat{x}$ where q is the fundamental charge and $\hat{x}$ is the position operator.

It is possible to separate this integral using the Wigner-Eckart theorem [69–71] and to extract from it a purely geometrical term, leaving a term which embodies other physical details of the transition:

$$\langle 5^2 P_{3/2}, F', m'_F | \hat{\mu} | 5^2 S_{1/2}, F, m_F \rangle = \langle J || \hat{\mu} || J' \rangle \times \mathrm{G}\left(I, J, F, m_F, J', F', m'_F, q\right) \quad (3.22)$$

where I is the nuclear spin. The 'reduced' matrix element, denoted by double bars $||$, depends on many details of the atom, including nuclear mass, and is not easily calculated; it can, however, be found experimentally from measurements of the upper-state lifetime, as described by Loudon [72, Eqn 2.57] and Demtröder [33]:

$$\Gamma = \frac{16\pi^3}{3\epsilon_0 h \lambda^3} \frac{2J+1}{2J'+1} \left| \langle J || \hat{\mu} || J' \rangle \right|^2 . \quad (3.23)$$

The second part is the product of geometrical terms:

$$\mathrm{G} = (-1)^{2F'+J+I+m_F} \sqrt{(2F'+1)(2F+1)(2J+1)} \times$$
$$\begin{pmatrix} F' & 1 & F \\ m'_F & q & -m_F \end{pmatrix} \begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix} \quad (3.24)$$

where the array-like symbol in brackets $(\ldots)$ is the Wigner 3-j symbol, and the similar term in curly brackets $\{\ldots\}$ is the Wigner 6-j symbol [73, §3.3]; both are closely related to the Clebsch-Gordan coefficients, and this is described in Edmonds [74].

This geometrical part describes the different coupling between hyperfine levels of definite angular momentum by a photon whose angular momentum is labelled by q as either left-circular σ_+ ($q = +1$), right-circular σ_- ($q = -1$), or linear π ($q = 0$). G naturally includes rules such as $m'_F = m_F + q$, which can be interpreted as conservation of angular momentum.

Linear polarisation has a very precise meaning in this context: linearly polarised light with the electric field oscillating parallel to the quantisation axis, provided by an external magnetic field, is labelled $q = 0$, but light with electric field orthogonal to the quantisation axis is described, and is seen by the atom, as a superposition of left- and

right- circularly polarised light. Any representation of two orthogonally linearly polarised waves in this circular basis must also be orthogonal and we hence label the possible polarisations as:

$$q \; = \; 0 : \; \pi_0 \tag{3.25}$$

$$q \; = \pm 1 : \; \sigma_\pm \tag{3.26}$$

$$\pi_\pm = \left(\sigma_+ \pm \sigma_-\right)/\sqrt{2}. \tag{3.27}$$

For single photon transitions, we evaluate this dipole matrix element and the transition rate is proportional to the modulus squared. However, for a two-photon Raman transition between sub-levels in the ground-state, there are multiple routes (via the sub-levels in the excited-state) which must be considered; this is discussed in Section 3.5

3.4 Rubidium

Rubidium ($Z = 37$) belongs to the alkali metals, all of which have one valence electron outside closed shells. At typical energies, the electrons in the closed shells do not participate, and only the outer electron need be considered. For each electron in each closed shell, there is another with opposite spin, and so the total spin angular momentum contribution of these electrons is zero; all of the angular momentum properties of the electron cloud are accounted for by the single outer electron [39, 69].

The relevant states of the atom, accounting for only this outer electron, are shown in Figure 3.2. The interactions which give rise to this structure have important consequences for our attempts to drive Raman transitions and so these are discussed below.

After LS coupling has been considered, one finds a single ground-state $5^2\mathrm{S}_{1/2}$ and two excited-states $5^2\mathrm{P}_{1/2}$ and $5^2\mathrm{P}_{3/2}$; the associated transitions are at 794.8nm and 780.2nm respectively, and are called the D1 and D2 lines. We work exclusively with the D2 line because of the closed transition between $\left| 5^2\mathrm{S}_{1/2}, F = 3 \right\rangle$ and $\left| 5^2\mathrm{P}_{3/2}, F' = 4 \right\rangle$, which we
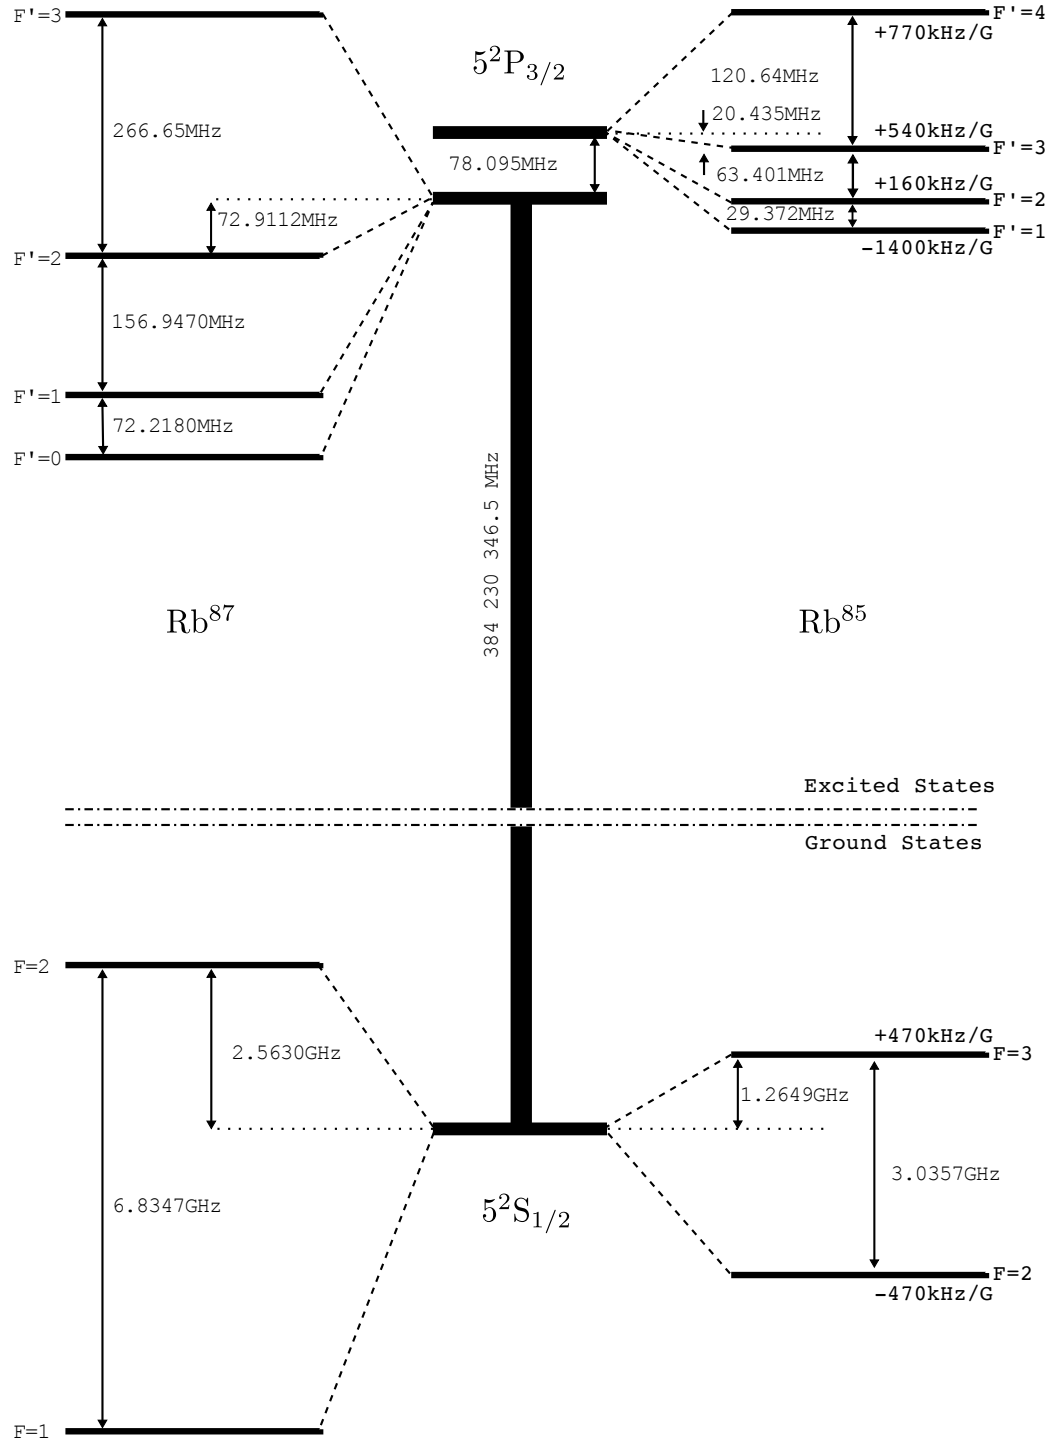
FIGURE 3.2: The relevant energy levels of rubidium 87 and 85 isotopes [73, 75]. The Zeeman shifts are given for the levels in rubidium 85, in terms of kilohertz per Gauss, which should be multipled by the m_F quantum number to find the shifted frequency.

$$\begin{aligned}
\Delta J &= 0, \pm 1 \quad \text{excluding } J = 0 \text{ to } J' = 0 \\
\Delta F &= 0, \pm 1 \quad \text{excluding } F = 0 \text{ to } F' = 0 \\
\Delta m_F &= q
\end{aligned}$$

TABLE 3.1: Summary of selection rules for electric dipole transitions [67].

$5^2\mathrm{P}_{3/2}$	$\|F' = 2\rangle$		$\|F' = 3\rangle$	
$5^2\mathrm{S}_{1/2}$	$\|F = 2\rangle$	$\|F = 3\rangle$	$\|F = 2\rangle$	$\|F = 3\rangle$
$m_F = -2$	$\sqrt{7/54}$	$\sqrt{10/189}$	$\sqrt{2/135}$	$\sqrt{25/216}$
-1	$\sqrt{7/36}$	$\sqrt{2/63}$	$\sqrt{2/45}$	$\sqrt{5/36}$
$+0$	$\sqrt{7/36}$	$\sqrt{1/63}$	$\sqrt{4/45}$	$\sqrt{5/36}$
$+1$	$\sqrt{7/54}$	$\sqrt{1/189}$	$\sqrt{4/27}$	$\sqrt{25/216}$
$+2$			$\sqrt{2/9}$	$\sqrt{5/72}$

TABLE 3.2: Coupling strengths for the transitions shown in Figure 3.3, where Raman transitions are driven by equal circular polarisation σ_+, σ_+, and $\Delta m_F = 0$ [73].

use for cooling, and the reduced dipole matrix element for this transition is [73, 76]:

$$\langle J || \hat{\mu} || J' \rangle = 4.227 a_0 q . \tag{3.28}$$

The $5^2\mathrm{S}_{1/2}$ level is split by the hyperfine interaction into two levels, $|F = 2\rangle$ and $|F = 3\rangle$, separated by 3035732440Hz [77] and $5^2\mathrm{P}_{3/2}$ is split into $|F' = 1 \ldots 4\rangle$, each of which is further split in to $2F + 1$ magnetic sub-levels. This structure is illustrated in Figure 3.2.

A zero of the geometrical coefficient G in Equation 3.24 corresponds to a forbidden transitions, and the rules which must be satisfied if this coefficient is to be non-zero are called selection rules; these are summarised in Table 3.1. We see that for most combinations of polarisations q_1, q_2 and m_F states, there are two upper states $|F' = (2, 3), m_F + q_1\rangle$ via which the transition $|F = (2, 3), m_F\rangle$ to $|F = (3, 2), m_F + q_1 - q_2\rangle$ is driven, and this is shown in Figure 3.3. If we are to understand Raman transitions in our system, we must account for these two possible routes.
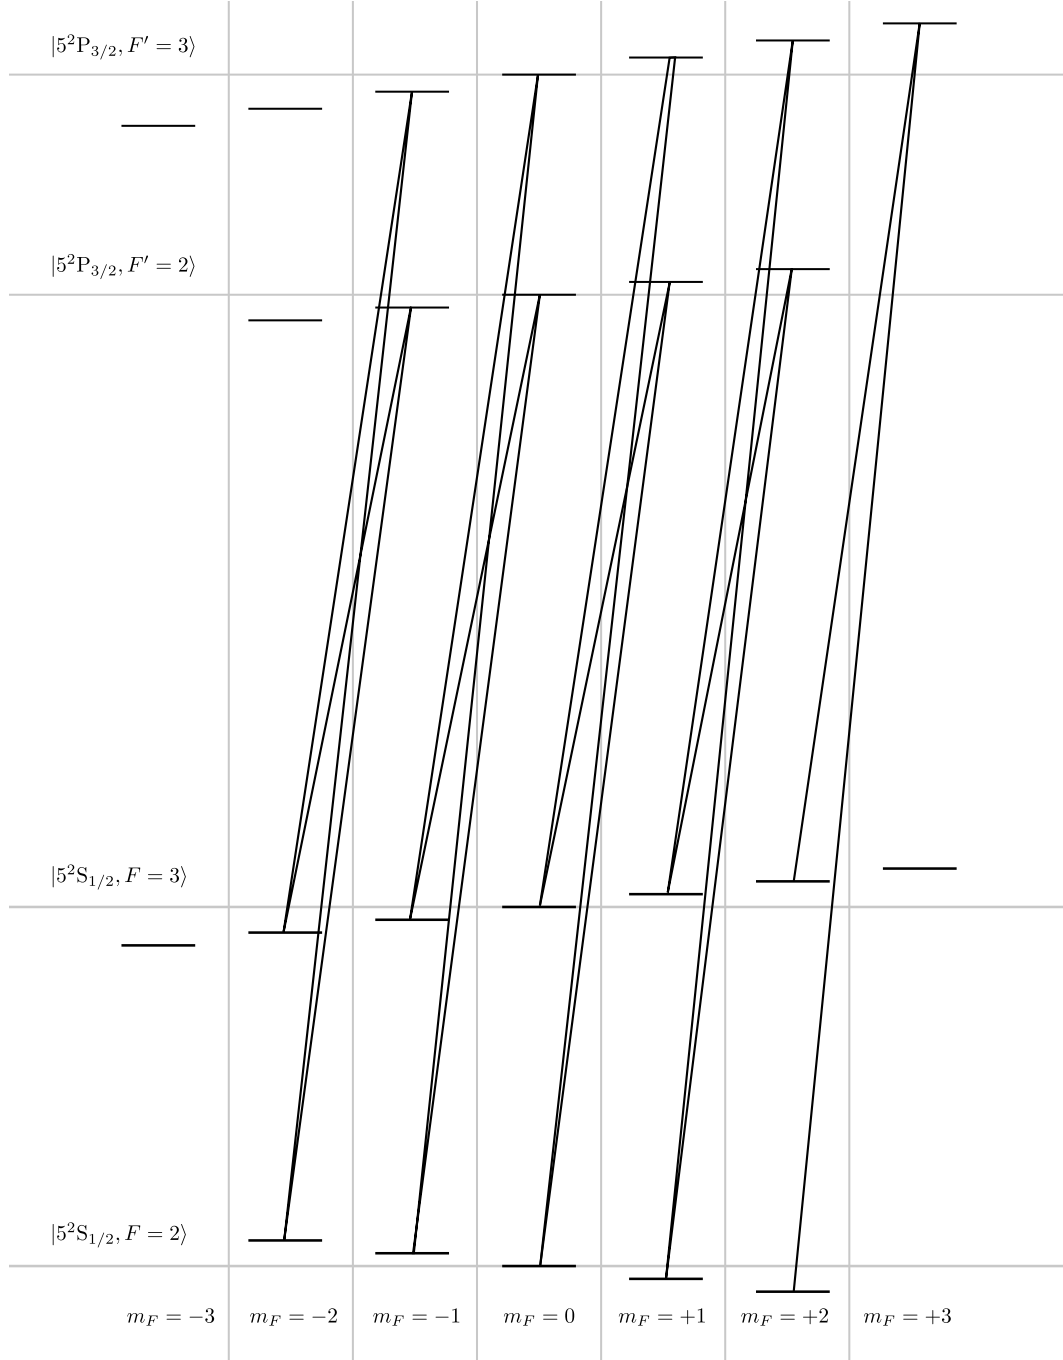
FIGURE 3.3: The states in Rubidium 85 via which our Raman transitions, using equal circular polarisations σ_+,σ_+, are driven. The separation between upper states is 63MHz and between lower states is the much larger 3036MHz, so the frequency differences are not shown to scale. The separation between $5^2\mathrm{P}_{3/2}$ and $5^2\mathrm{S}_{1/2}$ is much larger still, at 384THz. The Zeeman-shifts are approximately to scale, and importantly have opposite sign in the lower hyperfine states. Note that for the extreme $m_F = +2$ state, there is only one upper level which couples both states. States $\left|5^2\mathrm{P}_{3/2}, F' = (1,4)\right\rangle$ are not shown, as the rule $\Delta F = 0$ ensures each of these couples only to a single lower state. The coupling strengths are listed in Table 3.2.

3.5   Multiple Raman Routes

The three-level treatment in Section 3.1 is, unfortunately, insufficient to describe our system; we must extend it to include four levels. As discussed Section 3.4, the structure of our atom unavoidably means there are often two upper states via which the Raman transition may be driven. The exceptions are the extreme m_F states $\left|5^2S_{1/2}, 2, \pm 2\right\rangle$ coupled to $5^2P_{3/2}$ via $\sigma_{\pm}$; each of these couples to a single upper level $\left|5^2P_{3/2}, 3, \pm 3\right\rangle$.

To fully describe the situation we must include four levels, but the extension of the three-level treatment to four levels is not straightforward. Four-level systems [78] and the interference of the two routes between states in situations like our own [79] have been studied, along with some more general transformations of the topology of such general systems to ones more easily analysed [80]. However, the separation between the two upper levels in our system is small compared with the typical detuning Δ, and we expect that, in this limit, the separation will not be resolved and overall the two levels will behave effectively as one. We hence write the effective Rabi frequency of this effective three-level system as in Equation 3.20, but now with an additional factor χ:

$$\Omega_{\text{Raman}} = \frac{\Omega_P^2 + \Omega_S^2}{4\Delta}\chi. \tag{3.29}$$

The parameter χ depends on the coupling strength of the individual transitions and so depends on the initial and final m_F states and the polarisations of fields. We anticipate that as $\Delta \to \pm\infty$, χ will tend to a constant. This is illustrated in Figure 3.4.

The four level atom may be difficult to tackle analytically, but it poses no additional problems for numerical simulation, especially if we are content to restrict our analysis to the experimentally relevant situation where we begin in some superposition of the ground states. The four-level Hamiltonian is

$$\boldsymbol{H} = \begin{pmatrix} \omega_0 & 0 & g_{2,0}\Omega_P \cos\omega_P t & g_{3,0}\Omega_P \cos\omega_P t \\ 0 & \omega_1 & g_{2,1}\Omega_S \cos\omega_S t & g_{3,1}\Omega_S \cos\omega_S t \\ g_{2,0}\Omega_P \cos\omega_P t & g_{2,1}\Omega_S \cos\omega_S t & \omega_2 & 0 \\ g_{3,0}\Omega_P \cos\omega_P t & g_{3,1}\Omega_S \cos\omega_S t & 0 & \omega_3 \end{pmatrix} \tag{3.30}$$
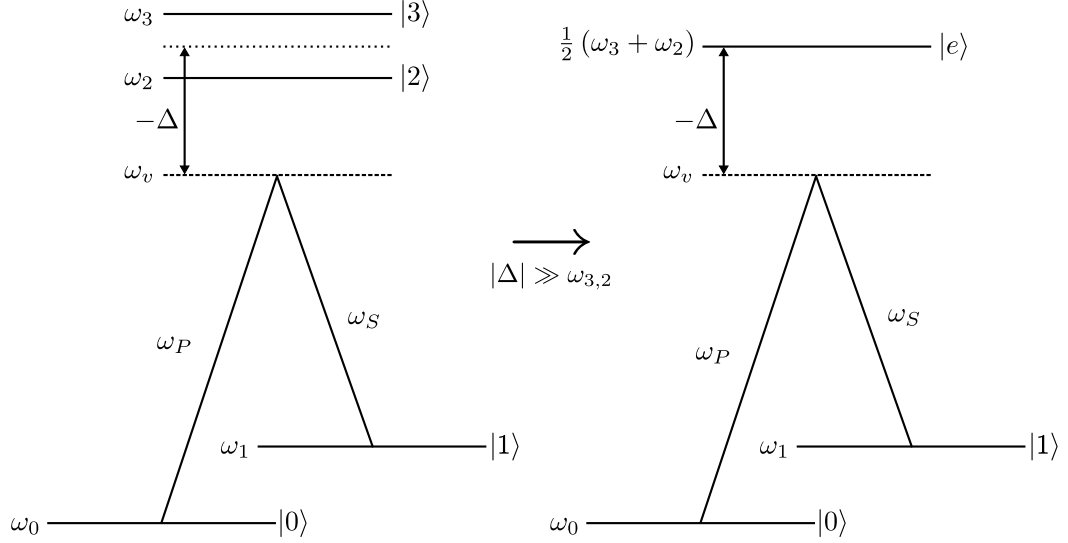
FIGURE 3.4: As Δ, measured relative to the mid-point between levels $|2\rangle$ and $|3\rangle$, becomes much larger than the separation between these levels, they behave as an effective single level $|e\rangle$ and we can approximate the four level system to a three-level system with an effective two-level Rabi coupling between the ground states.

where $g_{i,f}$ are the geometrical part of the matrix element G (Equation 3.24) for the states $|i\rangle$ and $|f\rangle$. We have separated out these geometrical terms and the Rabi frequencies $\Omega_{P,S}$ now only contain the reduced matrix element:

$$\Omega_{P,S} = \frac{|\langle J||\hat{\mu}||J'\rangle|}{\hbar} E_{P,S} \ . \tag{3.31}$$

After moving to the interaction picture and making the rotating wave approximation (as in the three-level treatment) this becomes

$$\tilde{\boldsymbol{H}} = \begin{pmatrix} 0 & 0 & \frac{1}{2}g_{2,0}\Omega_P & \frac{1}{2}g_{3,0}\Omega_P \\ 0 & 0 & \frac{1}{2}g_{2,1}\Omega_S & \frac{1}{2}g_{3,1}\Omega_S \\ \frac{1}{2}g_{2,0}\Omega_P & \frac{1}{2}g_{2,1}\Omega_S & \Delta_2 & 0 \\ \frac{1}{2}g_{3,0}\Omega_P & \frac{1}{2}g_{3,1}\Omega_S & 0 & \Delta_3 \end{pmatrix} \tag{3.32}$$

where $\Delta_{2,3}$ are the detunings from two the states $|2,3\rangle$ via which $|0\rangle$ and $|1\rangle$ are coupled.

For a range of detunings Δ, with $\Omega_P = \Omega_S = 10^{-6}\Delta$ to satisfy the assumption that we are far detuned, we calculate the eigenvalues and eigenvectors of the matrix $\tilde{\boldsymbol{H}}$. The eigenvectors tell us how the bare-states are mixed and the eigenvalues tell us at what

frequencies the phases of these superpositions oscillate. If there is no mixing of the bare states, then there are no oscillations of population in our system. Parameters for the distinct combinations of polarisations (q_P, q_S) are shown in Appendix A. For each case where the eigenstates are mixtures of the bare states, two eigenstates can be expressed as a rotation of the two bare ground states, and the other two are simply equal to the upper states. If we parameterise this rotation by an angle θ, we have equations of motion very similar to those found for the three-level system:

$$\tilde{c}_0(t) \; = \; \cos\theta v_0(0)e^{-i\lambda_0 t} - \sin\theta v_1(0)e^{-i\lambda_1 t} \tag{3.33}$$

$$\tilde{c}_1(t) \; = \; \cos\theta v_1(0)e^{-i\lambda_1 t} + \sin\theta v_0(0)e^{-i\lambda_0 t} \tag{3.34}$$

$$\tilde{c}_2(t) \; = \; 0 \tag{3.35}$$

$$\tilde{c}_3(t) \; = \; 0 \tag{3.36}$$

where v_i and λ_i are the eigenvectors and eigenvalues of $\tilde{\boldsymbol{H}}$. The population in $|0\rangle$ hence oscillates at frequency $(\lambda_1 - \lambda_0)$ with amplitude, as previously:

$$\cos\theta \sin\theta v_0(0)v_1(0) \leq \sin 2\theta \; . \tag{3.37}$$

The rotation matrix, angle, maximum modulation depth (for equal Ω_P and Ω_S) and the parameter χ relating the naïvely calculated effective Rabi frequency to that found via this numerical method, are shown in Appendix A.

This approach reproduces the result of the analytical approach to the three-level atom, and was also extended to include more than two upper-levels. This extension is necessary to understand transitions driven by linearly polarised light with electric field perpendicular to the quantisation axis, where there are a total of six states. Transitions are only driven by orthogonal linear polarisations, and the table of parameters is included in Appendix A.

Experimentally, we do not yet have independent control over the polarisations of our Raman beams, so we must use equal circular polarisations. From these calculations, we see that for equal intensities we can expect $> 90\%$ modulation depth, and an effective Rabi frequency $0.18 \times \Omega^2/(2\Delta)$ where $\Omega_P = \Omega_S = \Omega$. This can be expressed in terms

of the more useful experimental parameter, intensity. This two-photon effective Rabi frequency is

$$\Omega_{\text{Raman}} = \chi \frac{|\langle J||\hat{\mu}||J'\rangle|^2}{c\epsilon_0\hbar^2}\frac{I}{\Delta} \tag{3.38}$$

where I is the intensity of each Raman beams. For the rubidium 85 D2 transition $\langle J||\hat{\mu}||J'\rangle = 4.227a_0q$ and this becomes

$$\Omega_{\text{Raman}} = \chi \frac{(2\pi \times 1.05\text{MHz})^2}{\Delta}I / \left(\text{Wm}^{-2}\right). \tag{3.39}$$

The maximum modulation depths reported in Appendix A are not symmetrical about $m_F = 0$. This is perhaps unexpected, and an experiment against which the model can be tested was reported by Dotsenko *et al* [81]. The authors describe an experiment where population is prepared in a particular lower hyperfine ground state and then transferred by a Raman pulse whose frequency difference is scanned across the Zeeman-shifted m_F levels. They use caesium rather than rubidium, but the model is easily adapted; coupling strengths are from Steck [82]. The optical powers were 120μW and 360μW in a Gaussian beam with a 70μm waist. The detuning of $-2\pi \times 150$GHz firmly places the system in the far detuned regime, and we can use the model to predict the maximum fraction which will be transferred for each m_F state. This comparison is shown in Figure 3.5. The authors modelled the upper two-levels as having an effective coupling strength, and the clear downwards slope for high m_F states, which is predicted by this model, was attributed to optical pumping.

3.6 Raman Transitions in a MOT

If we wish to resolve the individual m_F levels, we must apply a magnetic field to lift the Zeeman degeneracy. Moreover, if the atoms are in a region of zero magnetic field, the absence of a quantisation axis implies that the interaction of the atoms with circularly polarised light will be the same as for linear polarisation, and we cannot expect our equally circularly polarised fields to drive Raman transitions.
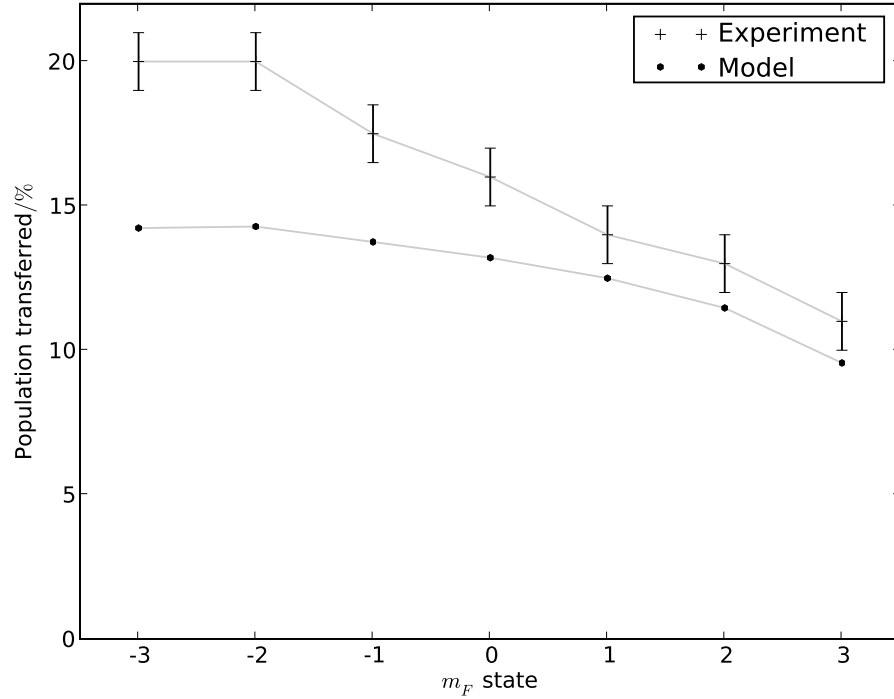
FIGURE 3.5: A comparison of the measured population transferred by Raman transitions in caesium, reported in [81], with that predicted by the effective three level model. Naïvely, we expect equal peak heights of 14%, corresponding to the population prepared equally across the Zeeman sub-levels. The clear downwards trend for high m_F states, previously attributed to optical pumping, is reproduced, and the difference can be accounted for by a 10% uncertainty in Rabi frequencies and a small $\sim 2\%$ offset from optical pumping.

The dynamics of a magneto-optical trap ensure that atoms accumulate around zero magnetic field and if we apply a static field, we will simply shift the position of this zero. Elsewhere, atoms are typically collected in a MOT and then loaded into a dipole trap where they can be held in position independently of the magnetic field [81, 83]. We would rather not introduce this additional stage, and so in the following sections I describe two methods which displace the cloud relative to the magnetic field zero. Of the two methods, we find the oscillating magnetic field method to be more experimentally useful, but the shift due to imbalanced intensities is also relevant to the experiment.

3.6.1 Oscillating Magnetic Field

An additional external magnetic field applied to the MOT shifts the position of the magnetic field zero, and within certain limits, the cloud follows. However, we require only that the atoms are in a magnetic field during any coherent manipulation; we can imagine turning on a magnetic field in the short time between extinguishing the MOT beams and beginning the coherent manipulation experiments. This approach is possible [84], but switching significant currents through large inductive loads is not experimentally trivial, especially if we are to repeat the entire pulse sequence quickly ($\sim 100\text{Hz}$) to average experimental noise.

An alternative to this step-change is to apply a smoothly oscillating magnetic field. We reason that the dynamics of the trap are such that although the cloud will follow a slowing moving magnetic field zero, the atoms will be unable to follow a sufficiently high frequency oscillation and will respond only to its time-average. By extinguishing the MOT beams and performing our coherent manipulation experiments near the maximum of such an oscillation, we can ensure there is a well defined quantisation axis and a steady-state MOT.

3.6.2 Imbalanced intensities

The equations derived in Section 2.3 for the force on an atom make a crucial assumption: that the intensities of the counter-propagating beams are equal. If we now relax this constraint, we can engineer the position dependent force to be zero at some position offset from the origin and, crucially, offset from the position where the magnetic field is zero.

We add the forces from two counter-propagating beams as in Equation 2.29, but now allow the intensities to be unequal:

$$F_{\text{total}} = (1 + m)F(\alpha\Gamma + \delta\Gamma) - (1 - m)F(\alpha\Gamma - \delta\Gamma) \tag{3.40}$$
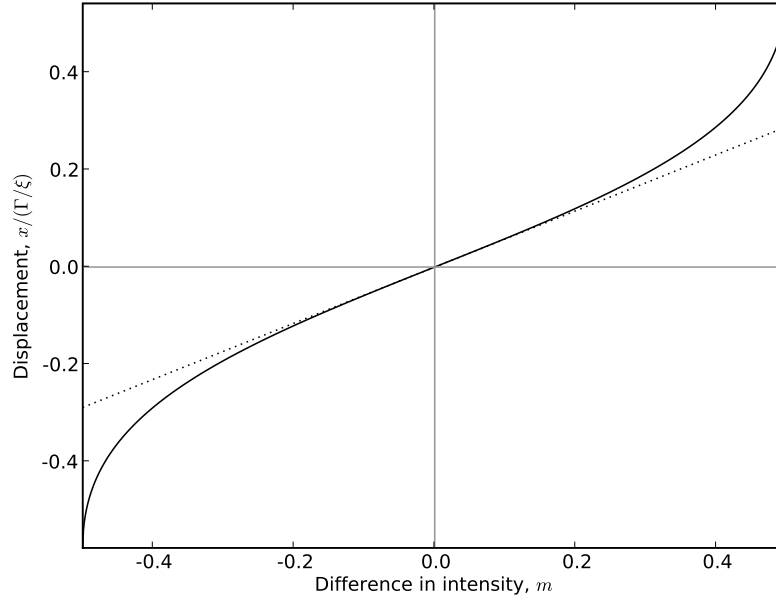
FIGURE 3.6: Plot of how the displacement of the position where forces are balanced depends on the imbalance in intensities, parameterised by m, for detuning $\alpha = \alpha_{\max}^{-}$. The linear approximation for small m (dotted) is compared with the actual curve (solid).

with $m = 0$ for equal intensities. The Taylor expansion now gives a non-zero force at the origin:

$$F_{\text{total}}(\alpha\Gamma) \simeq F_0 \frac{2m}{1 + (2\alpha)^2} - F_0 \frac{16\alpha}{[1 + (2\alpha)^2]^2}\delta \tag{3.41}$$

where we are considering only the spatial contribution $\delta = x/(\Gamma/\zeta)$. Using the gradient at the origin, which to first order is unaffected by this imbalance, we estimate the position where the force *is* zero:

$$x = (\Gamma/\zeta) \times m \frac{1 + (2\alpha)^2}{8\alpha} \ . \tag{3.42}$$

This linear approximation, compared with the actual curve, is shown in Figure 3.6. A result of this is that spatial structure of a laser beam leads to an equilibrium position which varies over the cross-section, and hence introduces unwanted structure to the cloud. As described in Sections 5.5.1 and 7.1, we found spatial filtering using a short stretch of single mode optical fibre greatly improved the shape of our cloud.

Chapter 4

Adiabatic Manipulation Theory

In Chapter 3 we showed how the ground hyperfine states of trapped atoms can be used to emulate a nearly ideal two level system. We now turn our attention to how we can manipulate this system, and in particular how we may perform the delicate operation of creating a superposition using imperfect experimental apparatus.

The work presented in this chapter originally appeared as a paper in Physical Review A [85], to which I was the main contributor. This work grew out of the perceived need to robustly create and manipulate superpositions of the ground hyperfine levels in our cold rubidium sample. Other work in this area generally uses STIRAP [86–88] and related techniques to create superpositions, but these inherently require more than two-levels. We are using rubidium to emulate a two-level system, and for our work to be generally applicable to two-level like systems, we required an alternative to STIRAP; we need to make superpositions in a simple two-level system, while ensuring the pulse is robust against experimental imperfections. Since publication, the Bakos group has published a paper on similar work [89].

The remainder of this chapter is an adapted version of the paper.

4.1 Introduction

The central elements of an optical interferometer are the mirrors and beam-splitters that deflect, divide and recombine the incident light. In an atom interferometer [90, 91], the corresponding elements are those that respectively switch prepared atoms between a pair of quantum states, and convert between either of the two eigenstates and equal superpositions of defined relative phase. The simplest implementations couple the two states using an optical field which, when resonant, results in Rabi oscillations. With judicious control of the optical intensity and duration, so that the interaction lasts just a half or quarter Rabi cycle, simple laser pulses achieve the required operations. The Rabi frequency depends, however, on the laser intensity and apparent frequency, which vary with the position and Doppler shift of the atoms in the laser beam. While simple interferometers are possible, variations in the prepared state can be significant.

Rapid Adiabatic Passage (AP) [86, 88, 92, 93] – which is quick in comparison with any incoherent processes (such as spontaneous emission) but slow in comparison with the Rabi oscillation – depends on off-resonant coupling of the two states and achieves a steady-state superposition that depends on the detuning between the laser and the atomic resonance. By varying the detuning slowly, population can be transferred with arbitrary fidelity from one state to the other. In the common Feynman, Vernon and Hellwarth pseudo-polarisation representation [94] of a two-level superposition as a 'state vector', the superposition precesses about an adjustable 'field vector' defined by the phase, intensity and detuning of the optical field. Provided that the field vector is varied sufficiently slowly, the angle between it and the precessing state vector remains fixed. Control of the azimuthal angle of the state vector around the field vector is lost during such an operation, but if the vectors are initially aligned, the azimuthal angle is inconsequential. A less than perfectly adiabatic operation, however, introduces an uncontrolled change in the angle between the vectors and hence increases the parameter-space to which the state is mapped; this renders the process irreversible and non-adiabatic.

Given initial alignment of the state and field vectors and a sufficiently slow variation of the latter, the superposition can be made to follow an essentially arbitrary path:

the state vector remains at all times parallel to the field vector, and the evolution is deterministic and adiabatic. For the mirrors of an atom interferometer, this path usually leads from one eigenstate to the other – it causes the complete inversion of the atomic population and, since the initial and final field vectors correspond to large detuning and zero intensity, the process is extremely insensitive to experimental variations. For the beam-splitters, however, the path instead ends or begins with an equal superposition: the final field must be exactly resonant, and the process is somewhat less robust.

In practice, the slow variation of the field vector desired for adiabaticity must be weighed against the need to complete the operation (and indeed the interferometry itself) before spontaneous emission and other processes of decoherence can occur. There may also be constraints from the finite duration of the atoms within the apparatus, or the accrual of, for example, velocity-dependent phases. Simple implementations of rapid adiabatic passage, such as linear chirps at constant intensity, satisfy the adiabaticity criterion to varying extents at different stages of the operation, and the overall fidelity can be far from optimum for the particular experimental constraints. Furthermore, the common method of creating superpositions – simply by truncating a pulse that would otherwise have caused complete population inversion (e.g. reference [95]) – is far from adiabatic during the final extinction and can be highly sensitive to experimental variations. In this chapter, we therefore consider the design and optimisation of mirror and beam-splitter pulses that maintain uniform adiabaticity, and hence reduce the sensitivity of such processes to variations in experimental parameters.

Our attempt to control adiabaticity throughout a pulse is not the first [96, 97] – a well known pulse for which the adiabaticity is constant is the tangent frequency sweep [98], for which one uses a constant intensity (or some other appropriate measure of coupling strength) while the frequency difference from resonance follows an inverse-tangent sweep in time. We shall see that, as is demonstrated by this example, the necessary condition that the detuning from resonance relative to the coupling strength be large at the beginning and end of the operation means any 'pulse' employing constant intensity throughout inescapably requires a large, perhaps impractically so, frequency range. Furthermore,

we are unaware of any attempt to derive a pulse of constant adiabaticity for the more challenging case of fractional, rather than full, adiabatic passage.

We extend the idea of constant adiabaticity for constant intensity to any form of controlled adiabaticity for, crucially, a smoothly modulated intensity. As well as reducing considerably the required frequency range, the inclusion of the extinction in the pulse design now allows us to consider pulses for half, as well as full, adiabatic passage.

4.2 Adiabatic Passage and rotations

Our technique is designed to mimic, in the language of quantum computing, a qubit rotation, and so it is important to assess to the differences between adiabatic evolution and a rotation of the state vector in the Feynman representation [94]. We approach this by considering the rotation of the basis necessary to move from bare states to dressed states, followed by the adiabatic evolution of these dressed states, and finally the rotation from dressed states back to bare states [30]. The result is that, for adiabatic following, only the phase of the overlap of the state with the (time-dependent) dressed eigenstates changes in time. This corresponds to the familiar result that the angle between state and field vectors remains constant (e.g. reference [67]), but is derived here independently of this geometrical representation. This result is used to assess the suitability of adiabatic beam-splitter and mirror pulses for interferometry.

4.2.1 Adiabatic Evolution of Dressed States

Consider a two-level system with coupling strength described by the Rabi frequency Ω, which varies (slowly) in time. This coupling is detuned from resonance by Δ, which is also permitted to vary in time. The bare states of a system, $|0\rangle$ and $|1\rangle$, can be rotated to form the dressed states, $|-\rangle$ and $|+\rangle$ by

$$\begin{pmatrix} |-\rangle \\ |+\rangle \end{pmatrix} = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2)e^{i\rho} \\ \sin(\theta/2)e^{-i\rho} & \cos(\theta/2) \end{pmatrix} \begin{pmatrix} |0\rangle \\ |1\rangle \end{pmatrix} \tag{4.1}$$

where $\tan\theta = \Omega/\Delta$ and ρ is the relative phase of system and perturbation. We use the FM frame [97], for which ρ is constant throughout the pulse, and we choose $\rho = 0$. The coefficients ($A_{0,1}$ and $A_{-,+}$) transform in the same way, and their evolution is described by

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} A_- \\ A_+ \end{pmatrix} = \frac{i}{2}\begin{pmatrix} \tilde{\Omega} & i\dot{\theta} \\ -i\dot{\theta} & -\tilde{\Omega} \end{pmatrix}\begin{pmatrix} A_- \\ A_+ \end{pmatrix} \tag{4.2}$$

where (as before) $\tilde{\Omega} = \sqrt{\Omega^2 + \Delta^2}$ and an over-dot indicates the time-derivative. In the adiabatic approximation ($\dot{\theta} \to 0$) [30] this integrates to

$$\begin{pmatrix} A_-(t) \\ A_+(t) \end{pmatrix} = \begin{pmatrix} e^{+i\phi/2} & 0 \\ 0 & e^{-i\phi/2} \end{pmatrix}\begin{pmatrix} A_-(0) \\ A_+(0) \end{pmatrix} \tag{4.3}$$

$$\text{where } \phi = \int_0^t \tilde{\Omega}(t')\mathrm{d}t' \tag{4.4}$$

We use R_t to represent the rotation matrix of Equation 4.1, with the subscript t reflecting the time-dependence inherited from θ. B and D represent the (time-dependent) column vectors of bare and dressed states respectively, with the equivalent kets, $|B\rangle$ and $|D\rangle$, where appropriate. Finally, $T_{b,a}$ represents the adiabatic evolution of the dressed states in Equation 4.3 between times a and b. In this representation,

$$D_t = R_t \cdot B_t$$
$$D_t = T_{t,0} \cdot D_0 \tag{4.5}$$

The final bare state amplitudes are hence obtained by

$$B_t = R_t^{-1} \cdot T_{t,0} \cdot R_0 \cdot B_0 \tag{4.6}$$

4.2.2 Time evolution of overlap with a test state

To illustrate, we use a test state, $|S\rangle$, and find the overlap $\sigma_t = \langle S|B_t\rangle = S^\dagger B_t$. If $S = (1,0)^T$, $B_0 = (1,0)^T$ and $\theta_t = \theta_0 = \theta$ then, using equation 4.6,

$$\sigma_t = \cos(\phi/2) + i\sin(\phi/2)\cos(\theta) \tag{4.7}$$

For on-resonance excitation, $\cos(\theta) = 0$ and, using $\tilde{\Omega} = \Omega$ in equation 4.4, we recover Rabi flopping: $|\langle S|B_t\rangle|^2 = \cos^2(\Omega t/2)$. The oscillation depth reduces for $\cos(\theta) \neq 0$ and the procedure is similar for the overlap of B_t with any state in the bare-states representation.

A more useful result is found for the special case of projection along a dressed eigenstate. If U represents the dressed eigenstate in the dressed-state basis (i.e. $U = (1,0)^T$ or $(0,1)^T$) the projection is simply

$$\begin{aligned}
\sigma_t &= U^\dagger D_t \\
&= U^\dagger T_{t,0} R_0 B_0
\end{aligned}$$

The equivalent S vector is $S_t = R_t^{-1} U$. In this basis, the product $U^\dagger T_{t,0}$ is trivial: $\left(e^{i\phi/2}, 0\right)$ or $\left(0, e^{-i\phi/2}\right)$. That is, the only effect of time evolution on the projection, σ_t, is a phase evolution and $|\sigma_t|^2$ is independent of time.

So, adiabatic manipulation keeps the overlap of the state with the dressed eigenstates constant, and manipulation of the bare-state amplitudes can be achieved by altering the dressed eigenstates. When the initial state of an atom does not fully overlap with an initial dressed eigenstate, adiabatic manipulation introduces, via an uncontrolled phase, uncertainty in the (measurable) overlap with bare eigenstates. This behaviour corresponds to the state vector remaining at a constant angle from the field vector, around which it rapidly rotates; see figure 4.1.

To clarify this behaviour, the next section compares an adiabatic beam-splitter with a traditional $\pi/2$ pulse.

4.2.3 Comparison of AP and Rabi Pulses

Consider first the classic Rabi $\pi/2$ pulse which, as neither frequency or intensity is changing, can be treated as adiabatic (except for the initial switch on and final extinction). Excitation is on-resonance ($\theta = \pi/2$) and a well defined rotation is induced ($\phi = \pi/2$).
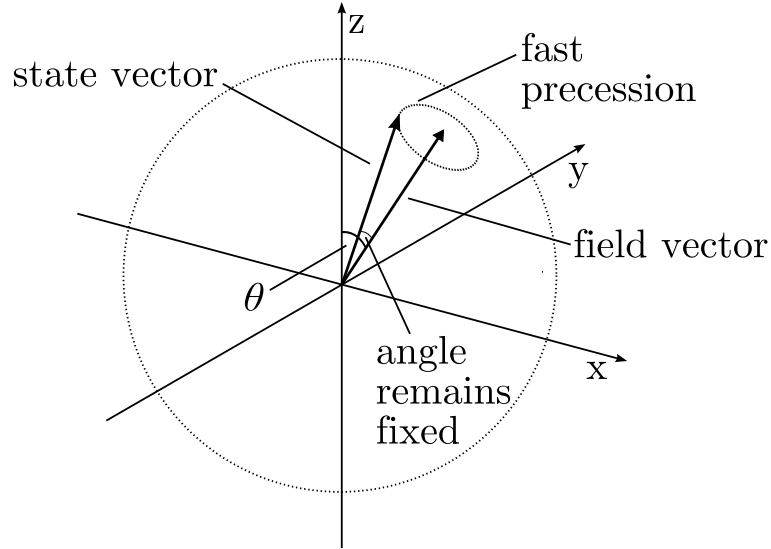
FIGURE 4.1: Adiabatic following in the Feynman representation.

Operation	$\pi/2$ Rabi pulses	AP based pulses
Initialise	$\lvert 0\rangle$	$\lvert 0\rangle$
First Pulse	$\lvert 0\rangle - i\lvert 1\rangle$	$e^{i\phi_1/2}(\lvert 0\rangle - \lvert 1\rangle)$
Free Evolution	$e^{i\gamma/2}\lvert 0\rangle - ie^{-i\gamma/2}\lvert 1\rangle$	$e^{i\phi_1/2}(e^{+i\gamma/2}\lvert 0\rangle - e^{-i\gamma/2}\lvert 1\rangle)$
Second Pulse	$i\sin\gamma/2\lvert 0\rangle - i\cos\gamma/2\lvert 1\rangle$	$e^{i\phi_-/2}i\sin\gamma/2\lvert 0\rangle - e^{i\phi_+/2}i\cos\gamma/2\lvert 1\rangle$
		where $\phi_- = \phi_1 - \phi_2$ and $\phi_+ = \phi_1 + \phi_2$

TABLE 4.1: Comparison of Rabi and AP based '$\pi/2$' pulses used to implement a simple interferometer. The state shown is that of the system after the stated operation. The state is initialised, the 'beam-splitter' pulse is applied, the system is permitted to evolve freely and accrue some phase, and finally the 'recombiner' pulse is applied to map this accrued phase to a (measurable) probability for each eigenstate.

The effect of two such pulses, with a time-delay between them, acting on a system initially in a bare eigenstate is shown in table 4.1. The accrued phase γ is mapped to the probability of inversion; this is, of course, a simple interferometer.

Next, consider the effect of two half AP pulses used in the same way. The first half AP pulse sweeps from far off resonance to on resonance ($\theta = 0 \rightarrow \pi/2$), and induces a large and uncertain number of precessions around the field vector; this is recorded in the phase ϕ_1. The second pulse now sweeps from on resonance to far off resonance ($\theta = \pi/2 \rightarrow \pi$), with a different large number of revolutions, ϕ_2. This progress is also shown in table 4.1.

For AP pulses, the precession around the field vector ($\phi_{-,+}$) is unknown and so, while the probability of inversion is identical to that obtained using the Rabi pulses, any

subsequent operation will be sensitive to this uncertain phase. Hence, we conclude that this manipulation technique may be limited to 'interferometers' containing two 'beam-splitters' only: a splitter, and a recombiner – the inverse of a beam-splitter in which the field vector sweeps adiabatically from horizontal to vertical.

In addition to beam-splitters, mirrors are usually required to build an interferometer. As before, these may be implemented with controlled area Rabi pulses (π-pulses), or full AP pulses. The Rabi version is a true rotation, but is experimentally fragile; the AP version, while robust, does not (in general) preserve the phase information recorded in γ. For example, the state following a half AP pulse and a full AP pulse is

$$e^{i\phi_1/2} \left(e^{-(\gamma+\phi_M)/2} |0\rangle - e^{+(\gamma+\phi_M)/2} |1\rangle \right)$$

where ϕ_M is the unknown phase of the full AP pulse.

All of these operations, including those of consecutive pulses, can be understood purely geometrically in the Feynman representation – one need only remember that, for any adiabatic operation, the angle between the state and the field vectors remains constant while the large and uncertain number of precessions loses all information about the azimuthal angle of the former around the latter.

4.3 Controlled Adiabaticity Pulse

Our analysis addresses the coupling of two non-degenerate quantum states, labelled $|0\rangle$ and $|1\rangle$, by a near-resonant optical field that is described at any time by its amplitude and frequency, which we write in the form of a coupling strength (or Rabi frequency) $\Omega(t)$ and detuning from resonance $\Delta(t)$. For the example of an electric dipole transition in an atom, the coupling strength, in terms of the time dependent amplitude of the electric field $E(t)$, the atomic transition frequency ω_0, and the dipole matrix element

$e \langle 1 | \hat{x} | 0 \rangle$, is [30]

$$\begin{aligned} \Omega(t) &= \frac{e \langle 1 | \hat{x} | 0 \rangle}{\hbar} E(t) \\ \Delta(t) &= \omega(t) - \omega_0 \end{aligned}$$

where $\omega(t)$ is the time-dependent excitation frequency.

In the Feynman pseudo-polarisation representation [94] the quadrature sum of these parameters $\tilde{\Omega} = \sqrt{\Omega^2 + \Delta^2}$ is the rate of precession of the state vector around the field vector, and the angle θ between field vector and the vertical (eigenstate) is given by $\tan \theta = \Omega / \Delta^1$. These parameters are the same as occur in the dressed states treatment, corresponding respectively to the energy splitting of the dressed states and the mapping of bare states on to dressed states [30, 39].

The condition for adiabaticity is that the rate of precession $\tilde{\Omega}$ be much greater than the rate of rotation of the field vector $\dot{\theta} \equiv \mathrm{d}\theta/\mathrm{d}t$ [30, 31]. A common measure is the adiabatic parameter $Q(t)$ [99] which, together for later convenience with its reciprocal $\epsilon(t)$, is defined by

$$Q(t) \equiv 1/\epsilon(t) \equiv \tilde{\Omega}/\dot{\theta} \tag{4.8}$$

When $Q \to \infty$ ($\epsilon \to 0$) the process is adiabatic. It follows that the time-dependent field parameters Ω and Δ are linked by this adiabatic parameter.

Our strategy is to take an experimentally straightforward amplitude Ω, constrain the adiabatic parameter ϵ to a given function, and hence derive the required frequency chirp Δ. The converse (where Δ is set and Ω calculated) is similarly possible, and follows the same procedure.

4.3.1 General Properties of the Pulse

We begin by noting that the relationships between θ and Ω & Δ may also be written as

$$\sin \theta = \left(\frac{\Omega}{\tilde{\Omega}} \right) ; \; \cos \theta = \left(\frac{\Delta}{\tilde{\Omega}} \right) ; \; \tan \theta = \left(\frac{\Omega}{\Delta} \right) \tag{4.9}$$

---

[1] We operate in the frequency modulation (FM) frame [97] which rotates at the instantaneous frequency of the pulse [99], and in which the direction of a field vector in the xy plane is fixed.

from which, using the definition of ϵ, we construct a differential equation:

$$\dot{\theta} = \epsilon\tilde{\Omega}$$

$$\Rightarrow \quad \dot{\theta} \,\sin\theta = \epsilon\tilde{\Omega}\left(\frac{\Omega}{\tilde{\Omega}}\right) = \epsilon\Omega \qquad (4.10)$$

with no specific constraint yet applied to the adiabatic parameter ϵ. We integrate from the time t_h at which detuning is zero and the field vector is horizontal ($\Delta(t_h) = 0$) to find an expression for $\cos\theta$ which, for brevity, we label as $\Gamma(t)$:

$$
\begin{aligned}
\left[\cos\theta(t')\right]_{t_h}^{t} &= -\int_{t_h}^{t} \dot{\theta}(t')\sin\theta(t')\mathrm{d}t' \\
&= \int_{t}^{t_h} \epsilon(t')\Omega(t')\mathrm{d}t' \qquad (4.11) \\
&\equiv \Gamma(t)
\end{aligned}
$$

and since $\cos\theta(t_h) = 0$, $\Gamma(t) = \cos\theta(t)$. We now find Δ in terms of Ω:

$$\Delta(t) = \frac{\Omega}{\tan\theta(t)} = \frac{\pm\Omega\Gamma}{\sqrt{1-\Gamma^2}} \qquad (4.12)$$

In addition to t_h, for which $\theta(t_h) = \pi/2$, we introduce t_v for which $\theta(t_v) = 0$, the latter corresponding to a field vector aligned with the vertical (eigenstate). With this definition, the integral in Equation 4.11 becomes

$$\int_{t_v}^{t_h} \epsilon(t)\Omega(t)\mathrm{d}t = 1 \qquad (4.13)$$

This integral is a general result for adiabatic following. It implies the constraint that, given $\Omega(t)$, a decrease in $\epsilon(t)$ (due, for example, to a change in some parameter affecting the detuning function $\Delta(t)$) necessitates an increase in $\epsilon(t)$ elsewhere. Examples of this behaviour occur later when we consider modifications to derived pulse forms, and their imperfect experimental implementations.

4.3.2 Cosine-squared envelope

As an example of smooth modulation, we consider an experimentally realizable, cosine-squared pulse envelope and constant adiabaticity:

$$\Omega(t) = \Omega_0 \cos^2\left(\frac{\pi t}{2\tau}\right); \quad -\tau \le t \le +\tau$$

$$\epsilon(t) = \epsilon_0$$

Additionally, if we choose the initial field vector to be vertical (that is, aligned with the eigenstate), $t_v = -\tau$ and the expression in Equation 4.13 becomes:

$$\frac{\epsilon_0 \Omega_0}{2}\left[t + \frac{\tau \sin\left(\frac{\pi t}{\tau}\right)}{\pi}\right]_{t_v = -\tau}^{t_h} = 1 \tag{4.14}$$

These general relationships can now be used to construct specific pulses, such as those necessary to invert a state or to create a superposition simply by choosing the time t_h at which resonance should occur.

4.3.2.1 Full Adiabatic Passage (Mirror)

For our 'full AP' pulse, we begin, in the Feynman representation, with the field vector aligned with one pole (eigenstate) and sweep it to the other. Resonance occurs mid-way, at $t = 0$, and hence the relevant condition is $t_h = 0$. With this condition, Equation 4.14 becomes $\epsilon_0 \Omega_0 \tau = 2$, and

$$\Gamma(t) = \frac{t}{\tau} + \frac{1}{\pi}\sin\left(\frac{\pi t}{\tau}\right) \tag{4.15}$$

from which $\Delta(t)$ follows using Equation 4.12, and is illustrated in figure 4.2. A sharp initial rise in detuning peaks at $t_{max} \simeq \mp 0.575\tau$ at a detuning of $\Delta(t_{max}) \simeq \pm 0.727\Omega_0$. The frequency modulation is symmetrical, and the maximum phase excursion – the maximum of the integral of detuning – is $\Phi_{max} = 1/\epsilon_0$; overall, the excursion is zero.

The shape of this detuning function is similar to that from a self-phase modulated pulse and the possibility of using an SPM pulse to perform the mirror operation is discussed further in Section 4.3.3.1.
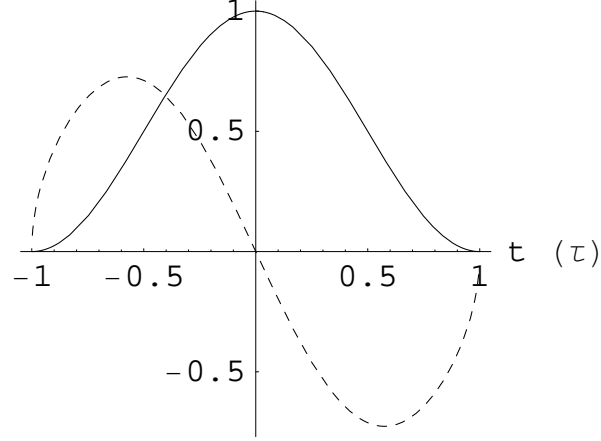
FIGURE 4.2: Coupling strength Ω (Ω_0) (solid) and detuning Δ (Ω_0) (dashed) for full AP 'mirror' pulse of duration 2τ. The frequency detuning is adapted throughout to maintain constant adiabaticity as the coupling strength follows an experimentally convenient $\cos^2$ profile.

4.3.2.2 Half Adiabatic Passage (Beam-splitters)

Using an identical procedure, we can derive a pulse which performs a 'half AP' operation. For this, the detuning must end on zero, and hence the condition is $t_h = \tau$. The relationship in Equation 4.14 is then simply $\epsilon_0 \Omega_0 \tau = 1$ and, using this,

$$\Gamma(t) = \frac{1}{2}\left[1 - \frac{t}{\tau} - \frac{1}{\pi}\sin\left(\frac{\pi t}{\tau}\right) \right] \tag{4.16}$$

The corresponding detuning function, again obtained from Equation 4.12, is illustrated in figure 4.3. Again, there is a sharp initial rise in detuning, which peaks at $t_{max} \simeq -0.506\tau$ at $\Delta(t_{max}) \simeq 1.09\Omega_0$. The total phase excursion is $\Phi_{max} \simeq 0.866/\epsilon_0$.

In contrast to the Rabi $\pi/2$ pulse, which we consider further in section 4.3.4, the beam-splitter and recombiner pulses based on AP are not identical. The recombiner can be regarded as the continuation of the beam-splitter that would complete the adiabatic inversion, or simply its time-reversal; in both cases, the recombiner pulse sweeps from on-resonance to far off-resonance.
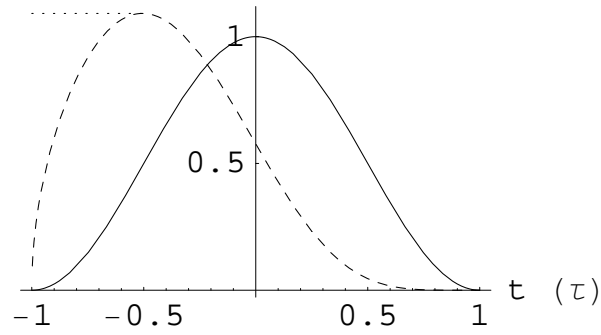
FIGURE 4.3: Coupling strength Ω (Ω_0) (solid) and detuning Δ (Ω_0) (dashed) for half AP 'beam-splitter' pulse of duration 2τ. As before, the frequency detuning is adapted throughout to maintain constant adiabaticity and the form of $\Delta(t)$ during the extinction ensures the superposition is approached adiabatically. The constant detuning (dotted) preceding the peak of the ideal function is a pragmatic modification, and is discussed in the text.

4.3.3 Comparison with traditional chirp schemes

Having determined the constant-adiabaticity frequency modulation for the mirror and beam-splitter operations, we find it instructive to compare them with established techniques.

4.3.3.1 Full Adiabatic Passage

The simplest method for performing adiabatic inversion might consist of a constant coupling amplitude $\Omega(t) = \Omega_0$ and a frequency sweep from well above to well below resonance. Unfortunately, for the initial and final field vectors to be aligned with the eigenstates, $\Delta(t)$ must sweep over a range much larger than the coupling strength: $\Delta_{max} \gg \Omega_0$. Inescapably, this approach requires an unreasonably large maximum detuning, even before any considerations of adiabaticity.

The situation is improved significantly by moving to a smooth modulation, such as the cosine-squared envelope previously described. For the simplest form of frequency chirp (a linear temporal chirp) the adiabatic parameter does vary during the pulse, but for the case of $\Delta_0 \simeq 2\Omega_0$ (which gives the same chirp rate at the middle of the pulse as our constant adiabaticity case), $\epsilon(t)$ does not significantly exceed that of the derived constant adiabaticity pulse. For small/large chirp rates, there are significant peaks in
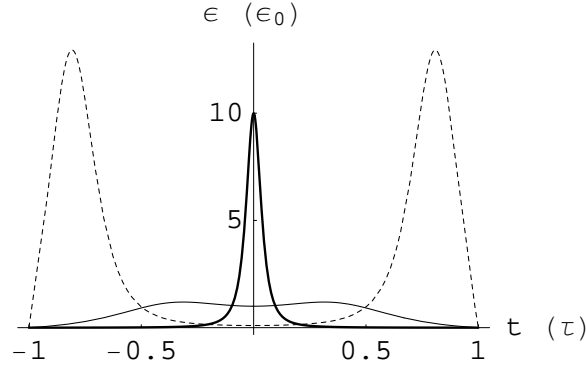
FIGURE 4.4: Variation of the adiabatic parameter $\epsilon(t)$, relative to the ideal $\epsilon_0 = 2(\Omega_0\tau)^{-1}$, during a linear temporally chirped pulse $\Delta(t) = \Delta_0 \times (t/\tau)$ with a smooth cosine-squared intensity modulation $\Omega(t) = \Omega_0 \cos^2\left[(\pi t)/(2\tau)\right]$. The three cases correspond to $\Delta_0 = 2\Omega_0$ (solid), $\Delta_0 = 20\Omega_0$ (thick), and $\Delta_0 = 0.2\Omega_0$ (dashed).

adiabaticity at the ends/middle of the pulse, as shown in figure 4.4. (Such increases in $\epsilon(t)$ are unavoidable, as embodied by equation 4.13.) For these cases, the state will not adiabatically follow the field and the pulse will not be experimentally useful. However, adiabatic inversion has been studied extensively, and there are many smoothly modulated pulse types which are robust [100].

We thus conclude that, while the precise pulse shape and chirp form are of limited consequence, a smoothly modulated pulse is important for full AP operations and makes far more efficient use of the available experimental resources. The form of figure 4.2 suggests that appropriate pulses might be provided by, for example, self-phase modulation of an ultrashort laser pulse; several such implementations have been considered in detail by Goswami and Warren [101].

4.3.3.2 Half Adiabatic Passage

The remarkable robustness of adiabatic inversion is not, unfortunately, shared by half adiabatic passage, for which the form of approach to resonance, and simultaneous extinction, is crucial. The simple truncation of a pulse designed for full AP will be far from adiabatic and sensitive to frequency errors which result in a non-zero final detuning. Conversely, pulses with no distinct end, such as a Gaussian intensity envelope, may approach resonance adiabatically but remain there long enough for small perturbations

to accrue. The optimum pulse must have a definite end, but should approach resonance in a controlled manner.

4.3.4 Practical Beam-Splitter pulses

We now focus on the practicalities of the half AP pulse, derived in section 4.3.2.2 and shown in figure 4.3. No real pulse will match the required pulse form; the dominant imperfections are likely to be variations in the Rabi frequency, due to intensity variations across a laser beam profile, and frequency shifts due, for example, to the various Doppler shifts for atoms within a thermal cloud. We consider these effects by introducing a coupling strength scaling factor α, and a frequency offset f, in the natural units of the problem, Ω_0. For our practical pulse, we therefore model the coupling strength and detuning functions by:

$$\Omega_e(t) = \alpha\Omega(t) \tag{4.17}$$

$$\Delta_e(t) = \Delta(t) + f\Omega_0 \tag{4.18}$$

The ideal pulse is given by $\alpha = 1$ and $f = 0$. A non-ideal pulse will result in a non-equal superposition and/or an error in the relative phase of the two components. We shall measure the fidelity of the superposition and phase errors using the definitions

$$q \equiv 1 - 2\left|\frac{1}{2} - |\langle 0|\psi\rangle|^2\right| \tag{4.19}$$

$$\sin(\chi) \equiv \text{Im}\frac{\langle 1|\psi\rangle}{\langle 0|\psi\rangle} \times \left|\frac{\langle 0|\psi\rangle}{\langle 1|\psi\rangle}\right| \tag{4.20}$$

where $|\psi\rangle$ is the state of the system after the pulse. Before proceeding, we make one pragmatic modification. Because the derived pulse of figure 4.3 begins with zero detuning, the initial direction of the field vector is extremely sensitive to errors in Δ and Ω. As the condition for field vector alignment with the eigenstate is $\Delta(-\tau) \gg \Omega(-\tau)$, we amend the form of $\Delta_e(t)$ so that the detuning begins at, and initially maintains, its peak value:

$$\Delta_m(t) = \begin{cases} \Delta_e(t_{max}) = 1.092\Omega_0 + f\Omega_0 & t < t_{max} \\ \Delta_e(t) & t \geq t_{max} \end{cases} \tag{4.21}$$

This modification guarantees the initial field vector alignment (provided f is greater than $\simeq -1$) at the expense of a small ($< 10\%$) variation in the adiabatic parameter $\epsilon(t)$, which is unlikely to be of experimental significance. More important is the behaviour as Δ and Ω approach zero at the end of the pulse, which is unaffected by this modification. Other modifications are possible, and this example is intended to illustrate the ability to meet both the constraints on ϵ and practical experimental requirements.

4.3.4.1 Realistic Rabi $\pi/2$ pulses

As a benchmark, we first consider the effect of experimental variations on the effect of a controlled area Rabi pulse characterised by a constant, resonant frequency and steady intensity which is described by

$$
\begin{aligned}
\Omega_e(t) &= \alpha\Omega_0 \\
\Delta_e(t) &= f\Omega_0 \\
\Omega_0\tau &= \pi/2
\end{aligned}
$$

The error in superposition composition is calculated analytically, and is shown, as a function of the coupling and detuning variations α and f, in figure 4.5(a). The plot shows numerous regions where an equal superposition is created (white). Along the axis $f = 0$ we see Rabi oscillations, with an equal superposition every odd integer multiple of a $\pi/2$ pulse. The behaviour away from the $f = 0$ axis may be understood by noting that, in the Feynman representation, the state vector is rotated about a vector that is inclined to the equator. The phase error of the superposition, shown for convenience by its sine in figure 4.5(b), manifests an essentially similar structure. In both cases, the key characteristic is relative tolerance of variations in frequency, but sensitivity to variations in coupling strength.
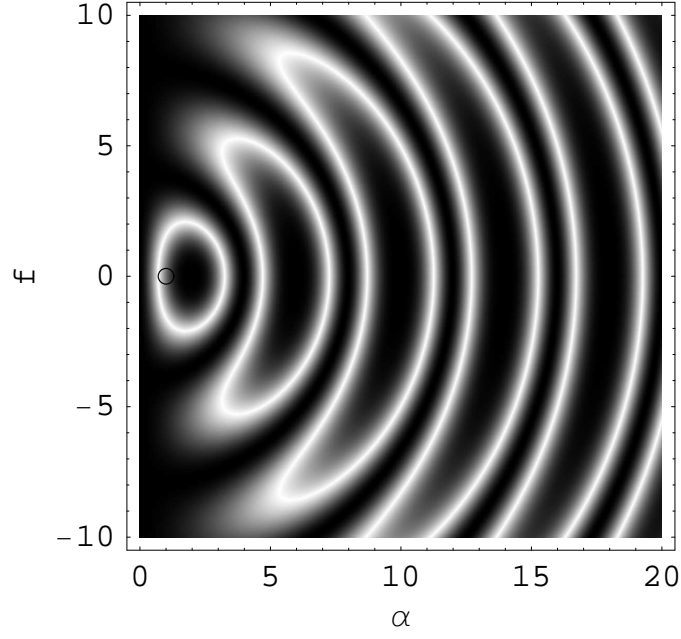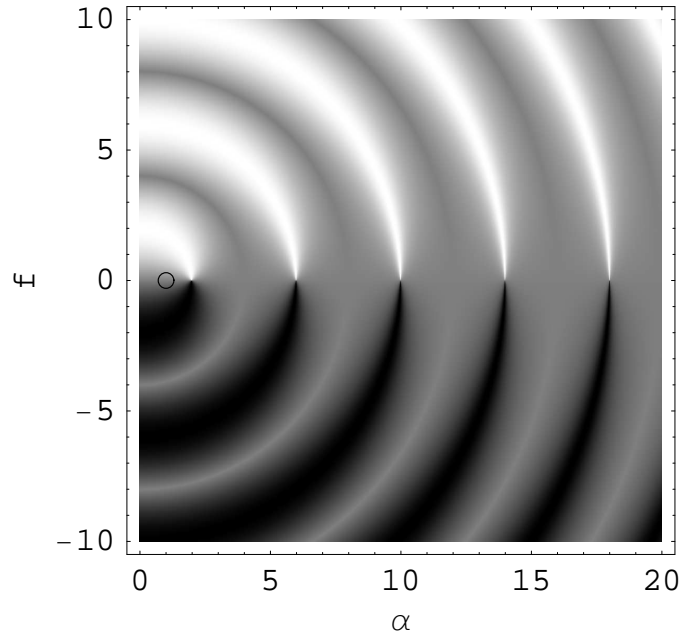
(a) Fidelity q

(b) Relative phase $\sin \chi$

FIGURE 4.5: (a) Fidelity q and (b) phase relative to the ideal $\sin(\chi)$ of the imperfect Rabi $\pi/2$ pulse, calculated analytically, with imperfections parametrised by a scaling of the coupling strength α and an offset of the resonance frequency $f\Omega_0$. Around the ideal pulse, circled at $\alpha = 1$ and $f = 0$, the fidelity follows the familiar Rabi oscillations for increasing coupling strength, and is relatively insensitive to frequency offsets; phase shows a similar structure. White is high and black is low.
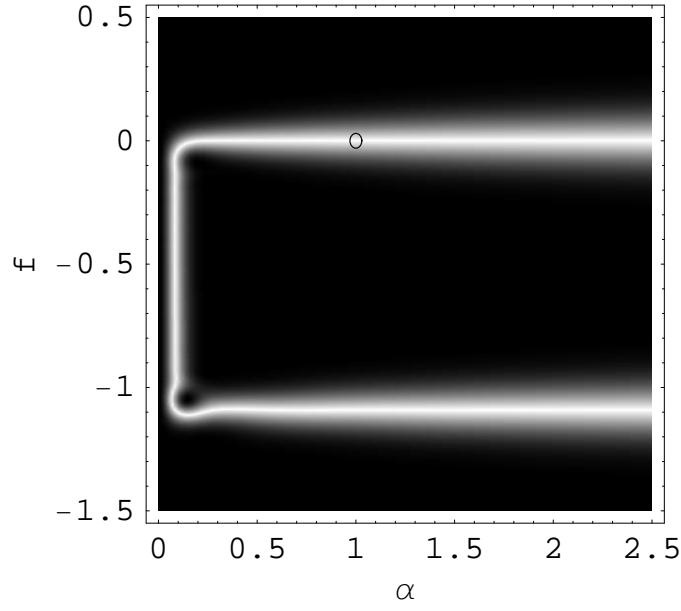
(a) Fidelity q
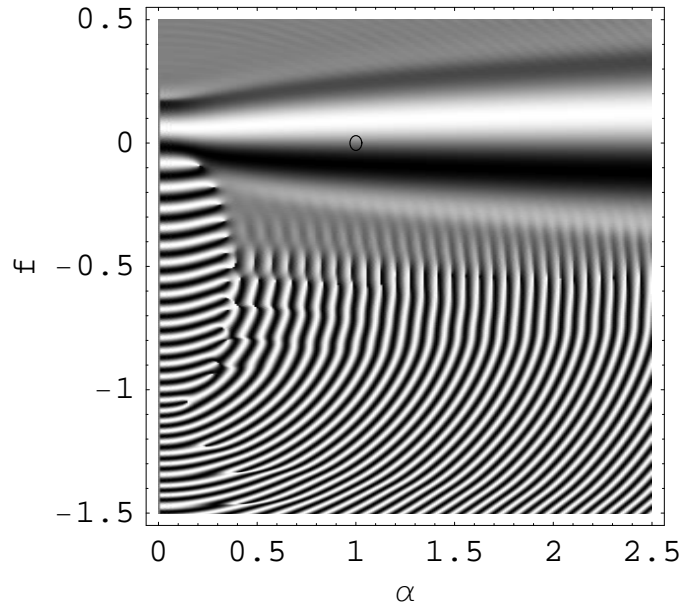
(b) Relative phase $\sin \chi$

FIGURE 4.6: The equivalent plot for an imperfect 'half AP' pulse, based on the constant adiabaticity pulse, with a modification to avoid initial sensitivity to an offset in detuning, calculated by numerical integration with $\epsilon_0 \sim 10^{-3}$. In contrast to the Rabi controlled area pulse, a large fraction of the parameter space for this pulse leaves the state unaffected; only the smaller, interesting region is shown. (a) Fidelity q shows as, expected, a narrow branch around $f = 0$ for which, above a minimum α, an equal superposition is created with (b) relative phase $\sin \chi$ unaffected by changing α but strongly dependent on f. The unexpected lower branch, centred on $f = -1.09$, for which fidelity is high and phase is rapidly varying, is discussed in the text. White is high and black is low.

4.3.4.2   Realistic half AP pulse

We now consider the effect of experimental imperfections on our modified constant-adiabaticity pulses, whose ideal form is illustrated in figure 4.3, modified according to Equation 4.21, and has imperfections described by equations 4.17 and 4.18. Our results, shown in figure 4.6, are determined by numerical integration of the Schrödinger equation (specifically, the dressed states evolution), using the second order trapezoidal method [102], for a range of α and f. A value of $\epsilon_0 \sim 10^{-3}$ was chosen, but most of the features are not critically dependent on this parameter.

For imperfect pulses, $(\alpha, f) \neq (1, 0)$, we cannot assume the adiabatic condition holds throughout; the adiabatic parameter will not remain constant and, for $f \neq 0$, the pulse will be strongly non-adiabatic towards the end where $\Omega_e \rightarrow 0$ and the offset in detuning becomes significant. The field vector in this case adiabatically approaches the equator while the rate of precession reduces, but the vector then swings violently (non-adiabatically) back to point towards the pole before the pulse is fully extinguished. We shall see later, in section 4.3.4.4, that this is similar to the second crossing in SCRAP. Although this increase in ϵ will in many cases be rapid and too short-lived for the atom to respond significantly, it is clear that the practical optimum will combine the best adiabaticity during the pulse with the greatest tolerance to experimental uncertainties at the beginning and end.

The error in superposition for the half AP pulse is shown in figure 4.6(a). As expected, there is a narrow region, close to the $f = 0$ axis, around which an equal superposition is created. Above a minimum coupling strength, the evolution is adiabatic and a further increase in coupling strength merely increases this adiabaticity.

If we assume that most imperfections accrue during the final extinction of the pulse we are able to offer an approximate description for this branch. We take the perturbation to begin at a time t_f when the detuning $\Delta_m(t_f) \sim \Omega_e(t_f)$ – that is, the field vector lies at $45°$ to its ideal direction – and the detuning function is dominated by the frequency error, $\Delta_m(t_f) \simeq f\Omega_0 \sim \alpha\Omega_0 \cos^2\left[(\pi t_f)/(2\tau)\right]$. We estimate the subsequent precession of the state vector by approximating the field vector to a constant for the remainder of the

pulse, and, by imposing that this precession ϕ must be small, deduce the following:

$$
\begin{aligned}
\phi &= \int_{t_f}^{\tau} \tilde{\Omega}(t)\mathrm{d}t \\
&\sim \Omega(t_f) \times (\tau - t_f) \\
&\sim f\Omega_0\tau\frac{2}{\pi}\left(\frac{\pi}{2} - \mathrm{ArcCos}\sqrt{\frac{f}{\alpha}}\right) \\
&\sim \Omega_0\tau\sqrt{\frac{f^3}{\alpha}} \ll 1
\end{aligned}
$$

The region of high fidelity in figure 4.6(a) is hence described by

$$
\sqrt{\frac{f^3}{\alpha}} \ll \epsilon_0 \tag{4.22}
$$

There is a further region, also achieving an equal superposition, centred around a negative detuning offset of $f = -1.09$, corresponding to the case in which the initial detuning $\Delta_m(-\tau)$ is zero. For this pulse, the field vector begins on the equator, at right angles to the state vector, before moving adiabatically towards the pole. The angle between state and field vectors remains constant during this evolution, but the azimuthal angle of the state vector relative to the field vector precesses rapidly. When the field vector comes to rest, the state is in an equal superposition, but the phase is undetermined. Superficially, this appears similar to the 'recombiner' pulse, but the detuning in this branch is not such that constant adiabaticity is maintained, and no consideration has been given to the phase of the initial (azimuthal) angle of the field vector along the equator. Figure 4.6(a) also shows a vertical (constant α) band joining the upper and lower branches: this is dependent on the adiabatic parameter ϵ_0, and is not experimentally useful.

The phase error for the half AP pulse is shown in figure 4.6(b). This shows quite different behaviour for the on- and off-resonance branches, with a region of slowly varying phase surrounding the on-resonance branch while, in stark contrast, the off-resonance branch falls in a region of rapidly varying phase. The upper branch is therefore useful; the lower branch is not.

4.3.4.3 Comparison of sensitivity to imperfections

The Rabi half AP pulses differ markedly in their tolerance to experimental variations, with the Rabi pulse being relatively tolerant of variations in frequency but sensitive to coupling strength, while the AP pulse offers tolerance to coupling strength at the expense of an increased frequency dependence. Specifically, while the Rabi pulse dependence on intensity variations about the ideal shows the familiar oscillation, the AP pulse is essentially unaffected by changes in coupling strength once the adiabatic regime is reached (for the example illustrated in figure 4.6, this is when $\alpha > 0.2$). This insensitivity to coupling strength is the familiar advantage of the adiabatic approach. The change in phase with coupling strength, provided there is no frequency error, is in both cases zero. The dependence of superposition phase on the frequency error, however, is (locally) zero for the Rabi pulse but strong for the AP pulse.

To illustrate the relative sensitivities of the two approaches to experimental variations, we consider a typical example of coherent atomic manipulation. Extending the example in section 4.3 of an atomic dipole, we consider the transition $5S_{1/2} \rightarrow 5P_{3/2}$ in a laser-cooled sample of ^{85}Rb, at a temperature of (typically) $< 100\mu$K [103], for which the Doppler width is approximately 150kHz. If the coupling is provided by a 100mW laser, focused to a spot $w_0 = 1$mm, the (maximum) Rabi frequency is $\Omega_1 \simeq 2\pi \times 350$MHz. The two hyperfine states of the $5S_{1/2}$ state can be coupled using two such beams (slightly detuned from the single photon resonance, and with a frequency difference equal to the hyperfine splitting) to drive a resonantly enhanced Raman transition [81, 104]. The coupling strength for this two-photon process is [67]

$$\Omega_2 = \frac{\Omega_1 \times \Omega_1}{2\delta} \simeq 2\pi \times 60\text{MHz}$$

for $\delta \gg \Omega_1$, where we have taken a typical single photon detuning of $\delta = 2\pi \times 1$GHz and the subscript refers to the number of photons involved in the process.

For this example, the insensitivity to coupling strength vastly increases the usable area of the beam. If we can accept an error of 1% in the superposition $q > 0.99$, the Rabi approach works over only a small region near the centre of the Gaussian beam profile,

extending to approximately 10% of the beam waist, and encompassing 2% of the power. This also assumes there is no jitter in the pulse duration. In contrast, the adiabatic approach, for which we can allow α to be as low as 0.2 (for the typical parameter $\epsilon \sim 10^{-3}$ used in these calculations, with a smaller adiabatic parameter permitting a further excursion towards $\alpha = 0$), extends the useful region to nearly 90% of the beam waist, encompassing 80% of the power.

The dependence of superposition phase on the frequency offset f, for $\alpha = 1$, is small for Rabi pulses, $\partial\chi/\partial f = 1$, while the AP pulse rests on a steep slope with $\partial\chi/\partial f \simeq 18$. In our example, the range of f will be determined by Doppler shifts for the thermal sample, with a typical value of

$$f \simeq \pm\frac{2\pi \times 300\text{kHz}}{\Omega_2} \simeq \pm 5 \times 10^{-3} \tag{4.23}$$

(The worst case of counter-propagating Raman beams is chosen, for which the total Doppler shift is the sum of that for each beam. The co-propagating arrangement is, to first order, Doppler insensitive.) Hence, at $\alpha = 1$, the phase uncertainty is $\delta\chi \simeq 5°$, compared with the Rabi case of $\delta\chi \simeq 0.3°$. While the cumulative error arising from this uncertainty may be unacceptable for complex interferometers, the increased frequency sensitivity will, for many applications, be less significant than the benefits accompanying the insensitivity to coupling strength.

Although the frequency and intensity errors will often be systematic, we have also examined the case of small random variations and their effect on a simple two-pulse interferometer: the dependence on α may generally be neglected, but variations in frequency lead to a reduction in the fringe visibility. If f is normally distributed about zero with a width σ_f, the visibility shows an approximately Gaussian dependence on σ_f, falling to $\frac{1}{2}$ at $\sigma_f \sim 0.01$, for the typical adiabatic parameter $\epsilon_0 \sim 10^{-3}$.

The analysis above shows the Rabi approach to be favourable if detuning is uncertain, and the AP approach to be best if coupling strength is uncertain. The key observation is that in many experimental situations (such as cold atom and ion experiments), the frequency can be controlled far better than the coupling strength.

4.3.4.4   Relation to other adiabatic techniques

The description adopted here, in terms of the parameters of detuning and coupling strength, can be applied to many implementations of adiabatic passage which involve time-delayed pulse pairs to achieve similar effects. A particularly close analogue is found in Stark-Chirped Rapid Adiabatic Passage (SCRAP) [105, 106], where two pulses impinge on a two-level system such that one (on-resonance) couples the states while the other (far off-resonance) provides a time-dependent detuning via the AC Stark shift. For full population transfer, the Stark (or detuning) pulse is delayed slightly relative to the coupling pulse, resulting in adiabatic following from the initial to the final eigenstate, and then a non-adiabatic return of the field to the starting position, *without being followed* by the state. SCRAP has also been adapted to create superpositions [107].

Additionally, there appears to be an interesting similarity between the pulse derived here and recent work by Vitanov *et al*[108], where Stimulated Rapid Adiabatic Passage (STIRAP) is recast in terms of a two-level system; see also [109]. The authors show a detuning and a coupling pulse displaced in time (Figure 2 in [108]) which, while an accurate model of their system, shows a strongly non-uniform adiabatic parameter at early and late times. This recasting of STIRAP also suggests that it may be possible to describe fractional-STIRAP [110, 111] in similar terms to the example pulse in this chapter.

Finally, we observe that simply constraining the adiabatic parameter is perhaps naïve. The adiabatic parameter varies significantly in the time displaced Gaussians of Vitanov's work, but it does so only when the coupling strength is weak, and hence does not adversely affect the operation. A more complete approach to finding an optimum adiabatic pulse might proceed by establishing some measure of deviation from ideal adiabatic behaviour, and then finding the function which minimises it, while also ensuring tolerance of experimental uncertainties, and accounting for experimentally limited parameters such as the maximum practical detuning and the finite time allowed for the pulse.

4.4   Conclusions

We have addressed the use of adiabatic passage for mirrors and beam-splitters in atom interferometry and the possibility of tailoring the combination of frequency and amplitude modulation to control the adiabaticity during the pulse. We have presented a straightforward analytical approach which leads to expressions of direct use for the frequency modulation of an experimentally realizable pulse and which can readily be applied to experimental implementations. In comparison with the linear temporal chirp and inverse tangent sweep, our smooth modulation and tailored chirp achieves better fidelity within experimental constraints. By observing the adiabaticity throughout the pulse, our approach also offers insight into related processes such as SCRAP and STI-RAP.

We have examined the sensitivity of these pulses to experimental variations in frequency and intensity of the excitation. In particular, for the experimentally challenging case of the beam-splitter, we have compared our tailored approach with the simple half Rabi cycle. Although the half Rabi pulse has a greater tolerance of variations in frequency, the chirp proves far more robust against variations in intensity, which are likely to be more problematic experimentally and result from variation across a beam profile as well laser intensity fluctuations. The effect of an imperfect adiabatic beam-splitter is understood simply using the Feynman representation and considering the deviation of the field vector during the final extinction of the pulse.

The inclusion of the extinction in the design procedure proves crucial: to our knowledge, this is the first application of controlled adiabaticity to the design of a half-adiabatic-passage pulse, and the first example of a pulse which can robustly create superpositions in a two-level system with a single exciting field.

Chapter 5

Magneto Optical Trap Apparatus

The Magneto-Optical Trap has become the staple of cold atomic physics, and while implementations are common-place, they are far from trivial. Our approach is similar to most but differs in some details, such as our method of applying the necessary well-controlled red detuning.

An overview of the lasers for the MOT is shown in Figure 5.1. The principle components of any MOT setup are a narrow linewidth laser frequency locked several linewidths to the red of a closed cooling transition and an anti-Helmholtz magnetic field; this is described in Section 2.4.3. In addition, as with most MOTs, we require a repump laser, as the cooling transition is only approximately closed. Our experiment requires that we be able to independently and rapidly control the amplitude of these two lasers; for this we use acousto-optical modulators (AOMs), and the frequency shift these introduce has implications for the laser locking. These two amplitude modulated beams must be sent, with the correct polarisation, to a vacuum chamber where they converge on the zero of the anti-Helmholtz magnetic field. Imbalances in intensity can perturb the trapping potential, disrupting the shape and velocity distribution of the cloud as described in Section 3.6.2. We use a short stretch of single mode optical fibre to ensure the beams sent to the MOT are clean, and that the repump and trapping lasers overlap. The density we require can be achieved by loading from background gas rather than actively loading the trapping region, and we use a rubidium getter, with electrical-feedthroughs to the vacuum chamber, to provide a controllable background rubidium pressure.

110MHz AOM

Repump laser

110MHz AOM

Single mode
optical fibre

to vacuum chamber

Molasses laser
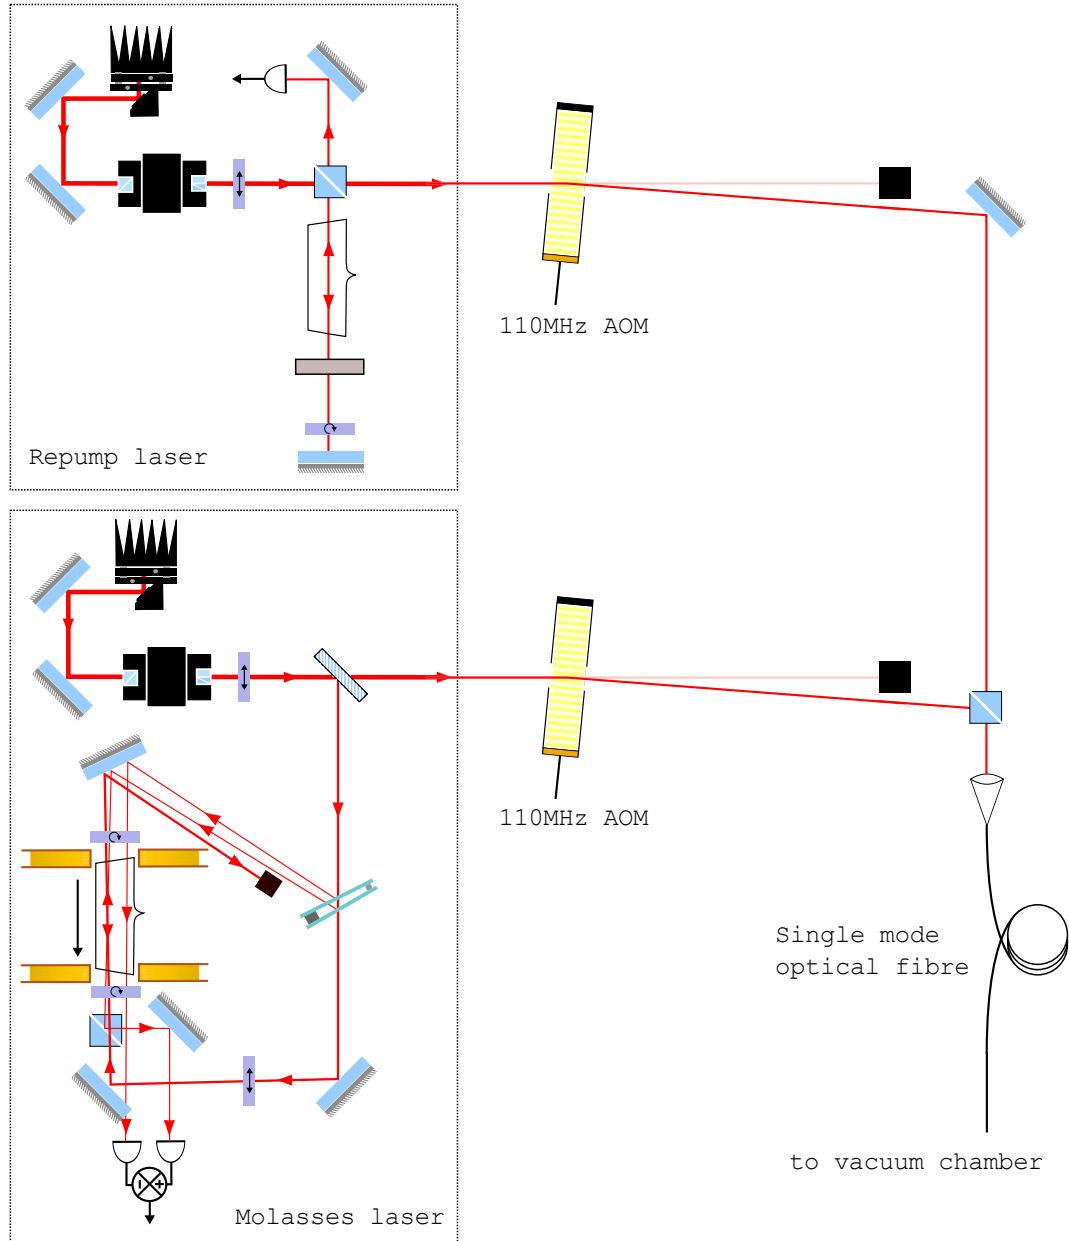
FIGURE 5.1: An overview of the MOT apparatus. Two locked lasers are independently amplitude modulated using AOMs, are combined using a polarising beam-splitter cube, and are sent to the vacuum chamber using a single-mode optical fibre. The repump and molasses laser are described in Figures 5.2 and 5.3.
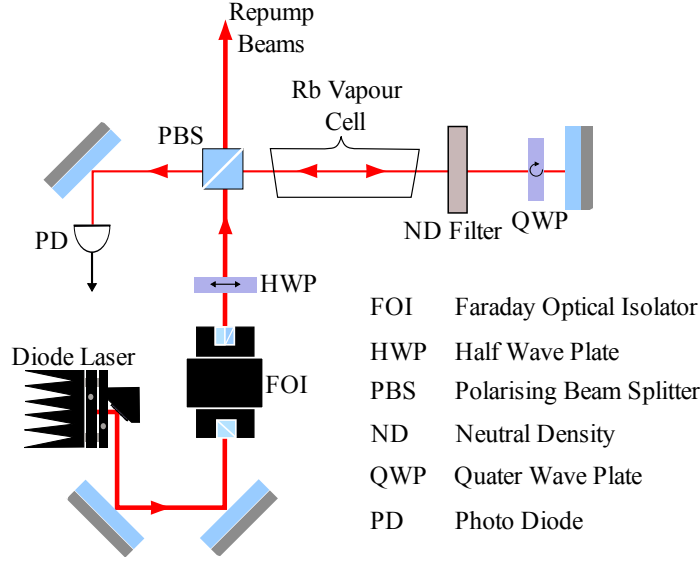
FIGURE 5.2: The simple SAS setup used to obtain a signal to which the repump laser is locked. Using the half-wave plate we send the majority of light through the polarising beam-splitter cube, with a small fraction, sufficient to perturb the population as described in Section 5.1.2, making one pass through the vapour cell, before being attenuated by the filter, and then retro-reflected. After two passes through the quarter-wave plate, the light has the correct polarisation to pass through the beam-splitter cube, and falls on the photodiode.

5.1 Lasers and Spectroscopy

Of the trappable atomic species, rubidium was chosen primarily because semiconductor diode lasers are available at a wavelength where it has a strong dipole transition; it is a fortuitous accident that the wavelength chosen as industry standard for Compact-Disc players so closely matches a transition in rubidium. We use Sharp GH0781JA2C laser diodes, which typically emit light with a wavelength of 784nm. This varies with temperature, and 780nm, the D2 line in rubidium, can be reached by cooling the device by a few degrees, to typically around $17°C$. The wavelength must, however, be controlled to much greater accuracy than the nearest nanometer; the typical frequency-scale relative to which the laser must be stable is the natural linewidth of the transition, as described in Section 2.2.1, and is $\sim 6MHz$. For comparison, a change of 1nm around 780nm is, in frequency, $\sim 493GHz$.

Temperature stabilisation cannot provide sufficient stabilisation to create a MOT. The high refractive index ($\simeq 3.3$) of the semiconductor gain medium causes large Fresnel reflection ($\simeq 29\%$) at the facets, creating a strong mode structure within the chip.
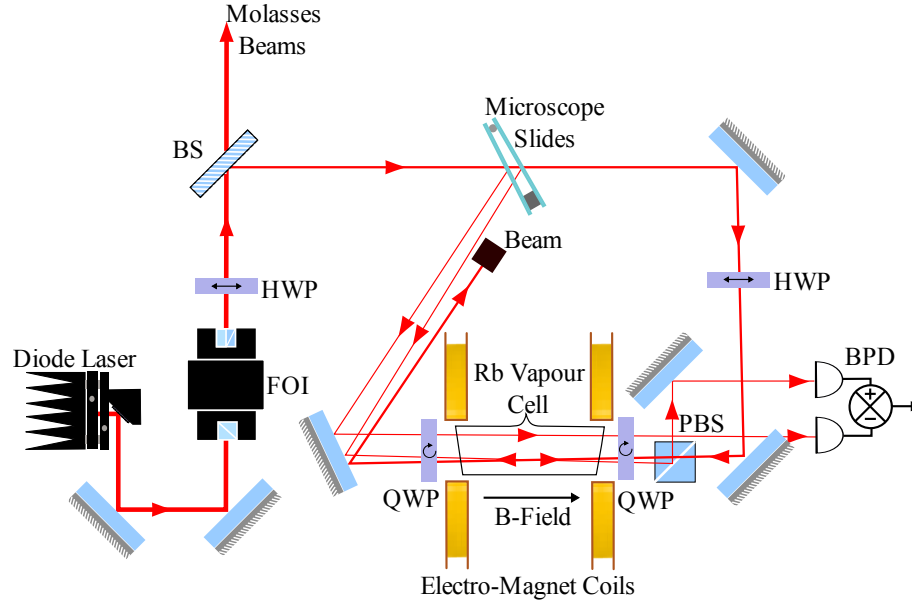
FIGURE 5.3: The modified SAS setup incorporating quarter-waveplates to optically pump to an extreme m_F state, and magnetic field-coils to Zeeman shift this level. In addition to the repump optics, we have a non-polarising beam-splitter (BS) and a Balanced photodiode (BPD). Despite superficial similarity, this is not a DAVL lock, which is described in Section 5.1.3. This, and Figure 5.2, are used with permission from Matthew Himsworth.

Typically light from only one of these facets emerges from the device, and the frequency stability can be vastly improved by adding a frequency selective element, such as a diffraction grating, after this output facet. The angle of light reflected from a diffraction grating depends on its wavelength, and so by placing the grating close to the output facet, we selectively feed-back only a narrow frequency range, which we can control via mechanical adjustment using a piezoelectric transducer. (Reflections from the grating create additional mode structure, as do, to a much lesser extent, reflections from the collimating lens which is between output facet and grating.)

A significant problem with this design is that changes to the grating angle (which are intended to affect the frequency) cannot avoid affecting the direction of the output beam. Significant adjustments can necessitate tedious re-alignment of the 6 cooling beams if the paths to the MOT are long. We found this to be a problem with free-space optics to the MOT, but by fibre-coupling we have reduced the free-space path length and have decoupled the alignment into the MOT (from a fibre-output coupler) from alignment into the fibre, the latter of which can be optimised simply using a power-meter. The next version of these home-built lasers will include a mirror which moves with the grating and

hence changes the small angular deviation to a small and hopefully insignificant beam displacement.

5.1.1   Stabilisation

A sufficiently narrow linewidth can be achieved using such an external-cavity diode-laser, but we additionally require that the laser frequency be held a specified red-detuning from the atomic transition. The 'optimum' detuning depends on many factors[1], but is expected to be of order a few linewidths; this is discussed in Section 2.4. In Rubidium, this hence requires a long-term stability much better than 6MHz on a frequency $\sim 3.8 \times 10^8$MHz. This is possible only if the frequency is monitored relative to an absolute reference, and clearly the ideal reference in this case is a signal derived from the transition itself.

One such method of deriving a signal is Doppler-free spectroscopy of a dilute atomic Rubidium vapour. The FWHM Doppler broadening of the D2 transition in a room temperature sample of rubidium is ~ 500MHz (see Section 2.2.2); a reference derived from the Doppler broadened spectrum is long-term stable, but is unable to provide the necessary resolution.[2] Using 'Doppler-free' spectroscopy, we can resolve resolve the finer details of the spectrum.

5.1.2   Saturated Absorption Spectroscopy

Our chosen method of Doppler-free spectroscopy is Saturated Absorption Spectroscopy, described in Foot [67]. For this, one uses two beams from the same laser, the frequency of which can be scanned over the atomic resonance, sent counter-propagating through a dilute atomic vapour. We consider a two-level atom, and one of these beams, the 'pump', should be sufficiently intense to significantly perturb the atomic populations: $I_{\text{pump}} \gtrsim I_{\text{saturation}}$. The other, the 'probe', should be weak such that it does not further perturb the change in populations caused by the pump: $I_{\text{probe}} \ll I_{\text{saturation}}$.

---

[1]These factors including what is meant by 'optimal'. The detuning which maximised the number of atoms is not necessarily the same as that which minimises the temperature.

[2]Locks based on the Doppler broadened spectrum are possible if extreme care is taken to account for experimental offsets; this is discussed in Section 5.1.3.

For a given laser detuning, the pump beam will perturb the populations of the velocity class of atoms whose Doppler shift is such that it brings the transition to resonance; this depletion of the ground state population is known as spectral hole burning, and an example of this hole for a laser detuning equal to half the Doppler width is shown in Figure 5.4. The finite width of this hole is due to the radiative (and other homogeneous) broadening of the transition.

The trick which allows us to perform Doppler-free spectroscopy is that, by sending the probe beam counter propagating, it sees the opposite Doppler shift and so interacts with a different subset of the atoms. For the majority of the frequency scan, the probe beam interrogates an unperturbed atomic distribution. It is only when the laser is on-resonance, and the Doppler shift of the atoms with which the beams interact is zero, that both beams interact with the same class and the probe beam sees the perturbed population distribution [67, §8].

The absorption of a laser by a two-level atom is the overlap integral between the laser linewidth and the transition profile. If we choose for simplicity a laser linewidth far narrower than the transition, we can approximate the laser linewidth to a delta function and the absorption is the Lorentzian $\sigma(\delta) = \frac{1}{\pi}\frac{\Gamma/2}{\delta^2 + (\Gamma/2)^2}$.

Extending the results of Section 2.3 to a Doppler broadened distribution, the population inversion for a given laser detuning δ_L scaled by the fraction of population at a given Doppler detuning δ_D is: $w(\delta_L, \delta_D) = D(\delta_D)w(\delta_L - \delta_D)$ where $D(\delta) := \frac{1}{\Delta_D\sqrt{\pi}}e^{-(\delta/\Delta_D)^2}$. This is illustrated for a fixed laser detuning $\delta_L = \Delta_D/2$ and $s = s_0$ in Figure 5.4.

To find the transmission of a probe beam through this perturbed atomic vapour, we note that the probe beam sees the opposite Doppler shift, and we integrate over the velocity distribution:

$$T_{\text{probe}}(\delta_L) = \int_{-\infty}^{+\infty} D(\delta_D)w(\delta_L - \delta_D)\sigma(\delta_L + \delta_D)\mathrm{d}\delta_D \ . \tag{5.1}$$

This integral is illustrated in Figure 5.5.

The Doppler-free signal of Saturated Absorption Spectroscopy occurs when the probe beam interacts with a velocity class of atoms whose population has been significantly
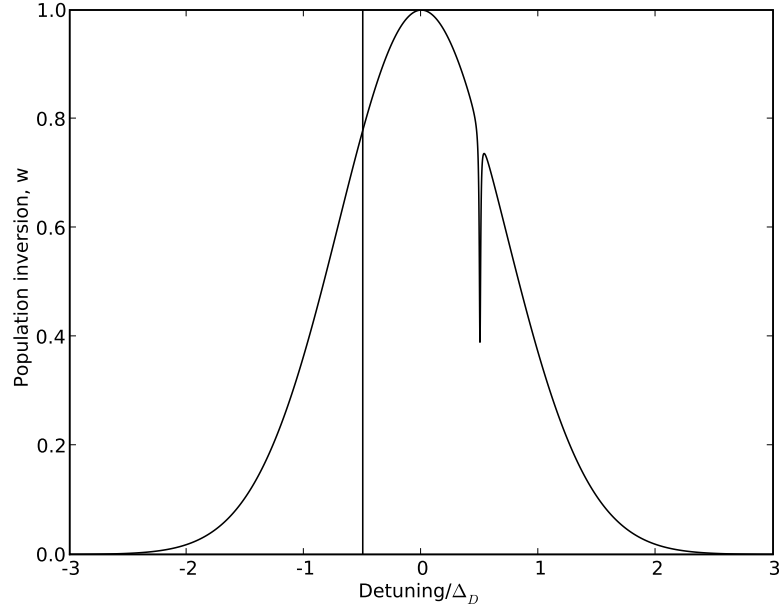
FIGURE 5.4: Illustration of spectral hole burning. An intense probe beam perturbs the population distribution of the velocity class whose Doppler shift brings the atoms to resonance. Here, the laser is blue-detuned by half the Doppler width and the natural linewidth of the atomic transition is 1% of the Doppler linewidth. The frequency with which the weak counter-propagating probe beam interacts is marked by the vertical line, and it is clear that, except within the natural-linewidth of the resonance, the probe beam will see an unperturbed population distribution.
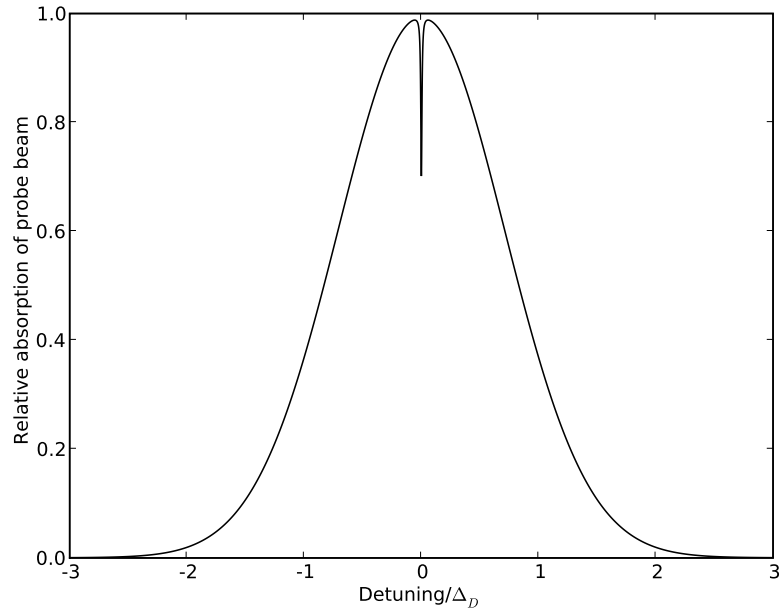
FIGURE 5.5: A typical Saturated Absorption Spectroscopy signal with the same conditions as Figure 5.4.

perturbed by the pump beam. If we now introduce a third level to our model of the atom, we can expect to see two independent Doppler-free resonances, but additionally, there will be a resonance for the non-stationary atoms whose Doppler shift brings one transition into resonance with the pump beam, and the other into resonance with the probe. These so-called 'cross-over' resonances are an artefact of the SAS technique; their position is mid-way between the atomic transitions and their width is the quadrature sum of the linewidths of the parent transitions.

The problem of multi-level atoms, and specifically of rubidium, was discussed by Maguire *et al* [112]; the effect of optical pumping, which we exploit in as discussed in Section 5.2 and which affects the amplitude of the cross-over peaks, was discussed by Smith & Hughes [113].

The spectra obtained for the upper and lower ground hyperfine levels of the two rubidium isotopes are shown in Figures 5.6, 5.7, 5.8 and 5.9. In each case, the real transitions are labelled by the quantum number of the upper state F', and the cross-overs by the F' numbers of the two upper states involved. These spectra were the subject of a detailed study by Siddons *et al* [75].
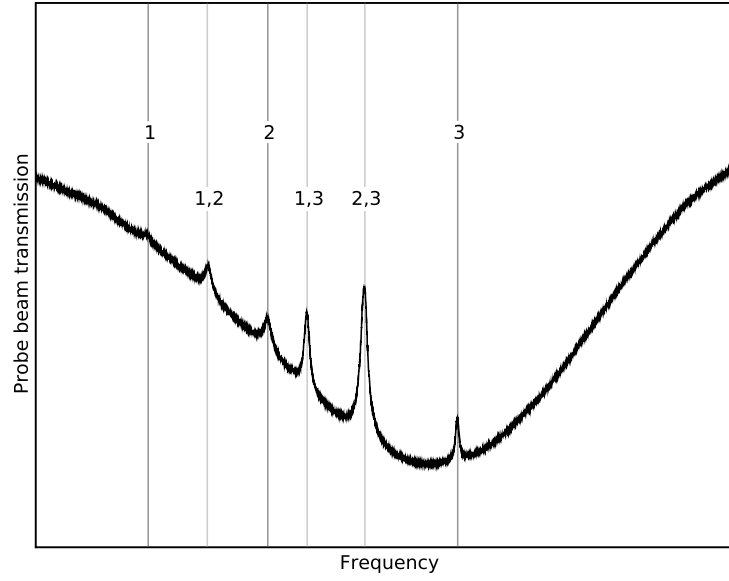
FIGURE 5.6: SAS spectrum for transitions from the upper hyperfine ground states of rubidium 87. The dips in absorption are shown against a large Doppler background. The difference between $F' = 3$ and $F' = 2$ is useful for calibration.
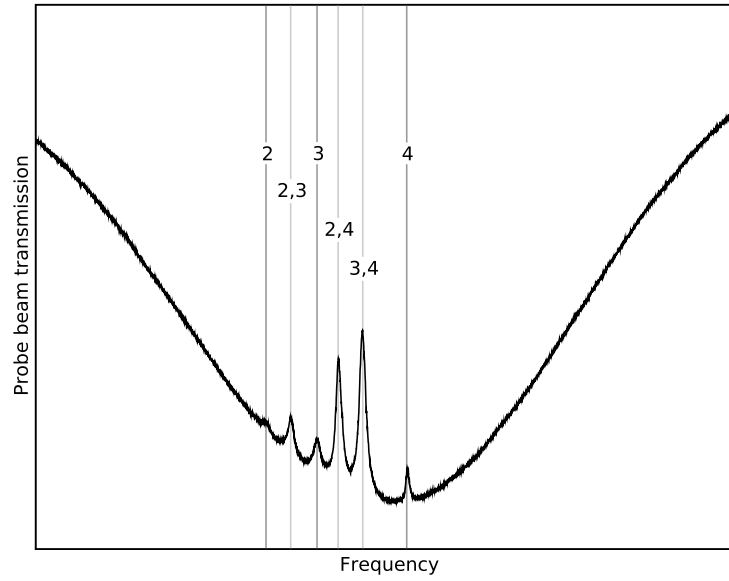
FIGURE 5.7: SAS spectrum for transitions from the upper hyperfine ground states of rubidium 85. The molasses laser is locked to the $F' = 4$ peak, which can be Zeeman shifted. (See Section 5.2.)
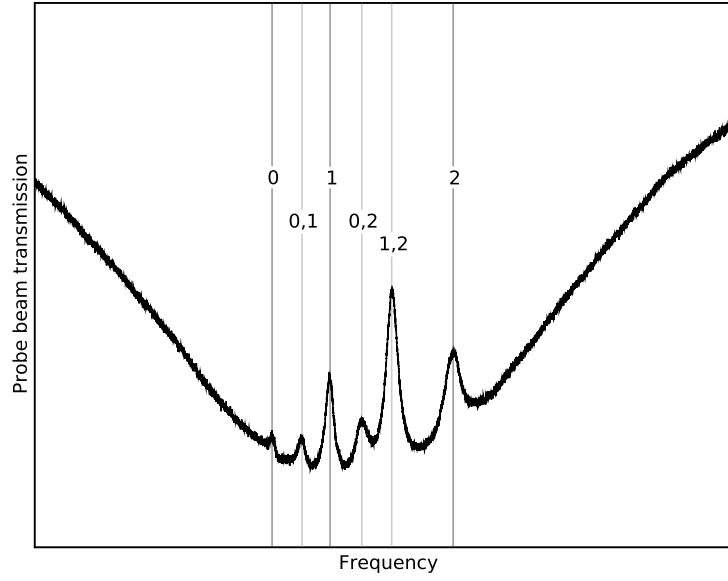
FIGURE 5.8: SAS spectrum for transitions from the lower hyperfine ground state of rubidium 87. This spectrum is included for completeness, but is generally not used in the experiment.
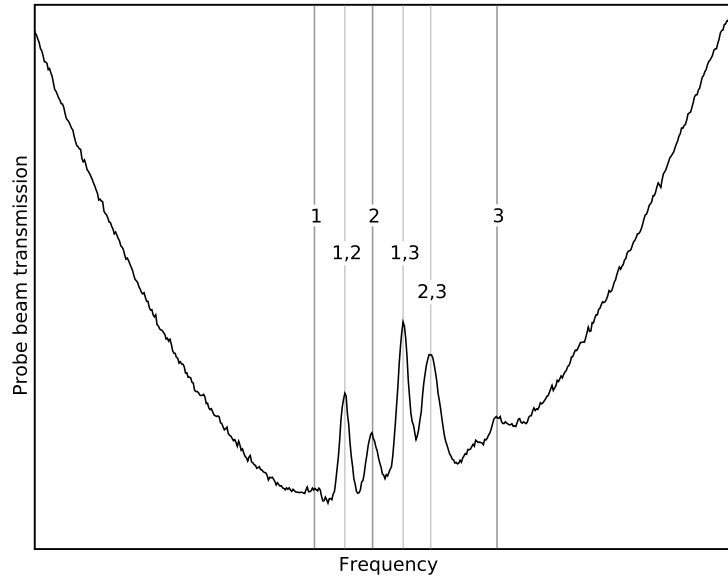
FIGURE 5.9: SAS spectrum for transitions from the lower hyperfine ground state of rubidium 85. The repump is locked to this spectrum, but it is the weakest of the four, and we often simply lock to the edge of the Doppler broadened background.

5.1.3   Alternative Technique: DAVLL

Despite the extremely broad width of the feature, laser locking based on Doppler broadened spectra can be successful. Care must be taken to account for systematic offsets, and the frequency is typically compared with (but not locked to) a Doppler-free signal. One may lock to the side of a Doppler broadened spectrum, or, using a Dichroic Atomic Vapour Laser Lock or DAVLL [114], one can lock to any point typically several hundred MHz around the absorption peak. The 'capture range', the range over which the signal remains linear, is also several hundred MHz wide, compared with the much smaller natural linewidth scale features of Saturated Absorption Spectroscopy.

Consider a rubidium vapour cell with an applied constant magnetic field. This field lifts the Zeeman degeneracy, and hence, for a given frequency, the absorption of left- and right- circularly polarised light will be different. By using a linearly polarised probe beam and measuring the difference between the absorption of the left- and right-circularly polarised components, one obtains a dispersion signal which is zero at the resonance and is approximately as wide as the Doppler broadened spectrum. This is polarisation spectroscopy in a vapour cell, and is similar to the Hänsch-Couillaud locking scheme [115].

The much larger capture range of DAVLL compared with Doppler-free signals is an obvious advantage, but there are unfortunately many sources of systematic offset, including voltage off-sets in the electronics and bi-refringence in the glass of the vapour cell. A DAVLL lock is under-construction in the group, and to confirm the laser frequency, a SAS signal is also measured, although it is not used for locking.

It is possible to combine the DAVLL technique with Saturated Absorption Spectroscopy. A strong pump beam creates the typical SAS dip in the Doppler broadened background for the right- and left- circularly polarised part of the spectrum, the difference of which is a sharp dispersion curve with a width of order the radiative lifetime; this is Doppler-free DAVLL [116].

5.1.4 Laser locking

Given the spectroscopic signal derived from the vapour cell, we must derive some signal which we can use as a measure of the deviation of frequency from some set-point to then feed back to and hence lock the laser. We typically wish to lock a laser to the maximum of an absorption feature, but locking to a maximum is not trivial. When locking a device, the sign of the error signal must correspond to the direction in which the signal has drifted, but for an extreme, a drift in either direction reduces the signal and the electronics cannot tell in which direction the correction should be applied. We can lock to the rising or falling edge of such a feature but we cannot lock arbitrarily close to the centre.

From the absorption maximum, we must derive a signal which is zero about our desired lock point, has a different sign either side of this point, and is, preferably, linear; the obvious candidate is the differential of the absorption feature. For DAVLL and other polarisation spectroscopy techniques, this derived signal is measured directly, but for SAS, we must perform the differential in electronics. In principle, we sample the signal at several nearby points, and take the difference; this can be achieved by dithering the signal, or using frequency modulation.

At first glance, deriving a signal from SAS with electronics may seem unnecessary when a similar signal is available easily using DAVLL and similar techniques. However, there is a strong advantage in that the linewidth scale feature of SAS can be frequency shifted, without additional broadening, and a laser locked to this shifted signal. This is described further in Section 5.2, and is of particular use in our experiment.

5.1.5 Frequency Modulation Spectroscopy

To introduce the technique of Frequency Modulated Spectroscopy, let us begin with dithering. Suppose that we can measure a signal proportional to the function $f(x)$, and we choose to apply a modulation to the free-parameter about an offset: $x(t) = x_0 + \beta \sin(\omega t)$. If the magnitude of this modulation is small, we can accurately approximate

the resulting function by a Taylor expansion:

$$f(x(t)) \simeq f(x_0) + \left.\frac{\partial f}{\partial x}\right|_{x=x_0} \beta \sin(\omega t) + \ldots \tag{5.2}$$

If we now 'mix' this signal with the same modulation (plus some phase delay ϕ) and then apply a low-pass filter such that frequencies at ω or above are rejected, we find

$$S(t) \simeq \left.\frac{\partial f}{\partial x}\right|_{x=x_0} \beta \sin(\phi) \tag{5.3}$$

and we may adjust ϕ to maximise this signal.

Hence we are left with a signal proportional to the differential of the function. Additionally, by continuing the Taylor expansion, it is clear that by mixing the signal with a harmonic of the modulation frequency we can extract higher order differentials.

This simple picture holds for low frequencies, but it fails as the modulation frequency becomes large. In this regime, we view the modulation in frequency space and, as described in more detail in Section 6.3.5, see that there are sidebands either side of the carrier. These additional frequency components probe the function absorption feature at $\pm \omega$ relative to the carrier and, by demodulating as before, we can measure the differential absorption of these components [117].

5.1.6 Beat Frequencies

With the lasers so locked, we set out to measure their stability and the technique described in this section is also used to characterise components of the manipulation laser apparatus.

A particularly practical method for characterising the linewidth of a laser, or indeed the stability of any oscillator, is to compare it with another, either a similar oscillator or a stable reference; we have two similar oscillators available as two lasers (repump and cooling transitions) are required to operate the MOT. With these two lasers locked to spectral features a few tens of MHz apart, the oscillation of intensity due to interference between the electric field generated by each of the lasers can be measured if the two
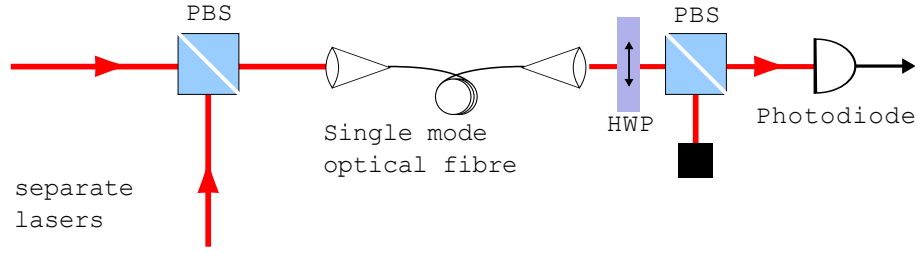
FIGURE 5.10: A typical setup used to record beat-frequency data. Spatial overlap of the beams is essential, and the short stretch of fibre provides the necessary spatial filtering. The electric fields must also be parallel, and the second beam-splitter cube ensures this condition. The half-wave plate can be used to tune the ratio of powers.

beams fall on the same photodiode. This relatively slow oscillation, compared with the frequency of the underlying electric field $\sim 384\text{THz}$, is a beat frequency.

The typical arrangement used for recording this beat-frequency data is shown in Figure 5.10. It is essential that the beams overlap and that they have the same polarisation. The beams are combined using a polarising beam-splitter cube, and then are converted to the same polarisation (unavoidably losing half the power) using additional polarising optics. Spatial filtering is provided by a short stretch of single mode optical fibre, and this ensures the electric fields of the two beams overlap perfectly. For these measurements there is plenty of optical power available, and we can afford to discard significant fractions of the power in order to obtain equal polarisations and a clean spatial mode.

The oscillation is recorded using a digital storage oscilloscope, and the standard way to analyse such data is to take the Fast Fourier Transform and identify the frequency components; this is shown in Figure 5.11. A peak corresponding to the frequency difference between the lasers dominates, and the width of this peak characterises the laser stability, but other peaks corresponding to various frequencies in the system are clearly visible and are identified in the figure.

A more informative plot of the frequency characteristics of a laser, especially one measured in this way, is the Allan variance [118]. This two-sample variance is

$$\sigma^2(\tau) = \frac{1}{2}\langle [\nu_{i+1} - \nu_i]^2 \rangle; \ \nu_i = \frac{c_i}{2\tau} \tag{5.4}$$
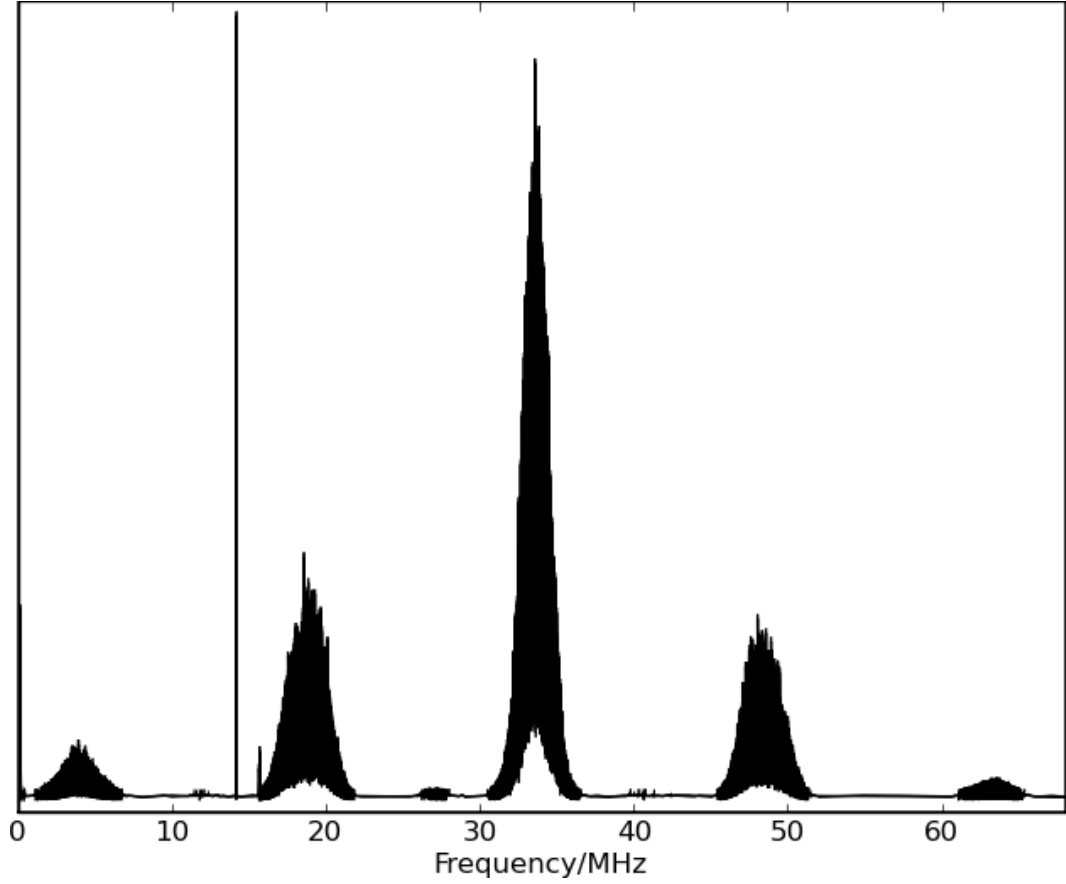
FIGURE 5.11: The Fast Fourier Transform of the beat frequency data. The broad central feature at 33.5MHz is the beat note and the first and second order sidebands are also visible. The narrow features at 14.15MHz and 15.65MHz are the modulation frequencies of the individual lasers.

where c_i is the number of zero crossings which occur between $t = i\tau$ and $t = (i+1)\tau$, and $\langle \ldots \rangle$ denotes the average [119]. We plot the Allan deviation, the square root of the variance, as a fraction of the optical carrier frequency ~ 384THz. This measure was originally formulated to characterise clocks which typically output regular ticks with a certain jitter, rather than a sinusoid, and a clock-like signal can be derived from our data by marking the positions of each of the zero-crossings. With this record of the time for each zero crossing, the Allan variance (or any other statistical measure of clock frequency) can easily be calculated [120, 121]. The Allan variance is shown in Figure 5.12.
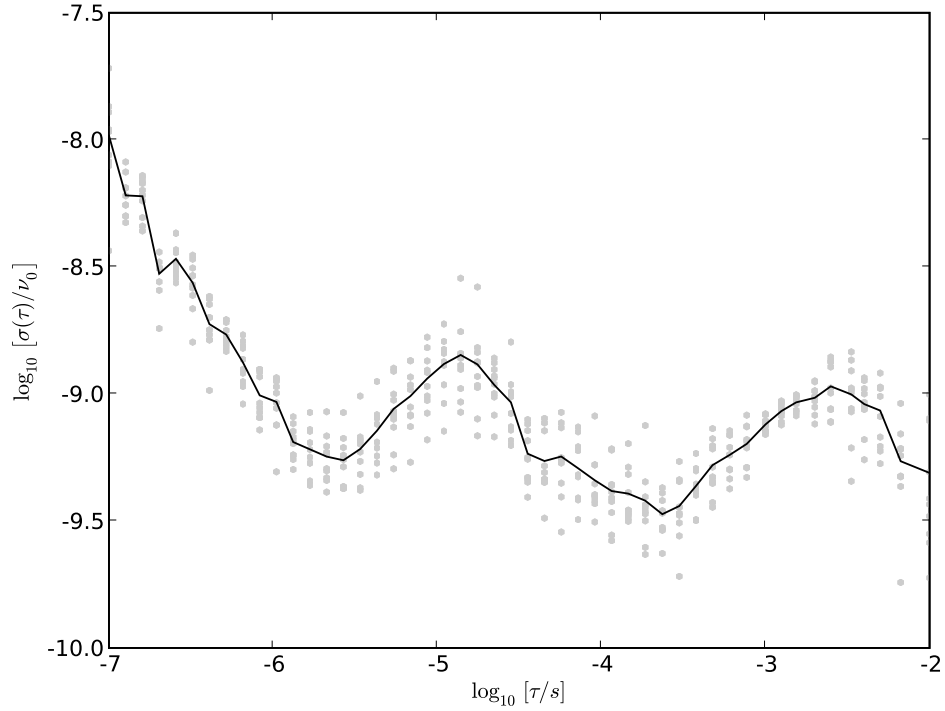
FIGURE 5.12: The Allan variance of the beat frequency between our two locked diode lasers, averaged over 10 data sets; the points for each set are shown in grey. The initial decrease is $1/f$ noise, which occurs too fast for the locking electronic to correct; the two lumps are so-called servo-bumps and are due to the response of the current feedback (100kHz) and piezo feedback (500Hz).

5.2 Zeeman-shifted Laser Locking

We now see how we can apply the precise red-detuning necessary for a MOT without sacrificing laser stability.

For a real multi-level atom, the ideal two-level system can be emulated by using a closed-transition; that is, a transition for which the upper-level can only decay to a single ground-state. The selection rule $\Delta F = \{0, \pm 1\}$ means we may use the transition $\left|5^2S_{1/2}, F=2\right\rangle \rightarrow \left|5^2P_{3/2}, F=1\right\rangle$ or $\left|5^2S_{1/2}, F=3\right\rangle \rightarrow \left|5^2P_{3/2}, F=4\right\rangle$. However, even exactly on-resonance, the overlap of the Lorentzian profile of adjacent upper F levels with the laser is not zero, and if the atom is pumped to one of these states, it has a finite chance of decaying to the wrong ground-state. Given we must red-detune by a few linewidths, the large separation between $\left|5^2P_{3/2}, 3\right\rangle$ and $\left|5^2P_{3/2}, 4\right\rangle$ (see Section 3.4) means this overlap is minimised for the latter of the two possible cooling transitions.

There are a variety of methods by which the necessary red-detuning can be applied. Typically, the laser is locked to the transition and the frequency shifted via an AOM. However, the method we use shifts the frequency of the *transition*, and we lock the laser to this shifted resonance.

When a quantisation axis is present, the Zeeman degeneracy is lifted and using circularly polarised light it is possible to pump the atom population to an extreme m_F state. With the population so prepared, only a single transition participates. Hence, by using circularly polarised pump and probe beams for SAS, it is possible to approximate in a real atom the idealised two-level situation. Furthermore, this quantisation axis can be provided by an external magnetic field, and it is straightforward to control the magnitude of the field on which, to first order, the Zeeman frequency shift linearly depends.

We use a pair of magnetic field coils, custom designed to apply a field uniform to within 1% for the volume swept out by the SAS beams, and apply a DC field; by adjusting quarter-waveplates, the single Doppler-free peak, corresponding to the cooling transition, is selected and the laser is locked to this feature. By adjusting the current in these Zeeman-shifting magnetic field coils, we are able to adjust the locked lasers, relative to the transition, reproducibly and with sub-MHz accuracy.

5.2.1 Zeeman shifting coils

In principle, the best geometry to apply a uniform magnetic field is the Helmholtz configuration, but unfortunately the maximum radius is limited by the height of the cell above the optical bench, giving a radius too small to enclose the cell if the Helmholtz condition is enforced. Instead, the optimum coil geometry, considering practical constraints such as post height and wire thickness, was found using a numerical model of the coils (a summation of the field from regularly spaced concentric current loops) and discussion with the mechanical workshop.

We use two coils, each with inner-radius 7cm, width 95mm, and 25×12 loops, with a separation of 8cm between the pair. We calculate the field in three-dimensions by taking

the field for a single loop and summing over the many loops making up the coils, and we find that over the narrow beam width of 2mm, the field varies by only a few percent.

5.3 Magnetic Fields

The spatially dependent force of a MOT is derived from the Zeeman shift due to a spatially dependent magnetic field. That the divergence of the magnetic field is zero, $\boldsymbol{\nabla} \cdot \boldsymbol{B} = 0$, means we cannot construct a field with a maximum in free-space, but we can construct a field with a minimum and this is achieved with the anti-Helmholtz configuration.

The position of the zero thus created is perturbed by external magnetic fields, notably that of the ion pump, and to allow us to correct for these influences and otherwise control the field, compensation coils have been added to the basic two-coils of the anti-Helmholtz pair.

5.3.1 Anti-Helmholtz coils

The Helmholtz configuration consists of two identical circular magnetic field coils placed co-axially such that their separation is equal to their radii and with currents flowing such that the sign of the fields is the same. This configuration optimises the uniformity of the field near the centre of the pair. The *anti*-Helmholtz configuration is similar except that the currents flow such that the fields oppose and the separation is chosen to optimise the linearity of the field about the centre.

The on-axis field from a current loop can be found from the Biot-Savart law:

$$B_z(z) = \frac{\mu_0}{4\pi} \frac{2\pi R^2}{(R^2 + z^2)^{3/2}} I.$$
(5.5)

The separation for Helmholtz condition is found by adding the field from two such loops, transposed forwards and backwards along the z axis by $s/2$, and finding the separation which sets the rate of change of field gradient about $z = 0$ to zero: $s = R/2$

The separation for the anti-Helmholtz condition is similarly found by subtracting the contribution and, since we are now interested in uniformity of the gradient, setting the third order differential to zero: $s = R\sqrt{3}/2$.

We use the relatively large coil inner radius of 103mm and $\simeq 430$ windings of Standard Wire Gauge 22 (0.710mm) copper wire. In the anti-Helmholtz condition, these provide a field gradient of ~ 5G/cm along the axis.

Given the uniform gradient along the z axis near the origin, the condition $\boldsymbol{\nabla} \cdot \boldsymbol{B} = 0$ and the radial symmetry of the problem allow us to deduce that the field gradient along the radial direction is half and negative that along the axial direction. In cylindrical co-ordinates, where r is the radial component,

$$\boldsymbol{B} = \left| \frac{\partial \boldsymbol{B}}{\partial z} \right|_{\text{origin}} \sqrt{(r/2)^2 + z^2} \times \frac{z\hat{\boldsymbol{z}} - r\hat{\boldsymbol{r}}}{\sqrt{r^2 + z^2}} \, . \tag{5.6}$$

Near the origin, surfaces of equal magnetic field magnitude are ellipsoids half as high as they are wide. The MOT region, discussed in Section 2.4.6, is hence also ellipsoidal, but extends beyond the range where the field depends linearly on displacement.

5.3.2   Stray fields and Compensation Coils

The exceptionally narrow width of the Raman transition encourages us to consider the perturbations which might shift the resonance and so hinder our search or otherwise influence our results. One such perturbation is that due to stray magnetic fields, a major source of which is the large magnets which are an essential part of the ion pump; these magnets are not only particularly strong ~ 1T, but are large ~ 20cm, meaning their strength will decrease more slowly with distance, and are close ~ 20cm to the cloud. One might therefore expect the atoms to be subjected to a non-zero magnetic field. This, fortunately, is not the case. The dynamics of the MOT ensure that, for equal powers (see Section 3.6.2), the cloud forms around zero magnetic field, regardless of where this zero occurs relative to the geometry of the chamber, ion pump, and other apparatus. This holds provided the stray magnetic field is approximately constant over

the width of the cloud, and while the displacement it causes to the position of the zero can be accommodated by alignment of the beams.

There is, however, virtue in applying an external magnetic field, especially if it can be changed in time, as was discussed in Section 3.6.1. We have four identical coils which are used in pairs and aligned along the axis of the two horizontal MOT beams, perpendicular to the axis of the anti-Helmholtz pair. The geometry of the vacuum chamber restricts the size of the coils and we cannot fulfil the Helmholtz condition. We require that the field be uniform over only a small region and, using 22 (axial) by 15 (radial) windings of 0.6mm copper wire, with inner radius 67mm and separation 260mm, the calculated field varies by less than 0.5% over the width of the beams.

The coils have an impedance of 13Ω and for the frequency range we use we can drive the coils with an audio amplifier. We use a 2 channel Velleman Power Amplifier producing $100\mathrm{W_{rms}}$ per channel (into 8Ω) which, driving two coils in parallel, is approximately $\pm7\mathrm{A}$ or $\pm45\mathrm{G}$ at the centre of the pair.

5.4 Fast switching

For many experiments it is necessary to quickly extinguish and otherwise control the intensity of the beams. The device we use is an Acousto-Optical Modulator and although this is discussed in detail in Section 6.4, its inclusion has a significant impact on the MOT implementation. An AOM modulates the intensity of an acoustic wave from which an optical beam is diffracted. This acoustic wave is travelling, and so the light acquires a frequency shift equal to the modulation frequency of 110MHz. Clearly we must account for this frequency shift if we are to keep to the correct red-detuning for Doppler cooling. Fortunately, the Zeeman shifting arrangement permits a sufficiently large adjustment of the locking transition. In our implementation, we choose a negative AOM shift -110MHz and compensate with $+92$MHz Zeeman shift to give, overall -18MHz detuning, or three linewidths to the red. Diffraction efficiencies can exceed 90% and any imperfections in the extinction ratio are due to imperfect extinction of the RF signal which is, fortunately, negligible.

Two alternative approaches, neither of which introduces an unwanted frequency shift, are Pockels cells and mechanical shutters. The former is conceptually an electrically switchable waveplate and is discussed in Section 6.6, but has not been used for amplitude modulation because of the relatively high voltages involved and the imperfect extinction ratio. The latter requires tight-focusing of the optical beam to achieve any respectable speed (as this reduces the distance the shutter must move), but has been used successfully elsewhere [122, 123]. Although there is no loss of power through such a device (and they can be significantly cheaper), it cannot compete with AOMs on either speed or reliability.

5.5 Vacuum Chamber and Optics

The layout of the free-space optics used to send the three orthogonal pairs of counter-propagating beams into the vacuum chamber of the MOT is shown in Figure 5.13. The main part of the vacuum chamber is a large form cut by computer-aided manufacturing, including precision knife edges, from a single block of stainless-steel, and sits on a raised platform above the optical table. There are a total of 14 ports, conforming to standard sizes, and each is sealed with a copper gasket and anti-reflection coated glass window, or a blank metal flange.

Light emerges from a 7.1mm FC/APC fibre collimator, and is then guided by free-space optics. A half-wave plate is used to alter the polarisations and send one third of the molasses laser power vertically. The remaining two thirds is sent via a periscope to the raised platform, and one half is sent along each horizontal axis. All free-space optics are $2''$ to allow for the possibility of larger beams in the future.

The vacuum windows are arranged on the surface of a sphere and the surface normal at the centre of each passes through the geometrical centre of the chamber. There is an unavoidable region of distortion around the circumference of each of the vacuum windows, but in this geometry we can make good use of even the smallest ports.
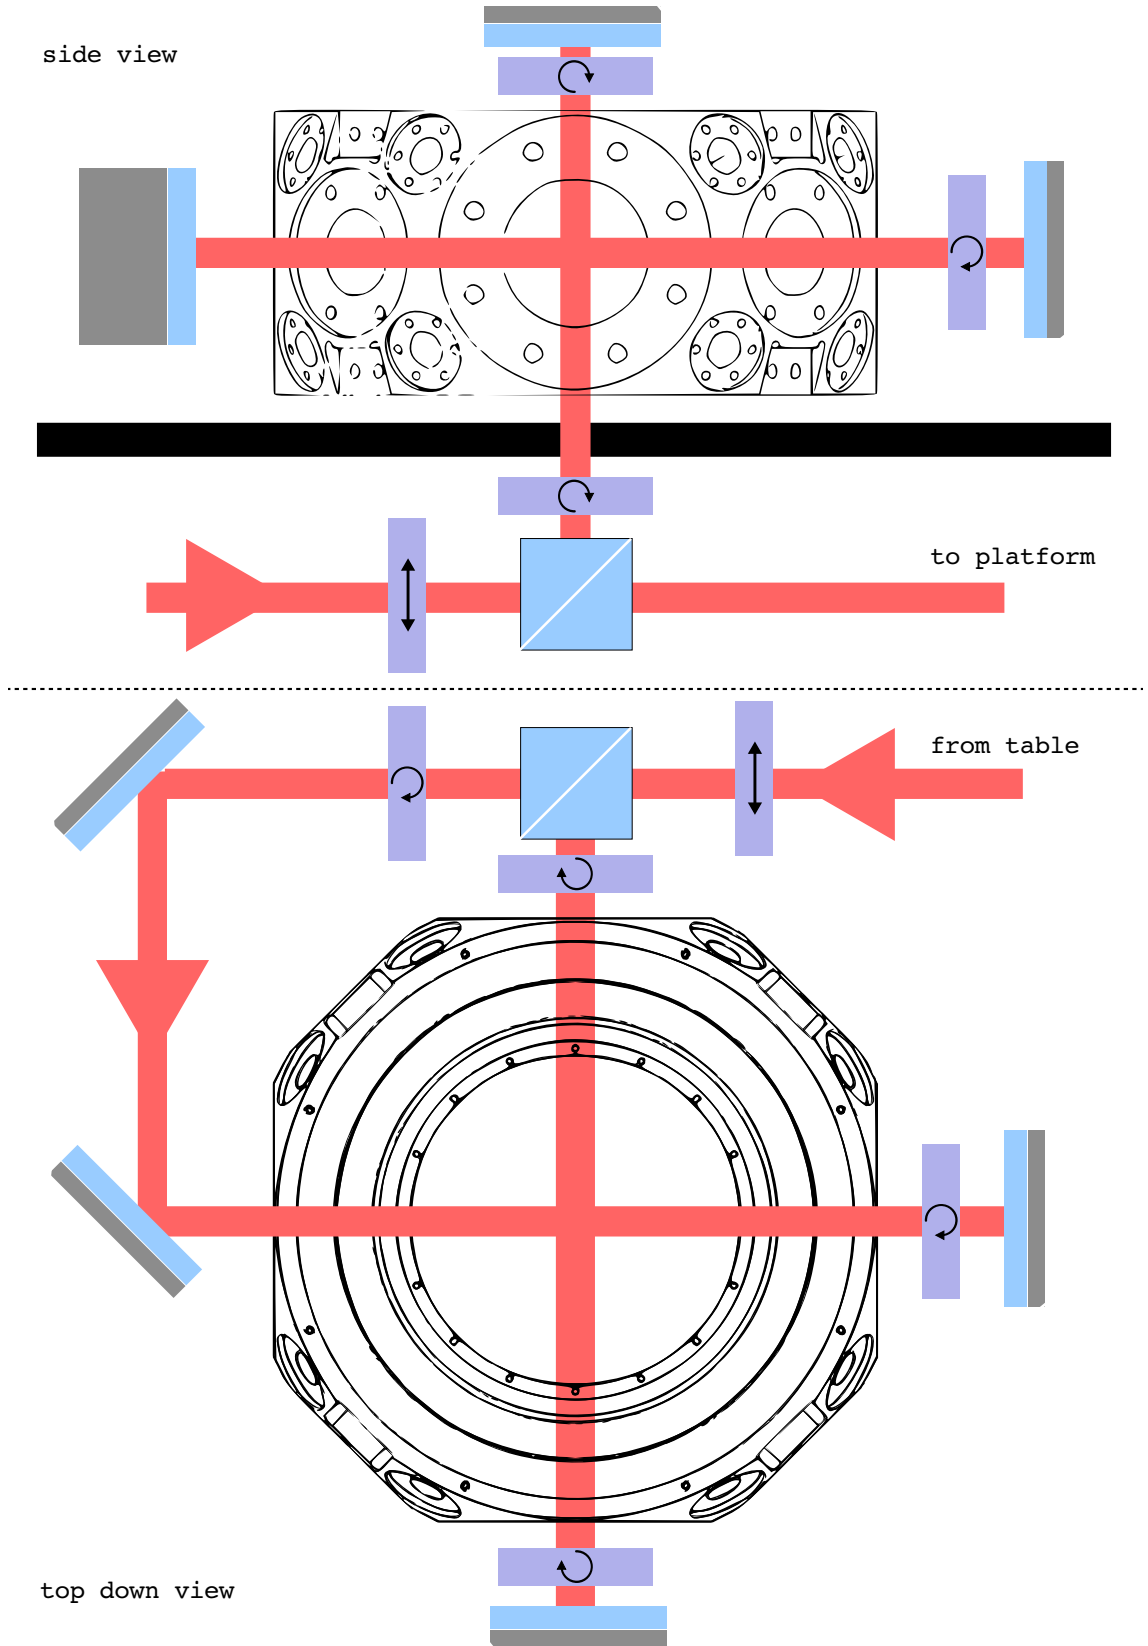
FIGURE 5.13: The principle optics used to produce counter-propagating beams along each of the three Cartesian axes. The chamber is on a platform above the optical table, and the vertical beam is split by a cube directly underneath the cloud. Each beam is directed into the chamber and then retro-reflected. The half-wave plates ($\updownarrow$) control the ratio of powers split by following cube, and the quarter-wave plates ($\circlearrowleft$) ensure the correct circular polarisation for the trapping beams.
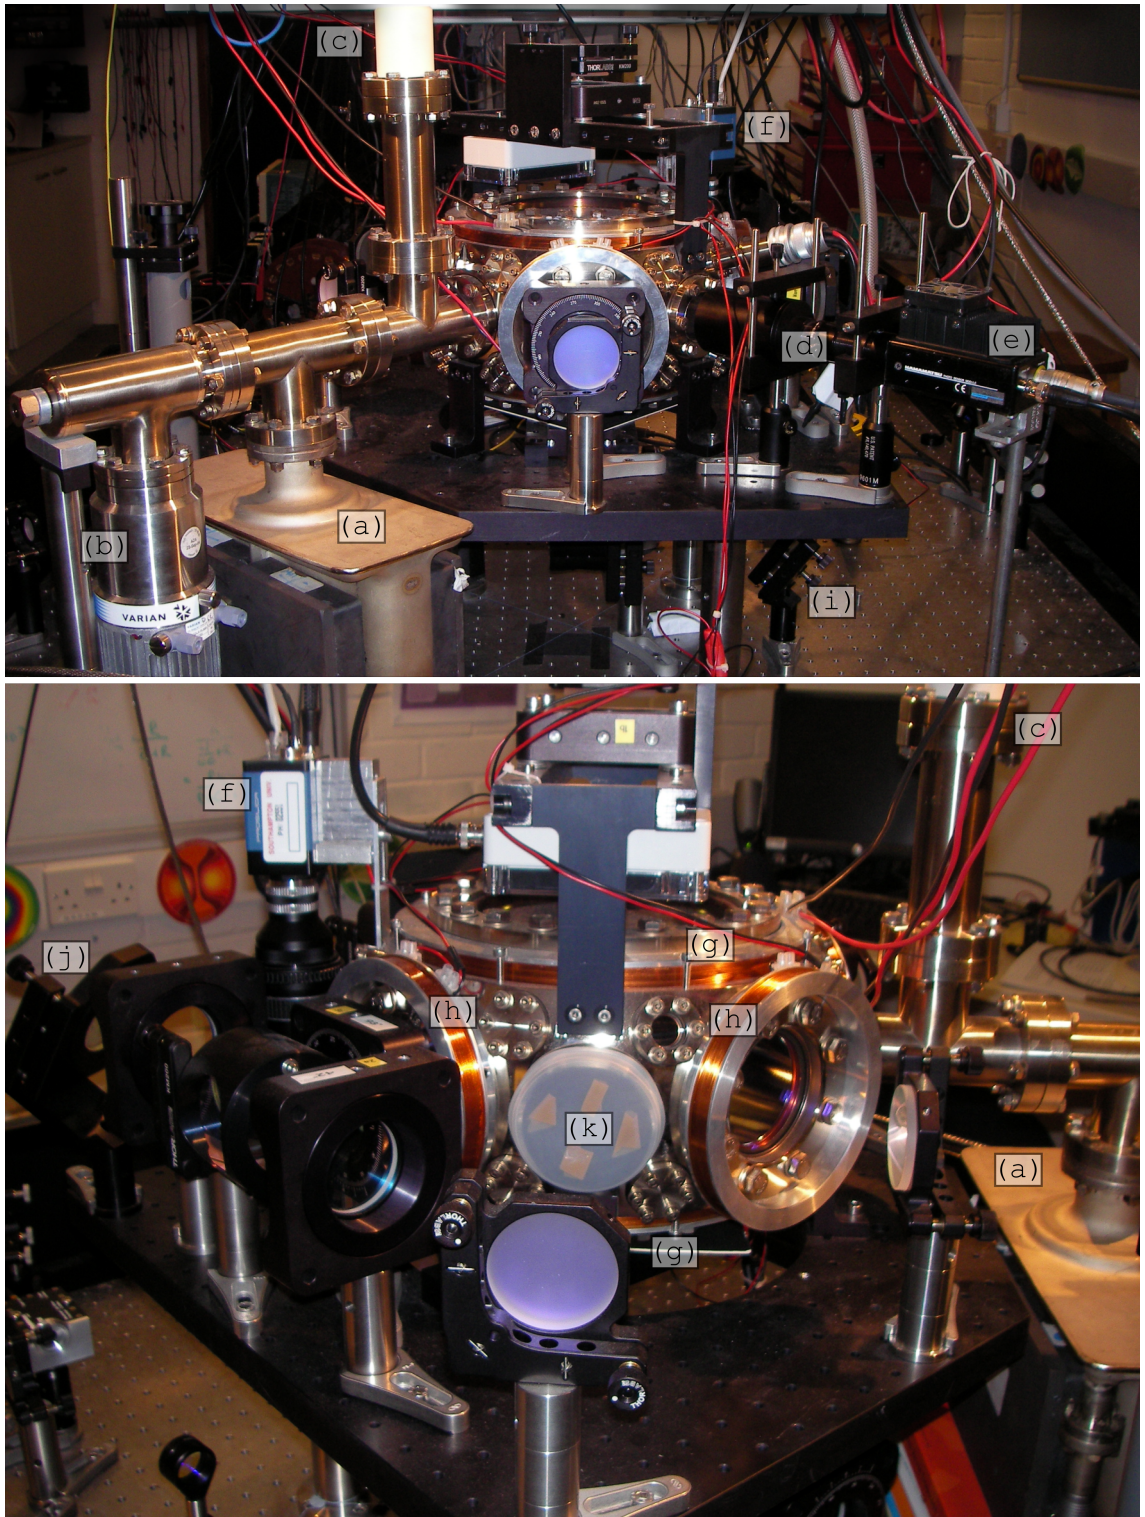
FIGURE 5.14: Front and side views of the vacuum chamber and optics used to create the magneto-optical trap. The following components are visible: ion (a) turbo (b) and sorption (c) pumps; collection optics (d), photo-multiplier tube (e), and fast camera (f); anti-Helmholtz coils (g) and compensation coils (h); mirrors to direct light upwards (i) to the platform, and horizontally along it (j); and a covered viewport opposite to the collection optics (k).

5.5.1 Fibre Coupling

The beam paths from the lasers to the MOT cloud are several metres long, and if the lasers need to be adjusted, even a small change in the angle of the grating at the output can displace the beams sufficiently that they no longer overlap. To remove this problem, we use a short stretch of single mode optical fibre. Any change in the beam angle may decrease coupling into the fibre, but we know any light emerging from the fibre is still aligned correctly into the MOT. We can maximise this power using a power meter and by adjusting the first two mirrors immediately after the laser.

It is important to use angled-cleaved fibres. The air/glass interface of the fibres is not anti-reflection coated, so we expected 4% reflection from both input and output facets. With well aligned flat cleaved fibre collimators, this reflection is directed back to the diode, and the small fraction which is not extinguished by the isolator is still sufficient to affect the stability. If instead the fibre cleaved is angled, coupling into the fibre is maximised when light is not normal to the interface, and the reflected light does not make its way back to the laser. We use an FC/APC air-spaced doublet collimator, manufactured specifically for 780nm, and with a $1/e^2$ waist 7.1mm.

5.5.2 Digression: Waveplates

Waveplates are well understood, but we have noticed an interesting property of quarter-wave plates and circular polarisation which greatly simplifies the setup of a magneto-optical trap, and we describe this here.

'Waveplate' is the name given to a slab of birefringent material selected such that the optical path difference between the fast the slow axes is (modulo 2π) a known phase delay. In the case of a 'half-waveplate', this delay is π and for a 'quarter-waveplate' it is $\pi/2$. The effect of a waveplate on an incident beam can be found by considering the relative delay of the components of the incident polarisation aligned with each of the axes. For example, a linearly polarised incident beam aligned at 45° to the axes of a half-waveplate can be analysed as two equal in-phase components which accrue a
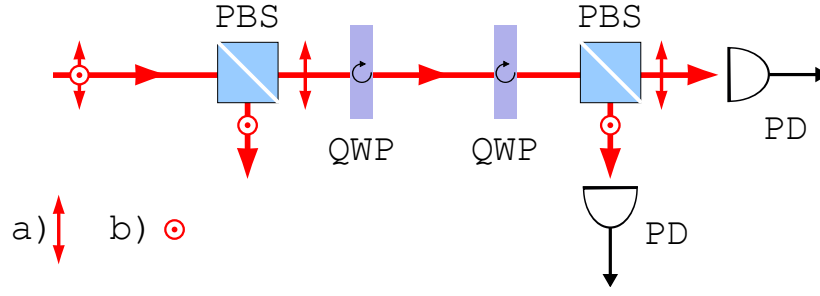
FIGURE 5.15: The optical setup used to find the alignment of the fast or slow axis of our waveplates, with a) horizontal and b) vertical polarisations, as indicated.

relative delay of π, yielding an output polarisation orthogonal to the input. More generally, we represent polarisation by a two-element column vector, and represent optical components, including waveplates, by two-by-two matrices; this is Jones calculus.

Unfortunately the markings on commercially bought waveplates indicating the fast and slow axes are not always present or reliable, and so it is necessary to determine this alignment experimentally. Using a polarising beam-splitter cube and a known input polarisation, one can quickly find the axis of a half-waveplate; the typical arrangement is shown in Figure 5.15. The waveplate is rotated to minimise transmission through the second cube; the minimum one can expect with a half-waveplate is $\simeq 0$, and with a quarter-waveplate $\simeq 1/2$. Unfortunately, with this method one cannot distinguish between fast and slow axes. For the half-wave plate this is inconsequential, but for the quarter-wave plate, this difference determines whether the output is right- or left-circularly polarised. However, it is possible to compare the axes of quarter-waveplates relative to each other. By inserting a second quarter-waveplate in addition to one for which the angle which gives minimum transmission through the second cube has already been found, one can rotate the second plate such that the phase delay of one component begun by the first plate is continued by the second, emulating a half-waveplate and approaching zero transmission through the second cube. In this case, the fast and slow axes of the two quarter-waveplates are parallel. This can be repeated for any number of quarter-waveplates so that, although their alignments are not known absolutely, they are all known relative to one master and hence to each other.

In our Magneto-Optical Trap, the circular polarisations are produced by using quarter-waveplates to convert linear polarisation, and retro-reflection including a double-pass

through a quarter-waveplate. This retro-reflecting arrangement has the unexpected advantage that, if one begins with circular polarisation, transmission through a quarter-waveplate, retro-reflection and transmission through the same waveplate will yield the opposite circular polarisation *regardless* or the orientation of this waveplate. We represent the incident circular polarisation v_{in}, waveplate Q, and rotation matrix $R(\theta)$ as follows:

$$v_{\text{in}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ \pm i \end{pmatrix} ; \ Q = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} ; \ R(\theta) = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{5.7}$$

The 'double pass' through the waveplate—consisting of one pass, retro-reflection, and a second pass—is, in the orientation of the waveplate, simply $Q \cdot Q$. We can rotate this object to the lab-frame via $R(-\theta) \cdot Q \cdot Q \cdot R(+\theta)$. This transformation applied to a circular polarisation yields

$$v_{\text{out}} = R(-\theta) \cdot Q \cdot Q \cdot R(+\theta) \cdot \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ \pm i \end{pmatrix} = \frac{e^{\pm i 2\theta}}{\sqrt{2}} \begin{pmatrix} 1 \\ \mp i \end{pmatrix} \tag{5.8}$$

which is the opposite circular polarisation, with only a global phase delay dependent waveplate angle.

When using these quarter-waveplates to produce the circular polarisation necessary for a MOT, we cannot know whether we are creating left- or right- circularly polarised light. However, since the magnetic field gradient, discussed in Section 5.3.1, has opposite sign along the vertical direction compared with the horizontal directions, by ensuring the circular polarisations of vertical and horizontal beams are opposite, we know that up to a change in sign of the field gradient, the polarisations are correct. This approach reduces the number of combinations from $2^6 = 64$ (magnetic field fixed and all waveplates set independently) to $2^3 = 8$ (orientation of waveplate before retro-reflector acknowledged to be inconsequential) to $2^1 = 2$; these two combinations are easily checked by reversing the current in the magnetic field coils. With so many other criteria which must simultaneously be met, this represents a significant simplification.

Chapter 6

Manipulation Laser Apparatus

The Raman beams to be used for coherent manipulation of the ground states of our trapped atoms must adhere to rather stringent constraints. The beams beams must be phase coherent, literally meaning the relative phase between them must behave deterministically over, for our purposes, the entire duration of not just an individual pulse, but all the pulses making up one interferometer sequence. This is feasible with separate narrow-linewidth (say $\lesssim 100$Hz) lasers locked to an external reference, or even with one laser injection-locked to a frequency-shifted version of the other, but we have chosen to use the more suitable approach of deriving both Raman beams from the same laser source. We favour this approach because we are able to directly synthesise the microwave frequency corresponding to the hyperfine splitting and so limit the 'linewidth' of the Raman excitation (which, in the case of two separate locked lasers is the quadrature sum of the individual linewidths) to be due entirely to the stability of our microwave source [124, 125]. Additionally, using standard microwave electronics, we can control not only the frequency of this difference, but the relative phase between successive pulses. One final practical advantage is that the detuning of the Raman beams from the single photon resonance which enhances the transition probability is tuned by adjusting one laser, and is decoupled from the difference frequency between the beams.

The following sections describe the apparatus and our implementation used to create these Raman beams.

6.1 Overview

The main frequency difference is provided by an electro-optical modulator, which modulates the phase of an incident optical beam at the applied RF frequency. This phase modulation creates 'sidebands' above and below the optical carrier frequency, spaced by the modulation frequency. If we modulate at the frequency splitting of the hyperfine levels, the combination of central carrier frequency and one sideband can resonantly drive the Raman transition. However, the positive and negative sidebands created by this phase modulation have opposite sign, and hence drive the Raman transition with opposite sign and exactly cancel. To break this symmetry, we use an acousto-optical modulator before the electro-optical modulator, which frequency shifts a controllable fraction of the beam power and, furthermore, deflects this frequency shifted fraction, spatially separating it from the unshifted remainder. By running the EOM at a frequency less than the hyperfine splitting, such that the sum of EOM and AOM frequencies equals this splitting, the combination of AOM shifted beam and *one* of the EOM sidebands drives the Raman transition, while all other combinations (carrier and either sideband; carrier and AOM shifted beam; AOM shifted beam and the other sideband) miss the two-photon resonance.

The spatially separated beams can be controlled separately and, specifically, their polarisations can be chosen independently. We set the polarisations to be orthogonal and recombine the beams using a polarising beam-splitter cube so that they are co-linear. This composite beam is then amplitude modulated using another AOM, the frequency shift from which is the same for each frequency component and hence leaves the crucial frequency difference unaffected. Note that this arrangement sacrifices independent amplitude modulation of the AOM and EOM beams but, as we intend to emulate a two-level system, it would be cheating to independently modulate the beams (as in, e.g. STIRAP, see Chapter 4) as this would make use of the multi-level nature of the atom.

This is the basic setup used to drive the Raman transition, and is illustrated in Figure 6.1. Certain modifications have been made to improve the situation—such as removing the carrier from the EOM line—and the generation of the RF signal to drive the EOM has yet to be described; these details, and descriptions of how the various modulators operate, is the subject of the following sections.
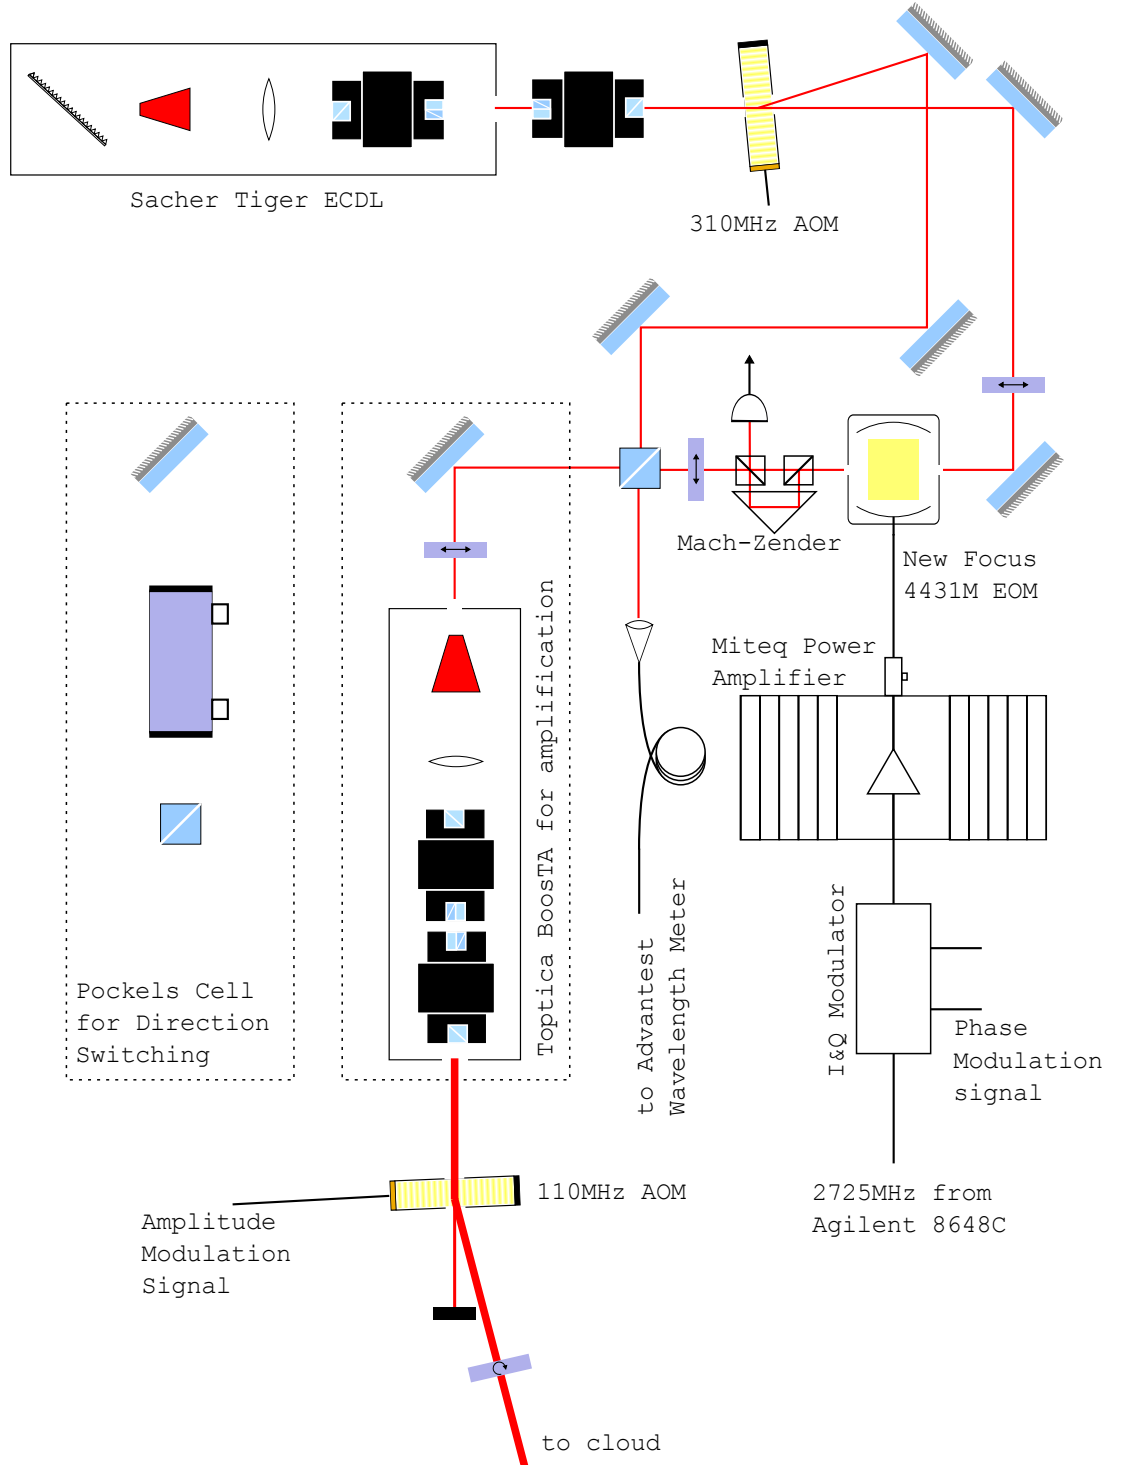
FIGURE 6.1: Overview of the apparatus used to generate the fields necessary to drive a two-photon Raman transition in rubidium 85. The dotted boxes enclose optional parts of the system which we cannot use simultaneously. The Pockels cell allows us to change the direction of the effective photon driving the effective two level atom, but for the current range of experiments, we have used the additional power available when using the BoosTA for amplification.

6.2 Laser Light Source

The laser chosen as the main source for the manipulation laser setup was a Sacher Tiger 1W external cavity diode laser. This commercial device has some interesting differences compared with typical ECDLs (such as those described in Section 5.1) and promises some significant advantages, but unfortunately the device is let down by comparatively less sophisticated mechanical construction, making the coarse alignment of the grating and cavity unreliable and difficult to adjust. The most striking difference is that the diode, unlike most, emits from both facets. This allows the grating to be placed on the back facet so that there is no possibility of changes to the grating angle affecting the direction of the output beam.

The second significant difference is to manage problems arising from the unusually high power output for a single diode. To ensure the optical intensity in the diode is low, the majority of the light makes only a single pass; the output facet is anti-reflection coated (only $\sim 1\%$ is reflected) and so the power circulating within the cavity is little more than the output power. This allows the device to be tapered, since the optical power is relatively low near the back of the device, and it need only be wide near the output, so that this high power is distributed over a larger area; this allows a larger power output for a given size of diode while maintaining a single spatial mode.

We have found the laser to be extremely unstable unless a 30dB isolator, in addition to the 30dB included inside the device, is used. Without this, even the relatively small reflection from the first lens in an optical chain causes sufficient back-reflection to perturb the device. With this additional isolator, the mode-hop-free tuning range is many GHz.

The intolerance of the device to optical feedback is perhaps an unavoidable feature of the high single-pass gain tapered diode approach, but the mechanical construction is, unfortunately, less than innovative. The simplistic mount is shown in Figure 6.2. Once aligned, the device behaves well, but any alignment performed at the factory is significantly perturbed in transit, and the simple mechanical construction means that any torque applied to a screw designed to control one axis inevitably affects the other, and adjustment must be undertaken with care.
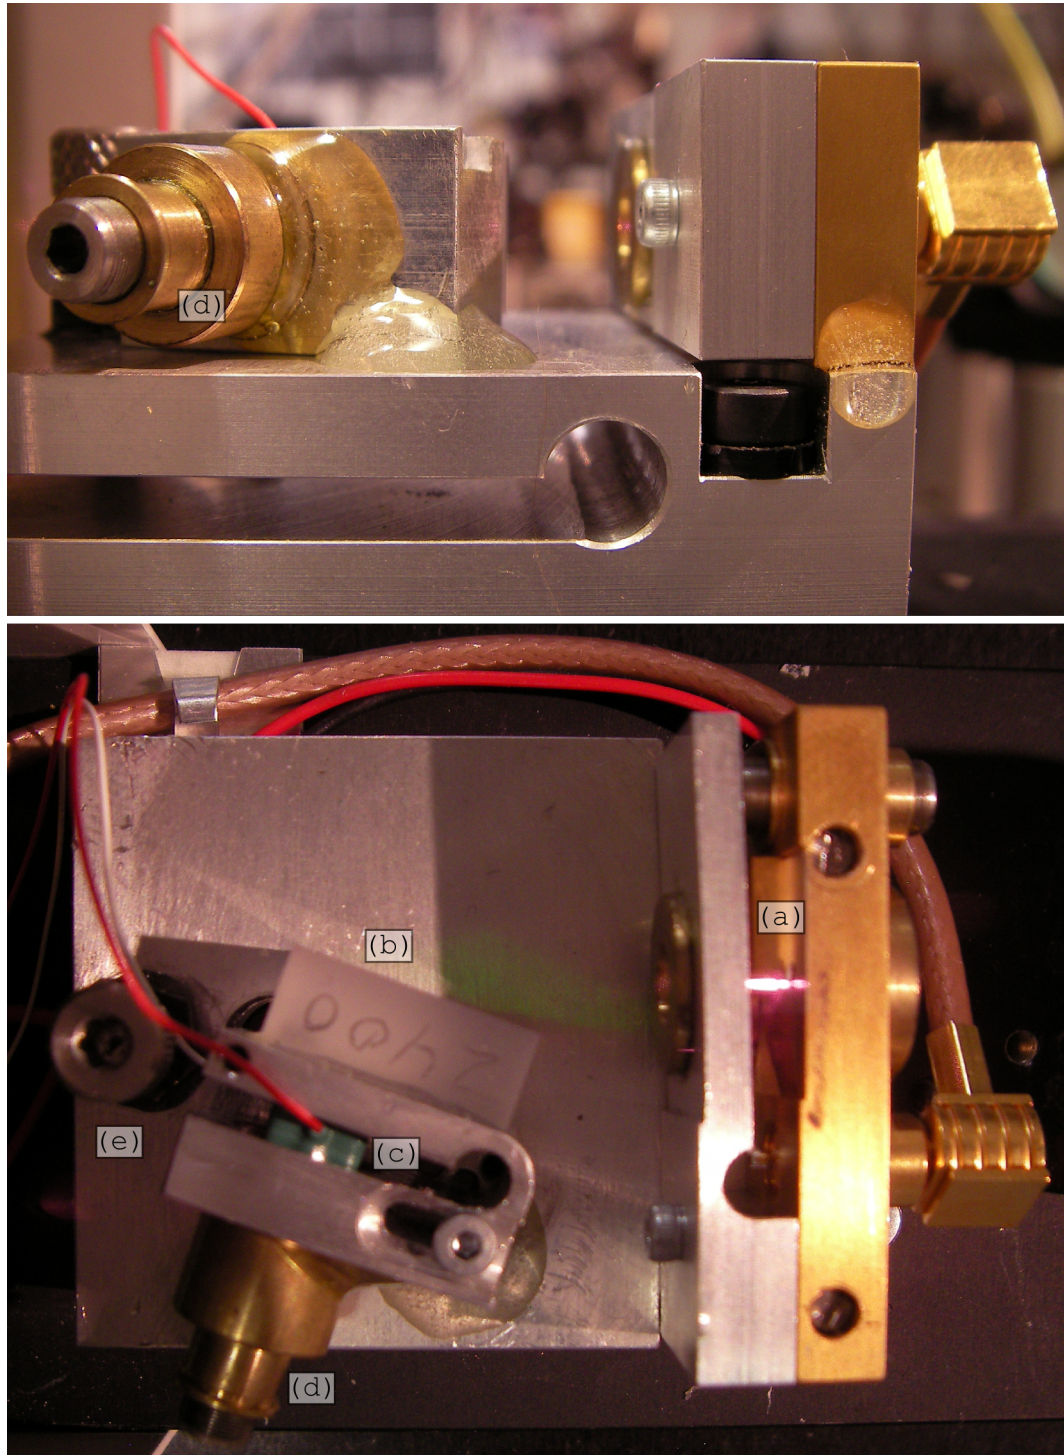
FIGURE 6.2: Internal optics of the Tiger Sacher 1W ECDL. Light emitted from the back facet of the diode (a) is reflected by a grating (b) whose angle can be controlled electronically by a piezo-electric stack (c) and more coarsely by a screw (d). The rotation of the grating about the orthogonal axis is via the cavity alignment screw (e).

6.2.1 Frequency Stability

The detuning of the main laser from which the Raman beams is derived influences the coupling strength (see Section 3.4), the hyperfine ground state splitting via the lightshift (see Section 8.3.2), and the possibility of populating the radiative upper state. Despite having good short-term stability, as in any similar device, the frequency drifts over minutes and hours and, for our experiments to be reproducible, we must account for or control this drift. Unlike the lasers used to produce the MOT (see Section 5.1), we must detune by many hundreds of MHz to where there is no absorption in rubidium 85, and hence to where there is no spectral feature to which we can lock.

The first method we consider to monitor the wavelength is simply to use a wavelength meter. We use a Advantest Q8326 which can provide 10^{-4}nm $\sim$ 50MHz precision, and can be calibrated against lasers locked to a rubidium transition, as shown in Figure 6.3. Additionally, this value can be read by a computer, which can then implement a feedback loop. We emulate a PI circuit—a feedback loop where terms proportional to the error signal and the time integral of the error signal are summed to create a feedback signal— and actuate the feedback by changing the computer-controllable piezo voltage on the laser. This setup was tried, but we found the 50MHz precision inadequate; the change in piezo voltage required to correct for the minimum drift detectable by the wavelength meter was often sufficient to cause the laser to mode-hop.

An alternative method is to lock to a spectral feature not in rubidium 85, but in the other abundant isotope rubidium 87. The Doppler broadened spectrum of this isotope extends to $\sim$ 3GHz above the transition in rubidium 85, and while there are suitable frequencies in this range, we cannot escape the region where lightshifts are significant (see Section 8.3.2) and population of the the upper-state is a danger.

We have considered frequency locking to one of the MOT lasers. By measuring the beat-frequency (see Section 5.1.6) we can derive a signal we can process in electronics and from which we can derive an error signal. However, for the GHz frequency shifts we wish to apply, the detectors and electronics are not trivial.
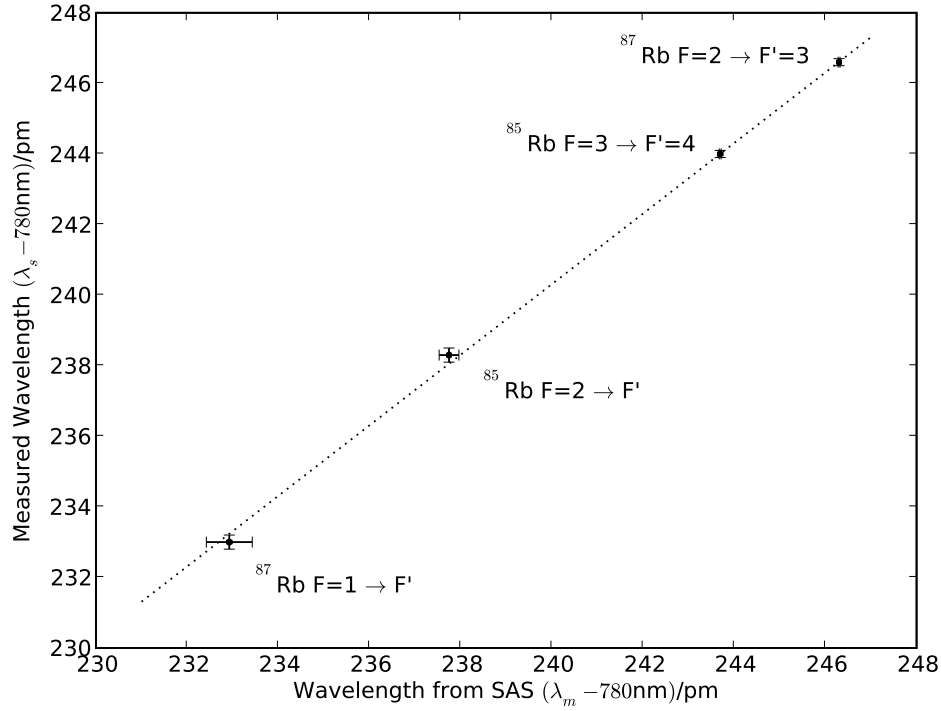
FIGURE 6.3: Readings of the Advantest wavelength meter for light of known frequency, found using Saturated Absorption Spectroscopy, discussed in Section 5.1.2. The wavemeter reading is offset by +0.3pm.

A final method which uses existing equipment is to use the Optical Spectrum Analyser in conjunction with the wavelength meter. The 50MHz precision of the wavelength meter overlaps with the 2GHz range of the OSA, and so, if we know the approximate frequency by using the wavelength meter, we can find the frequency much more precisely by measuring the combined spectrum of the Tiger and a reference laser using the OSA. Our precision is limited by the 12MHz resolution of the OSA, and the accuracy with which we know the frequency of the reference. From this, either a specifically built circuit or fairly simple computer processing, we can derive an error signal to feedback to the laser piezo.

Measurements of the short-term frequency stability of the laser, which the slow piezo lock cannot influence, is reported in Section 6.3.6. Overall, it is found, as specified in the device datasheet, to have a free-running linewidth of under 1MHz over the typical ~ 10ms duration of our interferometry experiments.

6.2.2   Amplification

A crucial part of the experiment so far, although one we hope to eventually do without, is the Toptica BoosTA. This device is similar to the Sacher Tiger, in that it is a high powered tapered laser diode transparent on both facets, but it has no grating feedback. Instead, this device allows one to seed the gain medium from an external laser, and provides single-pass amplification up to $100\times$ with a maximum output power approaching 1W. Unlike the Tiger, stability is not an issue as the BoosTA, up to some unavoidable amplification noise factor, reproduces the stability of the seeding laser. For a typical seeding power of 30mW, amplified spontaneous emission, that is, amplification under single pass of photons spontaneously emitted by gain medium, is typically 50dB below the amplified signal.

Currently, we use the Tiger as the laser source, which is frequency modulated using the AOM and EOM as described below, before this much weaker combined beam is used to seed the BoosTA. The BoosTA faithfully reproduces all frequency components of this beam, but can amplify only one polarisation component. Therefore, while before the BoosTA we are able to separate the composite beam using polarisation, for reasons described in Section 6.6, we are unable to after using the BoosTA for amplification. This ultimately limits the range of experiments we can attempt, but for early work the power increase it provides has been invaluable.

We hope that by increasing the efficiency of various parts of the modulation optics, we can greatly increase the optical power and so drive Raman transitions without the BoosTA. This would allow us to attempt a wider range of experiments, and specifically experiments to coherently manipulate the atomic momenta.

6.3   Microwave Frequency Shift

The majority of the 3.036GHz frequency shift is provided by an Electro-Optical Modulator driven by a series of radio-frequency components. There are various challenges in

creating a signal to drive the EOM. The signal must be accurately known around a relatively high RF frequency, must be phase coherently modulated around this centre, and must be at a relatively high power. The EOM, and the technological solutions devised to drive it, are discussed below.

6.3.1 Frequency Source

At the heart of the electronics for the Raman beam apparatus is an Agilent 8648C frequency synthesiser. The stability of this source is extremely important if we are to perform experiments using the very narrow hyperfine transition. The natural width of this transition is negligible, and calculations using typical parameters suggest we can expect a Raman frequency of several kHz; the uncertainty in the synthesised frequency must be less than this if we are to reliably 'dial-up' the transition frequency.

The power output of the Agilent is more than sufficient to drive the carrier driven I&Q modulator (see Section 6.3.2.1), but the 2.7GHz signal produced by this device is derived from an internal 10MHz reference oscillator, about which we cannot be so confident. The short-time performance is excellent (see Section 6.3.6), but a concerning imperfection is not the 'linewidth' of the RF signal, but the slow drift of the frequency over time. The quoted drift is $< \pm2$ppm/year and, by the end of the experiment, the Agilent was approximately 3 years old; temperature variations (within the operating range of $25°C \pm 5°C$) are expected to affect the frequency by $< \pm1$ppm and variations in the line voltage by an additional $< \pm0.5$ppm (for the typical expected voltage variation of 5%) [126]. Overall, the uncertainty in the 10MHz frequency reference is approximately $< \pm7$ppm which suggests a worst-case uncertainty in the derived 2.7GHz EOM drive signal of nearly 20kHz.

The majority of this drift is slow 'aging' drift, so, once the narrow transition frequency is found, the correction required to the frequency is not expected to change. However, the much faster temperature and line-voltage contributions could account for 4kHz of uncertainty, and this is unacceptable. The Agilent has the facility to adjust this reference frequency, but we first need a method to accurately measure this imperfection.

6.3.1.1   Frequency Referencing

One approach to monitor the stability of the Agilent is to buy a more stable frequency reference, such as a rubidium reference[1], but a more interesting approach and economical is to use the accurate pulse-per-second signal provided by the Global Positioning System and to count the number of oscillations from the Agilent in a given time; this method is not only cheaper, but provides better long-term stability and is directly traceable back to a primary standard.

For counting, we use several of the ubiquitous 8254 chips, each of which contains 3 independent 16-bit counters. We use a device (Acces I/O Products USB-CTR-15) which incorporates 5 such chips and facilitates programming of these chips over USB. Programming includes loading an initial value to the counter and setting the mode, as described below. Three pins (Clock, Gate, and Output) of each counter are available for external connections. The chip can be operated in various modes, and the two we employ are described below.

- Mode 2. The counter is loaded with a specific value which it decrements on each clock pulse, whilst the gate is high. When zero is reached, the output is high, and on the next clock pulse the counter resets.

- Mode 3. The counter is loaded with a specific value and again it decrements. However, in this mode, the output is high during the first half of the count and low for the second half. Internally, this is implemented by the counter holding output high and decrementing from the load value by two for each clock count, and then repeating with the output held low.

An arrangement of counters connected via their external pins is shown in Figure 6.4. If we wish to count the 10MHz signal for N seconds, we load $2N$ to the first counter which is in mode 3, and so holds the output high for N seconds. This output is connected to the gate of a second counter which will therefore count the 10MHz signal connected to its clock input for N seconds; a third counter is concatenated in the standard manner.

---

[1]e.g. Stanford Research Systems PRS10 which provides an absolute accuracy of 0.1 Hz at 9.2 GHz [81]
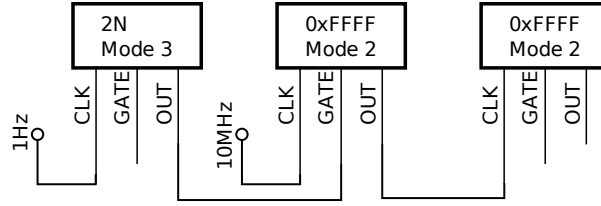
FIGURE 6.4: Connection of 3 counters (not necessarily within the same 8254 chip), numbered left to right, used to count a 10MHz signal for N seconds; the 1Hz signal is derived from a GPS receiver. Counters 2 and 3 are initially loaded with the maximum possible for a 16-bit counter of $2^{16} - 1 = $ 0xFFFF $= 65\,535$. The modes are described in the text.

Together, counters 2 and 3 have sufficient range to count for 0xFFFFFFFF/10MHz $\simeq$ 429s. After N seconds have elapsed, we have N seconds during which counting is paused when we may interrogate the chip via USB. During this time we read the value of the counters, which we later compare to Ns $\times$ 10MHz, and reset them to 0xFFFF. Additionally, and not shown in the schematic, we connect the output of counter 1 to an inverter and duplicate the arrangement of counters 2 and 3 so that, during the N second dead-time of one chip, we can continue to count the 10MHz signal.

For development and testing, we counted the internal 10MHz clock included in the device and used a 1Hz derived from this same reference; as hoped, we counted exactly 10 Million/second, with no dead-time[2] during the switch-over from one counter pair to another.

The GPS receiver is the Garmin OEM model 18x LVC. After the device has obtained a fix of its position a pulse, whose rising edge is quoted to be within 1μs of the GPS second, is produced every second. It is difficult to maintain this absolute time accuracy given the various delays, including propagation delays in cables, which can accrue, but any such delay is constant and the pulse to pulse jitter should not exceed the stated uncertainty in the timing-signal.

If we require a frequency shift on the 2.7GHz signal of at most a few 100Hz, we require the reference to be accurate to 10^{-7}. If we count at rate f for time τ, with timing-jitter $\delta\tau$, the relative uncertainty is $\alpha = (f\delta\tau)/(f\tau)$ which implies we must count for $\tau = \delta\tau/\alpha$, or approximately 10s for the GPS signal. Additionally, by averaging over

---

[2]The dead-time is much less than one clock cycle, making it unlikely the device will miss a count.
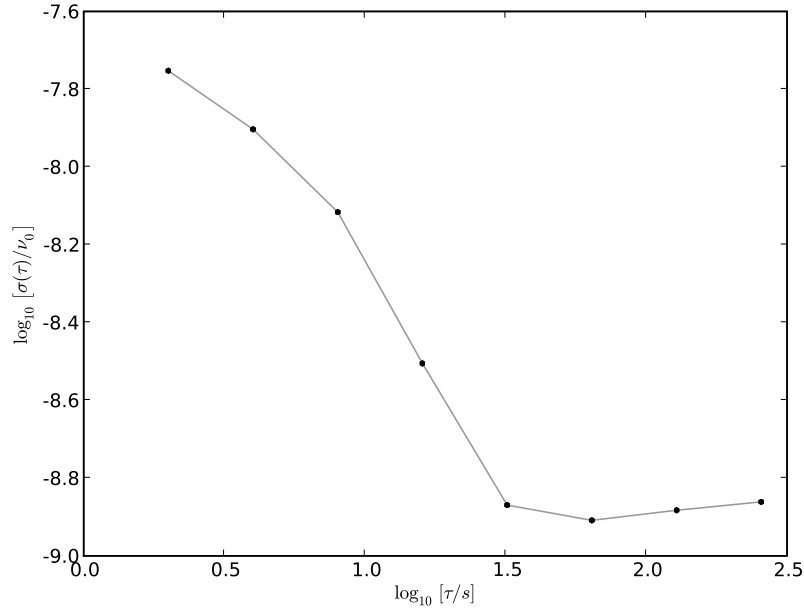
FIGURE 6.5: Allan variance plot showing the stability of the Agilent 10MHz reference compared with the GPS time-signal for time-scales from 2s to 256s.

successive counts, the uncertainty can be reduced by $\sqrt{N}$ where N is the number of independent counts, if we assume the signal does not drift during this time.

From the Allan variance plot in Figure 6.5, we see the oscillator is increasingly stable over long-time scales, up to a limit of 3ppb at ~ 200s, at which the absolute frequency (which can be adjusted) is offset from 10MHz by 439ppb [3]. The larger instability at short (< 32s) time-scales is most likely due to the instabilities in the pps output of the GPS receiver; it is likely the Agilent maintains the impressive 3ppb stability to much shorter time-scales than suggested by the plot.

An alternative time-reference which could be used is a time obtained via Network Time Protocol over the Internet, or a Local Area Network if a NTP server is available more locally. An NTP server will reference its clock against a reliable source, such as the GPS signal, and so may be considered long-term stable, but despite the attempts of the protocol to reduce timing uncertainty due to the latency of the network connection, the jitter is still several milliseconds. Taking the typical case of 10ms, we must count for 3 hours, assuming we can extract an accurate time-signal from the computer. This

---

[3]ppb here means parts-per-billion, or $\times 10^{-9}$.

would require concatenating a third counter, but appears practical if no GPS signal is available.

A final approach might be to use the MSF Time signal from NPL, broadcast from Cumbria with a carrier of 60kHz within 10^{-12} [127]. One could compare a count of this accurate frequency reference with the local 10MHz oscillator.

6.3.2 Phase Modulation

With only a fixed frequency source, regardless of how stable, we are severely limited in the types of manipulation we can perform. Should we wish to implement Ramsey interferometry we can do so only by purposely detuning the frequency so that, between the two $\pi/2$ pulses, the Agilent and atom accrue some phase difference. However, it is not possible, even in principle, to implement any kind of adiabatic manipulation, such as described in Chapter 4. In general we would like to modulate the frequency in a phase coherent way *during* the pulse.

The Frequency Modulation and Phase Modulation inputs on the Agilent may seem like the obvious choice, but neither can provide the *single* sideband modulation we require. To achieve this we use an I&Q modulator, whose two modulation inputs are driven by the two independent channels of a 300M samples per second arbitrary signal generator. These devices are described below.

6.3.2.1 I&Q Modulator

To modulate the 2.7GHz carrier, we use an I&Q modulator, standing for 'in-phase' and 'quadrature', also known as a Single Sideband Modulator. This device takes the carrier frequency generated by the Agilent, splits it, and phase shifts one path by 90°. These paths are then recombined with amplitudes determined by the two modulation inputs. By choosing that these I and Q inputs are sine and cosine functions of a some common phase ϕ, simple trigonometric identities show the output is a single frequency component phase shifted by this amount. For an input $\sin(\omega t)$, split to $\sin(\omega t)$ and
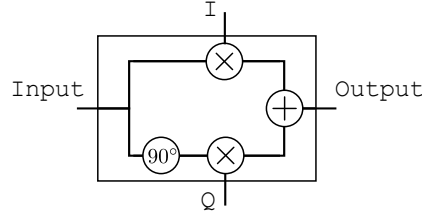
FIGURE 6.6: Block diagram showing the components of an I&Q or Single-sideband modulator. By choosing $I = \cos\phi(t)$ and $Q = \sin\phi(t)$, we produce the single frequency output $\sin(\omega t + \phi(t))$.

$\sin(\omega t + \frac{\pi}{2}) = \cos(\omega t)$, and I and Q inputs $\cos\phi$ and $\sin\phi$, the output is $\sin(\omega t) \times \cos\phi + \cos(\omega t) \times \sin\phi = \sin(\omega t + \phi)$.

A block diagram of the device is shown in Figure 6.6. We have used a Miteq I&Q Modulator SDM0104LC1CDQ, which requires no external power supply and derives the power necessary to drive the active components from the carrier which consequently must carry between 10dBm and 13dBm. The output power after modulation is considerably less than this relatively high input (approximately -10dB less), but this is still sufficiently strong that it can be used by the amplifier to drive the EOM. We must, however, take care to ensure the output power is in the correct range so as to drive the EOM yet not damage the fixed gain power amplifier.

6.3.2.2 Arbitrary Signal Generator

The power requirements of the device to drive the modulation inputs of the I&Q, at a maximum 5dBm or 560mV amplitude (in 50Ω), are not so demanding, but we must be able to control the amplitude of these inputs independently. For this, we use a Tabor 3362 arbitrary signal generator, which can synthesise a pre-loaded arbitrary waveform at 300M samples per second with 12bit resolution and up to 16M points. This device, like many others in the experiment, uses the Agilent (see Section 6.3.1) as a frequency reference, and can trigger or be triggered by other signal generators so that any modulation (to within 100ns) occurs at the correct time in the overall sequence.

We choose the outputs of the signal generator to be an amplitude multiplied by the cosine and sine respectively of a common phase. For the simple example of a Ramsey

interferometry, which comprises two separated $\pi/2$ pulses, this phase is constant during each pulse, but different from one to the next, and we can use the dead-time between the pulses to make the change. For a chirped pulse, such as those described in Chapter 4, the phase is the integral of the detuning, and varies continuously during the pulse.

6.3.3 Power Amplifier

Although the amplification required to boost the Agilent, or Agilent and I&Q modulator, output to the required level is not itself particularly large, the absolute power is. Fortunately, the relatively high AOM frequency shift of 310MHz, when subtracted from the hyperfine ground state splitting of Rubidium-85, brings the RF frequency just within the range covered by standard microwave communications. It is not uncommon for applications in this field to require high-power signals, and so suitable amplifiers in this range are significantly cheaper than those operating even 100MHz higher.

We use a Miteq Power Amplifier AMF-6B-027029-40-37P, which is designed to provide a gain of 40dB, a maximum power of +37dBm $\simeq$ 5W, and to handle without damage an input power of 0dBm. The bandwidth of this device is easily sufficient to faithfully reproduce the carrier plus any modulation the 300MS/s signal generator can provide.

6.3.4 RF Power Meter

To measure the RF power at the various stages of the chain we use a Mini Circuits ZX47-40-S+, which outputs a DC voltage in response to an RF signal and is characterised up to 8GHz and 27dBm. We assume the Agilent signal generator produces a known power, and we use this to calibrate the power meter. Additionally, to extend our measurement range to include the high powers created by the power amplifier, we use a Miteq CD-202-402-30S -30dB directional coupler. This device uses RF transmission lines to extract a small fraction of the power on a line without disturbing the remainder; we can measure this extracted fraction and hence obtain a measure of the power sent to the EOM while an experiment is running.
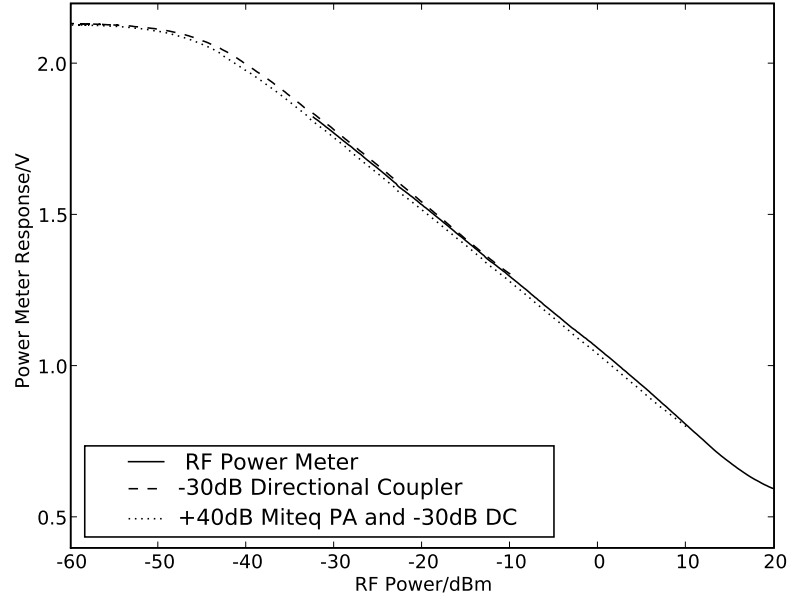
FIGURE 6.7: Calibration of the Mini Circuits ZX47-40-S+ RF Power Meter, using the Agilent 8648C as a reference. The plots are shifted according to the quoted 30dB attenuation and 40dB amplification; the curves are expected to overlap.

The output DC voltage from the RF power meter is shown in Figure 6.7, using the RF power reported by the Agilent as a reference. All measurements are at the typical EOM frequency of 2726MHz. These measurements show the power amplifier continues to provide +40dB gain, even up to the maximum safe input power of 0dBm, and hence provides +40dBm = 10W, which is twice the quoted maximum power. This is useful for some characterisation measurements (see Section 6.3.5), but there is the possibility of damage if the device is operated continuously at these powers.

6.3.5 Electro-Optical Modulator

The radio frequency electrical signal in the coaxial transmission line interacts with the light via an Electro-Optical Modulator, a device which incorporates a crystal whose refractive index varies in response to an electric field. Our chosen EOM is a New Focus 4431M, which is a $MgO:LiNbO_3$ (magnesium oxide doped lithium niobate) crystal in a resonant microwave cavity. The resonant cavity imposes a finite bandwidth, but ensures the amplitude of voltage oscillations across the crystal is far larger than could otherwise

be achieved. Our EOM has a quality factor $Q = 273$, which gives a bandwidth 0.4% of the carrier, and an RF power in the cavity $P_c = P_i Q$, where $P_i = 4$W is the input power. The expression for the phase modulation depth in this situation uses several parameters, notably the crystal dimensions, which are not known. However, we can use measurements of the modulation depth for other parameters to estimate the response of the device in our experiment.

With no cavity, voltage is related to RF power by $V_{\text{RMS}} = \sqrt{RP}$ where R is the transmission line impedance. In a resonant cavity, the circulating power is increased by a factor Q, and the voltage, using now an effective impedance, by a factor $\sqrt{Q}$. Overall, and using the voltage amplitude rather than RMS, $V_{\text{AMP}} = \sqrt{2R_{\text{eff}}QP_i}$. The datasheet for our device provides a plot of modulation depth $m = \Delta\phi/V_{\text{AMP}}$ for a variety of $\Delta\phi$, for a typical cavity $Q = 200$; from this we can estimate the effective impedance, $R_{\text{eff}} \simeq 0.25\Omega$.

We are now able to calculate the voltage across the crystal. The final ingredient is the characteristic voltage V_π at which the crystal induces a phase shift of π; this depends linearly on wavelength and is $V_\pi(1.06\mu\text{m}) = 63\text{V} \implies V_\pi(780\text{nm}) = 46\text{V}$. Finally, the phase modulation amplitude for an applied power is

$$\phi_{\text{AMP}} = \frac{\pi}{V_\pi(\lambda)}\sqrt{2R_{\text{eff}}QP_i} \tag{6.1}$$

which, for our cavity $Q = 273$ and $P_i = 5$W is $\phi_{\text{AMP}} = 1.77$rad.

Now we know the depth of this sinusoidal phase modulation, we can calculate the spectrum of the output optical beam. If we represent the incident optical beam by a complex exponential at frequency ω, the output beam is $e^{i\omega t + i\phi_{\text{AMP}}\cos(\Omega t)}$ where Ω is the frequency of the modulation. Using the Jacobi-Anger identity, we can expand this in terms of Bessel functions:

$$e^{i\omega t}e^{i\phi_{\text{AMP}}\cos(\Omega t)} = e^{i\omega t}\sum_{n=-\infty}^{+\infty}i^n J_n\left(\phi_{\text{AMP}}\right)e^{in\Omega t} \tag{6.2}$$

where J_n is the n^{th} order Bessel function. Hence the spectrum is composed of an infinite array of 'sidebands' with frequency an integer times the modulation frequency away from the carrier and amplitudes given by the Bessel functions. The power in a sideband
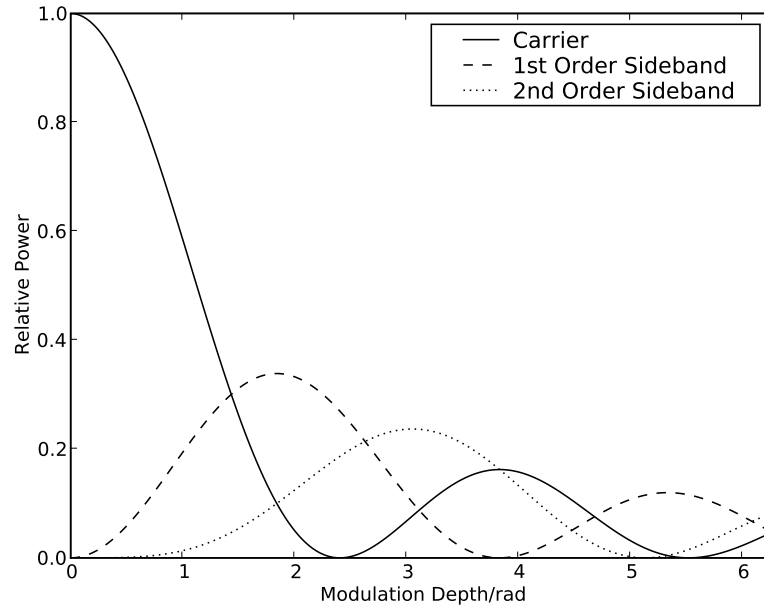
FIGURE 6.8: The power in the first few frequency components of the phase modulated beam as a function of the modulation depth.

is proportional to the square of this amplitude and the first few are plotted as a function of modulation depth in Figure 6.8.

Power in the first order sidebands is maximised for $\phi_{\mathrm{AMP}} = 1.84\mathrm{rad}$, which for our device is an RF power $P_i = 5.4\mathrm{W}$; this is only slightly above the maximum power of the Miteq Power Amplifier. A typical spectrum, measured with an Optical Spectrum Analyser, is shown in Figure 6.9.

The non-linearity exploited by the crystal is polarisation dependent, and the polarisation of the light must be linear and aligned with the crystal axis to achieve this phase modulation. Additionally, the wavefronts must be flat as they pass through the crystal or different parts of the wave will experience different phase delays; the small active aperture of the device makes it possible but not trivial to meet this requirement.

We can model the imperfect phase modulation of the EOM using a parameter β which is the fraction of the optical power which is modulated, with the corresponding fraction $1 - \beta$ passing through the EOM unaffected. The power in the n^{th} order sideband is hence $\beta \left| J_n \left(\phi_{\mathrm{AMP}} \right) \right|^2$, and the carrier, which is composed partly of the unmodulated fraction, is $(1 - \beta) + \beta \left| J_0 \left(\phi_{\mathrm{AMP}} \right) \right|^2$. Using this model, we measure the relative power
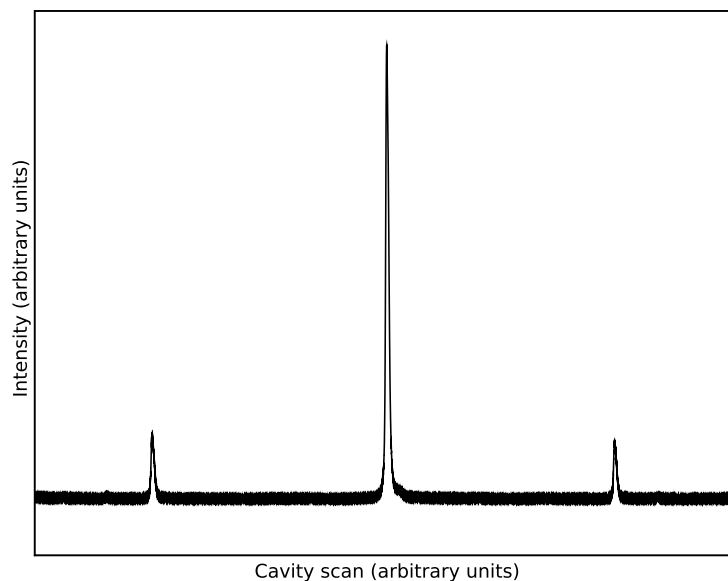
FIGURE 6.9: A typical spectrum after phase modulation by the EOM, as measured by the OSA. The central carrier and first order sidebands are clearly visible. One of the nearly discernible second order sidebands is shown more clearly in Figure 6.10. The modulation frequency is larger than the range of the OSA, so the relative frequency shift of the sideband is not faithfully mapped to position along the x axis.
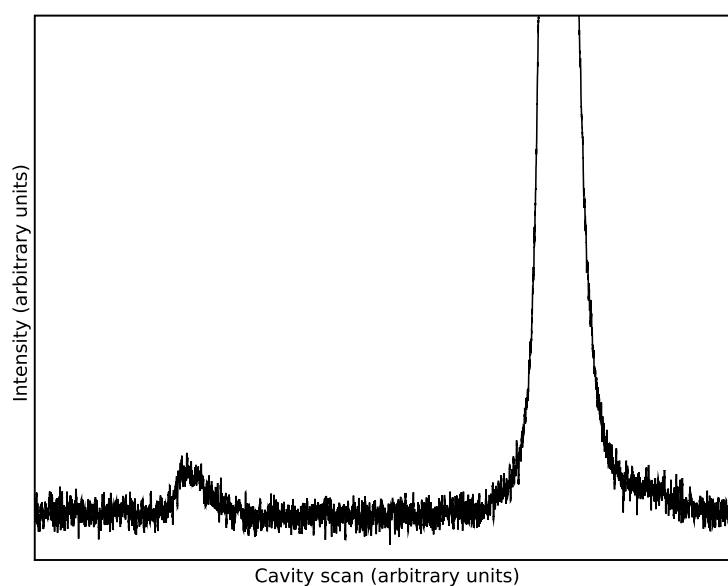
FIGURE 6.10: An expanded region of Figure 6.9 showing a second order sideband next to a much larger first order sideband.
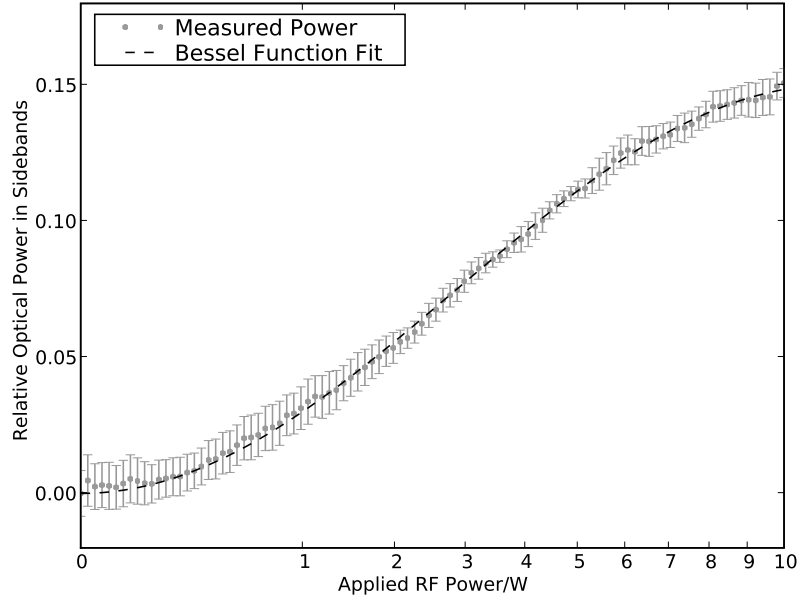
FIGURE 6.11: The fraction of power remaining in the carrier and placed into the first order sidebands, as a function of applied RF power, with Bessel function fits. Only $\simeq 44\%$ of the incident optical power is subject to the phase modulation.

in the carrier and the sidebands, using an Optical Spectrum Analyser, and fit to the first order sidebands using β as a free-parameter. Additionally, although we know the approximate modulation depth and are confident of its dependence on power, we allow this scale to stretch or contract to optimise the fit. This is shown in Figure 6.11. We find that only $\simeq 44\%$ of the incident optical power is affected by the phase modulation. Larger sidebands have been observed with device, and we suspect there is damage to the crystal.

There are various factors which could account for less than all the optical power being subjected to the phase modulation. The obvious candidate is imperfect alignment of the polarisation with the axis of the crystal, but care has been taken to ensure the input polarisation is linear (using a 30dB polarising beam-splitter cube), and the sideband power was maximised by rotating this polarisation using a half waveplate. If the wave-fronts of the beam are not flat, different parts of the beam will experience different phase delays; overall, this will appear as though only a portion of the beam is subjected to the phase modulation. Care has been taken to ensure the beam is collimated, but spatially dependent damage to the crystal could produce a similar effect. We are confident of the
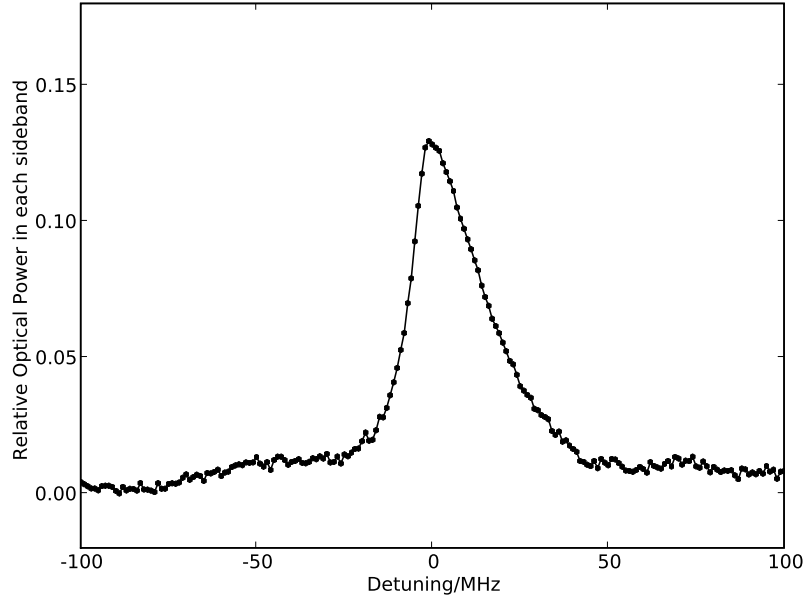
FIGURE 6.12: The EOM RF cavity resonance found by measuring the optical power in each of the 1st order sidebands at -3dBm as a function of detuning from the resonance at 2726MHz. FWHM $= (+15\text{MHz} - -7\text{MHz}) = 22\text{MHz}$ and $Q = 124$.

RF drive, and so fear the EOM is damaged; without specialist equipment, we shall have to return it to the manufacturer for diagnosis.

We may not be able to determine why not all optical power is affected by the modulation, but the Bessel function fit also tells us the dependence of modulation depth on RF power: $\phi_{\text{AMP}} = 0.54\sqrt{P_i}$. From this, using Equation 6.1, we can estimate the product of cavity Q and the effective impedance R_{eff} by which we modelled the cavity. Additionally, we can measure the Q of the EOM microwave cavity by measuring the width of the resonance peak as the drive frequency is changed. The plot, shown in Figure 6.12, is not symmetrical, but has a characteristic FWHM $= 22\text{MHz}$, which implies $Q \simeq 124$ and confirms $R_{\text{eff}} \simeq 0.25\Omega$. The EOM is shown in Figure 6.13.

6.3.6 Characterisation

It is important to ensure the frequency difference between the Raman beams remains coherent for the duration of an pulse sequence, even if we are simply trying to drive Rabi oscillations. We can estimate the coherence time based on the specifications of
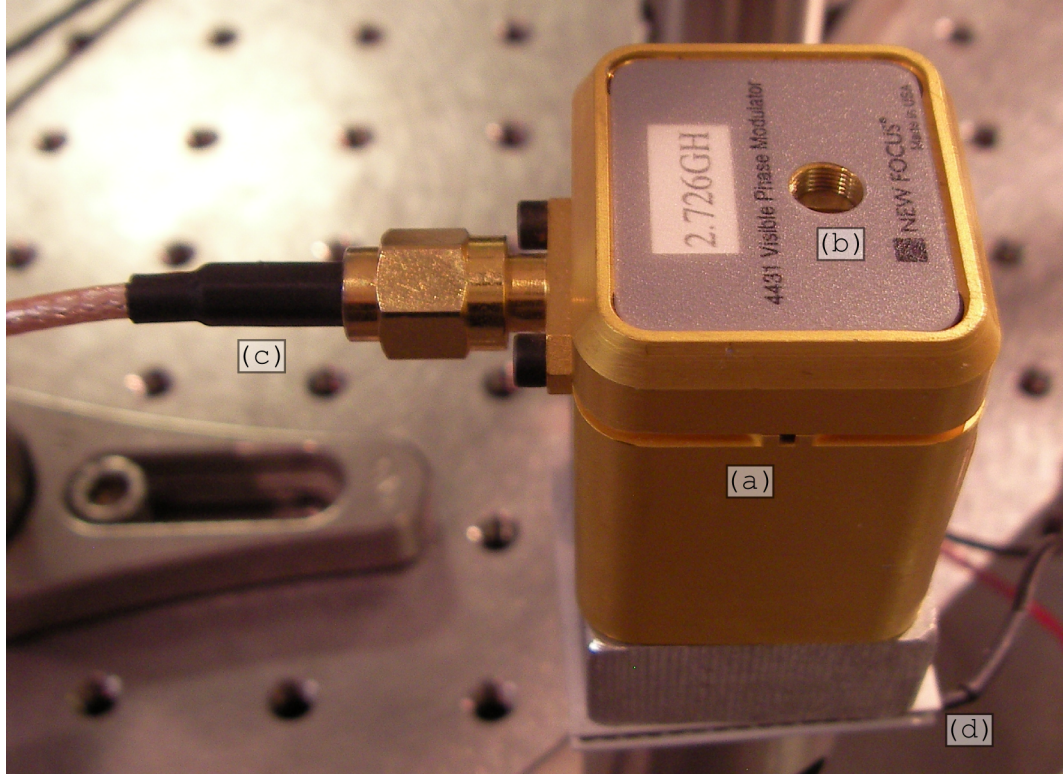
FIGURE 6.13: The New Focus 4431 Electro-Optical Modulator which provides the majority of the 3GHz shift required to drive our Raman transitions. Light passes through the crystal via the small 1mm aperture (a). The crystal is in a resonant microwave cavity, the centre frequency of which can be tuned using the adjustment screw (b) and which is driven via a semi-rigid SMA cable (c). Also visible is the Peltier-effect cooling (d) which is (optionally) used to cool the device.

the devices used to create the signal (phase noise of the Agilent oscillator and noise figure of the Miteq amplifier), but using the beat frequency technique of Section 5.1.6, we can directly measure it. Ideally, we could measure the 3GHz beat frequency between EOM sideband and AOM shifted beam, but as we do not have a photodetector which can respond to this frequency, we chose instead to compare the sidebands with a stable reference: namely, the trapping laser locked to a rubidium transition.

With the beams aligned as in Figure 5.10, we must first tune the Tiger frequency within < 100MHz of the rubidium feature to which we intend to lock the trapping laser. We scan the trapping laser across the spectral features by modulating its piezo voltage, and we observe the signal from the photodiode used to measure the beat frequency. When the laser frequencies are within the bandwidth of this detector, we observe a modulation; this is shown in Figure 6.14, and is a quick method to tune the Tiger frequency to resonance.
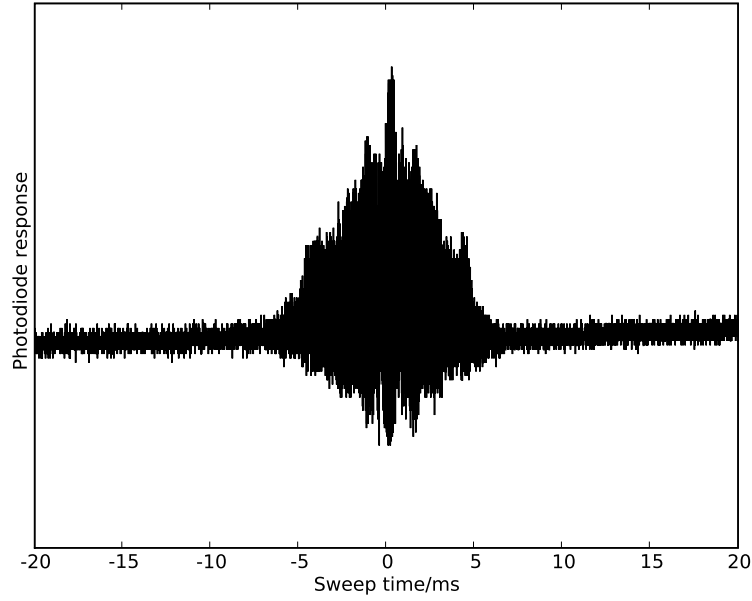
FIGURE 6.14: The AC coupled response as the piezo scan of the trapping laser sweeps the beat frequency with the Tiger across the bandwidth of the photodetector.

The beat frequency between the unmodulated Tiger and the trapping laser, and the beat frequency between an EOM sideband derived from the Tiger and trapping laser are shown in Figure 6.15. For short-times, where the main contribution is $1/f$ noise, the Allan variance of the sideband is larger than that of the carrier; both slopes are consistent with $1/f$ noise, and we conclude the increase is due to the additional noise introduced by the RF modulation. By treating the noise from the RF modulation and the noise from the laser as uncorrelated sources which add in quadrature, we can infer the contribution from the modulation. Unlike the Tiger, this modulation is long-term stable, and extrapolating the $1/f$ noise to the time-scales of our coherent experiments ($10\mu s < \tau \leq 10ms$), we conclude that the linewidth of the frequency difference over the relevant time-scales should not exceed $3\mathrm{GHz} \times 10^{-9} < 10\mathrm{Hz}$.

6.4 Acousto-Optical Modulators

Acousto-optical modulators are quite different from electro-optical modulators, and can be used to perform several modifications to the incident beam depending on the alignment and type of AOM.
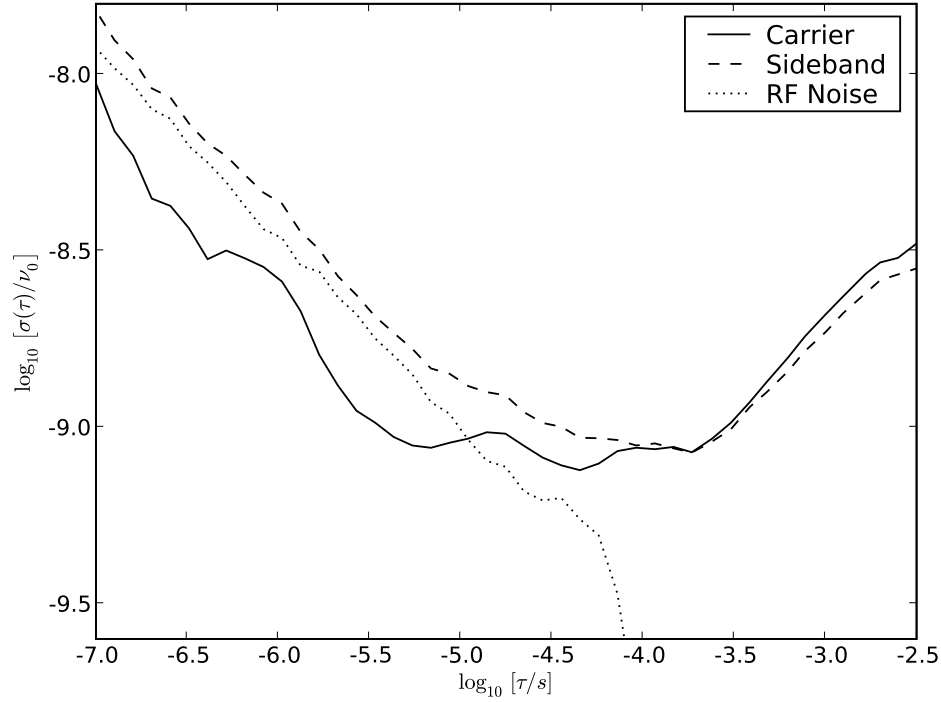
FIGURE 6.15: Allan variance plot showing the frequency stability of the Tiger carrier and an EOM sideband and the inferred noise from the RF modulation. The signal was measured via a beat frequency with a laser locked to a rubidium transition, as described in Section 5.1.6, and the drift of the unlocked Tiger dominates for long-time.

The AOM consists of a mechanical oscillator (typically a piezo-electric) attached to a compressible medium (such as a glass or crystal) through which an acoustic wave propagates; in our AOMs, this wave is absorbed once it has traversed the medium. The compressions and rarefactions associated with the acoustic wave cause a sinusoidal modulation of the refractive index; it is with these features that the incident optical wave will interact. All of our AOMs are so-called fixed frequency AOMs and, apart from comparatively slow amplitude modulation, we need only consider a single frequency.

The feature in the AOM is more than a moving mirror; it is a periodic modulation, and hence we should expect Bragg reflection. The light reflected from the moving mirror will also be Doppler shifted, but only the component of photon momentum normal to the surface of this mirror will be affected, and hence the angle of light reflected by this moving mirror will be different from the angle of incidence [33]. For the special case where the Bragg diffraction angle and this Doppler-shifted angle are equal, we enter the

so-called Bragg regime, and only one diffracted order is present.

The Bragg regime is experimentally relatively straightforward to achieve; the collimated optical beam is passed through the AOM whose angle is carefully adjusted (to within $\sim 0.1°$) while the diffraction pattern is observed in the far field (typically $> 10\text{cm}$). Collimation of the beam is crucially important because if the beam is diverging, the Poynting vector (i.e. the propagation direction) will vary across the beam and only part of the beam will be in the Bragg regime at one time. This is clearly visible in the diffraction pattern where one can observe part of the beam having been diffracted whilst the remainder is unaffected and appears in the undiffracted zeroth order.

We employ two very similar AOMs for creation and control of the Raman beams. The first operates at 310MHz and diffracts approximately 25%, the undiffracted 75% being sent to the EOM, while the second operates at 110MHz, although this frequency is not crucial, and is used for amplitude modulation.

6.5 Mach-Zender interferometer

If the BoosTA is used to amplify the light, the maximum power we can inject is finite and we would rather the power in the useful frequency components was not limited because some power was reserved for the relatively large carrier component. Additionally, the lightshift from the carrier, as described in Section 8.3.2, affects the Raman transition frequency, and removing an unnecessary frequency component simplifies this lightshift landscape.

We use an interferometer to separate the frequency components of the EOM beam, and hence to remove the high power carrier and leave only the sidebands; this was inspired by Dotsenko *et al* [81] and the device was meticulously assembled by André Xuereb. This use of a Mach-Zender interferometer is described in Haubrich *et al* [128], along with other techniques which can be used to combine beams of near-equal frequency.

The Mach-Zender interferometer is shown in Figure 6.16. Two 1cm non-polarising beam splitter cubes are glued to a metal base, and a right-angle prism reflects the path split
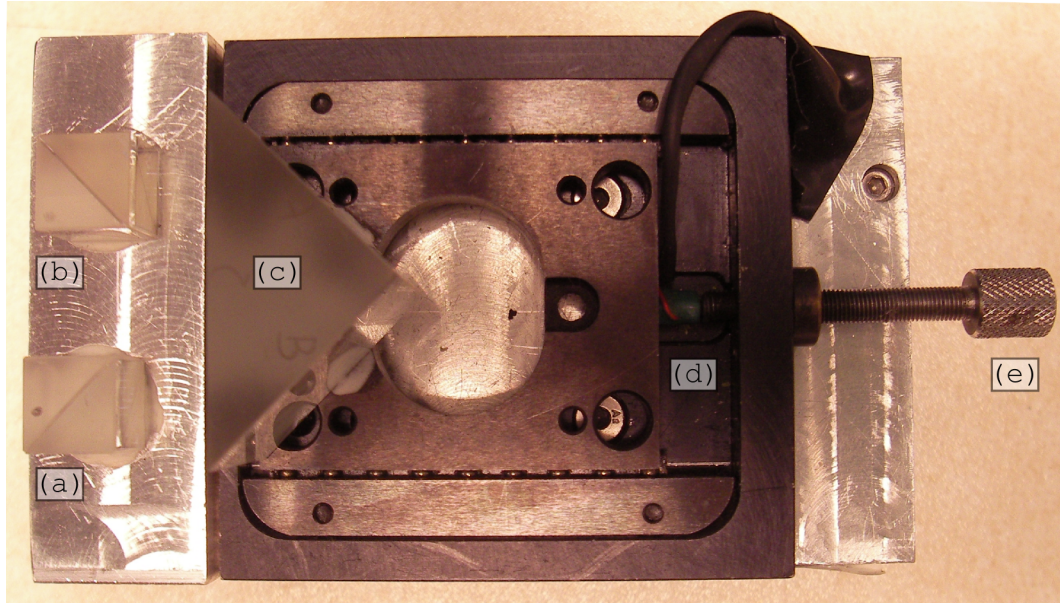
FIGURE 6.16: The Mach-Zender interferometer, showing the layout of optical components. Light enters a 1cm non-polarising beam-splitter cube (a), and the two paths are recombined in a second cube (b), with one having been retro-reflected by a right-angled prism (c) mounted on a translation stage. The optical delay is controlled electronically using a piezo-electric stack (d) and a micrometer screw (e).

by the first cube back to the second where they are recombined. Alignment of the component parts of this device is crucial, and without investing in multi-axis translation and rotation mounts for each cube, the best approach is to use very low expansion UV curing glue, align by hand, and then fix using a UV lamp. The prism is also fixed using this glue, and can be translated in one direction using a screw for coarse alignment and a piezo-electric stack for fine adjustments and feedback from a locking circuit.

The interferometer splits an input signal into two parts and recombines these after applying a time-delay to one path. We require that one path results in destructive interference for the carrier frequency, and that the other path results in constructive interference for the sidebands. Considering only the carrier and first order sidebands of the phase modulated beam, we require that $\omega\tau = 2\pi n$ where τ is the time-delay, ω is the carrier frequency, and n is an integer, and that $(\omega \pm \Omega)\tau = 2\pi n \pm \pi$ where Ω, as in Section 6.3.5, is the modulation frequency. Hence the path difference is $x = c\tau = c\pi/\Omega = 55$mm; this is an optical path length, and the refractive index of the glass shortens the actual distance. Structures on this length scale are easily assembled by hand, and over a centimetre scale range of path difference the carrier and sidebands exit

from different ports of the interferometer. However, *which* of the ports each leaves from depends on wavelength scale changes in the path difference and while the mechanical construction of the interferometer means it is passively stable, temperature changes and other causes of drift in the metal mean it drifts over several minutes, and it must be actively stabilised.

As in Dotsenko *et al* [81], we monitor the power from one of the ports and lock to a maximum. We reason that, especially with our imperfect modulator (see Section 6.3.5), the majority of the power will remain in the carrier, and hence by ensuring the maximum power exits one port, we are also ensuring the sidebands exit via the other. As discussed in Section 5.1.5, we cannot lock directly to a maximum; instead, we dither the piezo voltage at the mechanical resonance of 2kHz and mix the photodiode signal with this same frequency to obtain the derivative about the maximum. From this, an error signal is derived and a PI circuit feeds back to the piezo offset voltage. Elsewhere, optical single sideband sideband modulation has been shown using an interferometer similar to our setup [129], and with polarising optics [130].

Unfortunately we have had difficulty with the photodiode amplifier used to measure the dithered 2kHz signal, and the device cannot reliably be locked. Unlike a failure of the MOT laser locking, or a slow drift of the unlocked Tiger wavelength, a slow change in intensity due to the unlocked Mach-Zender is not monitored in our experiment. We are working to address this—both by improving the reliability of the locking circuit and by accurately and continuously monitoring the optical powers—but to ensure they were reproducible, some of the experiments described in Chapter 8 we performed without this component.

6.6   Direction Switching

If counter-propagating beams are used to drive the Raman transition, the photon in the effective two-level picture described in Section 3.1 has the difference in energy between the beams, but now has the (scalar) sum of their momenta. This configuration allows us to greatly influence the momentum of a cloud of atoms, but we must be able to choose
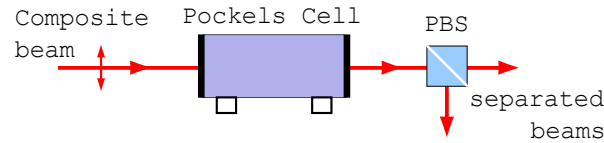
FIGURE 6.17: A Pockels cell can be used as an electrically switchable waveplate to rotate the polarisation by 90° and so control whether a beam is reflected or transmitted by a polarising beam-splitter cube. If the input beam contains two components with orthogonal polarisations, these will leave via opposite ports of the PBS.

the direction of our effective photon: we must be able to controllably switch which beam approaches from which direction.

With the two Raman beams combined using a polarising beam-splitter cube, the fields overlap spatially but have opposite polarisation. We can apply some common operation to the two overlapped beams, separate them using a second beam-splitter cube, and then send them counter-propagating into the vacuum chamber. If this common operation is rotation of the polarisation using a half-wave plate we can influence which of the Raman beams is reflected and which is transmitted by the second beam-splitter cube. We use a Pockels cell, which in this arrangement acts as an electrically switchable wave plate, and by changing the voltage on this device, we can change the direction of the effective photon during the short dead time between pulses in a sequence. The layout is shown in Figure 6.17.

A typical application of Pockels cells is to emulate a mirror of switchable reflectivity, for use in Q-switched lasers, and it is for this that most of the commercially available electronics is designed. This application requires the high voltage needed to activate the cell be maintained for a short time ($\sim 10\mu$s). We require that the voltage be held high or low, ideally indefinitely. The Pockels cell behaves as a capacitor, and so there is not current flow even when the voltage is high, and a simple circuit using high voltage FETs to control the required 300V is shown in Figure 6.18. The board itself is shown in Figure 6.19. The two MOSFETs are used as switchable resistors, and when one is high, the other is low. Current flow is only that which is needed to charge the 60pF capacitance of the Pockels cell, and the HV source is never connected directly to ground.

An unfortunate consequence of using polarisation to distinguish the beams is that we can no-longer use the BoosTA as this device only supports one polarisation. Although
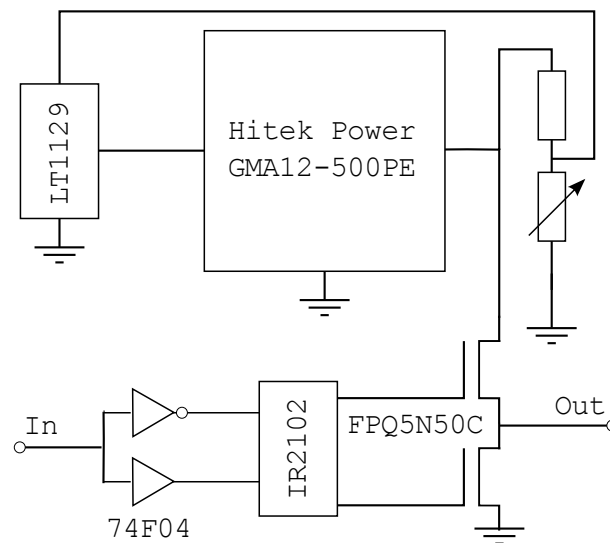
FIGURE 6.18: A simplified schematic of the electronics used to drive the Pockels Cell. The input is buffered by the 74F04 Hex inverter, which then drives the IR2102 MOSFET driver. The high voltage is provided by the Hitek Power DC to DC converter, which is designed to provide 500V, but which can be regulated by controlling the input voltage using the LT1129 adjustable regulator with feedback from the potential divider. The MOSFETs are FPQ5N50C which can withstand 500V and have on-resistance 1.4Ω.
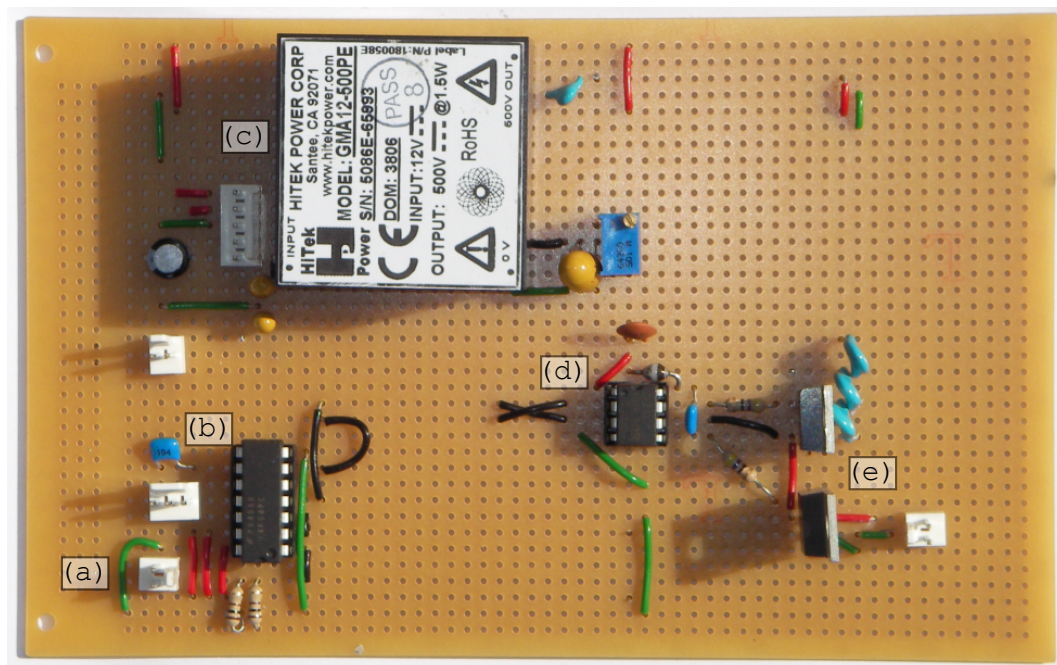
FIGURE 6.19: A photograph of the electronics used to drive the Pockels Cell. Molex connectors (white plastic with between 2 and 5 pins) are common, and are used for external connections including BNC jacks, power supplies, and active components, such as voltage regulators, which may require heat-sinking. The photograph shows the input signal connection (a), which is buffered and inverted by the hex inverter (b), which then controls the high voltage produced by the DC to DC converter (c), via the MOSFET driver (d) and the MOSFET pair (e).

it appears possible, we have yet to attain sufficient power to drive Raman transition acceptably fast without the use of this optical amplifier.

6.7 Computer Control

The majority of the instruments for the current suite of experiments are, by some means, computer controlled. A significant exception is the locking circuitry which it makes little sense to automatically adjust and which typically require manual intervention when the device they control falls out of lock. For all the other instruments, an interface – such as GPIB, USB, or Ethernet – allows us to connect the device to a computer and hence to automate control as though a person were using the instrument's front-panel, to upload commands or sequences which it would be impossible or impractical to enter via a front-panel (such as $\sim$ 1M point sequences of arbitrary waveform data), and to quickly retrieve measurements.

The task of instrument communication is greatly simplified by Virtual Instrument Software Architecture (VISA) which is an industry API[4] which shields a program from the details of the underlying physical connection and allows the developer to concentrate on what needs to be said, rather than how one should speak.

Standard Commands for Programmable Instruments (SCPI) is a standard which defines the structure of what may be said over the communications link managed by VISA. This is the high-level language with which we control the majority of our instruments.

The programming task involves primarily sending pre-defined sequences of commands, or commands generated according to some pre-defined rules, to a variety of instruments. Timing of commands is not crucial since we rely on the accurate clocks of the instruments for any time-critical experiments. It is also necessary to generate sequences extending to millions of points for uploading to arbitrary signal generators; on modern hardware this can be a quick operation if the correct tools are used. Finally, the programs should be readable and be quick to develop and modify.

---

[4]Application Programmers' Interface

LabView by National Instruments is extremely widely used to control instruments, and while it is unmatched for quick development of graphical user interfaces, it favours development of virtual front-panels to instruments, not automation, and is particularly cumbersome for sending scripted sequences of commands.

Our chosen language is Python. This high-level language has clear syntax and minimal core set of functions which can be extended with imported libraries; this allows us to tailor the language to our task. By importing numerical computation libraries, large arbitrary waveforms can be generated quickly, and VISA libraries can then be imported to upload this waveform to an instrument. We can build on the abstraction provided by VISA and construct functions specific to our instruments. The Tabor waveform generators require the programmer to specify the amplitude and offset of an arbitrary waveform, which is then encoded as 16-bit integers. With a specifically written function, we can pass a list of real-value voltages which are interpreted by the computer and uploaded to the instrument in the form it requires. This shields us from the details of how the waveform we wish to apply is represented by the instruments.

We can also communicate with the digital oscilloscope (Tektronix DPO4000) used to record the photodetector signal. We use HTTP over 100Mb/s Ethernet to download 10M 16-bit points in under 30s. This impressive speed was an enormous help in our search for the elusive Raman signal.

Chapter 7

MOT Characterisation

Our description of the experiments begins with those designed to characterise and understand the cold atomic sample in the magneto-optical trap. It is essential that we understand the environment and properties of our cloud before we attempt any coherent manipulation.

Three important properties of the cloud are the number of atoms and the spatial and momentum distributions of these atoms. The conceptually simplest method to measure the number of atoms is to measure the signal their fluorescence induces in a photodiode and, by calculating the rate of fluorescence for each atom and the solid angle subtended by the detector, calculate the number of atoms. Unfortunately this method is plagued by experimental uncertainty, ranging from calculating the fluorescence rate given that the cloud is optically thick, to the angle-dependent loss of light passing through the thick glass windows of the vacuum chamber. We can nonetheless obtain an estimate of the number of atoms, and fortunately the large uncertainty has little consequence for our intended experiments.

A more important property is the spatial distribution of the cloud. For experiments where we manipulate the atom momenta (see Section 9.2), we intend to measure the effect of our pulse sequence by mapping the momenta to position, and any perturbation in the initial spatial distribution must be considered.

The final property is the momentum distribution. As discussed in Section 2.4.4, this should be well described by a Maxwell-Boltzmann distribution with a characteristic temperature. By approximating the distribution to a Maxwell-Boltzmann, we endeavour to measure this characteristic temperature using the method of release and recapture.

7.1   Cloud Size and Shape

A sufficiently accurate measurement of the cloud size can be found by using an imaging system of known magnification and a camera of known pixel spacing. Alternatively, to avoid the cumulative uncertainties of this indirect calibration, we can image a structure of known size and then use the same optical system, including the camera, to image the cloud; we use this method to calibrate the imaging system.

We use a Watec WAT-902DMS3 CCD camera, with the unexceptional pixel count of 384×288 (8 bit gray-scale) but which has excellent sensitivity at 780nm. Care must be taken to ensure that an image of the cloud does not saturate any pixels.

By using the practically perfect beams emerging from the single-mode optical fibre, and by careful alignment, we are able to obtain a round and small-scale structure free cloud; this was not possible without the aggressive spatial filtering of the short stretch of fibre. The additional magnetic fields perturbing the trap are approximately constant over the cloud's dimensions, and do not cause unwanted structure. We learn in Section 8.3.6, however, that the cloud is still slightly offset from the magnetic field zero.

A false-colour image of the cloud is shown in Figure 7.1. Little structure is visible within the central peak, and cross-sections through the brightest part of the image, shown in Figure 7.2, allow us to calculate the width of this distribution. The recorded signal is due to the integrated fluorescence along the optical axis and if, as predicted in Section 2.4.5, the cloud were a uniform density ellipsoid, we would not see the long Gaussian-like tail far from the centre. We conclude that the cloud is not hard-edged, perhaps due to absorption of the trapping beams as they pass through the optically dense cloud (which was neglected previously by assuming the light was very far detuned) giving a larger confining force but also a larger scattering rate to atoms near the surface, and we
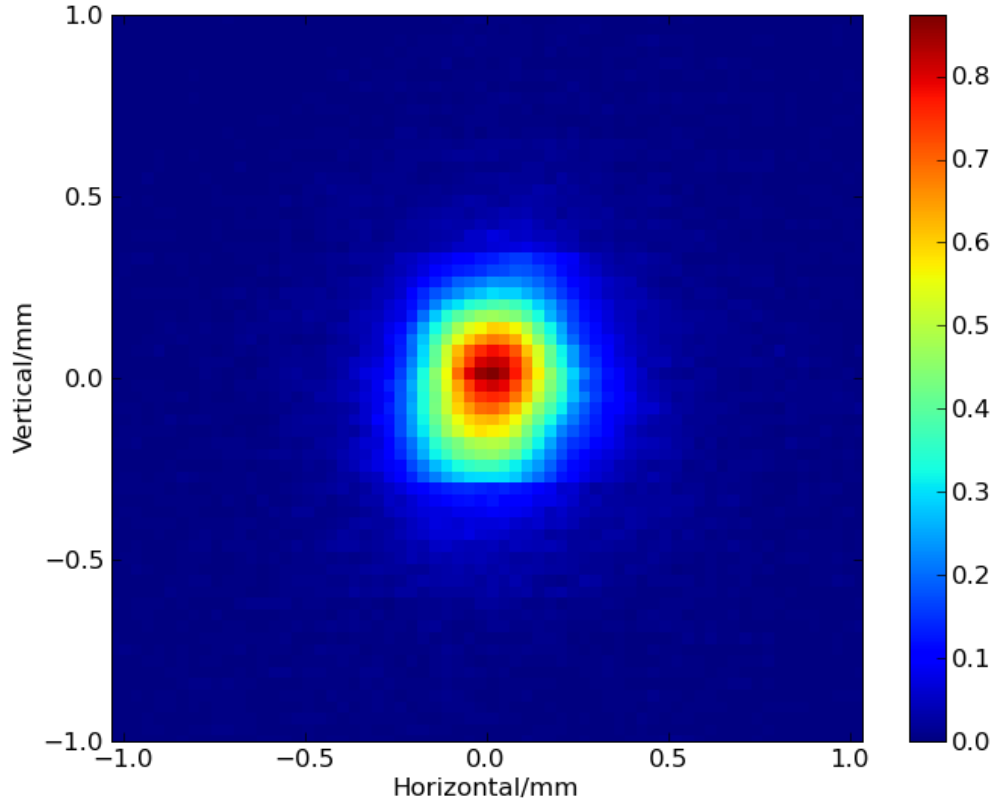
FIGURE 7.1: A false-colour image of the trapped atoms. The amplitude is linear with intensity. The transverse scale was calibrated by imaging a ruler using the same optical system; the short depth of focus ensures that the ruler and atom cloud are subject to the same magnification.

estimate the cloud size from the FWHM of the cross-section: the cloud is $(206 \pm 5)\mu$m across by $(245 \pm 10)\mu$m high.

As mentioned in Section 2.4.5, a number of interesting shapes have been observed in magneto-optical traps. By deliberately misaligning one of the horizontal retro-reflected beams, we introduce a shear and the atoms are confined to a toroidal rather than a spherical region. This misalignment can be carried out quickly and reversibly without losing a significant fraction of the atoms. The cloud before and after deliberate misalignment is shown in Figure 7.3.
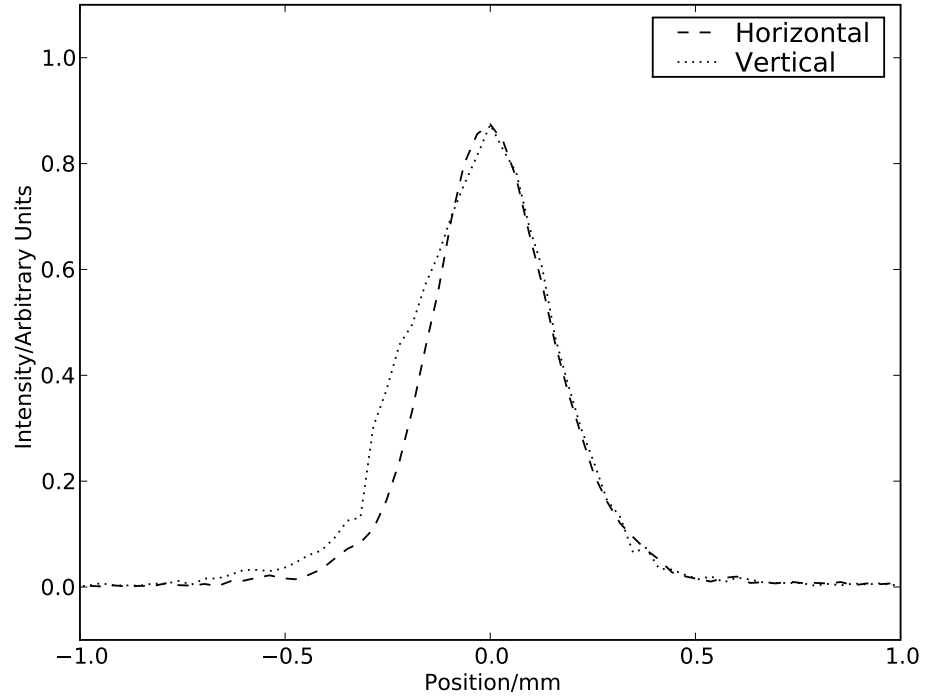
FIGURE 7.2: Cross-sections through the brightest part of the cloud image, shown in Figure 7.1. The distributions are approximately Gaussian and have widths $(206 \pm 5)\mu$m and $(245 \pm 10)\mu$m.
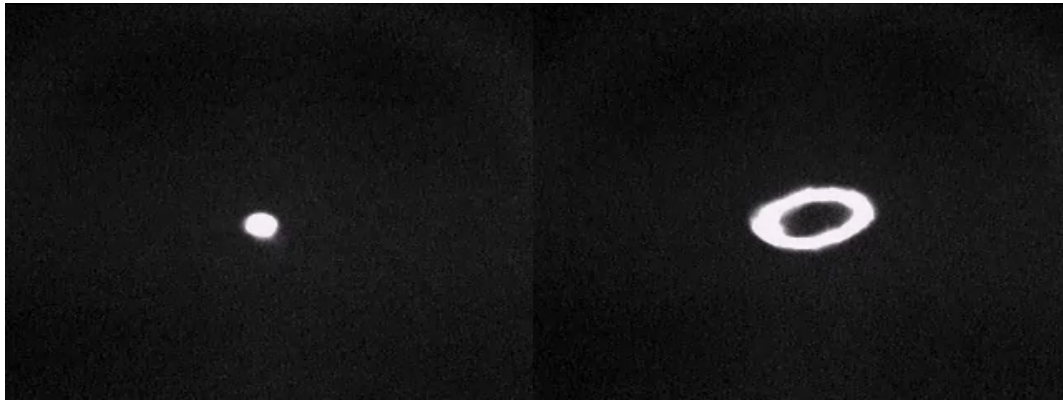
FIGURE 7.3: Images of the MOT cloud before (left) and after (right) misalignment of a horizontal retro-reflected beam. When the beams are not counter-propagating, atoms are confined to a ring with a clear absence of atoms near the centre.

7.2 Atom Number and Density

The number of atoms in the cloud can be found by measuring the fluorescence from the cloud using a photodiode and, by calculating the fluorescence rate for each atom using the intensity and detuning of the beams, estimating the number of atoms necessary to produce this signal. This method is simple, but is often troubled by experimental uncertainties.

Our beams emerge from a fibre output coupler and so have a well defined waist of 7.1mm (specified by the manufacturer datasheet), and the powers can also be measured relatively accurately:[1] repump power $P_{\mathrm{RP}} = (19.4 \pm 0.1)$mW and molasses power $P_{\mathrm{mol}} = (10.0 \pm 0.1)$mW. The detuning of the molasses laser is known to less than a linewidth: $\Delta = -2\pi \times (18 \pm 1)$MHz. We can now calculate the average fluorescence rate per atom as the upper state lifetime multiplied by the steady-state populating in the upper state. This population is typically estimated using the formula derived in Section 2.3 for the steady state population of a two-level atom:

$$\rho_{1,1} = \frac{1}{2}\frac{s}{s+1}. \tag{7.1}$$

However, the rubidium atom is more complicated than a two-level atom, and in order to calculate the fluorescence rate we should account for these additional levels. As described further in Section 8.2.1, we use a numerical model which treats each transition as a two-state system, and, while ignoring coherences, calculates the population dynamics for a number of applied fields. Using this model, and the measured powers and beam waists, we find the steady state population $\rho_{1,1} = (20 \pm 1)\%$; we have assumed the atoms are near the centre of the Gaussian beam profile. In contrast, the simplistic two-level model gives $\rho_{1,1} = 22\%$. The total fluorescence from our cloud is therefore

$$\begin{aligned} P_{\mathrm{cloud}} &= N \times hf \times (20 \pm 1)\%\,\Gamma \tag{7.2} \\ &= N \times (1.9 \pm 0.1)\,\mathrm{pW} \end{aligned}$$

where N is the number of atoms.

---

[1]We use a Newport 840-C with 818-UV head and 883-UV attenuator.
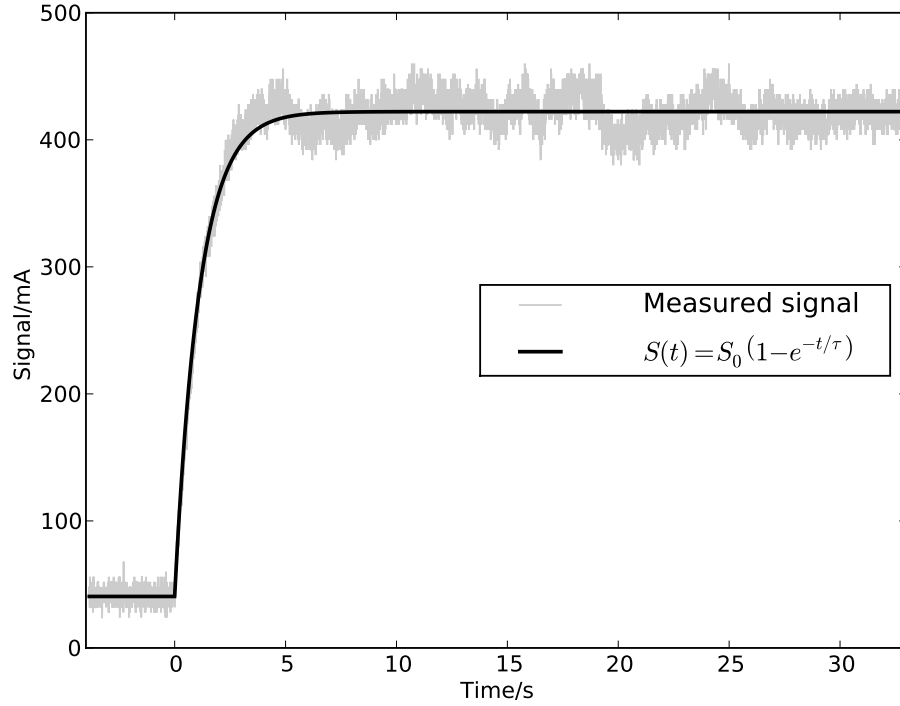
FIGURE 7.4: The measured photodiode signal as the MOT cloud forms. This measurements allows us to measure the fluorescence from the cloud on top of a large background from laser scatter. The parameters for the fit are: loading rate $\tau = 1.1\,\text{s}$ and amplitude $S_0 = (382 \pm 1)\,\text{mA}$. The offset is $41\,\text{mA}$.

Next we measure the fluorescence detected by a photodiode and a plot of this signal as the cloud forms is shown in Figure 7.4. The measured current is $(382 \pm 1)\text{mA}$ and from the photodiode responsivity $(85 \pm 5)\text{mA/W}$ and the amplifier gain 2×10^4, we estimate the power incident on the detector as

$$
\begin{aligned}
P_{\text{detector}} &= \frac{(382 \pm 1)\,\text{mA}}{(85 \pm 5)\text{mA/W}} \times \frac{1}{2 \times 10^4} \quad &(7.3)\\
&= (225 \pm 15)\,\mu\text{W} \quad &(7.4)
\end{aligned}
$$

The final part of this calculation is one that unfortunately introduces a large uncertainty. We must estimate the fraction of fluorescence from the atoms which falls on the active area of the photodiode. A fast response is not important, so we use a large area photodiode (Thorlabs DET110) and, crudely, the collection optics can be approximated by a 2″ lens at 160mm, which therefore collects 0.63% of the light. Using the ray-tracing model described in Section 7.3.2 we find this crude approximation to be valid here where the

atoms are near the focus and near the optical axis. Not so easily modelled is the loss due to scatter and reflection from the various optical components between the cloud and the detector: anti-reflection coated but thick vacuum view-port; narrow band (780 ± 1)nm filter (Chroma Z780/10X) used to remove any signal from non-laser light; and the anti-reflection coated but still imperfect surfaces of the lenses themselves. The measured transmission of the filter is $(91 \pm 2)\%$, and the transmission through the vacuum ports is quoted as $(93 \pm 2)\%$. The fraction of light which falls on the detector is therefore $0.63\% \times (91 \pm 2)\% \times (93 \pm 2)\% = (0.53 \pm 0.02)\%$ and our atom number is

$$
\begin{aligned}
N &= \frac{1}{(0.53 \pm 0.02)\%} \frac{(225 \pm 15)\,\mu\mathrm{W}}{(1.9 \pm 0.1)\,\mathrm{pW}} \\
&= (22 \pm 2) \times 10^9 \text{ atoms.}
\end{aligned}
\tag{7.5}
$$

This number is typical of atom traps loaded from background vapour. The uncertainty quoted is comparatively small, but this is for one instance of the MOT, and day to day the actual number varies by a factor of ten or so. Additionally, the method used to calculate the steady state population of the upper state assumes we can define a saturation intensity for each transition, but this is not strictly true when the transition is not closed. Fortunately, the atom number is not crucial for our later experiments, and the main purpose of this measurement is to ensure we understand the radiative processes in our atoms, and the characteristics of our collection optics.

If we model the density as a Gaussian distribution in each axis

$$
\rho(x, y, z) = \rho_0 e^{-\left[(x/x_0)^2 + (y/y_0)^2 + (z/z_0)^2\right]} ,
\tag{7.6}
$$

noting that we observe the integral along one axis (y) of this distribution and assuming the dimensions along x and y are the same $(x_0 = y_0)$, we can estimate the density from the measured atom number and full-width at half-maximum:

$$
\rho_0 x_0^2 z_0 \pi^{3/2} = N
\tag{7.7}
$$

where we have used $\int_{-\infty}^{+\infty} e^{-(a/a_0)^2} = a_0\sqrt{\pi}$. The density is $\rho_0 = (385 \pm 45) \times 10^9 \text{ mm}^{-3}$.

7.3 Temperature

The method of release and recapture involves in principle switching off the trapping (and cooling) potential, allowing the atoms to move ballistically for a known time, and then measuring the fraction of atoms which remain in the trapping volume when the potential is switched back on. The potential can be switched—the AOM controlled trapping and repumping beams can be quickly extinguished and the magnetic gradient alone has a small effect on the atoms[2]—but it is not so simple to measure the fraction of atoms remaining in the trapping volume, or even to accurately define what is meant by this volume. In reality, spontaneously emitted photons are collected by a system of lenses and imaged on to a detector; fluorescence from atoms in front of and behind the centre of the cloud relative to the imaging system hence also falls on the detector. One may proceed with geometrical models—one of which is outlined below—but the correct approach is to accurately model the physical system. A simple but very effective numerical model is presented and from it the temperature is deduced.

7.3.1 Geometrical Models

In the simplest geometrical model, we restrict ourselves to 1D and consider atoms uniformly distributed within a hard-edged trapping volume and a Maxwell-Boltzmann momentum distribution.

$$s_{\mathrm{i}}(x) = \begin{cases} 1/a & |x| \leq a/2 \\ 0 & |x| > a/2 \end{cases} \tag{7.8}$$

$$p_v(v) = \frac{1}{\sqrt{2\pi}} \sqrt{\frac{m}{k_{\mathrm{B}}T}} \exp\left[-\frac{mv^2}{2k_{\mathrm{B}}T}\right] \tag{7.9}$$

or, as a spatial spread after time τ,

$$p_x(x) = \frac{1}{\tau} \frac{1}{\sqrt{2\pi}} \sqrt{\frac{m}{k_{\mathrm{B}}T}} \exp\left[-\frac{m(x/\tau)^2}{2k_{\mathrm{B}}T}\right] \tag{7.10}$$

---

[2]Atoms in the extreme m_F state see a frequency shift 470kHz/G $\times$ 3 = 1.4MHz/G, or an acceleration of 6.6m/s^2 for our field gradient. For the typical free-expansion time of 10ms this is only 80μm, which is far smaller than the $\sim$ 1mm point-spread function.

The full-width at half-maximum of this point spread function hence grows proportionally to τ with a characteristic speed

$$v_{\text{PSF}} = \frac{\partial}{\partial \tau} \text{FWHM} = 2\sqrt{2 \ln 2} \sqrt{\frac{k_B T}{m}} \tag{7.11}$$

The spatial distribution after a time τ is found simply by convolving the initial distribution with the point spread function $p_x(x)$. To find the signal measured by some detector, we integrate this final distribution over the region ($|x| \leq b/2$) within which atoms contribute to the signal.

$$
\begin{aligned}
d(\tau) \quad &\propto \quad \int_{-b/2}^{+b/2} \left[s_{\text{i}} * p_{\text{x}} \right](x) \, \mathrm{d}x \tag{7.12} \\
&\propto \quad \int_{-b/2}^{+b/2} \int_{-a/2}^{+a/2} p_{\text{x}}(x - x') \, \mathrm{d}x' \, \mathrm{d}x \tag{7.13}
\end{aligned}
$$

Typical curves are shown in Figure 7.5.

This model can be extended to three dimensions, but to correctly model the contribution of fluorescence of each atom, the volume imaged on to the detector must extend in front of and behind the cloud centre relative to the detector; the imaging characteristics of the collecting optics become important if we are to build an accurate model which, ideally, requires only temperature as a fitting parameter.

7.3.2 Realistic Model

A realistic reproduction of the signal on a detector can be found by simulating the ballistic expansion of an initial distribution of N particles, each of which isotropically emits M photons, some of which, after propagating through a lens or system thereof, fall on the active area of a detector; the contribution of an atoms fluorescence is found directly. Such a numerical model has the advantage that it is possible to build in such experimental inconveniences as finite beam size or occlusion of the beam with little additional complexity.
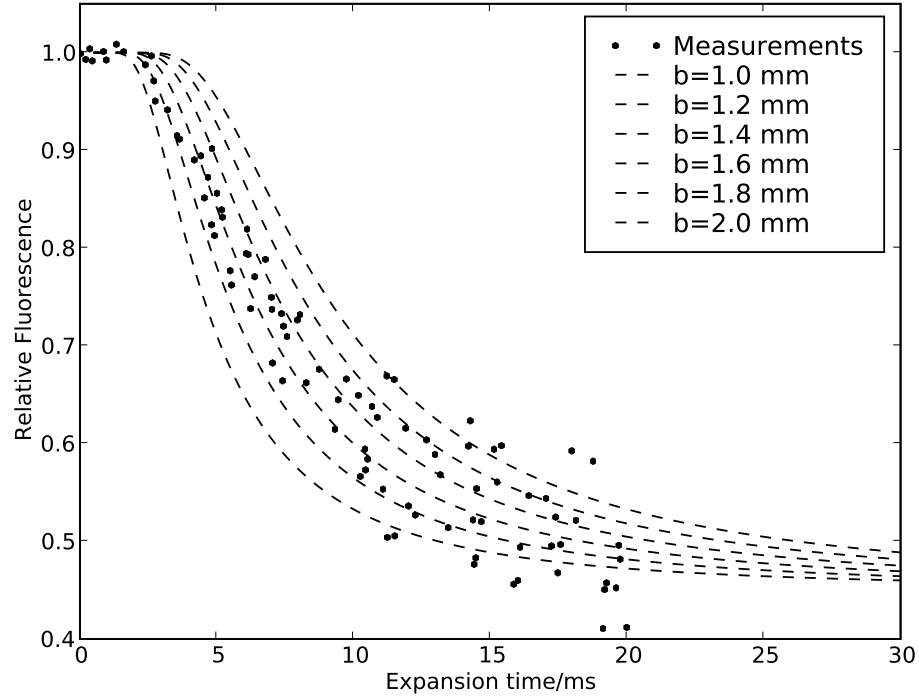
FIGURE 7.5: Simple geometrical model of the release and recapture measurement. To account for expansion in the two dimensions transverse to the axis of the imaging optics, the square of the function described in the text is plotted. A convincing fit to the data can be found for a wide range of temperatures by adjusting the free parameter b; a model with fewer free parameters is required.

It is possible to use the initial distribution of atoms as observed in Section 7.1, but as all atoms in the initial distribution are imaged on to the photodetector, small differences between the modelled and the real spatial distribution are not expected to severely influence the results. The characteristic speed (see Equation 7.11) for rubidium 85 atoms with the typical temperature $T \sim 100\mu$K is $v_{\mathrm{PSF}} \simeq 0.16$mm/ms. Atoms typically do not begin to leave the imaged region until after a few milliseconds, so we need only consider spatial density perturbations larger than a millimetre; this is larger than the cloud, and so we can safely treat the cloud as a uniform density sphere with a small radius. Far more important is the velocity distribution, which is modelled as a Maxwell-Boltzmann distribution (see Section 2.4.4) with temperature as the free parameter which we hope to fit to the data.

The emission of photons from each of these atoms is assumed to be isotropic, and hence the direction of each ray is chosen at random to be a uniformly distributed point on the

surface of a sphere. The model assumes that each photon will not be scattered between atoms and lens; this is unlikely to be true for the initially optically dense cloud, but any absorbed photon will be re-emitted and any overall delay will be far smaller than other time-scales in this situation.

7.3.3 Measurements

The sequence of pulses used for this measurement is a simple example of those discussed further in Section 8.2. The MOT beams are extinguished and, after a delay, the trapping laser is flashed on for a short time. The fluorescence of this read out pulse is recorded and is plotted, relative to the fluorescence during normal operation of the MOT, as a function of the delay time; see Figure 7.6. We note that the fluorescence measured by the trapping laser for short delays is the same as that recorded for the steady-state MOT. From this we conclude that the majority of the population is cycling on the cooling transition during normal MOT operation, and this is confirmed by numerical simulations described in Section 8.2.1. The measured fluorescence is due to spontaneous emission from the cooled atoms, but there is also a large component due to scattering from the vacuum system windows and the internal surfaces of the chamber. This appears as an offset, and causes the fluorescence to decay to a constant rather than zero.

Also shown in Figure 7.6 are the results of the simulation described in Section 7.3.2, the only free-parameters of which are temperature and the offset due to background fluorescence. We estimate the temperature to be $T = (80 \pm 20)\mu\text{K}$ and state with confidence that this is less than the Doppler-limit (Equation 2.38) for rubidium cooled via the D2 line of $T = 146\mu\text{K}$ [73]. (This sub-Doppler temperature measurement was made on a smaller and less dense cloud than that described in sections 7.1 and 7.2. The temperature of that cloud was significantly higher at $(225 \pm 10)\,\mu\text{K}$.)

7.3.4 Imaging Optics

The fluorescence of atoms in the magneto optical trap is captured by a system of lenses and imaged on to a photodetector. In general, these atoms move away from the focal plane, their fluorescence is occluded by various apertures, and the imaging is subject
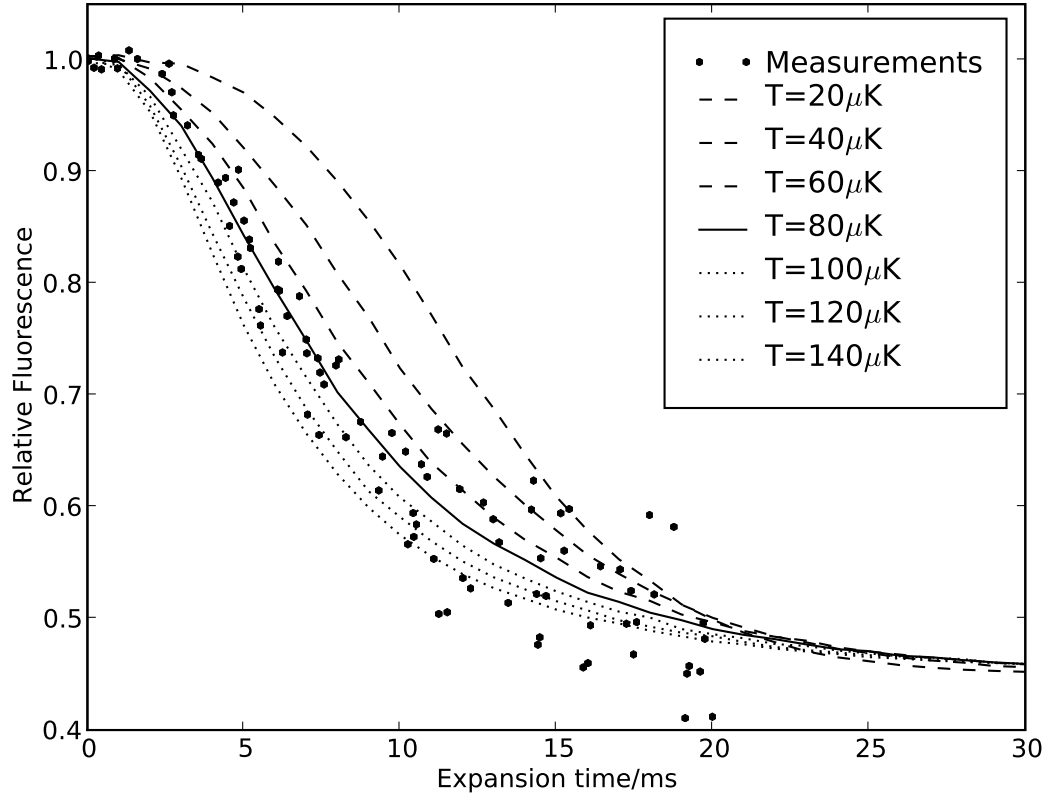
FIGURE 7.6: Release and recapture measurements and a numerical model fit. Various temperatures, the best fit of which is $T = 80\mu K$, are shown. The offset from zero is due to scatter from optical components.

to spherical aberrations. For some measurements, notably release and recapture temperature measurement discussed below, a model which ignores the imperfections in the imaging optics is insufficient.

The length-scales involved in the optics are very much larger than the wavelength and so it is reasonable to expect an accurate model can be constructed using geometrical ray-tracing. Objects such as a lenses and irises can be used to construct a realistic model of the imaging system, as illustrated in Figure 7.7. We expect spherical aberrations to be significant, so we shy away from the oft-used thin-lens approximation. The way in which these components are modelled is described below.

Rays are modelled as having a position and a velocity which is perturbed as it propagates through each optical element. Free space propagation of rays is trivial and refraction
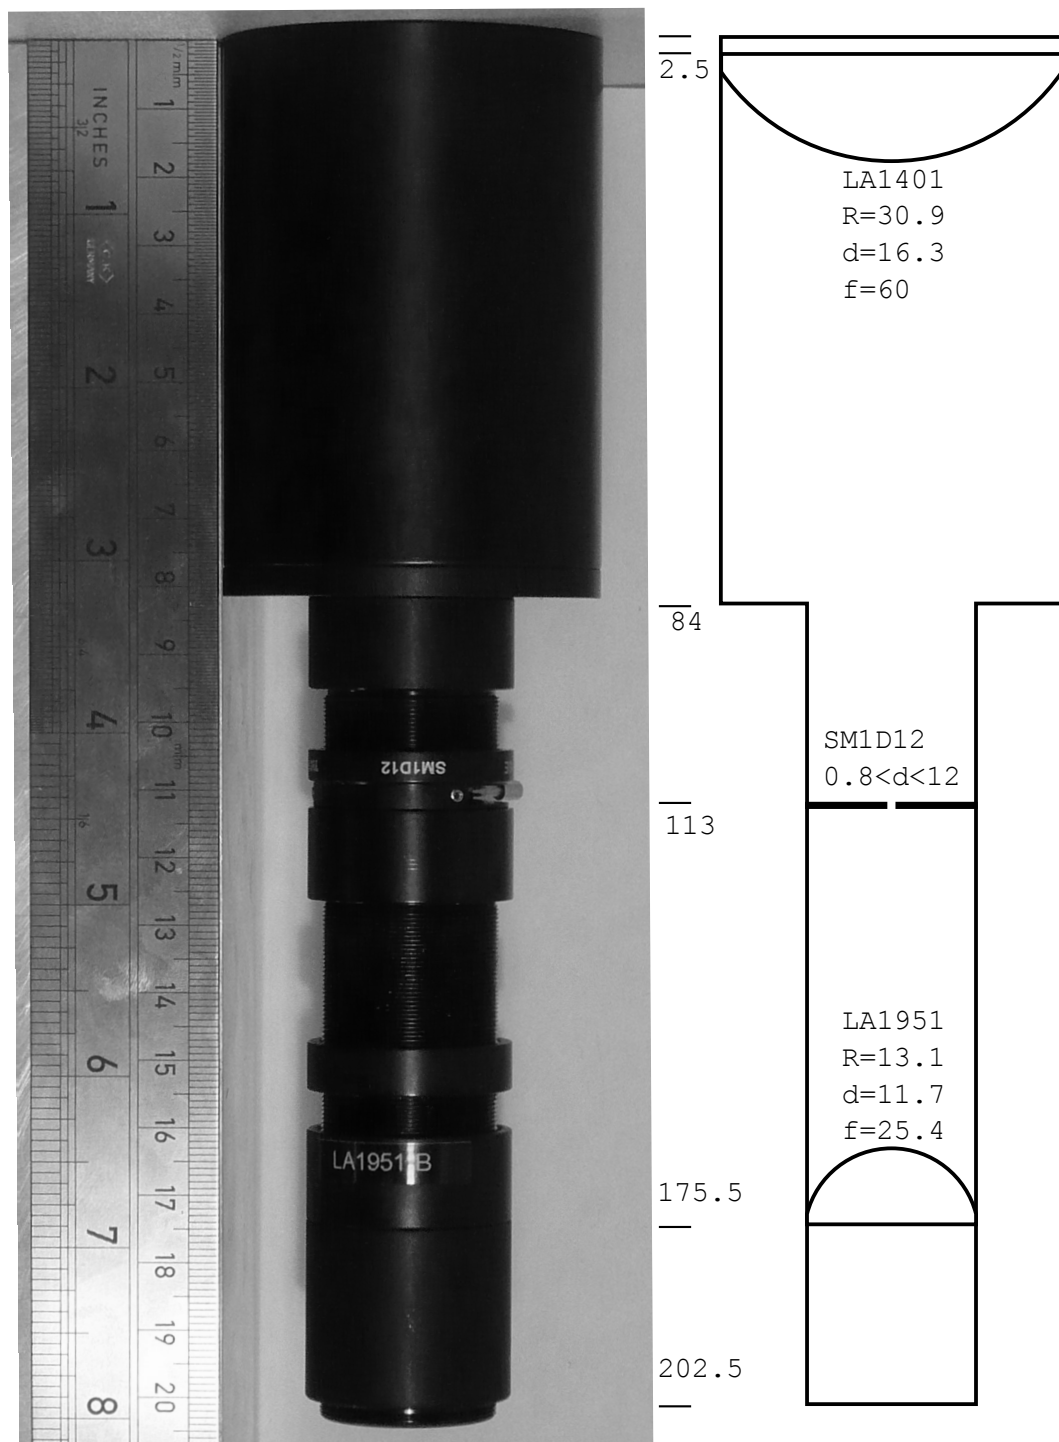
FIGURE 7.7: Photograph and schematic of the imaging optics. Dimensions are in millimetres and proportions, including lens curvatures, are to scale. For lenses, the Thorlabs part number, radius of curvature R, thickness d, and focal length f are shown.

from a curved surface can be modelled by finding the vector normal to the surface where the ray impacts, and applying Snell's law of refraction.

Given the surface normal and the ray velocity vectors, we rotate the co-ordinate system such that the surface normal is aligned with the x axis and the velocity vector is in the x, y plane with positive x. The problem is thus reduced to two-dimensions and the refracted angle is found easily via Snell's law, after which the co-ordinate system is rotated back to find the refracted velocity vector in the original three-dimensional co-ordinate system. The problem of modelling refraction is now reduced to finding the position at which a ray touches a surface, via free-space propagation, and the surface normal at that position.

For the plano-convex lens, the principle object from which our system is constructed, the point at which a ray hits the spherical surface is found by solving the three-dimensional vector equation $|\vec{x} + \vec{v}t - \vec{c}_c| = r_c$ for t where $\vec{c}_c, r_c$ are the centre and radius of curvature respectively; the position is then found by incrementing the position: $\vec{x} \rightarrow \vec{x} + \vec{v}t$. The surface normal is then simply $\vec{n} = \vec{x} - \vec{c}$, and the refracted angle is calculated from this vector and the refractive indices, as described above. The refracted ray is then propagated to the plane face of the lens and, using a surface normal $\vec{n} = \hat{z}$, the refracted angle at which the ray leaves the lens is found.

7.3.5 Aside: Ray Tracing Programs

The approach described above was implemented by specifically written code, but this situation is an example of the more general problem of ray-tracing, for which many excellent (and numerically efficient) applications already exist. It would be interesting therefore to investigate to what extent these programs could be used to simulate such physical situations. A good choice of program might be PovRay as this internally represents object geometries as constructions (combined via unions, differences, intersections...) based on primitives (planes, spheres, cuboids...) in contrast to the majority which use may these primitives for interface with a user but which ultimately represent the resulting objects as collections of vertices. The approach used by PovRay can be expected to be much more efficient and accurate for objects such as planoconvex lenses which can very naturally be represented by these primitives.

Chapter 8

Hyperfine Manipulation

The culmination of all that has been discussed so far is coherent manipulation of the hyperfine ground states of cold rubidium 85. The first stage of such manipulation is to prepare the population in a known state. This operation, the way in which it can be used as a calibration, and the closely related read operation, are discussed first. I then describe several experiments we have undertaken to drive Raman transitions and understand the environment of our atom cloud.

8.1   Digression: Previous Work

Before presenting the results of our experiments, we acknowledge work that has gone before. As well as placing the current work in context, this highlights the general features which we can hope to observe.

There has been much experimental work done to drive Raman transitions in trapped alkali metal atoms, and typically the atoms have been far more localised than those in our diffuse cloud. This is a requirement of our approach, because if we are to coherently manipulate the atom momentum, our particles must be free during the manipulation sequence. Atoms collected in a MOT are typically loaded into a dipole trap created by a tightly focused far detuned laser [81, 131–134]. By trapping the atoms independent of their internal state, one can take much more time over the measurements. In Dotsenko

142
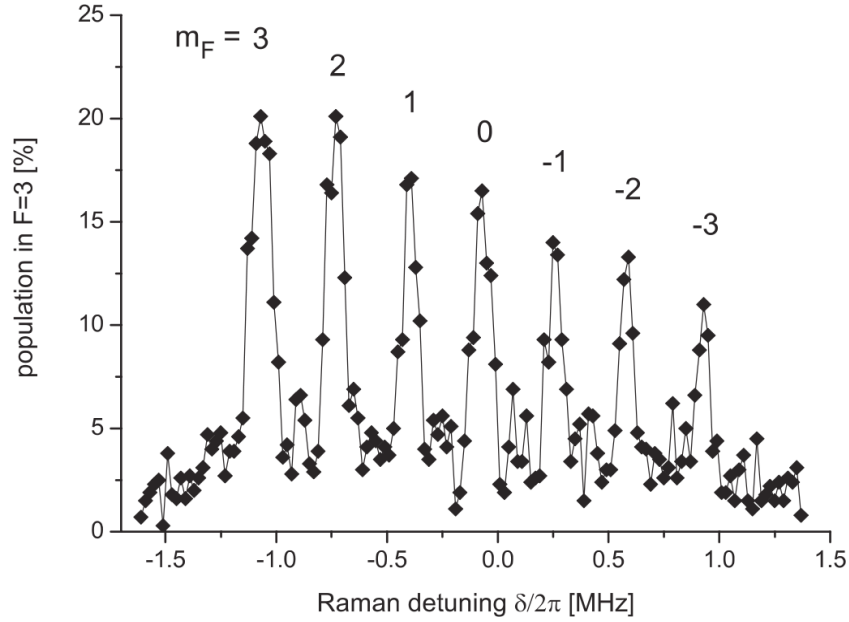
FIGURE 8.1: A Raman frequency scan clearly showing seven peaks corresponding to the Zeeman sub-levels of Caesium $|F = 3\rangle$. Transitions are driven by equal circular polarisations and $\Delta m_F = 0$ [81].

et al [81], atoms are coherently manipulated and then those in one hyperfine state removed from the dipole trap by a 'push-out' laser. The fluorescence from the remaining atoms is a measure of the population transfer, and the result of their Raman frequency scan is shown in Figure 8.1. Rabi oscillations between the hyperfine states of atoms so trapped have been observed at frequencies up to several megahertz, and an example is shown in Figure 8.2 [134]. Ramsey interferometry experiments have also confirmed that the coherence between the states is long-lived, and is not significantly perturbed by the confining potential.

An alternative approach is to use trapped ions. As point-like charges, these particles can be trapped independently of their internal state using electric and magnetic fields. A linear array of ions in a Paul trap [33, 135, 136] exhibits superb coherence, control, and addressability, and this approach has been singled out as the most promising scheme for a scalable quantum computer [137]. Single qubit operations are almost mundane in this field, and two-qubit operations were first demonstrated over a decade ago [138]. Particles can be cooled to their motional ground state [139] and excellent work is on going [140].

In our experiment, we have sacrificed the tight localisation of particles to keep open
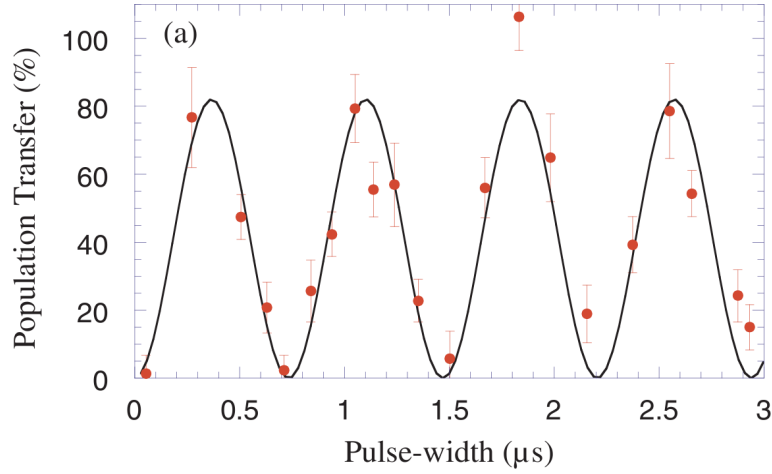
FIGURE 8.2: Fast MHz frequency Rabi oscillations between the hyperfine ground states of dipole trapped Caesium [134].

the possibility of manipulating momentum. Bakos *et al* [104, 141–143] have explored manipulation of atom momenta with laser pulses, but using single-photon transitions, not Raman transitions.

8.2 State Preparation and Read-out

Coherent experiments typically require state preparation, manipulation, and read out operations. The two-state system is embodied by the ground hyperfine levels of the atomic ensemble, and the various operations are provided by computer controlled laser pulses. These operations are broadly classed as either incoherent or coherent, the former covering preparation and read out and the latter covering everything from π pulses to adiabatic rapid passage and complex quantum computer-like pulse sequences.

The sequence begins by extinguishing the MOT beams; when the trap is operating, the atoms are cycling on a radiative transition which redistributes the population and quickly destroys any coherence. When these beams are switched off, the trapping potential is removed and the cloud begins to expand ballistically, subject only to gravity and the magnetic field gradient; all operations must complete before a significant fraction of atoms leave the region where Raman beams and MOT beams coincide. If we assume that all atoms imaged on the photodector are equally affected by the Raman beams, the
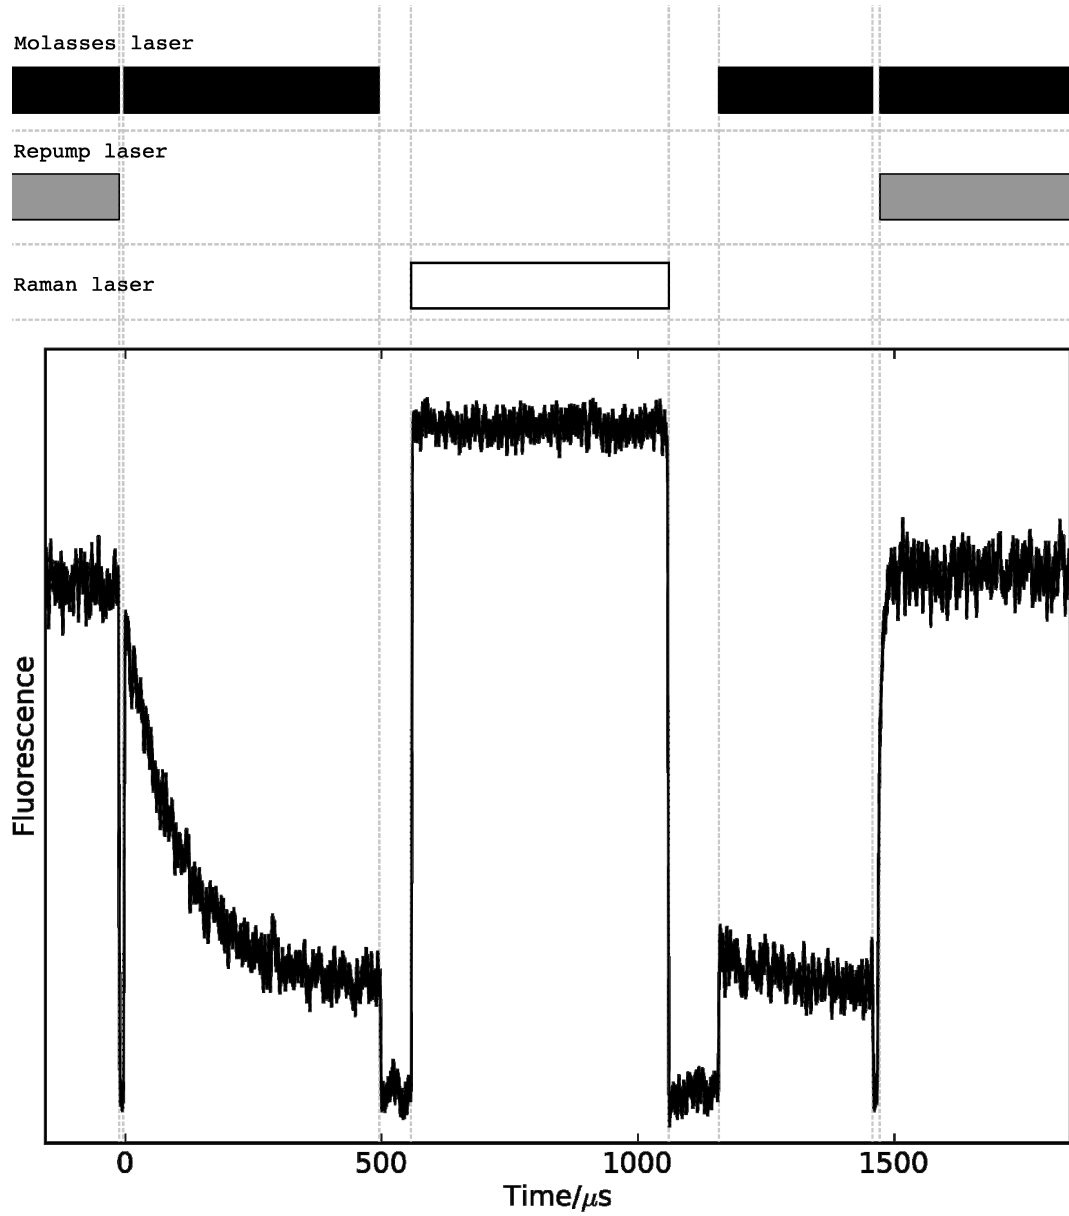
FIGURE 8.3: The fluorescence recorded during a typical pulse sequence. Initially the MOT operates normally, and then both trapping and repumping beams are extinguished. Here, as normally, the trapping beam is used for the prepare and read operations. Between these incoherent operations, a sequence of Raman pulses coherently manipulates the pseudo-two-level system of the ground-hyperfine levels $|F = 2\rangle$ and $|F = 3\rangle$. The manipulation laser is not absorbed by the atoms, but the large scatter from optical components is visible on the detector.

time-scale for which a 'significant' fraction remains can be found from the measurements in Section 7.3.3 and is $\simeq$ 4ms.

The typical pulse sequence is illustrated in Figure 8.3. The first operation prepares the atoms in a particular state; this is achieved by turning on *one* of the MOT beams. This couples one of the hyperfine states to the radiative upper state and, after spontaneous

emission, a fraction of the population will have decayed to the non-coupled level. This population remains in this non-coupled level as that which returned to the coupled level is once again coupled to the radiative state. Very quickly, all the population accumulates in the non-coupled state. Next we apply a pulse or sequence of pulses from the manipulation laser; as previously discussed, this laser coherently couples the two hyperfine ground states, and is parameterised by amplitude, frequency, and phase, all of which can be controlled during the pulse; the simplest example, shown in the figure, is of a single top-hat pulse at a fixed frequency. The final stage is the read operation for which, similarly to the prepare operation, we switch on one of the MOT beams. This pumps the population to a particular state, but by observing the fluorescence due to spontaneously emitted photons during this pumping process, we obtain a signal proportional to the fraction of population found in the hyperfine ground state coupled by the chosen MOT beam. We compare the fluorescence observed during the read operation to that recorded during the state preparation.

8.2.1 Optical Pumping Rates

We normally use the molasses laser to prepare the quantum system. This laser is locked approximately three linewidths below the $\left|5^2\mathrm{S}_{1/2}, 3\right\rangle$ to $\left|5^2\mathrm{P}_{3/2}, 4\right\rangle$ transition and during normal operation of the MOT most of the population is cycling on this transition. Approximating this to a two-level system as in Section 2.4, the transition rate is

$$R = \frac{s_0/2}{1 + s_0 + (2\Delta/\Gamma)^2}\Gamma.$$

(8.1)

The saturation parameter is the ratio of intensity to the saturation intensity and depends on the levels being coupled,

$$I_{\mathrm{SAT}} = \frac{c\epsilon_0\,(\hbar\Gamma)^2}{4\,|\langle f|\,\hat{\mu}_q\,|i\rangle|^2}$$

(8.2)

where $|i\rangle$, $|f\rangle$ are the initial and final states and q is the polarisation; for each pair of states, we calculate the absorption rate due to a laser or number of lasers, and the spontaneous emission rates according to the branching ratios. The dynamics are similar to those described in Section 2.1.2, and specifically, if we represent these couplings by a

matrix, to ensure population is conserved, the matrix must obey the constraint imposed on Γ: $\sum_n \Gamma_{n,m} = 0$ for each m.

For 2mW per MOT beam (σ_- and σ_+ along each of the 3 dimensions) with $1/e^2$ diameter 7.1mm, and starting with the population equally distributed amongst the m_F sub-levels of $\left|5^2\mathrm{S}_{1/2}, F = 3\right\rangle$, the population in $\left|5^2\mathrm{P}_{3/2}\right\rangle$, to which the fluorescence is proportional, is predicted to increase initially during the first 100ns to $\simeq 6\%$, and then to decrease with a half-life $\tau_{1/2} = 56\mu$s. This model neglects the motion of atoms through the spatially dependent polarisation and intensity created by the beams (see Section 2.5), can only estimate the intensity seen by the atoms at their position which is displaced from that of ideal MOT alignment (see Section 8.3.6), and by considering the populations, we have ignored any coherences (the off-diagonal elements of the density matrix) which might appear[1]. Nevertheless, an exponential fit to the data in Figure 8.3 gives a half-life of 69μs, which is 30% longer than predicted, and can be explained by uncertainties in the intensities and detunings. We can be confident that the underlying physics, if not the precise numbers, is modelled correctly. The model confirms that population is evenly distributed across the sub-levels of $\left|F = 2\right\rangle$ in the very long-term, but that after a few times the half-life, there is approximately 10% more population in states $m_F = \pm 1$ than in $m_F = 0$: $p(\pm 2) = 18\%$; $p(\pm 1) = 23\%$; $p(0) = 16\%$ after $50\tau_{1/2}$.

8.2.2 Analysis of Fluorescence Signals

A recorded trace, such as that shown in Figure 8.3, must be analysed in order to determine what fraction of population is transferred from $\left|F = 2\right\rangle$ back to $\left|F = 3\right\rangle$ by the manipulation sequence. As can be seen in the figure, this transferred population appears as an exponential decay of the read-out pulse, with the same time-scale as the prepare pulse, but a different amplitude. We are able to accurately mark the positions of these signals using a recording of the signal applied to the AOM used to switch the trapping laser (see Section 5.4) and hence we can separate these signals from the other features in the data. With these extracted exponentials, we proceed to fit, using the MINPACK libraries for least-squares fit, two exponentials with a common offset and decay rate, but

---

[1]For example, this model cannot reproduce Electromagnetically Induced Transparency [144, 145].

with different amplitudes. The ratio of these amplitudes is a measure of the signal we sought, and the absolute value of the prepare pulse amplitude can be used to discriminate results taken when the MOT was not present due, for example, to the trapping laser falling out of lock (see Section 5.1.4).

8.3 Raman signal

Despite testing all components of the MOT and manipulation laser apparatus, the characteristic Raman signal remained elusive for many months. Having invested the effort to make the majority of the apparatus computer controlled, we made good use of the automation and large data collection facilities during our search. We parameterised our experiments with values including the duration of a manipulation pulse and the frequency difference between the composite Raman beams. These values were set by the computer and the resulting fluorescence signal, recorded by a digital oscilloscope, downloaded for later processing.

When searching for an elusive signal, and especially one in which there is a large background, there is a danger that a slow drift can be mistaken for the signal. We must, for example, ensure any change in state population is due to Raman transitions and not to an increase in optical pumping because the laser frequency has drifted. Using computer control, we are able to choose the value of a parameter at random, within a certain range, and hence map slow drifts in a signal to uncorrelated noise. After several false positives with the simpler regular and sequential measurements, this random point picking has allowed us to be certain of the origin of the signals we observed.

During such lengthy searches, there is unfortunately a chance that some part of the experiment will falter. The Tiger wavelength drifts, as has been discussed previously (see Section 6.2.1), but more often, the MOT lasers fall out of lock. When this happens, the MOT cloud disappears and hence, the amplitude of the prepare signal found from the data, as described in Section 8.2.2, decreases significantly. This decrease may also occur slowly over several runs if it is the repump frequency which drifts, as this has a less drastic effect on atom number. When plotting the result of such lengthly runs,

therefore, we discard the points for which the amplitude of the prepare signal drops below a threshold.

We are left with a scatter of points which, while retaining all the features of the signal, is difficult to interpret by eye. We would like to plot a curve, but the irregular spacing of the data means smoothing algorithms must be used with care. We proceed by choosing a regularly spaced set of values for the independent variable, typically frequency, and for each point calculating a weighted mean and standard deviation for the vertical *and* horizontal positions. We choose a Gaussian weighting for the points, and calculate the values as follows for each of the regularly spaced frequencies:

$$\bar{a} = \frac{\sum_i p_i a_i}{\sum_i p_i} \quad \text{and} \quad \sigma_a = \sqrt{\frac{\sum_i p_i \left(a_i - \bar{a}\right)^2}{\sum_i p_i}} \tag{8.3}$$

for both the frequency x and the relative fluorescence y; i.e. $a = (x, y)$. p_i is our weighting function, here chosen to be $p_i = e^{-(x_i - x_0)^2/w^2}$ where w is a characteristic width, typically 3kHz, and is reported where this algorithm is used.

We now apply these tools to the scatter of points recorded for various experiments.

8.3.1 Narrow central resonance

We apply an external magnetic field, as described in Section 3.6.1, which lifts the Zeeman degeneracy of the m_F sub-levels. Using equal circular polarisation, we drive Raman transitions with $\Delta m_F = 0$ and, having prepared the population in $|F = 2, m_F = (-2 \ldots +2)\rangle$, we expect to observe 5 resonances. Initially we apply a large magnetic field and search near zero detuning for the $m_F = 0$ resonance, whose frequency is not affected by this field. We use a wavelength 780.262nm $< \lambda \le$ 780.268nm, corresponding to a detuning -13GHz $< \Delta \le -10$GHz. The optical power P is distributed across a Gaussian beam profile, and the central intensity is

$$I_0 = \frac{2P}{\pi w^2} \tag{8.4}$$
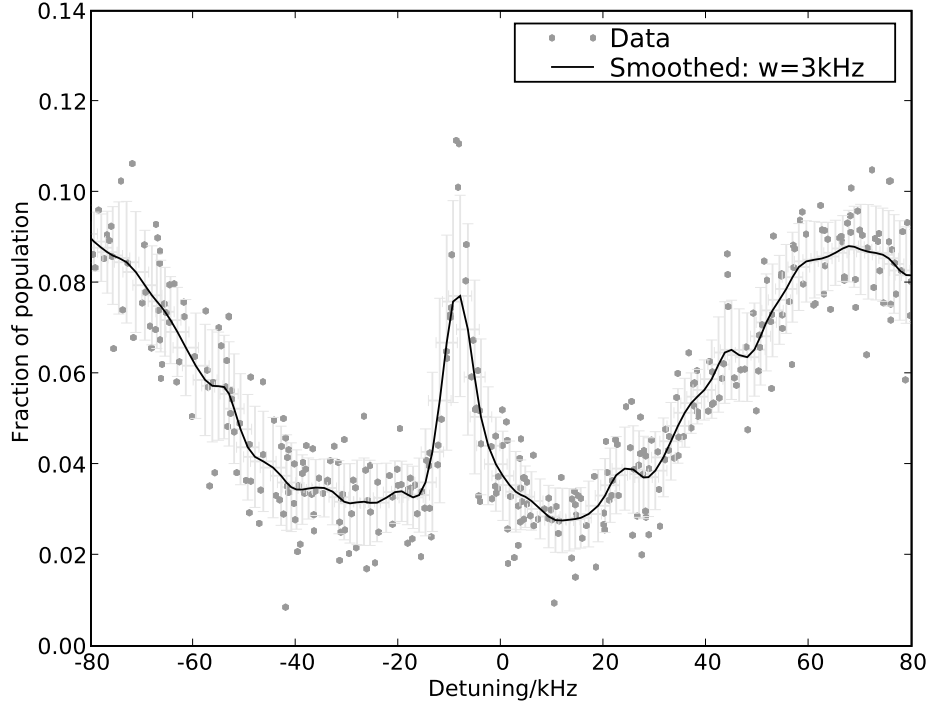
FIGURE 8.4: Narrow $m_F = 0$ Raman resonance, offset from the expected zero detuning by $-(8.5 \pm 0.5)$kHz and with a FWHM ~ 10kHz. Part of the broader $m_F \neq 0$ peaks are visible beyond ± 50kHz.

where w is the $1/e^2$ *radius* of the beam [146]. By tuning the Tiger to resonance, we can ensure the Raman beams overlap with the cloud, but it is difficult to ensure this alignment to better than 1mm. By using a large waist, we can ensure all atoms in the cloud are subject to the same intensity, even with this imperfect alignment. The intensity away from the centre falls as $I = I_0 e^{-(x/w)^2}$ and since $\partial I/\partial x|_{x=0} = 0$ and $\partial I/\partial w|_{w=0} = 0$, the combined uncertainty in intensity is just the uncertainty in the central intensity for *small* uncertainties in alignment. The intensity in each component beam is thus

$$\begin{aligned} I_{P,S} &= 2 \times (50 \pm 5)\text{mW}/\left(\pi \times [(5 \pm 2)\text{mm}/2]^2\right) \\ &= (510 \pm 410)\text{mW}\,\text{cm}^{-2}. \end{aligned}$$

The effective Rabi frequency, from Equation 3.38, is

$$\Omega_R = 2\pi \times (84 \pm 70)\text{kHz} \qquad (8.5)$$

The large uncertainty is due almost entirely to the uncertainty in the beam waist and resulting poorly known intensity. The power broadened width of the peak suggests a Rabi frequency at the lower end of this prediction, and perhaps that the assumption that the misalignment at the cloud was a small fraction of the beam waist is not valid. We use a pulse length $\tau = 500\mu$s and so we can expect, even in the lower frequency limit, to drive many Rabi cycles.

The results of this first narrow scan are shown in Figure 8.4. The plot is symmetrical about a clear central peak, which we identify as the Raman transition between $m_F = 0$ states. Much broader features, which we attribute to the $m_F \neq 0$ states, are visible at the edges of plot. The maximum relative fluorescence for this central peak is around 0.1. If we assume we prepare the population equally amongst the 5 m_F levels of the lower hyperfine ground state, we can expect an amplitude of 0.2, but by driving a large and uncertain number of oscillations, we might expect to only see the average of this maximum within our noisy data. There is therefore good agreement, to within our limited experimental accuracy, between the measured and the theoretical maximum peak heights.

The peak centre is offset from the expected zero detuning; we are confident of the EOM drive frequency (discussed in Section 6.3.1) and attribute this to the lightshift, which is discussed in the following section. The measured shift is (-8.5 ± 0.5)kHz and, for our detuning and intensities, the predicted lightshift is (10 ± 3)kHz.

8.3.2 Light Shifts

When the field from our laser is applied to the atoms in the trap, it is no longer sufficient to consider either in isolation; the dipole interaction couples the two quantum systems. For off-resonant coupling, where absorption is small, the main effect of the interaction is to shift the energy levels of the atom. There are various ways of looking at this effect, known as the light shift. One may consider the full multi-level Hamiltonian, composed of bare-atom and interaction parts, and observe that the eigenvalues are different when the coupling is present [147]; alternatively, when the magnitude of the shift is small, it is valid to take the light-shift found for a two-level system coupled by a single frequency, apply
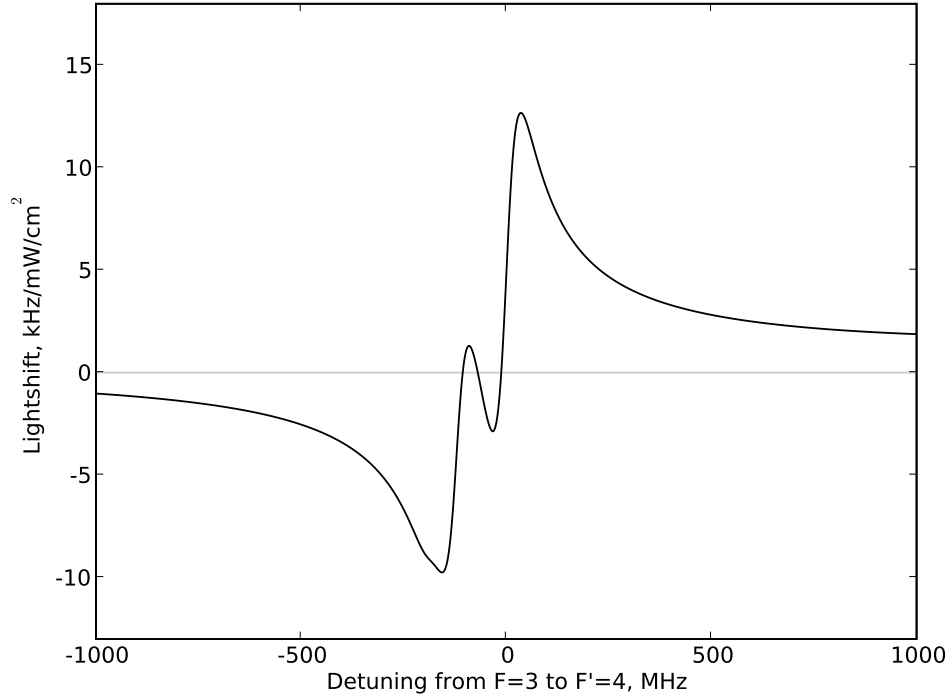
FIGURE 8.5: The lightshift induced between rubidium 85 hyperfine ground states by a single field, with detuning measured relative to the $|F = 2, m_F = 0\rangle$ to $|F' = 4, m'_F = 0\rangle$ transition. In the limit where detuning is greater than the Rabi frequency, the shift depends linearly on frequency. A similar dependence has been observed experimentally [149].

this perturbation to each pair of levels and, if there is more than one, each interacting field, and sum to find the overall light-shift [148].

This perturbative method can be applied to our system of cold (Doppler broadening neglected) rubidium interacting with the numerous frequency components produced by our manipulation laser apparatus. For detunings much larger than the Rabi frequency, the light shift depends linearly on intensity; a plot of light shift for a single applied field, assuming this linear dependence, is shown in Figure 8.5. Our manipulation laser apparatus applies four main frequency components, and the light shift including these EOM sidebands and AOM shifted beam, with the typical power ratio of 15% of the carrier, is shown in Figure 8.6; the AOM beam here has positive frequency shift. The calculated shifts are invalid near to single-photon resonances, but it is these regions we must avoid if we are to avoid significantly perturbing the atomic energy levels.
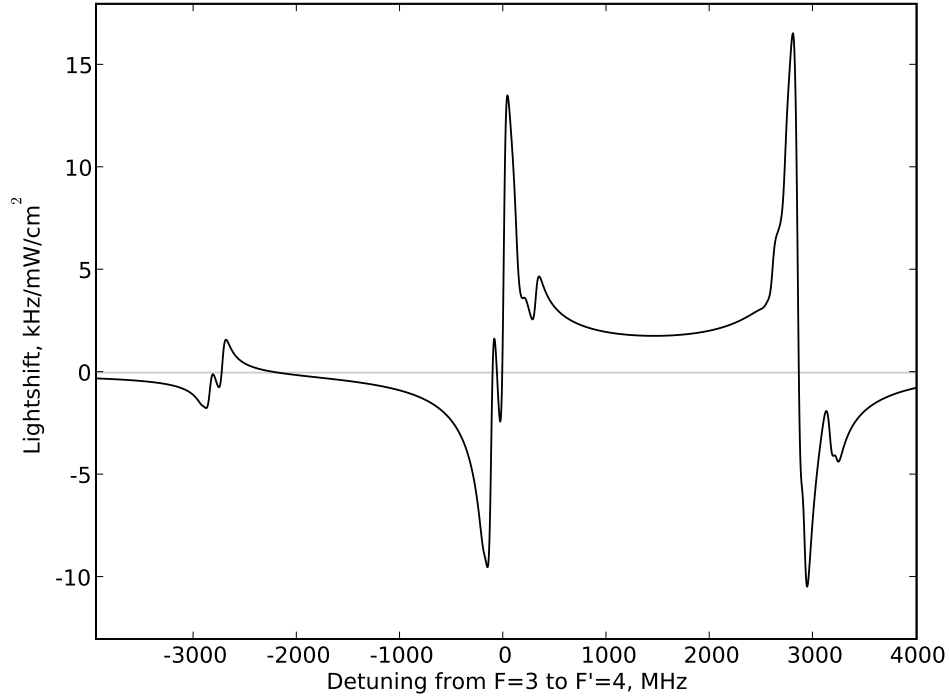
FIGURE 8.6: The total lightshift induced between rubidium 85 hyperfine ground states
for the fields produced by our manipulation laser apparatus. The EOM sidebands and
AOM beam are shifted relative to the carrier, and each derived component has 15% of
the carrier's power. Here, the AOM shift is positive, and overlaps with the EOM lower
sideband around a detuning of +3000MHz. The details are calculable, but in general,
we wish to avoid such detunings.

Clearly we wish to avoid regions where the light shift varies rapidly; the plateaux are

also undesirable but it may be necessary to use these to maximise the two photon Rabi

frequency (see Section 3.4). Unless we are to detune by many GHz, it is therefore

important to know the frequency of the main manipulation laser to within a few MHz,

more details of which are described in Section 6.2.1.

8.3.3   Light-shifted Resonance

To confirm that the shift observed in Section 8.3.1 is indeed due to the light shift, we

would like to measure the position of this central peak for various Raman detunings, but

unfortunately we do not yet have sufficient control over the Tiger wavelength to perform

this characterisation systematically. However, we can compare this small lightshift with

that which we observe when the Mach-Zender (see Section 6.5) is not used to remove the

carrier. This plot is shown in Figure 8.7; the observed lightshift is $-(460 \pm 40)$kHz. The Tiger wavelength $\lambda = (780.2505 \pm 0.0002)$nm places the lower sideband near to single-photon resonance, greatly increasing the lightshift, but also making this shift strongly dependent on the precise laser frequency.

The laser detuning from $\left|5^2\text{S}_{1/2}, 3, 0\right\rangle$ to $\left|5^2\text{P}_{3/2}, 4, 0\right\rangle$ was $\Delta = -2\pi(3.35 \pm 0.10)$GHz. Although we can measure the total power (170mW) and the relative power in each frequency component (Carrier:EOM:AOM $= 20 : 3 : 3$) relatively accurately, for reasons discussed in Section 8.3.1 we do not know the intensity accurately. The beam has a relatively large diameter (5 ± 2)mm, and we estimate the intensity at the centre as $I_0 = (1730 \pm 1400)\text{mW}\,\text{cm}^{-2}$. The lightshift at this detuning, with these power ratios, is $-(0.5 \pm 0.1)\text{kHz}/\left(\text{mW}\,\text{cm}^{-2}\right)$, and hence we calculate a lightshift to be $-(870 \pm 720)$kHz. The measured shift is within this uncertainty, and is less than the estimate. This discrepancy is consistent with misalignment of the beams, and we would have to be far more sure of the actual intensity falling on the atoms before we could make a meaningful comparison between theory and prediction.

8.3.4 Polarisation

The theory described in Section 3.5 and by which we understand the Raman process in our atom predicts that, for equal linear polarisations, Raman transitions will not be driven. Using identical parameters to those in the previous section, we performed a frequency scan with a comparable number of points and over the same region. The only change was the removal of a quarter-wave plate, which converted the linear polarisation of the BoosTA to circular polarisation. The scan, shown in Figure 8.8, shows no evidence of Raman transitions. We intend to perform this experiment with the Mach-Zender interferometer so as to reduce the lightshift and the noise from single photon pumping. However, this device is not yet fully reliable, and the best way to ensure the conditions were identical for both runs was to not use it.
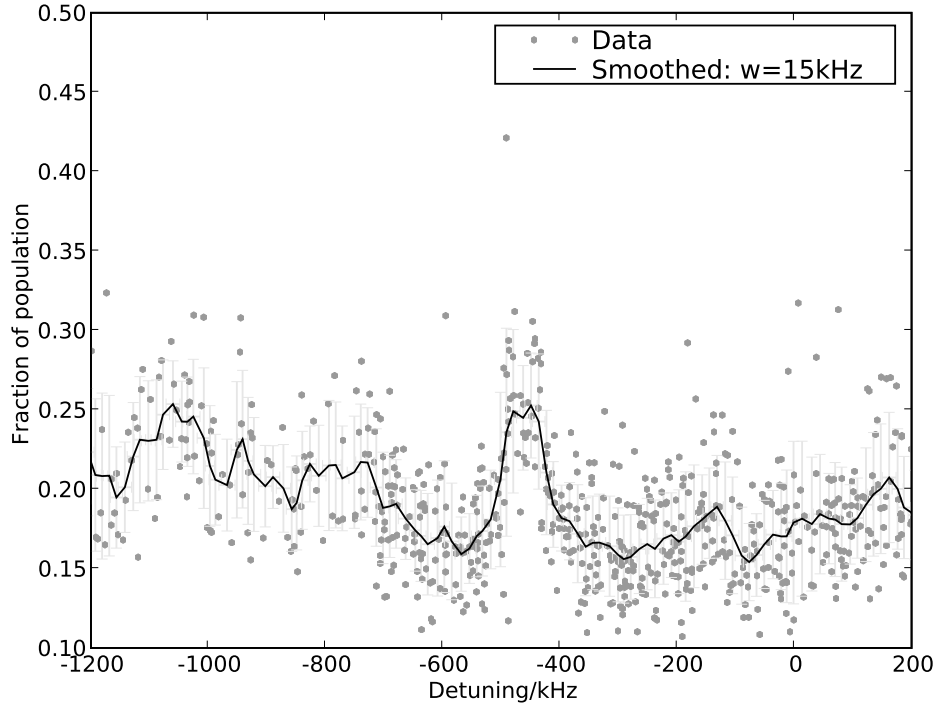
FIGURE 8.7: A Raman frequency scan conducted without the Mach-Zender interferometer. A significantly light-shifted $m_F = 0$ resonance is visible, with $m_F = -1$ and $m_F = -2$ resonances at lower frequency. The data are significantly more noisy, and offset from zero, because of optical pumping by the carrier; the smoothed line may lack some features of the underlying data, such as a peak around -800kHz. The $m_F = +1$ and $m_F = +2$ resonances are apparently not visible, or at least less than their negative counterparts, but the noise is such that we cannot confidently claim they are absent.

8.3.5   Power Broadened Resonances

By choosing a small laser detuning and ensuring the carrier is well separated by the Mach-Zender interferometer, we can maximise the Rabi frequency while keeping the light shift small. The result of such a scan is shown in Figure 8.9; a bias field was applied to lift the Zeeman degeneracy and five peaks are clearly visible. There is unfortunately significant noise due to single photon optical pumping, but this is an unavoidable consequence of decreasing the laser detuning; ideally we would like to increase the intensity and keep the detuning large. Nevertheless, we observe that the central peak which, unlike the others, is not affected by spatially or temporally varying magnetic fields, but which is still broader than that seen previously in Figure 8.4.

The pulse length of $500\mu s$ has not been changed, and so we attribute this broadening
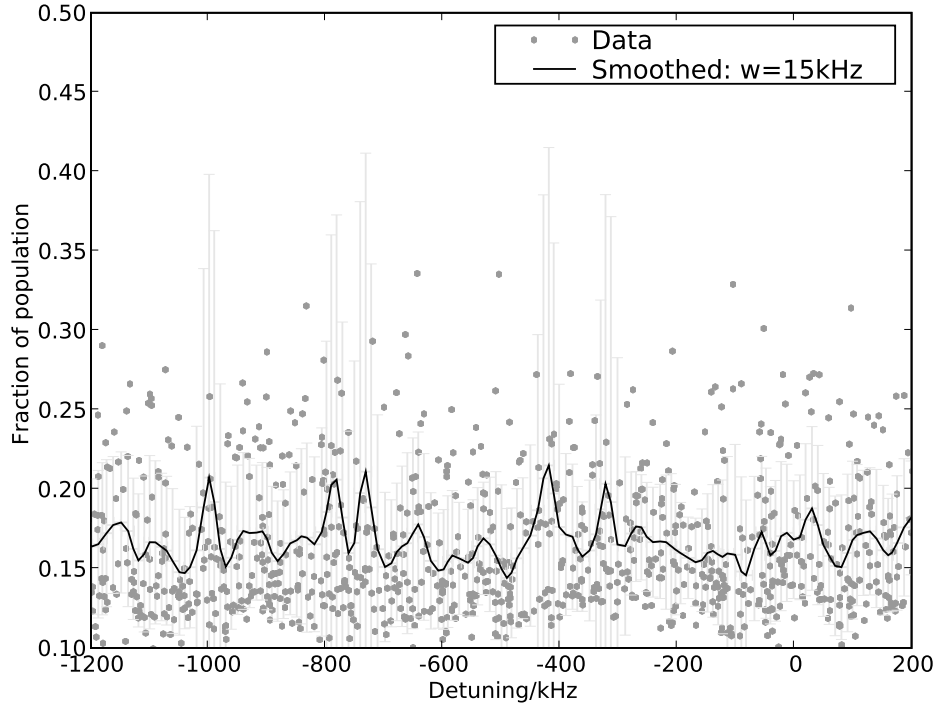
FIGURE 8.8: A scan with identical conditions to those for Figure 8.7 but with equal linear polarisations π_+,π_+. This is not expected to drive Raman transitions, and the clear central peak and distinct increase associated with $m_F < 0$ transitions, seen previously, are absent.

to power broadening. This mechanism can be understood in the Bloch vector picture described in Section 4.2. The depth of modulation depends not on the detuning, but on the ratio of detuning to coupling strength. Therefore, for a larger coupling strength, the modulation depth remains significant for larger detunings. The Tiger wavelength was $\lambda = (780.249 \pm 0.001)$nm, which is a detuning of $-(2.6 \pm 0.6)$GHz and places the upper EOM sideband near to single-photon resonances, and which accounts for the observed background optical pumping. The calculated lightshift in this region is positive, which agrees with the observation, but the magnitude varies quickly in this region, and the shift is $+(0.9\pm0.3)$kHz/mW/cm^2. We also have reason to doubt the exact results of this prediction so near to a resonance, as discussed in the previous section. The intensity is the same as in Section 8.3.1 $I = (510\pm410)$mW cm^{-2}, and we calculate the lightshift to be (460 ± 400)kHz. The results do not disagree, and we again attribute the over estimate to possible misalignment.
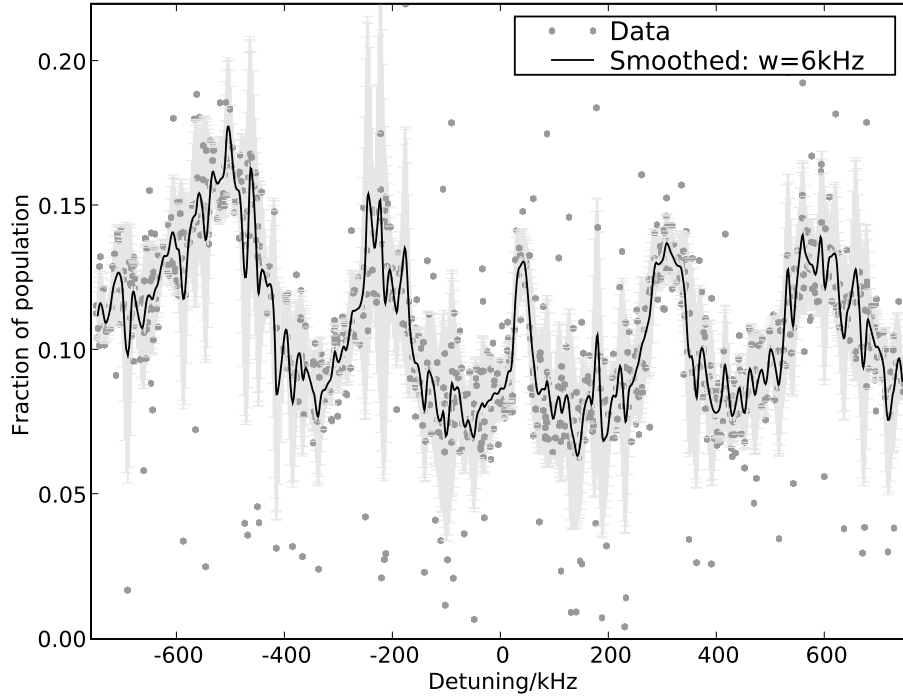
FIGURE 8.9: Five peaks corresponding to each of the $\left|5^2\mathrm{S}_{1/2}, F = 2\right\rangle$ Zeeman-split sublevels. The central $m_F = 0$ peak at $+(38 \pm 2)$kHz is power broadened to a FWHM of (40 ± 5)kHz.

We can also calculate the Rabi frequency: $\Omega_R = (410 \pm 350)$kHz. The width of the $m_F = 0$ transition in Figure 8.9 suggests that the transition is driven at 40kHz, which is just outside the lower bound of our calculation. Significantly higher Rabi frequencies have been reported [134], typically by focusing the beam to a radius smaller than our cloud.

The Mach-Zender interferometer, despite being locked, is still prone to long term drift. Such drifts affect the overall intensity as well as the ratio of powers. The numbers obtained from the theory are approximately correct, and we believe it would be prudent to ensure we have accurate experimental values to use to before we attempt to model the system in more detail. We can, at least, be confident that the increase in the width of the $m_F = 0$ transition is due to power broadening.
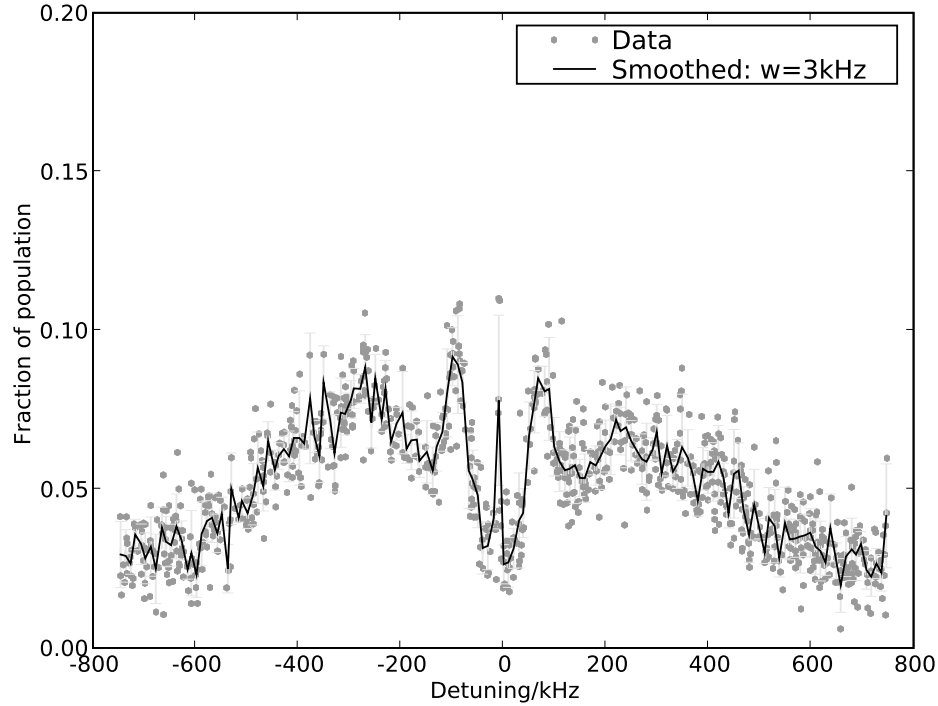
FIGURE 8.10: Raman frequency scan for no applied magnetic field. Features corresponding to each of the 5 possible transition are visible, including the very narrow (and hence under-sampled) $m_F = 0$ transition. The $m_F = \pm 1$ resonances are at $-(90 \pm 10)$kHz and $+(80 \pm 10)$kHz with a FWHM 40kHz. The $m_F = \pm 2$ resonances are at $-(270 \pm 20)$kHz and $+(250 \pm 20)$kHz with FWHM 130kHz.

8.3.6 Zeeman Shifted Resonances

As discussed in Section 3.6, we expect that, in the absence of an external magnetic field, atoms should be in zero magnetic field during the Raman operation and hence, as discussed in Section 3.6, the $m_F = 0$ transition will not be driven. We also expect that the $m_F \neq 0$ transitions will be degenerate, and will appear as a single peak. The results of a scan with no applied magnetic field are shown in Figure 8.10. Unexpectedly, there are 5 features visible, corresponding to $m_F = \pm 2$, $m_F = \pm 1$, and $m_F = 0$. The central peak is under-sampled by virtue of being far narrower than the other peaks, but its presence is confirmed by not just the high points, but the absence of low points around the centre; the region around zero detuning is magnified in Figure 8.11.

From the frequency offset of the peaks, we conclude that despite our best efforts at alignment, the intensities of the counter-propagating MOT beams are not equal and, as
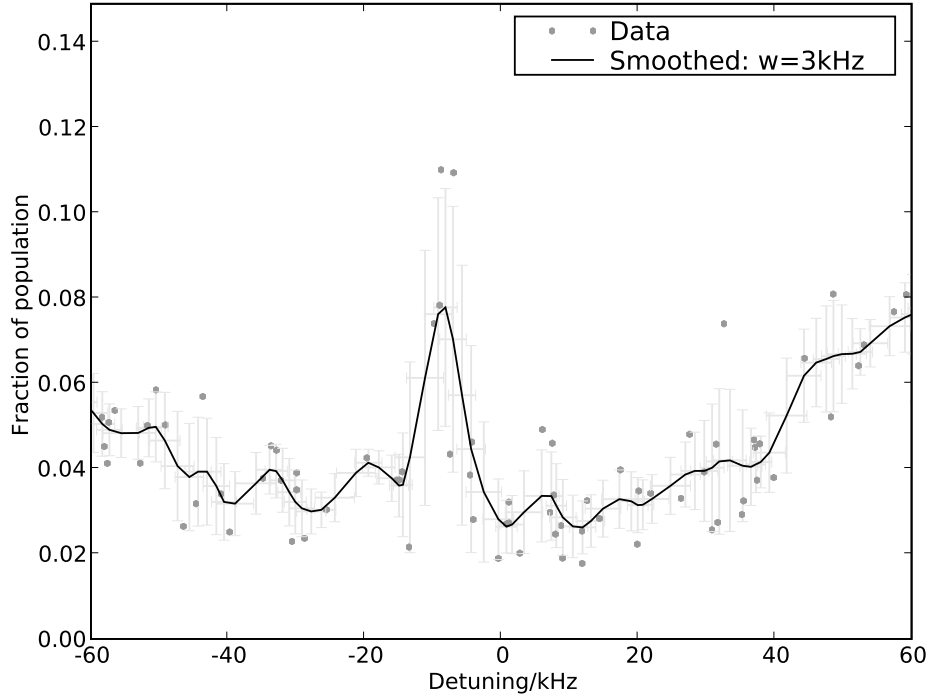
FIGURE 8.11: The region around zero detuning showing from Figure 8.10 showing the peak due to $m_F = 0$.

discussed in Section 3.6.2, the cloud forms around a point of non-zero magnetic field. We can also conclude, since peaks corresponding to m_F levels of opposite sign do not overlap, that the spatial extent of the cloud does not include the magnetic field zero. From the FWHM of the $m_F \neq 0$ peaks, we can estimate the spatial spread of the cloud. Using the horizontal field gradient 2.5G/cm, and the FWHM 130kHz for transitions between the $m_F = \pm 2$ states, we estimate the FWHM of the spatial distribution as $(280 \pm 50)\mu$m; this is approximately the same as that measured directly in Section 7.1.

We are able to control the externally applied magnetic field, and by running large frequency scans for different magnitudes, we can observe the frequency shift and confirm its dependence. Frequency scans for three applied fields are shown in Figures 8.12, 8.13, and 8.14.

The peaks, each of which can be easily associated with an m_F level, move to higher frequency shift for higher applied fields, as expected. The measured positions of these peaks are shown in Figure 8.15. The magnitude of the applied magnetic field is not well
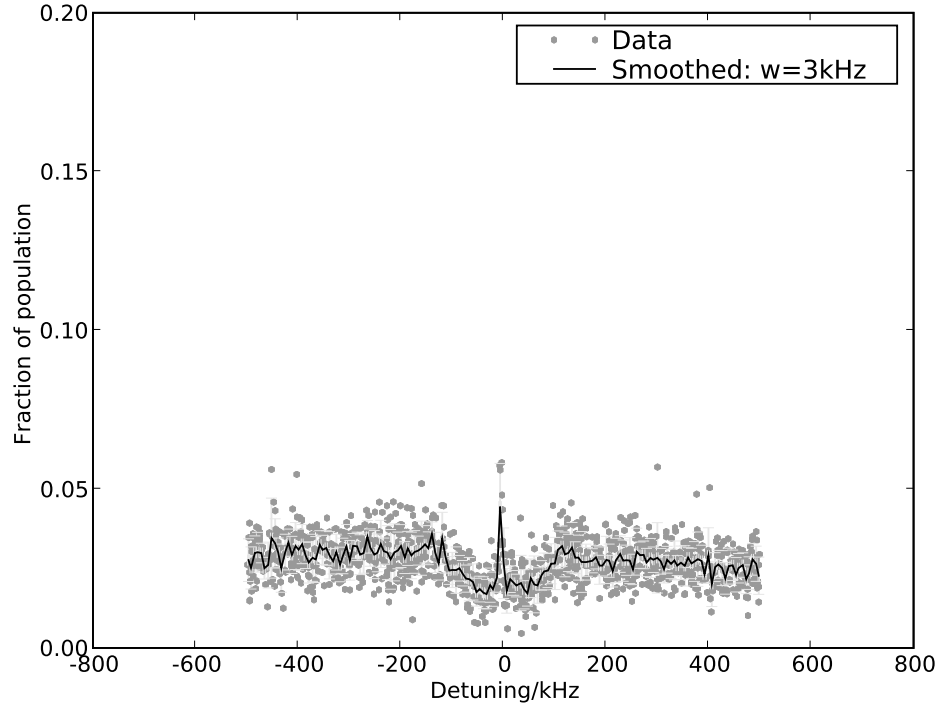
FIGURE 8.12: Applied voltage of $V_{\text{in}} = 200\text{mV}$ with a clear peak for $m_F = 0$, but otherwise indistinct features. The inferred magnetic field is $B = 0.39\text{G}$.

known, but by plotting the frequencies as a function of the applied voltage, on which the field linearly depends, we find three straight lines which intercept at a non-zero voltage $-V_0 = 351\text{mV}$, corresponding to the offset due to displacement of the cloud, and a non-zero frequency -6.5kHz, due to the light-shift. The measured frequency dependence on voltage is 336kHz/V which, if we use the known Zeeman dependence of 470kHz/G, allows us to estimate the offset magnetic field in which the cloud sits. The frequency offset is 118kHz, and the field is $B_0 = 0.25\text{G}$. Thus the field is $B(V) = 0.71\text{G/V} \times V + B_0$

The magnetic field offset implies, for our field gradient (see Section 5.3.1), that the cloud is displaced by $\simeq 1\text{mm}$; for our detuning ($\Delta = -3\Gamma$), this corresponds to a difference in intensities of $< 3\%$. (See Section 3.6.2.) This small imperfection can be accounted for by scatter from various components, or light being reflected from vacuum chamber windows (before the quarter waveplates) and hence some light of the wrong circular polarisation hitting the cloud.
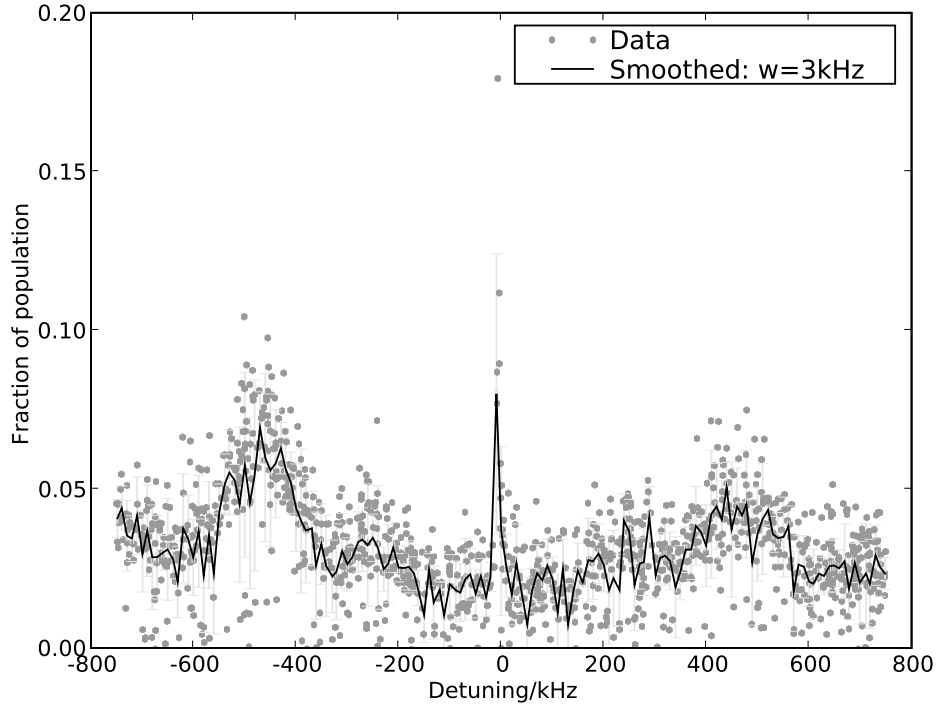
FIGURE 8.13: Applied voltage of $V_{\text{in}} = 400$mV with centre peak at $-(6 \pm 2)$kHz, $m_F = \pm 1$ at $-(260 \pm 20)$kHz and $+(250 \pm 20)$kHz, and $m_F = \pm 2$ at $-(470 \pm 40)$kHz and $+(460 \pm 40)$kHz. The inferred magnetic field is $B = 0.54$G.

8.4 Improved State Preparation

The prepare operation described in Section 8.2.1 and used in all experiments to date uses the trapping beam with polarisations constrained by the requirements of a MOT. This operation distributes the population roughly equally across the m_F sub-levels of $\left|5^2\text{S}_{1/2}, F = 2\right\rangle$, but we would far prefer to concentrate the majority of the population in a single m_F state and hence address a larger fraction of the population with our Raman pulses. This would be a convenience for the current suite of experiments where we observe the fluorescence, but is almost a requirement if we are to observe the results of attempts to manipulate the momentum of the cloud.

We are unable to address the individual m_F levels by frequency because the Zeeman-splitting is far less than the natural linewidth of the states. We can however use the different coupling strengths for various polarisations to design more useful preparation
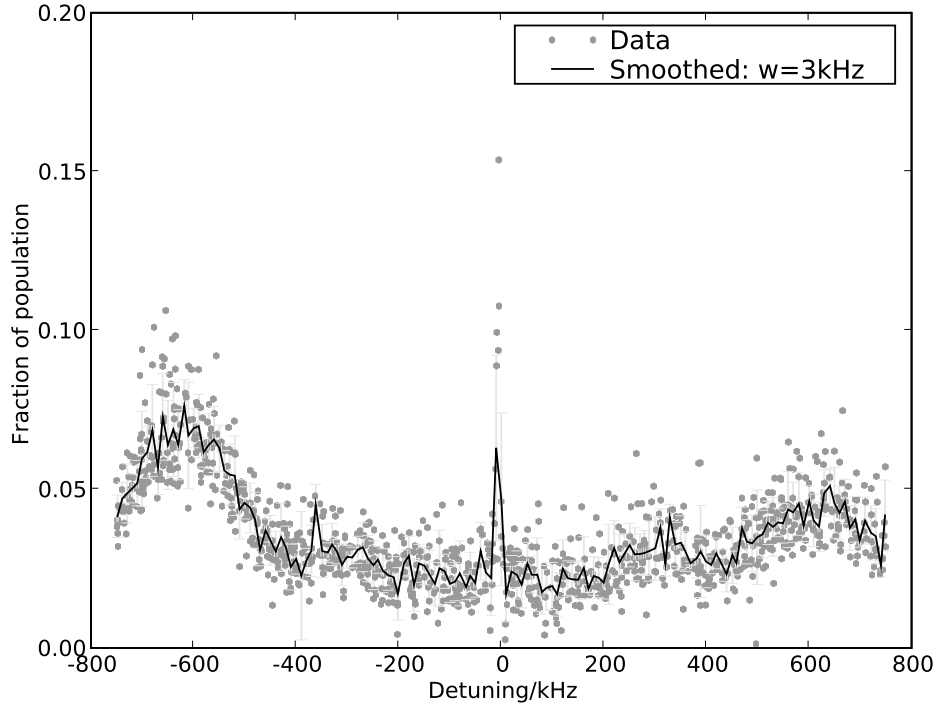
FIGURE 8.14: Applied voltage of $V_{\text{in}} = 600$mV with centre peak at $-(5 \pm 2)$kHz, $m_F = \pm 1$ at $-(310 \pm 20)$kHz and $+(300 \pm 20)$kHz, and $m_F = \pm 2$ at $-(630 \pm 40)$kHz and $+(620 \pm 40)$kHz. The inferred magnetic field is $B = 0.68$G.

schemes if we use separate and independently controlled optical paths for either the trapping or the repump beams.

A particularly useful feature of the coupling strengths is a zero for the transition $\left|5^2\text{S}_{1/2}, 3, 0\right\rangle$ to $\left|5^2\text{P}_{3/2}, 3, 0\right\rangle$ for π polarisation. If, therefore, we use *both* the repump and a frequency shifted (perhaps using an AOM) version of the trapping beam now resonant with $\left|5^2\text{S}_{1/2}, 3\right\rangle \rightarrow \left|5^2\text{P}_{3/2}, 3\right\rangle$, both with π polarisation, we optically pump populations in both $F = 2$ and $F = 3$ in all m_F sub-levels *except* $\left|F = 3, m_F = 0\right\rangle$ and hence it is here that population will accumulate. Using 1mW per beam, retro-reflected to ensure the average momentum absorbed is zero, in 7.1mm beams as previously, numerical simulation suggests we can prepare 98% of the population in $\left|F = 3, m_F = 0\right\rangle$ with a half-life 1.7μs. This is considerably simpler than some other spin-polarising processes [83], and while it does introduce heating, unlike some which simultaneously cool [150, 151], it relies only on the selection rule arising from the requirement that at least one of the final or initial states have non-zero angular momentum [152, Complement A_{XIII}], and
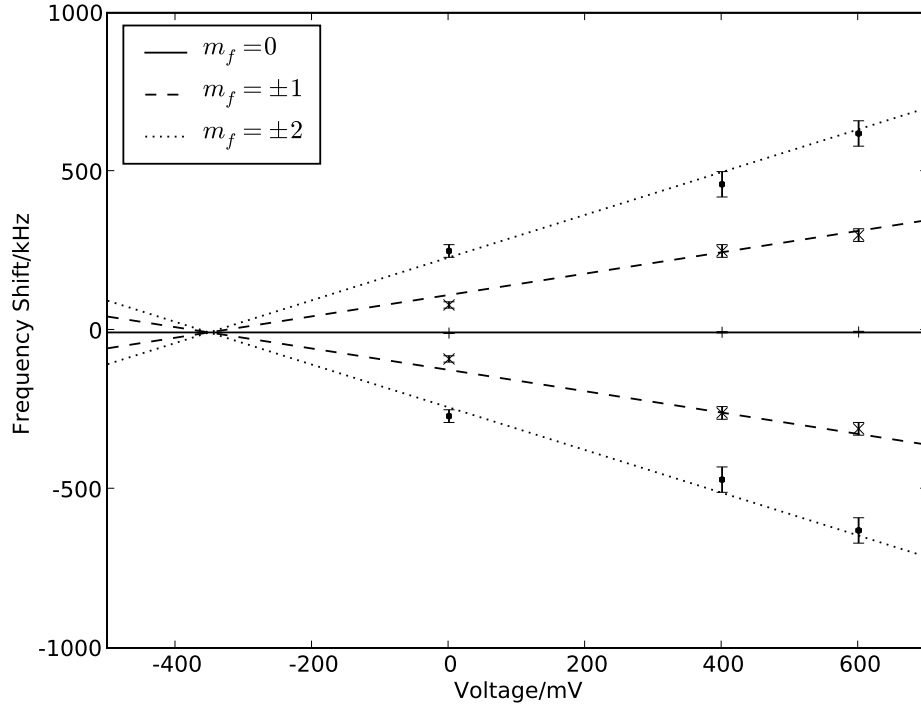
FIGURE 8.15: The energy level shift of the possible $\Delta m_F = 0$ Raman transitions as a function of the voltage applied to the audio amplifier, on which the applied magnetic field linearly depends. The linear fits have gradients proportional to the m_F level, and intercept at non-zero voltage and non-zero frequency, the origins of which are discussed in the text.

hence applies to all half-integer nuclear spin alkali metals [39, Appendix D]. A similar scheme using caesium has shown at least 60% of the population can be prepared in this way [81]. An alternative is to use the Raman transitions for the prepare operation [153].

Chapter 9

Conclusions

Throughout this thesis we have covered many of the steps necessary to demonstrate the coherent cooling schemes described in the introduction. We have constructed a magneto-optical trap, which provides cold atoms, and developed programmable apparatus to create the fields necessary to manipulate the effective two-level system provided by the ground hyperfine states of these cold atoms.

We have shown that we are able to drive Raman transitions between the hyperfine ground states of these atoms, which are cooled and collected by the magneto-optical trap, but which are free to move during the manipulation. Unlike atoms held in dipole traps or otherwise confined, we can hope to coherently manipulate the momentum of these free particles. We have shown that by applying an oscillating magnetic field and performing coherent operations around the maximum of this oscillation we can create a steady-state magneto-optical trap yet lift the Zeeman degeneracy sufficiently to address individual sub-levels.

The Raman fields are derived from a single narrow linewidth laser using a variety of acousto- and electro-optical modulators. We have characterised the components of this apparatus and have demonstrated that we can successfully implement simple pulse sequences; we foresee no barriers in extending these to more intricate modulation schemes. The derived fields were measured and confirmed to remain coherent for the duration of envisaged pulse sequences.

We have observed the lightshift of the narrow Zeeman insensitive $m_F = 0$ transition, and we have developed a model which allows us to calculate this shift. The model appears valid in the limit of far detuning, and we identify the need to more accurately measure and control the Raman laser detuning if we are to reliably tune to this resonance. We are confident that, with improved control over the Raman laser detuning, we will be able to reproducibly tune to the $m_F = 0$ resonance and hence begin to drive Rabi oscillations and to explore interferometric pulse sequences.

The prepare operation currently distributes population roughly equally across the Zeeman sublevels of the lower hyperfine state, and so a pulse can only influence one fifth of the population. The current suite of experiments would be greatly improved if we were able to prepare the majority of the population in the state $m_F = 0$, and a scheme to achieve this was proposed. The improved state preparation appears essential for future experiments where we hope to manipulate the momentum of the cloud.

Theoretical study of the various routes via which a Raman transition between hyperfine ground states can be driven has enabled us to predict the effective Rabi frequency for this effective two-level atom. Unexpectedly, the model also predicts that the ratio of coupling strengths will limit fraction of population which can be transferred. Evidence for this effect was found in previously published work by the Meschede group, and although our results are significantly more noisy, we do note a consistent weighting of our Raman peaks in favour of low m_F states; this is consistent with the predictions listed in Appendix A. A significant step in this method was finding the eigenvectors of the Hamiltonian numerically. We intend to formulate an analytical approach where any number of upper levels decouple from the two ground states in the limit of far detuning, and we expect that this may provide some physical insight.

9.1 Immediate Improvements

We note several difficulties found in our attempts to drive Raman transitions compared with what has been demonstrated in other implementations. The problem of controlling the bias field has been a thread running throughout this thesis, but the spatial extent of

the cloud is also a concern. Unlike atoms held in a dipole trap [81, 132–134] or ions in ion traps [133, 138–140, 154, 155], our atoms are free, and move ballistically during the prepare, manipulation, and read out pulses. Additionally, our particles are distributed over several hundred micrometers, and if these are all to experience the same field, the intensity of the our Raman beams must not vary significantly over this length scale. We have used an approximately Gaussian beam with a waist larger than the cloud, but this has left us with significantly less intensity, and hence a lower Rabi frequency, than would be possible if the particles were more localised. We must also find a method to ensure better alignment of the Raman beams on to the cloud, and take more care to accurately measure the beam waist. A more mundane problem is that we have chosen a large vacuum chamber, and the solid angle subtended by our collection optics is relatively small.

As a first improvement we can use a Toptica FiberDock, which is a fibre coupler specifically designed for the BoosTA. Using this we can expect approximately 700mW into a single mode fibre and for a 2.2mm output coupler an intensity $I \sim 9000\mathrm{mW\,cm^{-2}}$ per frequency component. The typical detuning frequency of 5GHz suggests an effective Rabi frequency of 450kHz. This is significantly larger than we have achieved so far, and in the few milliseconds before atoms leave the trapping region we can expect to drive hundreds of oscillations. This will allow us to perform simple coherence experiments—such as Rabi oscillations and Ramsey interferometry—but is still short of the requirements to fully explore the adiabatic regime. A more exotic technique is to engineer a top-hat beam profile so that we concentrate more of the power on the cloud while ensuring intensity does not vary significantly. Using a phase-only spatial light-modulator and a micron sized waveplate we can map a computer controlled phase profile to an intensity profile at the cloud [156], and so, in software, optimise the overlap of the Raman beams with the cloud.

We must also improve the Mach-Zender electronics so that the lock is reliable, and it may be necessary to interrupt the experiments while the Electro-Optical Modulator is repaired by the manufacturer. The Raman detuning must also be accurately measured

and controlled and we intend to implement the approach discussed in Section 6.2.1, using existing equipment.

9.2   Towards Coherent Manipulation of Momentum

Despite clear evidence that we are able to drive the Raman transition between the long-lived hyperfine ground-states, it is not necessarily a trivial step to drive Rabi oscillations. The manipulation laser apparatus is constructed so that despite being derived from a laser with a relatively large linewidth, the *difference* between the Raman beams should remain coherent for much longer. This appears to be confirmed by beat-frequency measurements described in Section 6.3.6, but we must observe Rabi oscillations before we can be certain.

Significant challenges remain if we are to coherently manipulate the momentum of the atoms and then observe this change. As well as addressing the issues discussed previously, particularly concentrating population in $m_F = 0$, we must improve the efficiency of the manipulation beam apparatus so that we have sufficient power to perform the operations *without* amplification by the BoosTA; this is discussed in Section 6.6. Switching times and extinction ratios of the Pockels cell device have been shown to be sufficient. We are still optimistic that with careful control of the Tiger frequency and improved phase modulation by the EOM we can perform at least simple interferometric sequences to affect the momentum of atoms in the cloud.

We must also observe the momentum distribution of the manipulated atoms. The release and recapture method described in Section 7.3 is insufficient because the distribution is no longer fully described by a Gaussian width. We hope to measure the momentum by allowing the atoms to expand ballistically for a known time before turning on a laser and taking a flash-photograph of their position. We have a fast camera (Prosilica GigE GE680) which has sufficient resolution (640x480), can take sufficiently short exposures ($\geq 25\mu s$), and, crucially, can be triggered externally in the same way as the signal generators used to create the pulses. Photographs have been taken in this manner, but we have so far used only one beam for the flash, which limits the 'exposure' time to

$\sim 50\mu$s as discussed in Section 8.2.1, and the signal to noise ratio has been small. We may also be able to measure the momentum distribution directly using recoil-induced resonance spectroscopy [157]. In either case, improved state preparation is essential as unless the population is concentrated in one m_F state the signal will be contaminated by atoms in the unaffected m_F states and the momentum distribution of those which are affected will be far harder to discern.

9.3 Cooling with the Dipole Force

As a final remark we return to the original motivation for this experiment—to cool species which cannot be cooled by traditional techniques—and consider a radically different approach. Although this work has concerned the scattering force, a particle also interacts with light via the dipole force. If a particle is seen as a refracting glass sphere, the down-stream light will be brought to a focus, and a second particle near this focus will be bound to the first [158]. By placing a mirror one half of a focal length down-stream, we bind the first particle not to a second, but to a *time-delayed* reflection of itself. The time-delay is crucial, as the force need no longer be conservative. Decay is now via leakage of light from cavity modes, not spontaneous emission of the atom [159, 160]; a two-level system is no longer required and the Doppler-limit is no longer relevant. This 'cavity cooling' has been demonstrated [161], and a large ESF funded collaboration is underway to study this effect theoretically and experimentally.

Appendix A

Raman Transition Coefficients

The following tables list the coupling we can expect for Raman transitions driven via the hyperfine levels of $5^2\mathrm{P}_{3/2}$ in Rubidium-85-like atoms, with various polarisation combinations. There are typically two sub-levels in $5^2\mathrm{P}_{3/2}$ via which the Raman transition can be driven, and in the far-detuned limit, these behave as one effective level. This section extends Section 3.3, which describes these parameters and the numerical methods used to obtain them.

In summary, $\left|5^2\mathrm{S}_{1/2}, 2\right\rangle$ is coupled to $\left|5^2\mathrm{P}_{3/2}\right\rangle$ via an electric field E_P of frequency ω_P and coupling strength $\hbar\Omega_P = \langle J||\hat{\mu}||J'\rangle E_P$, and similarly for $\left|5^2\mathrm{S}_{1/2}, 3\right\rangle$ via field E_S. We assume these two coupling strength are equal $\Omega_P = \Omega_S = \Omega$, and the effective Rabi-frequency in the limit of far detuning for this coupling between sub-levels in $\left|5^2\mathrm{S}_{1/2}, 2\right\rangle$ and $\left|5^2\mathrm{S}_{1/2}, 3\right\rangle$ is $\Omega_R = \chi\Omega^2/(2\Delta)$, where the parameter χ is found from the transition strengths between the Zeeman sub-levels.

The different modulation depths, especially the asymmetry, is perhaps unexpected; these calculations may explain the observation of differences in peak heights reported in [81], there attributed to optical pumping.

m_F	M	θ	depth	χ
-2	$\begin{pmatrix} +0.49 & +0.87 \\ -0.87 & +0.49 \end{pmatrix}$	0.338π	0.852	0.124
-1	$\begin{pmatrix} +0.85 & +0.53 \\ -0.53 & +0.85 \end{pmatrix}$	0.176π	0.895	0.176
$+0$	$\begin{pmatrix} +0.82 & +0.57 \\ -0.57 & +0.82 \end{pmatrix}$	0.191π	0.933	0.179
$+1$	$\begin{pmatrix} +0.78 & +0.63 \\ -0.63 & +0.78 \end{pmatrix}$	0.216π	0.977	0.161
$+2$	$\begin{pmatrix} +0.74 & +0.67 \\ -0.67 & +0.74 \end{pmatrix}$	0.235π	0.995	0.125

TABLE A.1: Parameters for transitions between $\left|5^2\mathrm{S}_{1/2}, 2, m_F\right\rangle$ and $\left|5^2\mathrm{S}_{1/2}, 3, m_F + 0\right\rangle$ driven by polarisations (σ_-, σ_-) via upper states $\left|5^2\mathrm{P}_{3/2}, (2,3), m_F - 1\right\rangle$.

m_F	M	θ	depth	χ
-2	$\begin{pmatrix} +0.70 & +0.72 \\ -0.72 & +0.70 \end{pmatrix}$	0.225π	0.999	0.215
-1	$\begin{pmatrix} +0.84 & +0.54 \\ -0.54 & +0.84 \end{pmatrix}$	0.183π	0.911	0.193
$+0$	$\begin{pmatrix} +0.90 & +0.43 \\ -0.43 & +0.90 \end{pmatrix}$	0.143π	0.781	0.174
$+1$	$\begin{pmatrix} +0.92 & +0.38 \\ -0.38 & +0.92 \end{pmatrix}$	0.125π	0.709	0.136
$+2$	$\begin{pmatrix} +0.89 & +0.46 \\ -0.46 & +0.89 \end{pmatrix}$	0.152π	0.815	0.068

TABLE A.2: Parameters for transitions between $\left|5^2\mathrm{S}_{1/2}, 2, m_F\right\rangle$ and $\left|5^2\mathrm{S}_{1/2}, 3, m_F - 1\right\rangle$ driven by polarisations (σ_-, π_0) via upper states $\left|5^2\mathrm{P}_{3/2}, (2,3), m_F - 1\right\rangle$.

m_F	M	θ	depth	χ
-2	$\begin{pmatrix} +0.83 & +0.55 \\ -0.55 & +0.83 \end{pmatrix}$	0.186π	0.919	0.234
-1	$\begin{pmatrix} +0.72 & +0.69 \\ -0.69 & +0.72 \end{pmatrix}$	0.243π	0.999	0.176
$+0$	$\begin{pmatrix} +0.75 & +0.66 \\ -0.66 & +0.75 \end{pmatrix}$	0.228π	0.991	0.137
$+1$	$\begin{pmatrix} +0.76 & +0.65 \\ -0.65 & +0.76 \end{pmatrix}$	0.226π	0.989	0.097
$+2$	$\begin{pmatrix} +0.97 & +0.24 \\ -0.24 & +0.97 \end{pmatrix}$	0.077π	0.464	0.120

TABLE A.3: Parameters for transitions between $\left|5^2\mathrm{S}_{1/2}, 2, m_F\right\rangle$ and $\left|5^2\mathrm{S}_{1/2}, 3, m_F - 1\right\rangle$ driven by polarisations (π_0, σ_+) via upper states $\left|5^2\mathrm{P}_{3/2}, (2,3), m_F + 0\right\rangle$.

m_F	M	θ	depth	χ
-2	$\begin{pmatrix} +0.79 & +0.61 \\ -0.61 & +0.79 \end{pmatrix}$	0.210π	0.968	0.128
-1	$\begin{pmatrix} +0.82 & +0.58 \\ -0.58 & +0.82 \end{pmatrix}$	0.195π	0.942	0.167
$+0$	$\begin{pmatrix} +0.82 & +0.57 \\ -0.57 & +0.82 \end{pmatrix}$	0.192π	0.934	0.179
$+1$	$\begin{pmatrix} +0.82 & +0.58 \\ -0.58 & +0.82 \end{pmatrix}$	0.196π	0.942	0.167
$+2$	$\begin{pmatrix} +0.79 & +0.61 \\ -0.61 & +0.79 \end{pmatrix}$	0.210π	0.968	0.128

TABLE A.4: Parameters for transitions between $\left|5^2\mathrm{S}_{1/2}, 2, m_F\right\rangle$ and $\left|5^2\mathrm{S}_{1/2}, 3, m_F\right\rangle$ driven by polarisations (π_+, π_-) via upper states $\left|5^2\mathrm{P}_{3/2}, (2,3), m_F \pm 1\right\rangle$.

Appendix B

Conference Talks and Posters

In addition to the paper **Fractional adiabatic passage in two-level systems**[85], the work described in this thesis was communicated via the following talks and posters.

Date			Title and venue
8th	July	2008	**New optical method of cooling atoms and molecules**
			The Rank Prize Funds Mini-Symposium
			Pulse or Pull of Optical Momentum
19th	September	2005	**Laser Cooling and Trapping of Molecules**
			ORC Photonics Day
			University of Southampton
1st	February	2008	**Cavity Cooling of Molecules**
6th	July	2007	**The coldest stuff in Southampton**
24th	February	2006	**Coherent Manipulation of Rubidium**
			Quantum Light and Matter Group Seminars
			University of Southampton

Date			Title and venue
			Coherent Manipulation of Rubidium
21st	September	2007	School of Physics and Astronomy
			University of Southampton
			Adiabatic Beam-splitter Pulses
24th	June	2007	18th International Conference on Laser Spectroscopy
			Telluride, Colorado
			Cold atoms near nano-structured surfaces
7th	April	2008	Cold Quantum Matter (EuroQUAM)
			Inauguration Conference, Barcelona
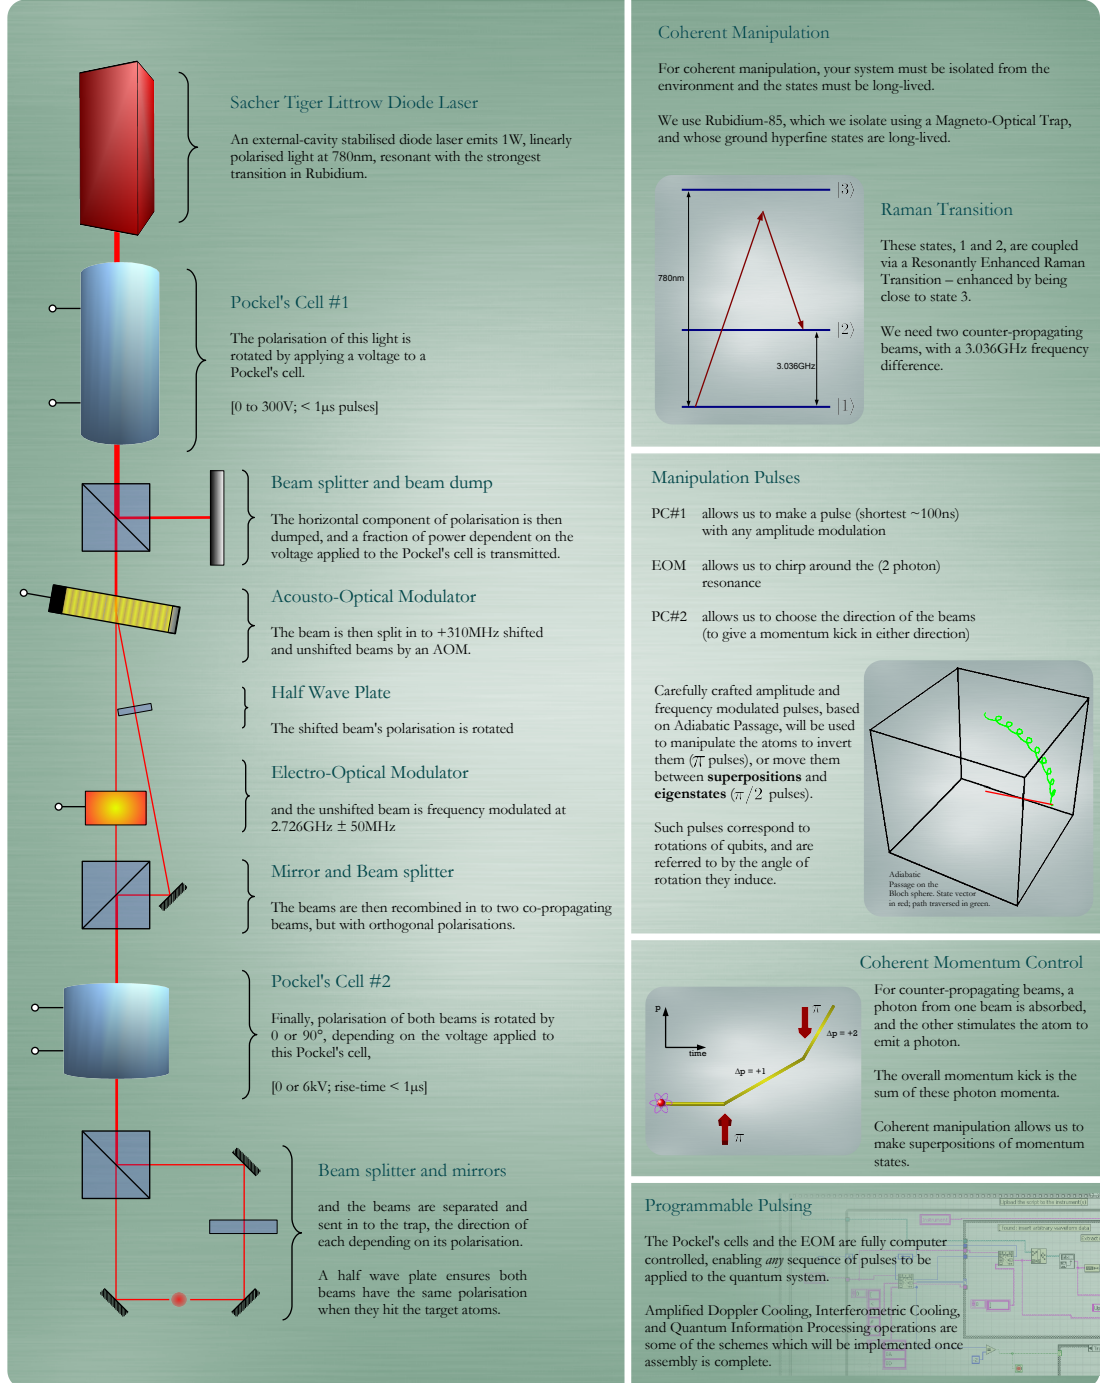
FIGURE B.1: An overview of the experimental apparatus and the planned experiments.

FIGURE B.2: A representation of the adiabatic beam-splitter pulse reported in [85].

FIGURE B.3: An outline of a cooling mechanism using the dipole force and of an experiment (under construction) to investigate the predicted effects.

Bibliography

[1] T. W. Hänsch and A. L. Schawlow. Cooling of gases by laser radiation. *Opt. Comm.*, 13(1):68, 1975.

[2] D. J. Wineland and H. Dehmelt. Proposed $10^{14}\Delta\nu < \nu$ Laser Spectroscopy on Ti Mono-Ion Oscillator III. *Bull. Am. Phys. Soc*, 20:637, 1975.

[3] D. J. Wineland, R Drullinger, and F. L. Walls. Radiation-Pressure Cooling of Bound Resonant Absorbers. *Phys. Rev. Lett.*, 40(25):1639–1642, June 1978. doi: 10.1103/physrevlett.40.1639.

[4] A. Ashkin and J. P. Gordon. Cooling and trapping of atoms by resonance radiation pressure. *Opt. Lett.*, 4(6):161, 1979.

[5] Steven Chu, J. E. Bjorkholm, A. Ashkin, J. P. Gordon, and L. W. Hollberg. Proposal for optically cooling atoms to temperatures of the order of 10^{-6} K. *Opt. Lett.*, 11(2):73, 1986.

[6] Paul D. Lett, William D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. Optical molasses. *J. Opt. Soc. Am. B*, 6(11):2084, 1989.

[7] Steven Chu, L. Hollberg, J. E. Bjorkholm, Alex Cable, and A. Ashkin. Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Phys. Rev. Lett.*, 55:48–51, 1985.

[8] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Phys. Rev. Lett.*, 59:2631–2634, 1987.

[9] 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, July 1988. doi: 10. 1103/physrevlett.61.169.

[10] J. Dalibard and C. N. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6(11): 2023, 1989.

[11] Steven Chu. The manipulation of neutral particles. *Rev. Mod. Phys.*, 70(3):685–706, 1998.

[12] C. N. Cohen-Tannoudji and William D. Phillips. New mechanisms for laser cooling. *Physics Today*, 43:33–40, October 1990.

[13] Mark Kasevich and Steven Chu. Atomic interferometry using stimulated Raman transitions. *Phys. Rev. Lett.*, 67(2):181–184, July 1991. doi: 10.1103/physrevlett. 67.181.

[14] Michael A. Lombardi, T. P. Heavner, and S. R. Jefferts. NIST primary frequency standards and the realization of the SI second. *The Journal of Measurement Science*, 2(4):74–89, December 2007.

[15] W. H. Oskay, S. A. Diddams, E. A. Donley, T. M. Fortier, T. P. Heavner, L. Hollberg, W. M. Itano, S. R. Jefferts, M. J. Delaney, K. Kim, F. Levi, T. E. Parker, and J. C. Bergquist. Single-Atom Optical Clock with High Accuracy. *Phys. Rev. Lett.*, 97(2):020801, 2006. doi: 10.1103/physrevlett.97.020801.

[16] M. H. Anderson, J. R. Ensher, M. R. Matthews, Carl E. Wieman, and E. A. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269:198–201, July 1995. doi: 10.1126/science.269.5221.19.

[17] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein Condensation in a Gas of Sodium Atoms. *Phys. Rev. Lett.*, 75(22):3969–3973, November 1995. doi: 10.1103/physrevlett.75. 3969.

[18] J. T. Bahns, W. C. Stwalley, and Phillip L. Gould. Laser cooling of molecules: A sequential scheme for rotation, translation, and vibration. *J. Chem. Phys.*, 104 (24):9689–9697, 1996. doi: 10.1063/1.471731.

[19] A. Fioretti, D. Comparat, A. Crubellier, O. Dulieu, F. Masnou-Seeuws, and P. Pillet. Formation of Cold $Cs2$ Molecules through Photoassociation. *Phys. Rev. Lett.*, 80(20):4402–4405, May 1998. doi: 10.1103/physrevlett.80.4402.

[20] A. N. Nikolov, E. E. Eyler, X. T. Wang, J. Li, H. Wang, W. C. Stwalley, and Phillip L. Gould. Observation of Ultracold Ground-State Potassium Molecules. *Phys. Rev. Lett.*, 82(4):703–706, January 1999. doi: 10.1103/physrevlett.82.703.

[21] C. Gabbanini, A. Fioretti, A. Lucchesini, S. Gozzini, and M. Mazzoni. Cold Rubidium Molecules Formed in a Magneto-Optical Trap. *Phys. Rev. Lett.*, 84(13): 2814–2817, March 2000. doi: 10.1103/physrevlett.84.2814.

[22] S. Inouye, M. R. Andrews, J. Stenger, H. J. Miesner, D. M. Stamper-Kurn, and W. Ketterle. Observation of Feshbach resonances in a Bose-Einstein condensate. *Nature*, 392(6672):151–154, March 1998. doi: 10.1038/32354.

[23] B. Deh, C. Marzok, C Zimmermann, and Ph. W. Courteille. Feshbach resonances in mixtures of ultracold ^{6}Li and ^{87}Rb gases. *Phys. Rev. A*, 77(1):010701, 2008. doi: 10.1103/physreva.77.010701.

[24] P. W. H. Pinkse, P. T. Junglen, T. Rieger, S. A. Rangwala, and G. Rempe. Filtering slow polar molecules from a thermal gas. *Quantum Electronics Conference, 2003. EQEC '03. European*, page 271, June 2003.

[25] T. G. M. Freegarde, Geoff Daniell, and D. Segal. Coherent amplification in laser cooling and trapping. *Phys. Rev. A*, 73(3):033409, 2006.

[26] M. Weitz and T. W. Hänsch. Frequency-independent laser cooling based on interferometry. *Europhys. Lett.*, 49(3):302–308, 2000. doi: 10.1209/epl/i2000-00149-4.

[27] T. G. M. Freegarde and D. Segal. Algorithmic cooling in a momentum state quantum computer. *Phys. Rev. Lett.*, 91(3):037904, 2003.

[28] Victor I. Balykin, V. G. Minogin, and Vladilen S. Letokhov. Electromagnetic trapping of cold atoms. *Reports on Progress in Physics*, 63(9):1429–1510, 2000.

[29] Harold J. Metcalf and P. van der Straten. Laser cooling and trapping of atoms. *J. Opt. Soc. Am. B*, 20(5):887, 2003.

[30] Bruce W. Shore. *Simple Atoms and Fields.* The Theory of Coherent Atomic Excitation. John Wiley & Sons, 1989. ISBN 0-471-61398-3.

[31] C. N. Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Quantum Mechanics.* John Wiley & Sons, 1977. ISBN 0-471-16433-X.

[32] Bruce W. Shore. *Multilevel Atoms and Incoherence.* The Theory of Coherent Atomic Excitation. John Wiley & Sons, 1990. ISBN 0-471-52416-6.

[33] Wolfgang Demtröder. *Laser Spectroscopy: Basic Concepts and Instrumentation.* Springer-Verlag, 2003. ISBN 3-540-65225-6.

[34] D. Meschede. *Optics, Light and Lasers.* Wiley-VCH, 2004. ISBN 3-527-40364-7.

[35] L Allen and J. H. Eberly. *Optical resonance and two-level atoms.* Courier Dover Publications, Inc., New York, 1987. ISBN 0-486-65533-4.

[36] J. Dalibard and Yvan Castin. Laser Cooling from the Semi-Classical to the Quantum Regime, 1992.

[37] J. P. Gordon and A. Ashkin. Motion of atoms in a radiation trap. *Phys. Rev. A*, 21(5):1606–1617, May 1980. doi: 10.1103/physreva.21.1606.

[38] Vladilen S. Letokhov, V. G. Minogin, and B. D. Pavlik. Cooling and trapping of atoms and molecules by a resonant laser field. *Opt. Comm.*, 19(1):72–75, October 1976.

[39] Harold J. Metcalf and P. van der Straten. *Laser Cooling and Trapping.* Graduate Texts in Contemporary Physics. Springer, 1999. ISBN 0-387-98747-9.

[40] Hannes Risken. *The Fokker-Planck Equation.* Springer, 1992. ISBN 354061530X.

[41] D. J. Wineland and W. M. Itano. Laser cooling of atoms. *Phys. Rev. A*, 20(4): 1521–1540, October 1979. doi: 10.1103/physreva.20.1521.

[42] Y. Castin, H. Wallis, and J. Dalibard. Limit of Doppler cooling. *J. Opt. Soc. Am. B*, 6(11):2046, 1989.

[43] Mary L. Boas. *Mathematical Methods in the Physical Sciences*. John Wiley & Sons, 1983.

[44] Thad Walker, David Sesko, and Carl E. Wieman. Collective behavior of optically trapped neutral atoms. *Phys. Rev. Lett.*, 64(4):408–411, January 1990. doi: 10.1103/physrevlett.64.408.

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

[46] V. S. Bagnato, L. G. Marcassa, M. Oria, G. I. Surdutovich, R. Vitlina, and S. C. Zilio. Spatial distribution of atoms in a magneto-optical trap. *Phys. Rev. A*, 48 (5):3771–3775, November 1993. doi: 10.1103/physreva.48.3771.

[47] G. Muller A. Hope, D. Haubrich and D. Meschede. Neutral Cesium Atoms in Strong Magnetic-Quadrupole Fields at Sub-Doppler Temperatures. *Europhys. Lett.*, 22(9):669–674, 1993.

[48] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, Christopher 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, August 1995. doi: 10.1103/physreva.52.1423.

[49] P. Jeffrey Ungar, David S. Weiss, E. Riis, and Steven Chu. Optical molasses and multilevel atoms: theory. *J. Opt. Soc. Am. B*, 6(11):2058–2071, 1989.

[50] John Weiner and P. T. Ho. *Light-Matter Interaction*. Wiley-Interscience, 2003.

[51] J. Lawall, F. Bardou, B. Saubamea, K. Shimizu, M. Leduc, A. Aspect, and C. N. Cohen-Tannoudji. Two-Dimensional Subrecoil Laser Cooling. *Phys. Rev. Lett.*, 73 (14):1915–1918, October 1994. doi: 10.1103/physrevlett.73.1915.

[52] J. Lawall, S. Kulin, B. Saubamea, N. Bigelow, M. Leduc, and C. N. Cohen-Tannoudji. Three-Dimensional Laser Cooling of Helium Beyond the Single-Photon Recoil Limit. *Phys. Rev. Lett.*, 75(23):4194–4197, December 1995. doi: 10.1103/physrevlett.75.4194.

[53] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. N. Cohen-Tannoudji. Laser Cooling below the One-Photon Recoil Energy by Velocity-Selective Coherent Population Trapping. *Phys. Rev. Lett.*, 61(7):826–829, 1988. doi: 10.1103/physrevlett.61.826.

[54] J. Hack, L. Liu, M. Olshanii, and Harold J. Metcalf. Velocity-selective coherent population trapping of two-level atoms. *Phys. Rev. A*, 62:013405, 2000.

[55] Mark Kasevich and Steven Chu. Laser Cooling below a Photon Recoil with Three-Level Atoms. *Phys. Rev. Lett.*, 69(12):1741–1744, 1992.

[56] J. Reichel, O. Morice, G. M. Tino, and Christophe Salomon. Subrecoil Raman Cooling of Cesium Atoms. *Europhys. Lett.*, 28(7):477–482, 1994.

[57] J. Reichel, F. Bardou, Maxime Ben Dahan, E. Peik, S. Rand, Christophe Salomon, and C. N. Cohen-Tannoudji. Raman Cooling of Cesium below 3 nK: New Approach Inspired by Lévy Flight Statistics. *Phys. Rev. Lett.*, 75(25):4575–4578, December 1995. doi: 10.1103/physrevlett.75.4575.

[58] H. J. Lee, Charles S Adams, Mark Kasevich, and Steven Chu. Raman Cooling of Atoms in an Optical Dipole Trap. *Phys. Rev. Lett.*, 76(15):2658, 1996.

[59] Nir Davidson, Heun Jin Lee, Mark Kasevich, and Steven Chu. Raman cooling of atoms in two and three dimensions. *Phys. Rev. Lett.*, 72(20):3158–3161, May 1994. doi: 10.1103/physrevlett.72.3158.

[60] Naoto Masuhara, John M. Doyle, Jon C. Sandberg, Daniel Kleppner, Thomas J. Greytak, Harald F. Hess, and Greg P. Kochanski. Evaporative Cooling of Spin-Polarized Atomic Hydrogen. *Phys. Rev. Lett.*, 61(8):935–938, August 1988. doi: 10.1103/physrevlett.61.935.

[61] Nikolay V. Vitanov and S. Stenholm. Analytic properties and effective two-level problems in stimulated Raman adiabatic passage. *Phys. Rev. A*, 55:648–660, 1997.

[62] Nikolay V. Vitanov. Analytic model of a three-state system driven by two laser pulses on two-photon resonance. *J. Phys. B*, 31(4):709–725, 1998.

[63] A. S. Shumovsky, E. I. Aliskenderov, F. LeKien, and N. D. Vinh. On a Jaynes-Cummings type model with multiphoton transitions . *J. Phys. A: Math. Gen.*, 19: 3607–3617, December 1986.

[64] Bruce W. Shore and Peter L. Knight. The Jaynes-Cummings Model. *J. Mod. Opt.*, 40(7):1195–1238, July 1993.

[65] Ying Wu. Effective Raman theory for a three-level atom in the Λ configuration. *Phys. Rev. A*, 54(2):1586–1592, August 1996. doi: 10.1103/physreva.54.1586.

[66] Bruce W. Shore. Definition of virtual levels. *Am. J. Phys*, 47(3):262, 1978.

[67] Christopher J. Foot. *Atomic Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2005. ISBN 0-19-850695-3.

[68] A. F. Linskens, I. Holleman, N. Dam, and J. Reuss. Two-photon Rabi oscillations. *Phys. Rev. A*, 54(6):4854, 1996.

[69] G. K. Woodgate. *Elementary Atomic Structure*. Oxford Scientific Publications, 2000. ISBN 0-19-851156-6.

[70] D. M. Brink and G. R. Satchler. *Angular momentum*. Clarendon Press, 1994. ISBN 978-0-19-851759-7.

[71] J L Meunier. A simple demonstration of the Wigner-Eckart theorem. *Euro. J. Phys.*, 8(2):114–116, 1987.

[72] Rodney Loudon. *The quantum theory of light*. Oxford science publications. Clarendon Press · Oxford, 1983. ISBN 019851155-8.

[73] Daniel A. Steck. Rubidium 85 D Line Data, 2008. *http://steck.us/alkalidata/*

[74] A. R. Edmonds. *Angular Momentum in Quantum Mechanics*. Princeton Landmarks in Physics. Princeton University Press, 1996.

[75] Paul Siddons, Charles S Adams, Chang Ge, and Ifan G. Hughes. Absolute absorption on rubidium D lines: comparison between theory and experiment. *J. Phys. B*, 41(15):155004 (10, 2008.

[76] B. E. Schultz, H. Ming, G. A. Noble, and W. A. van Wijngaarden. Measurement of the Rb D2 transition linewidth at ultralow temperature. *Euro. Phys. J. D*, 48: 171–176, July 2008. doi: 10.1140/epjd/e2008-00109-0.

[77] E. Arimondo, M. Inguscio, and P. Violino. Experimental determinations of the hyperfine structure in the alkali atoms. *Rev. Mod. Phys.*, 49(1):31, 1977.

[78] Z. Deng. Separated subsystems in four-level atom. *Opt. Comm.*, 48:284–286, December 1983. doi: 10.1016/0030-4018(83)90216-x.

[79] S. J. Buckle, S. M. Barnett, Peter L. Knight, M. A. Lauder, and D. T. Pegg. Atomic Interferometers: Phase-dependence in Multilevel Atomic Transitions. *J. Mod. Opt.*, 33:1129–1140, September 1986.

[80] E. Kyrölä and M. Lindberg. Serial and parallel multilevel systems. *Phys. Rev. A*, 35(10):4207–4225, May 1987. doi: 10.1103/physreva.35.4207.

[81] I. Dotsenko, W. Alt, S. Kuhr, D. Schrader, M. Muller, Y. Miroshnychenko, V. Gomer, A. Rauschenbeutel, and D. Meschede. Application of electro-optically generated light fields for Raman spectroscopy of trapped cesium atoms. *Appl. Phys. B*, 78:711–717, 2004.

[82] Daniel A. Steck. Cesium D Line Data, 2008. *http://steck.us/alkalidata/*

[83] J. I. Cirac, M. Lewenstein, and P. Zoller. Collective laser cooling of trapped atoms. *Europhys. Lett.*, 35(9):647–652, 1996.

[84] C. J. Dedman, K. G. H. Baldwin, and M. Colla. Fast switching of magnetic fields in a magneto-optic trap. *Rev. Sci. Inst.*, 72:4055–4058, November 2001. doi: 10.1063/1.1408935.

[85] James E. Bateman and T. G. M. Freegarde. Fractional adiabatic passage in two-level systems: Mirrors and beam splitters for atomic interferometry. *Phys. Rev. A*, 76(1):013416, 2007. doi: 10.1103/physreva.76.013416.

[86] Nikolay V. Vitanov, T. Halfmann, Bruce W. Shore, and K. Bergmann. Laser-induced population transfer by adiabatic passage techniques. *Ann. Rev. Phys. Chem.*, 52:763–809, 2001.

[87] K. Bergmann, H. Theuer, and Bruce W. Shore. Coherent population transfer among quantum states of atoms and molecules. *Rev. Mod. Phys.*, 70(3):1003–1025, 1998.

[88] Bruce W. Shore, Klaas Bergmann, A. Kuhn, S. Schiemann, J. Oreg, and J. H. Eberly. Laser-induced population transfer in multistate systems: A comparative study. *Phys. Rev. A*, 45(7):5297, 1992.

[89] G. P. Djotyan, J. S. Bakos, G. Demeter, Zs. Sörlei, J. Szigeti, and D. Dzsotjan. Creation of a coherent superposition of quantum states by a single frequency-chirped short laser pulse. *J. Opt. Soc. Am. B*, 25(2):166–174, 2008.

[90] M. Weitz, B. C. Young, and Steven Chu. Atomic Interferometer Based On Adiabatic Population Transfer. *Phys. Rev. Lett.*, 73(19):2563–2566, 1994.

[91] J Baudon, R Mathevet, and J Robert. Atomic interferometry. *J. Phys. B*, 32(15):R173–R195, 1999.

[92] J. S. Melinger, Suketu R. Gandhi, A. Hariharan, Debabrata Goswami, and Warren S. Warren. Adiabatic population transfer with frequency-swept laser pulses. *J. Chem. Phys.*, 101(8):6439, 1994.

[93] V. S. Malinovsky and J. L. Krause. General theory of population transfer by adiabatic rapid passage with intense, chirped laser pulses. *Euro. Phys. J. D*, 14:147–155, 2001.

[94] Richard P. Feynman, Jr Frank L. Vernon, and Robert W. Hellwarth. Geometrical Representation of the Schrödinger Equation for Solving Maser Problems. *J. App. Phys.*, 28(1):49–52, 1957.

[95] P. D. Featonby, G. S. Summy, Jocelyn L. Martin, Wu Huang, K. P. Zetie, Christopher J. Foot, and K. Burnett. Adiabatic transfer for atomic interferometry. *Phys. Rev. A*, 53(1):373–380, January 1996.

[96] Lorenz Mitschang and Herbert Rinneberg. Broadband population inversion by a frequency-swept pulse beyond the adiabatic approximation. *J. Chem. Phys.*, 118 (12):5496–5505, 2003.

[97] J. Baum, R. Tycko, and A. Pines. Broadband and adiabatic inversion of a two-level system by phase-modulated pulses. *Phys. Rev. A*, 32(6):3435, 1985.

[98] C. J. Hardy, W. A. Edelstein, and D. Vatis. Efficient Adiabatic Fast Passage for NMR Population Inversion in the Presence of Radiofrequency Field Inhomogeneity and Frequency Offsets. *J. Mag. Res.*, 66:470–482, 1986.

[99] Daniel Rosenfeld and Yuval Zur. Is the sech/tanh Adiabatic Pulse Really Adiabatic? *J. Mag. Res.*, 132(1):102–108, May 1998.

[100] T. Lu, X. Miao, and Harold J. Metcalf. Bloch theorem on the Bloch sphere. *Phys. Rev. A*, 71(6):061405, 2005.

[101] Debabrata Goswami and Warren S. Warren. Effects of pulses with simple phase and frequency modulations. *Phys. Rev. A*, 50(6):5190–5196, December 1994. doi: 10.1103/physreva.50.5190.

[102] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1992. ISBN 0-521-43108-5.

[103] C. D. Wallace, T. P. Dinneen, A. Kumarakrishnan, Phillip 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.

[104] G. P. Djotyan, J. S. Bakos, G. Demeter, P. N. Ignácz, M. Á. Kedves, Zs. Sörlei, J. Szigeti, and Z. L. Tóth. Coherent population transfer in Rb atoms by frequency-chirped laser pulses. *Phys. Rev. A*, 68(5):053409–8, 2003.

[105] L. P. Yatsenko, Bruce W. Shore, T. Halfmann, K. Bergmann, and A. Vardi. Source of metastable H(2s) atoms using the Stark chirped rapid-adiabatic-passage technique. *Phys. Rev. A*, 60(6):R4237–R4240, December 1999. doi: 10.1103/physreva. 60.r4237.

[106] T. Rickes, L. P. Yatsenko, S. Steuerwald, T. Halfmann, Bruce W. Shore, Nikolay V. Vitanov, and Klaas Bergmann. Efficient adiabatic population transfer by two-photon excitation assisted by a laser-induced Stark shift. *J. Chem. Phys.*, 113(2): 534–546, July 2000. doi: 10.1063/1.481829.

[107] L. P. Yatsenko, Nikolay V. Vitanov, Bruce W. Shore, T. Rickes, and K. Bergmann. Creation of coherent superpositions using Stark-chirped rapid adiabatic passage. *Opt. Comm.*, 204(1-6):413–423, 2002.

[108] Nikolay V. Vitanov and Bruce W. Shore. Stimulated Raman adiabatic passage in a two-state system. *Phys. Rev. A*, 73:053402, 2006.

[109] Ying Wu. Effective Raman theory for a three-level atom in the Λ configuration. *Phys. Rev. A*, 54(2):1586, 1996.

[110] P. Marte, P. Zoller, and J. L. Hall. Coherent atomic mirrors and beam-splitters by adiabatic passage in multilevel systems. *Phys. Rev. A*, 44(7):R4118–R4121, 1991.

[111] Nikolay V. Vitanov, K. A. Suominen, and Bruce W. Shore. Creation of coherent atomic superpositions by fractional stimulated Raman adiabatic passage. *J. Phys B*, 32(18):4535–4546, 1999.

[112] L P Maguire, R M W van Bijnen, E Mese, and R. E. Scholten. Theoretical calculation of saturated absorption spectra for multi-level atoms. *J. Phys. B*, 39 (12):2709–2720, 2006.

[113] David A. Smith and Ifan G. Hughes. The role of hyperfine pumping in multilevel systems exhibiting saturated absorption. *Am. J. Phys.*, 72(5):631–637, 2004. doi: 10.1119/1.1652039.

[114] Kristan Lee Corwin, Zheng-Tian Lu, Carter F. Hand, Ryan J. Epstein, and Carl E. Wieman. Frequency-Stabilized Diode Laser with the Zeeman Shift in an Atomic Vapor. *Appl. Opt.*, 37(15):3295–3298, 1998.

[115] T. W. Hänsch and B. Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. *Opt. Comm.*, 35:441–444, December 1980. doi: 10.1016/0030-4018(80)90069-3.

[116] T. Petelski, M. Fattori, G. Lamporesi, J. Stuhler, and G. M. Tino. Doppler-free spectroscopy using magnetically induced dichroism of atomic vapor: a new scheme for laser frequency locking. *Euro. Phys. J D*, 22(2):279–283, February 2003.

[117] Eric D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *Am. J. Phys*, 69(1):79, 2000.

[118] S Kunze, S Wolf, and G. Rempe. Measurement of fast frequency fluctuations:Allan variance of a grating-stabilized diode laser. *Opt. Comm.*, 128:269–274, July 1996.

[119] D. W. Allan. Statistics of atomic frequency standards. *Proceedings of the IEEE*, 54(2):221–230, 1966.

[120] D. A. Howe, D. W. Allan, and J. A. Barnes. Properties of Signal Sources and Measurement Methods. *Thirty Fifth Annual Frequency Control Symposium. 1981*, pages 669–716, 1981.

[121] J Rutman and F. L. Walls. Characterization of frequency stability in precision frequency sources. *Proceedings of the IEEE*, 79(7):952–960, July 1991. doi: 10. 1109/5.84972.

[122] Kilian Singer, Selim Jochim, Marcel Mudrich, Allard Mosk, and M Weidermüller. Low-cost mechanical shutter for light beams. *Rev. Sci. Inst.*, 73(12):4402–4404, 2002. doi: 10.1063/1.1520728.

[123] L P Maguire, S. Szilagyi, and R. E. Scholten. High performance laser shutter using a hard disk drive voice-coil actuator. *Rev. Sci. Inst.*, 75(9):3077–3079, 2004. doi: 10.1063/1.1786331.

[124] J. E. Thomas, P. R. Hemmer, S. Ezekiel, C. C. Leiby, R. H. Picard, and C. R. Willis. Observation of Ramsey Fringes Using a Stimulated, Resonance Raman Transition in a Sodium Atomic Beam. *Phys. Rev. Lett.*, 48(13):867–870, March 1982. doi: 10.1103/physrevlett.48.867.

[125] Mark Kasevich and Steven Chu. Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer. *Appl. Phys. B*, 54:321–332, May 1992. doi: 10.1007/bf00325375.

[126] Agilent Technologies. *Agilent 8648ABCD.* Agilent

[127] National Physical Laboratory. *NPL Time & Frequency Services: MSF 60kHz Time and Date Code.* National Physical Laboratory, 2007.

[128] D. Haubrich, M. Dornseifer, and R. Wynands. Lossless beam combiners for nearly equal laser frequencies. *Rev. Sci. Inst.*, 71(2):338–340, 2000. doi: 10.1063/1. 1150204.

[129] A. Loayssa, D. Benito, and M. J. Garde. Single-sideband suppressed-carrier modulation using a single-electrode electrooptic modulator. *Photonics Technology Letters, IEEE*, 13(8):869–871, Aug 2001. doi: 10.1109/68.935831.

[130] B. W. Barr, S. H. Huttner, J. R. Taylor, B. Sorazu, M. V. Plissi, and K. A. Strain. Optical modulation techniques for length sensing and control of optical cavities. *Appl. Opt.*, 46(31):7739–7745, 2007.

[131] Nicolas Schlosser, Georges Reymond, Igor Protsenko, and Philippe Grangier. Sub-possonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841): 1024–1027, June 2001. doi: 10.1038/35082512.

[132] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and Dieter Meschede. Neutral Atom Quantum Register. *Phys. Rev. Lett.*, 93(15):150501, October 2004. doi: 10.1103/physrevlett.93.150501.

[133] Dieter Meschede and Arno Rauschenbeute. Manipulating Single Atoms. *ATOMIC PHYSICS 19: XIX International Conference on Atomic Physics; ICAP 2004*, 770 (1):359–367, 2005. doi: 10.1063/1.1928870.

[134] D. D. Yavuz, P. B. Kulatunga, E. Urban, T. A. Johnson, N. Proite, T. Henage, T. G. Walker, and M. Saffman. Fast Ground State Manipulation of Neutral Atoms in Microscopic Optical Traps. *Phys. Rev. Lett.*, 96(6):063001–, February 2006. doi: 10.1103/physrevlett.96.063001.

[135] Wolfgang Paul. Electromagnetic Traps for Charged and Neutral Particles. *Nobel Lecture*, December 1989.

[136] H. Dehmelt. Experiments with an isolated subatomic particle at rest. *Nobel Lecture*, December 1989.

[137] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Phys. Rev. Lett.*, 74(20):4091–4094, May 1995. doi: 10.1103/physrevlett.74.4091.

[138] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a Fundamental Quantum Logic Gate. *Phys. Rev. Lett.*, 75:4714, 1995.

[139] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and Phillip L. Gould. Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy. *Phys. Rev. Lett.*, 75(22):4011–4014, November 1995. doi: 10.1103/physrevlett.75.4011.

[140] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, Christian F Roos, J. Eschner, and R. Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422:408–411, March 2003.

[141] J. S. Bakos, G. P. Djotyan, G. Demeter, and Zs. Sörlei. Transient laser cooling of two-level quantum systems with narrow natural linewidths. *Phys. Rev. A*, 53(4): 2885–2888, 1996.

[142] G. Demeter, G. P. Djotyan, and J. S. Bakos. Deflection and splitting of atomic beams with counterpropagating, short, chirped laser pulses. *J. Opt. Soc. Am. B*, 15(1):16–24, 1998.

[143] G. Demeter, G. P. Djotyan, Zs. Sörlei, and J. S. Bakos. Mechanical effect of retroreflected frequency-chirped laser pulses on two-level atoms. *Phys. Rev. A*, 74 (1):013401, 2006. doi: 10.1103/physreva.74.013401.

[144] K. J. Boller, A. Imamoğlu, and S. E. Harris. Observation of Electromagnetically Induced Transparency. *Phys. Rev. Lett.*, 66(20):2593–2596, May 1991.

[145] M. D. Lukin. Colloquium: Trapping and manipulating photon states in atomic ensembles. *Rev. Mod. Phys.*, 75(2):457–472, April 2003. doi: 10.1103/revmodphys. 75.457.

[146] Anthony E. Siegman. *Lasers*. University Science Books, 1986. ISBN 0-935702-11-3.

[147] Y. Stalgies, I. Siemers, B. Appasamy, and P. E. Toschek. Light shift and Fano resonances in a single cold ion. *J. Opt. Soc. Am. B*, 15(10):2505–2514, 1998.

[148] Emeric de Clercq and P. Cerez. Evaluation of the light shift in a frequency standard based on Raman induced Ramsey resonance. *Opt. Comm.*, 45(2):91–94, March 1983.

[149] P. F. Liao and J. E. Bjorkholm. Direct Observation of Atomic Energy Level Shifts in Two-Photon Absorption. *Phys. Rev. Lett.*, 34(1):1–4, January 1975. doi: 10.1103/physrevlett.34.1.

[150] H. J. Lee and Steven Chu. Atomic spin polarization by Raman cooling. *Phys. Rev. A*, 57(4):2905–2908, April 1998. doi: 10.1103/physreva.57.2905.

[151] H. Perrin, A. Kuhn, I. Bouchoule, T. Pfau, and Christophe Salomon. Raman cooling of spin-polarized cesium atoms in a crossed dipole trap. *Europhys. Lett.*, 46(2):141–147, 1999.

[152] C. N. Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Quantum Mechanics.* John Wiley & Sons, 1977. ISBN 0-471-16435-6.

[153] A. D. Boozer, R. Miller, T. E. Northup, A. Boca, and H. J. Kimble. Optical pumping via incoherent Raman transitions. *Phys. Rev. A*, 76(6):063401, 2007. doi: 10.1103/physreva.76.063401.

[154] J. L. Sørensen, D. Møller, T. Iversen, J. B. Thomsen, F. Jensen, P. Staanum, D. Voigt, and M. Drewsen. Efficient coherent internal state transfer in trapped ions using stimulated Raman adiabatic passage. *New Journal of Physics*, 8:261–, November 2006. doi: 10.1088/1367-2630/8/11/261.

[155] Carl E. Wieman, D. E. Pritchard, and D. J. Wineland. Atom cooling, trapping, and quantum manipulation. *Rev. Mod. Phys. Supplement*, 71:253–, March 1999. doi: 10.1103/revmodphys.71.s253.

[156] Jesper Glückstad, Paul C. Mogensen, and René L. Eriksen. The generalised phase contrast method and its applications. *Danish Optical Society*, 16(1):49–54, 2001.

[157] 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. doi: 10.1103/physrevlett.72.3017.

[158] N. K. Metzger, K. Dholakia, and E. M. Wright. Observation of Bistability and Hysteresis in Optical Binding of Two Dielectric Spheres. *Phys. Rev. Lett.*, 96(6): 068102, 2006. doi: 10.1103/physrevlett.96.068102.

[159] Peter Horak, Gerald Hechenblaikner, Klaus M. Gheri, Herwig Stecher, and Helmut Ritsch. Cavity-Induced Atom Cooling in the Strong Coupling Regime. *Phys. Rev. Lett.*, 79(25):4974–4977, December 1997. doi: 10.1103/physrevlett.79.4974.

[160] Vladan Vuletić and Steven Chu. Laser Cooling of Atoms, Ions, or Molecules by Coherent Scattering. *Phys. Rev. Lett.*, 84(17):3787–3790, April 2000. doi: 10.1103/physrevlett.84.3787.

[161] P. Maunz, T. Puppe, I. Schuster, N. Syassen, P. W. H. Pinkse, and G. Rempe. Cavity cooling of a single atom. *Nature*, 428:50–52, March 2004. doi: 10.1038/nature02387.