UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL AND APPLIED SCIENCES

School of Physics & Astronomy

Phase and intensity control of lasers for atom interferometry

by

Rachel Gregory

Thesis for the degree of Doctor of Philosophy

November 2015

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
School of Physics & Astronomy

Doctor of Philosophy

PHASE AND INTENSITY CONTROL OF LASERS FOR ATOM INTERFEROMETRY

by Rachel Gregory

Atom interferometry is a highly versatile experimental technique for precision measurement, with applications in inertial sensing and tests of fundamental physics. Such interferometers rely on high fidelity transfer of atoms between internal states, which can be challenging when working with thermal atom clouds, where a range of resonance frequencies exists. Inhomogeneities in the spatial intensity distribution of the manipulation lasers can also give rise to a range of coupling strengths; both of these effects serve to reduce the fidelity of coherent manipulation operations.

Composite pulses offer one route to high fidelity coherent manipulations. In this thesis we investigate the performance of composite pulses in a thermal cloud of Rubidium 85. We find remarkable agreement between theory and experimental investigation of composite inversion pulses, and build on this to model the performance of such pulses in atom interferometers. We also find excellent agreement between our model and the first experimental demonstration of an application of interferometry in the cooling of atoms; this is the first step towards implementing this cooling method in complex structures like molecules.

Acknowledgements

First I must thank my supervisor, Dr Tim Freegarde, for taking me on as a student and always having a fresh way of approaching problems.

I am indebted to all the Quantum Controllers who worked on the experiment before me; in particular, Dr James Bateman and Dr Matt Himsworth, who were happy to spend an hour or two helping me to fix things. I owe all of my lab skillz to experimental wizard Dr Alex Dunning. Sorry I broke so many things. Thanks also to Dr Jo Rushton who patiently corrected my many electronics screwups.

The past few years would have been much less enjoyable without fellow Quantum Controllers Jonathan "J. Effective" Woods and Andy "Mary Berry" Chapman. Thanks for all the tea and for listening to me complain when the experiment broke for the nth time (where n is a *very* large number). Also on the listening-to-me-complain list is Sanja - thanks for eating my cake and shouting at me in the bouldering room.

To George, partner in crime, fabulous bird lady and almost definitely not a Russian spy, you're brilliant and I love you to bits. To Hamid, who kept me smiling through the writing of this thesis (and beyond), doost daram, joon.

I'm incredibly fortunate to have a wonderfully supportive family who have supported me in all my choices. Dad, I hope you enjoy reading this! (Let me know if you spot any mistakes.)

Above all, I would never have made it through the last four years without my number 1 cheerleader, knitting consultant, gardening advice bureau and best friend. Mum, I owe you everything.

dedicated to Michael

who always knew I'd be a scientist

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Rachel Gregory , declare that the thesis entitled *Phase and intensity control of lasers for atom interferometry* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. 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: [2],[3]

Signed:.....

Date:.....

Introduction

First proposed in 1975 [4], laser cooling of atoms is based on the exchange of momentum between atoms and photons; precisely tuning the frequency of a laser such that atoms moving in a direction opposite to the wavevector of the laser are more likely to scatter photons allows a velocity-dependent frictional force to be engineered. The first demonstration of the use of a laser to decelerate atoms was in 1982 [5]. With the addition of appropriate magnetic fields, atoms can be confined in addition to being cooled: this is a magneto-optical trap (MOT). The first MOT was reported in 1987 with sodium atoms [6].

In the years following its first experimental demonstration, the MOT has become the workhorse of a wide range of investigations into quantum phenomena. For example, in 1995 the first atomic Bose-Einstein condensates were produced within months of each other (Rubidium 87 [7] and Sodium 23 [8]); both experiments relied on laser cooling in combination with other methods to reach the μK (and below) temperatures required. Laser cooling is an essential first step in confining atoms in optical lattices; these experiments have applications in the study of many-body phenomena [9] [10].

The application which is of interest in this thesis is atom interferometry. Just as an optical interferometer takes a wavelike object (a beam of light), splits it, and measures the phase difference between the two paths taken when the beams are recombined, an atom interferometer takes a wavelike object (the wavepacket of an atom) and performs the same operations. Similarly, the requirement on a beam of light used for interferometry to be coherent and monochromatic has an analog in atom interferometry: in order to observe interference fringes, the centre of mass separation of the atomic wavepackets must be less than half of the coherence length [11]. This is possible in ultracold atomic gases which have a small spread in atomic velocities. Being massive, an atom is subject to forces that light does not experience, and the difference in these forces along the two paths within the interferometer manifests itself in the phase shift at the output

of the interferometer. This enables measurement of these forces to a very high precision [12], and applications exist in accelerometers [13], gyroscopes [14], tests of fundamental physics [15] [16] [17] [18] [19], measurement of atomic properties [20], and chip-based quantum information processing [21].

In the majority of atom interferometry experiments, stimulated Raman transitions between ground state hyperfine levels are used to transfer momentum to the atoms. This momentum transfer upon excitation means that the internal and external states of the atom are correlated, and so the two paths taken within the interferometer may be labelled by the internal state [22]. When acting as a beamsplitter, a Raman pulse creates a coherent superposition of two wavepackets taking separate paths; when acting as a mirror, a Raman pulse redirects the wavepackets towards each other. Additional augmentation pulses may be applied to increase the sensitivity of the interferometer [23].

An optical Mach Zehnder interferometer with only partially reflective mirrors would suffer a loss of signal at the output, due to some of the light being lost upon reflection. Similarly, atom interferometers require effective augmentation pulses, i.e. high fidelity transfer of population between internal states. This can be difficult in systems with a range of resonance frequencies - for example, when a range of velocities is present in the atomic ensemble, each atom will have a different Doppler-shifted resonance frequency, and an associated off-resonance error. One method of creating higher fidelity augmentation pulses is to replace the single mirror pulse with a composite pulse. A composite pulse is a series of consecutive pulses with different durations and phases, designed to compensate a particular type of error. Composite pulses as augmentation pulses were shown to increase the contrast of an atom interferometer [24], but only two composite pulses were investigated.

The main focus of this thesis is a systematic study of composite pulses in ultracold Rubidium 85. We show that very good fits to data from cold atom experiments in our system can be achieved through simulation, and use this model for further investigations. This involves testing other composite pulses through simulation including one pulse which we believe to be novel. We also investigate whether a correlation exists between the performance of a composite pulse as an inversion pulse (for transfer of atoms initialised in one internal state to the other internal state - i.e., initial and final states are not superposition states) and its performance as an augmentation pulse. Some of this work is published in [2].

The cooling of molecules (and indeed any element lacking a two level closed radiative transition) is much more challenging than the cooling of atoms with simple internal structures like Rubidium. This is because each electronic state is split into many vibrational states, each with a different energy, and many different lasers are required to maintain the cycle of absorption and spontaneous emission [25] [26]. Techniques such as optical pumping are often employed to reach low temperatures [27] [28].

Although atom interferometers are mainly used for precision measurement, it is also possible to use interferometry to manipulate the velocity distribution of an ensemble of atoms, by choosing an interferometer sequence such that the phase difference accumulated between the two paths is dependent on the velocity of the atom. Such a velocity sensitive excitation can be used in combination with spontaneous emission to cause cooling. This 'interferometric cooling' [29] is independent of the detuning of the laser from resonance and has potential applications in the cooling of species with unfavourable energy level structures which may otherwise require many lasers. Interferometric cooling was first proposed in 2000 and had not been experimentally demonstrated until recently. We show that the first demonstration of interferometric cooling in our experiment [3] can be well described by our model, and build on this agreement to simulate improvements to the cooling scheme.

The layout of this thesis is as follows. Chapter 1 introduces some background to the cooling and trapping of atoms. Chapter 2 is an overview of Raman transitions and presents a derivation of the equations upon which the model of our experiment is based. In Chapter 3, atom interferometry is introduced in more detail and interferometric cooling is described. In Chapter 4, we introduce composite pulses, the types of error for which they can compensate, and how to quantify their effectiveness.

Chapters 5 and 6 present the experimental details of, respectively, the magneto-optical trapping system and the Raman laser system. In Chapter 7, the use of a spatial light modulator for beamshaping of the Raman beams is investigated.

In Chapter 8 we present investigations of composite pulses in our system. Excellent agreement is established between our model and the experimental data, and this is used to justify further investigation of composite pulses through simulation. In particular, we simulate composite pulses which were not tested experimentally; we also simulate their use in atom interferometers. In Chapter 9 we show that our

model can be used to fit to experimental data from the first demonstration of interferometric cooling, and use the model to investigate as yet untested modifications to the cooling.

Several relevant derivations are presented in the Appendices, along with an example of the code used for simulations.

Part I

Theory

Chapter 1

Cooling and trapping of Rubidium 85

In experiments where the preservation of coherence of atomic states is important, cold, dilute atomic gases are ideal: the low density and low temperature (small spread of velocities) result in a smaller probability of coherence-destroying collisions on the timescale of the experiment. Low density is easily achievable, but cooling and confining the atoms requires consideration of the internal energy structure of the atoms and their interaction with light and magnetic fields.

This chapter outlines the atomic structure of Rubidium 85 and the rules governing dipole transitions within this structure. The principles of laser cooling and magneto-optical trapping are described and consideration is given to sub-Doppler cooling mechanisms.

1.1 Rb 85 atomic structure

1.1.1 Gross structure

Electrons bound to atoms can occupy different energy states; the energies and wavefunctions of these states are obtained by solving the time-independent Schrödinger equation for the atom. The simplest atom is Hydrogen with one proton and one electron. The electrostatic (Coulomb) potential experienced by the electron due to the nucleus is spherically symmetric and so the Schrödinger equation

$$-\frac{\hbar^2}{2m}\nabla^2\psi - \frac{Ze^2}{4\pi\epsilon_0 r}\psi = E\psi, \tag{1.1}$$

where Z is the atomic number, can be solved exactly (with $Z = 1$)[30]. The wavefunctions which solve this equation are product functions composed of a radial part depending on n and l and an angular part characterised by the quantum numbers l and m_l, and are usually written in the form $\psi(r, \theta, \phi) = R_{n,l}(r)Y_{l,m_l}(\theta, \phi)$. The parameter n is the principal quantum number, which can be thought of as a measure of how large the orbit of the electron is; as n increases, the probability of finding the electron at greater distances from the nucleus increases. n is always a positive integer. l is the orbital (or azimuthal) quantum number, which determines the shape of the orbit and the angular momentum of the electron; m_l is the projection of l onto the z axis. l is also a positive integer, and m_l is an integer restricted to the range $-l \leq m_l \leq l$. The angular eigenfunctions $Y_{l,m_l}(\theta, \phi)$ are the spherical harmonics and in Hydrogen, the radial eigenfunctions $R_{n,l}(r)$ are proportional to Laguerre polymonials. The electron also possesses a spin of 1/2, denoted by s, with projection on an arbitrary axis[1] $m_s = \pm 1/2$. The spin eigenstates are most concisely written as $|s, m_s\rangle$, and the eigenbasis of this system (Hamiltonian as eq. 1.1 plus spin) may be written as $|n, l, m_l, s, m_s\rangle$: an eigenstate is completely defined by listing its good quantum numbers.

In Hydrogen, the energy of the nth level $E_H(n)$ is given by

$$E_H(n) = hc\frac{R_\infty}{n^2}, \tag{1.2}$$

where R_∞ is the Rydberg constant. Note that neither l nor m_l feature in this equation; in Hydrogen, states of the same n but different l are degenerate in energy. The degeneracy with respect to m_l is a consequence of the spherically symmetric potential, and the degeneracy with respect to l is due to a symmetry of the two-body problem [31] which is not present in atoms with more than one electron.

The alkali metals (elements found in Group 1 of the Periodic Table) have one valence electron, which experiences a spherically symmetric potential to a very good approximation. An approximation to their gross structure energy levels $E_{AM}(n, l)$ can be adapted from equation 1.2

[1] This axis is usually defined by the direction of an external magnetic field.

$$E_{AM}(n,l) = hc\frac{R_{AM}}{(n-\delta_l)^2} \tag{1.3}$$

where these energies are no longer degenerate in l for a given n. δ_l is called the quantum defect and accounts for deviations from the Hydrogen case due to imperfect screening of the nuclear charge by the inner shell electrons. It is l-dependent because the spatial wavefunctions are l-dependent and the shape of these wavefunctions determines the extent of the screening.

To calculate these energy levels to a higher precision, it must be taken into account that the form of the potential is no longer spherically symmetric, due to the electrostatic interaction between the electrons. To solve the Schrödinger equation in these cases, the Central Field Approximation is employed, which separates the potential into two parts: a spherically symmetric part, and a residual part, enabling solutions to be found by using the solutions to the spherically symmetric potential as a basis and treating the residual potential as a perturbation. Here we do not concern ourselves with the details of this calculation (which may be found in [30]) and simply note that (i) we may still label the energy levels according to the good quantum numbers n, l, m_l, s, m_s, and (ii) only the radial wavefunction is affected; the angular wavefunctions are as in Hydrogen.

1.1.2 Fine structure

The next step is to consider the interaction of the spin of the electron with its orbital angular momentum; this gives rise to the fine structure, the scale of these splittings being of order $Z\alpha^2$ where $\alpha \approx 1/137$ is the fine structure constant. From this point we consider only alkali metal atoms with a single outer shell electron. The electron is orbiting relativistically in an electric field and the result of this is to set up a magnetic field. The perturbation Hamiltonian takes the form

$$\hat{H}_2 = -\hat{\boldsymbol{\mu}} \cdot \hat{\mathbf{B}}, \tag{1.4}$$

where the intrinsic magnetic moment of the electron due to its spin is $\hat{\boldsymbol{\mu}} = -g_s \frac{\mu_B}{\hbar}\hat{\boldsymbol{s}}$ and the magnetic field set up by the orbiting electron in the electrical potential $\phi(r) = U(r)/e$ is

$$\mathbf{B} = \frac{1}{emc^2}\frac{1}{|\mathbf{r}|}\frac{\partial U(r)}{\partial r}\hat{\mathbf{l}}. \tag{1.5}$$

Here, $\hat{\mathbf{l}} = \hat{\mathbf{r}} \times \hat{\mathbf{p}}$ is the orbital angular momentum of the electron, g_s is the electron spin g-factor, and μ_B is the Bohr magneton. For an alkali metal we can write $U(r) = -\frac{Ze^2}{4\pi\epsilon_0 r}$ and the spin-orbit operator becomes

$$-\hat{\boldsymbol{\mu}}_{\boldsymbol{s}}.\hat{\mathbf{B}} = \frac{\mu_0}{4\pi}\mu_B g_s Z \frac{\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}}{r^3\hbar^3}, \tag{1.6}$$

where a factor of two has been incorporated to account for the relativistic transformation between the electron and nucleus rest frames [32].

Calculating the energy shift due to this perturbation is not trivial since the eigenstates $|n, l, m_l, s, m_s\rangle$ of $\hat{H}_0$ are not eigenstates of the spin-orbit operator $\hat{H_{SO}}$. This is clear from the commutators $[\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}, \hat{l}_z] \neq 0$ and $[\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}, \hat{s}_z] \neq 0$: two non-commuting operators do not share a simultaneous eigenbasis, and so m_l and m_s are not good quantum numbers in fine structure. However, $\hat{\mathbf{l}}^2$ and $\hat{\mathbf{s}}^2$ do commute with $\hat{\mathbf{s}}.\hat{\mathbf{l}}$, as does the z component of the total angular momentum $\hat{j}_z$ (where $\hat{\mathbf{j}} = \hat{\mathbf{l}} + \hat{\mathbf{s}}$), and the total angular momentum squared $\hat{\mathbf{j}}^2$. The eigenvalues of $\hat{\mathbf{j}}^2$ and $\hat{j}_z$ are j and m_j respectively, and in combination with l and s they completely describe the eigenbasis of the spin-orbit operator $|l, s, j, m_j\rangle$. m_j takes values from $-j$ to j in integer steps, and angular momentum addition rules restrict the value of j to $|l - s| \leq j \leq |l + s|$.

Contructing this new basis now allows the spin-orbit energy shift ΔE_{SO} to be determined. To calculate the expectation value of $\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}$ it is useful to replace $\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}$ with the squared operators using $\hat{\mathbf{j}}^2 = \hat{\mathbf{l}}^2 + \hat{\mathbf{s}}^2 + 2\hat{\mathbf{s}} \cdot \hat{\mathbf{l}}$. We also note that the expectation values of the squared operators are $j(j+1), l(l+1), s(s+1)$, and combining this with the expectation value of the radial part of $\hat{H}_{SO}$ (denoted by β_{nl}), it is found that

$$\Delta E_{SO} = \beta_{nl}\frac{j(j+1) - l(l+1) - s(s+1)}{2}. \tag{1.7}$$

Finally we note that the previous eigenbasis may be expressed in terms of the new eigenbasis:

$$|l, s, j, m_j\rangle = \sum_{m_l, m_s} C(lsjm_j; m_l, m_s) \, |l, m_l, s, m_s\rangle \tag{1.8}$$

where the C coefficients are called Clebsch-Gordan coefficients: the elements of the unitary matrix transformation between the basis of total angular momentum states and the uncoupled basis[30].

1.1.3 Hyperfine structure

We now address nuclear effects on the energy level structure: specifically, the interaction of the nuclear spin with the total angular momentum of the electron. This hyperfine interaction between the magnetic field set up by the electron (proportional to $\hat{\mathbf{j}}$) and the nuclear spin magnetic moment $\hat{\mathbf{I}}$ can be written as $\hat{H}_{HFS} = A_J \hat{\mathbf{I}} \cdot \hat{\mathbf{J}}$, where the constant of proportionality A_J depends on the electron's total angular momentum J.

In the previous section we considered an interaction proportional to the dot product of two angular momenta. We found that the eigenvalues of the operators corresponding to the z projections of those angular momenta were no longer good quantum numbers, but the eigenvalues corresponding to the sum of the angular momenta and its z component were good quantum numbers. In the same way we construct the total operator for angular momentum of the atom $\hat{\mathbf{F}}$ and its z component $\hat{F}_z$, their eigenvalues being F and m_F, forming the new eigenbasis $|l, s, j, F, m_F\rangle$. As before, m_F takes values from $-F$ to F in integer steps, and angular momentum addition rules restrict the value of F to integers in the range $|I - J| \leq F \leq |I + J|$. Similarly, the hyperfine structure energy shift ΔE_{HFS} is

$$\Delta E_{HFS} = A_J \frac{F(F+1) - I(I+1) - J(J+1)}{2}. \tag{1.9}$$

1.1.4 Spectroscopic notation

Now that we have a set of quantum numbers that can completely specify an electronic state under the interactions considered so far, it is useful to introduce a shorthand for the state of the outer shell electron, known as spectroscopic notation. The state is described in the form

$$|\psi\rangle = nL_J, F = ... \tag{1.10}$$

where n and J are replaced with their numerical values, and L is usually replaced with the letter corresponding to the orbital angular momentum value S, P, D, F, etc. For example, the state of an electron in the 5th shell of an atom of nuclear spin 1/2, with $L = 1$, $J = 3/2$, $F = 2$ would be written $5P_{3/2}, F = 2$.

1.1.5 Zeeman effect

In the absence of an external magnetic field, there is no "preferred" orientation of the atom in space, and so the m_F sublevels of a particular state are degenerate. An external magnetic field breaks this symmetry and splits the sublevels in energy according to the orientation of the atom magnetic moment relative to the applied field. The way in which an applied magnetic field affects these levels is dependent on the strength of the interaction compared to the strength of the hyperfine interaction; we will only be concerned with the case where the external field interaction is weaker than $A_J\hat{\mathbf{I}}\cdot\hat{\mathbf{J}}$, and so it can be treated as a perturbation to the hyperfine structure. The magnetic field $\mathbf{B}$ interacts with the total magnetic moment of the atom $\mu_{atom} = -g_J\mu_B\mathbf{J} + g_I\mu_N\mathbf{I}$. Here, g_J (a function of J, L, and S) and g_I are the Landé g factors for the electron and the nucleus; a derivation of their form may be found in [33].

In this situation, the vector $\mathbf{F}$ precesses slowly around the applied magnetic field, which we take to be along the z axis. The vectors $\mathbf{I}$ and $\mathbf{J}$ rotate around $\mathbf{F}$ at a rate such that their z components are no longer well-defined; however, the z component m_F of their sum F is a good quantum number. Levels of the same F will be split by an amount dependent on the degree to which $\mathbf{F}$ is aligned with $\mathbf{B}$ - i.e., the projection of the total magnetic moment of the atom onto the external field. Since $\mu_N \ll \mu_B$, in determining the form of the interaction operator we can neglect the contribution of the projection[2] of $\mathbf{I}$ along $\mathbf{F}$ and find the perturbation Hamiltonian to be $H_{Zeeman} = g_F\mu_B B F_z$. The g factor $g_F \propto g_J$ is a function of F, I, and J, and the expectation value of this Hamiltonian is $\Delta E_{Zeeman} = g_F\mu_B B m_F$.

[2]However, we do not ignore $\mathbf{I}$ in the calculation of g_F. The nuclear magnetic moment is small compared to the electron magnetic moment, but the vector $\mathbf{I}$ is not small compared to the vector $\mathbf{J}$.

1.2 Dipole transitions

In this section we outline the rules for electric dipole transitions[3] between hyperfine sublevels of Rubidium 85. The atom has no permanent electric dipole moment; rather, one is induced by the incident electric field, and it is the interaction between this dipole and the field itself which gives rise to dipole transitions. Fermi's Golden Rule [34] states that the transition rate W_{if} from an initial state $|\phi_i\rangle$ to a final state $|\phi_f\rangle$, where the transition is induced by a perturbation Hamiltonian H_1, is proportional to the matrix element squared and the density of final states ρ_f:

$$W_{if} = \frac{2\pi}{\hbar}|\langle\phi_f| H_1 |\phi_i\rangle|^2 \rho_f. \tag{1.11}$$

In the case of dipole transitions (with no external magnetic field), the initial and final states which concern us are the eigenstates of the atom under the hyperfine interaction; since $\mathbf{B}_{ext} = 0$ the m_F levels are degenerate. The perturbation Hamiltonian is the dot product of the electric dipole operator $-e\mathbf{r}$ with the electric field of the incident radiation $-e\mathbf{r}\cdot\mathbf{E}$, where the electric field is determined by its amplitude E_0, oscillation frequency ω, and polarisation vector $\boldsymbol{\epsilon}$: $\mathbf{E} = \boldsymbol{\epsilon}\, E_0 \cos\omega t$.

We consider only the case where the amplitude of the electric field of the incident radiation does not vary significantly over the size of the atom; this allows us to factor out E_0 from our matrix elements. We are justified in this approximation because the wavelength of the radiation we use to induce transitions is much larger than the extent of the electron orbitals. Under this approximation, the matrix elements which concern us (for transitions between states of the same n) are $\langle l, s, j, I, F, m_F| \mathbf{r}\cdot\boldsymbol{\epsilon} |l', s', j', I', F', m'_F\rangle$.

The evaluation of these matrix elements is not trivial and the usual method is to take advantage of certain properties of tensor operators. Since we are dealing with eigenstates involving spherical harmonics, the spherical basis is a good choice for the operator $\mathbf{r}$. In this basis, $\mathbf{r}$ has components r_q (with $q = -1, 0, 1$) which may be written in terms of the spherical harmonics $Y_{1,-1}, Y_{1,0}$ and $Y_{1,1}$:

$$\mathbf{r} \propto Y_{1,-1}\frac{\hat{\mathbf{e}}_x + i\hat{\mathbf{e}}_y}{\sqrt{2}} + Y_{1,0}\hat{\mathbf{e}}_z + Y_{1,1}\frac{\hat{-\mathbf{e}}_x + i\hat{\mathbf{e}}_y}{\sqrt{2}}, \tag{1.12}$$

[3] Magnetic dipole transitions (caused by the interaction of the atom's magnetic moment with the oscillating magnetic field of the incident radiation) occur at an energy scale which is small enough to not be of interest to us; magnetic sublevels are split by frequencies in the radio frequency range, whereas we are interested in optical frequencies.

where $\hat{\mathbf{e}}_{x,y,z}$ are unit vectors along the Cartesian axes.

Similarly, the polarisation vector $\boldsymbol{\epsilon}$ with components $\epsilon_q, q = -1, 0, 1$ is written in the spherical basis as

$$\boldsymbol{\epsilon} = A_{\sigma^-}\frac{\hat{\mathbf{e}}_x - i\hat{\mathbf{e}}_y}{\sqrt{2}} + A_\pi \hat{\mathbf{e}}_z + A_{\sigma^+}\frac{-\hat{\mathbf{e}}_x - i\hat{\mathbf{e}}_y}{\sqrt{2}}, \tag{1.13}$$

where A_π depends on the component of the field along the z axis and $A_{\sigma^\pm}$, representing left- and right-circular polarisation, depend on the component in the xy plane.

Written in this form, $\mathbf{r}$ is a rank 1 spherical tensor and the matrix elements with respect to angular momentum eigenstates

$$\begin{aligned}\langle l, s, j, I, F, m_F | \, \mathbf{E}\cdot\mathbf{r} \, | l', s', j', I', F', m'_F\rangle = \\ E_0 \sum_q (-1)^q \epsilon_q \langle l, s, j, I, F, m_F | \, r_{-q} \, | l', s', j', I', F', m'_F\rangle\end{aligned} \tag{1.14}$$

can be reduced using the Wigner-Eckart Theorem[35] as follows

$$\begin{aligned}\langle l, s, j, I, F, m_F | \, er_q \, | l', s', j', I', F', m'_F\rangle = \\ (-1)^{F-m_F}(l, s, j, I, F||e\mathbf{r}||l', s', j', I', F') \begin{pmatrix} F & k & F' \\ -m_F & q & m'_F \end{pmatrix}.\end{aligned} \tag{1.15}$$

The term in brackets is a Wigner $3-j$ symbol, which may also be written in terms of Clebsch-Gordan coefficients; it is related to the integral of the product of three spherical harmonics. This term is only non-zero when $m_F = m'_F + q$, where $q = -1, 0, 1$ is the angular momentum carried by, respectively, right-circular, linear, and left-circular polarisations. Here we have factored out the angular momentum dependence of m_F, leaving the other angular momentum interactions in the reduced matrix element; we can repeat the reduction process to bring out the F dependence from this element:

$$(l,s,j,I,F||e\mathbf{r}||l',s',j',I',F') = (l,s,j||e\mathbf{r}||l',s',j')(-1)^{F'+J+1+I}\sqrt{(2F'+1)(2J+1)}\begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix}. \tag{1.16}$$

The term in brackets here is a Wigner $6-j$ symbol. $3-j$ and $6-j$ symbols can always be expressed in terms of square roots of rational numbers and are easily computed (see for example the Mathematica functions `ThreeJSymbol` and `SixJSymbol`). The reduced matrix element in equation 1.16 is empirically determined from measurements of the lifetime of the transition concerned (for Rubidium 85 numbers, see [1]).

1.2.1 Selection rules

Selection rules are restrictions on the allowed dipole transitions due to conservation of angular momentum: light carries at most one unit of angular momentum and so the dipole operator can only change the angular momentum of the atom by one unit. Equation 1.15 gave us two selection rules for dipole transitions between hyperfine levels: $\Delta m_F = 0, \pm 1$, $\Delta F = 0, \pm 1$ (no $F = 0 \rightarrow F = 0$ transitions allowed). In addition to this, changes in J are restricted to $\Delta J = 0, \pm 1$ (no $0 \rightarrow 0$) and the odd parity of the dipole operator requires that $\Delta l = \pm 1$.

1.3 Cooling and trapping

When an atom scatters a photon, a transfer of momentum occurs. Reducing the width of the momentum distribution of a cloud of atoms reduces its temperature, and so if the photon scattering can be controlled to some degree then we have influence over the temperature. Magneto-optical trapping is based on engineering the likelihood of a scattering event to depend on the velocity (for cooling) and the position (for trapping) of the atom. Many comprehensive reviews of the physics of magneto-optical trapping exist ([36], [37], [38]) and in this subsection we outline the aspects relevant to our experiment.

1.3.1 Laser cooling

Consider a laser slightly detuned in frequency from a closed transition incident on an atom in the ground state. If the component of the atom's velocity along the wavevector of the laser beam is such that the frequency of the light is Doppler-shifted into resonance, an absorption event occurs. A photon with momentum $\hbar\mathbf{k}$ is absorbed by the atom. By conservation of energy, the atom is now in the excited state, and by conservation of momentum, the atom's momentum is changed by $\hbar\mathbf{k}$. The atom will then spontaneously decay back to the ground state and emit a photon in a random direction. Over a large number of absorption-spontaneous emission cycles, the momentum change due to spontaneous emission will average to zero, and the net momentum change is in the direction of propagation of the laser beam. Since the laser is red detuned, the absorption event becomes less likely as the velocity of the atom towards the laser source decreases. If lasers are incident along all three axes, the momentum of the atom is decreased in all three directions.

The momentum of a photon resonant with this transition is small compared to the average momentum of the atoms before cooling, and so a large number of spontaneous emission events must occur to cool the atoms. The force on the atoms is equal to the rate of change of momentum, or equivalently the number of photons scattered per unit time multiplied by the momentum of one photon: $\mathbf{F} = \gamma\rho_{ee}\hbar\mathbf{k}$, where ρ_{ee} is the population of the excited state and γ is the spontaneous emission rate. The form of ρ_{ee} is derived in Appendix A, and here we note that the force due to one laser beam with wavevector $\mathbf{k}$ is

$$\mathbf{F} = \gamma \frac{\Omega^2}{\gamma^2 + 4\delta^2 + 2\Omega^2}\hbar\mathbf{k}, \tag{1.17}$$

where Ω is the Rabi frequency of the transition and δ is the detuning from resonance.

Here we only consider low laser intensities $\Omega \ll \gamma$ so that the transition is not saturated. If we neglect the Ω term in the denominator and take the resultant force on an atom with velocity $\mathbf{v}$ due to two counter propagating laser beams, accounting for the Doppler shift in the detuning term, we find

$$\mathbf{F} = \gamma \left(\frac{\Omega^2}{\gamma^2 + (4\delta + \mathbf{k}\cdot\mathbf{v})^2} - \frac{\Omega^2}{\gamma^2 + (4\delta - \mathbf{k}\cdot\mathbf{v})^2} \right) \hbar\mathbf{k}. \tag{1.18}$$

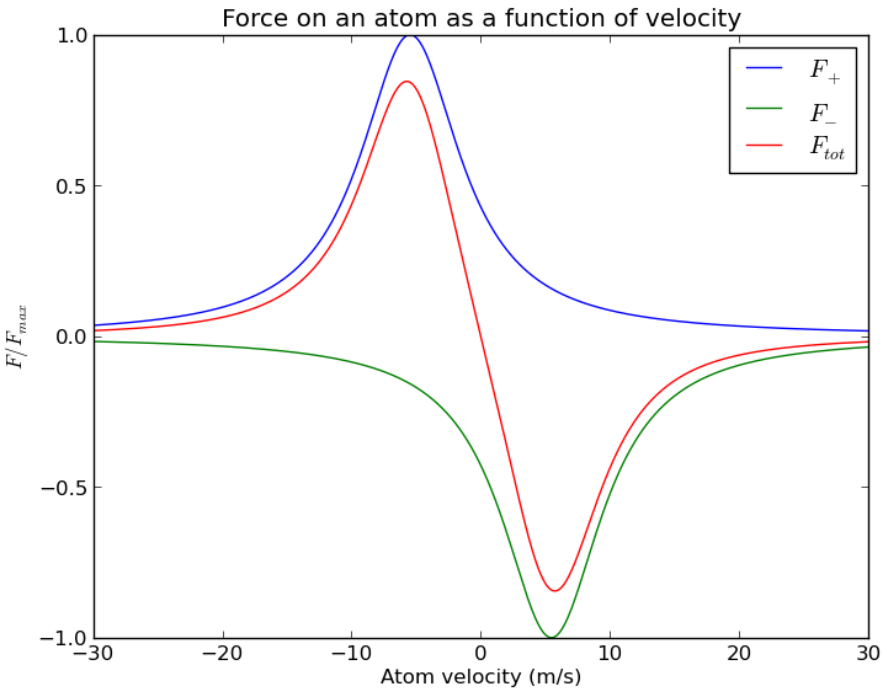

Figure 1.1: Force versus atomic velocity in one dimension. The blue and green lines are the forces due to each counterpropagating beam, and the red line is their sum. F_{max} is the maximum value of F_+ (which is equal to the maximum value of F_-).

When $\delta > 0$ the sign of this force is opposite to the sign of $\mathbf{v}$, and the force is zero when $\mathbf{v} = 0$. For small $\mathbf{v}$ the force can be approximated as

$$\mathbf{F} \approx \hbar\mathbf{k}\gamma\delta\Omega^2 \left(\frac{-16\mathbf{k}\cdot\mathbf{v}}{\gamma^2 + 4\delta^2} \right); \tag{1.19}$$

this is a frictional force which is antisymmetric in $\mathbf{v}$. A plot of this force as a fucntion of velocity is shown in Figure 1.1; the force is linear around the origin. A gas of atoms where each atom experiences this frictional force along all three axes will experience a decrease in temperature.

1.3.2 Magnetic trapping

The frictional force in equation 1.19 is independent of the position of the atoms and does not confine them to any region of space. In order to achieve trapping in addition to cooling, a position-dependent force is required. One way of providing

this is a static magnetic field in combination with careful choice of the polarisations of the MOT beams.

In 1.1.5 the Zeeman effect was described, where a magnetic field causes a shift in the hyperfine levels of $\Delta E_{Zeeman} = g_F \mu_B B m_F$. If a magnetic field can be produced which increases linearly from the centre of the trap, the hyperfine levels will be shifted accordingly and the detuning from resonance made to depend on position. The anti-Helmholtz coil configuration, where two coils with currents in opposing directions are placed coaxially, gives such a magnetic field [39], with the gradient in the z-direction being twice that in the x- and y-directions; the field is linear to a very good approximation for z positions much smaller than the radius and separation of the coils.

Figure 1.2 shows the effect of this B field for a simplified energy level scheme. As the MOT beams are red-detuned, the lower energy levels ($m_J = -1$) are shifted more into resonance with increasing $|z|$. According to angular momentum selection rules, a transition from $J = 0, m_J = 0$ to $J = 1, m_J = -1$ requires a photon with -1 unit of angular momentum. The MOT beam polarisations are chosen so that on either side of the zero magnetic field, σ^- transitions - those which involve a change in the angular momentum of the atom of -1 - are driven. This means that the beams are left-circularly polarised, i.e. from the point of view of the source, the electric field vector rotates in an anti-clockwise sense around the **k** vector of the beam.

This position-dependent force, in combination with the velocity-dependent molasses force, results in cooling and trapping.

1.3.3 Sub-Doppler cooling

When Doppler cooling was first realised, it was thought that the smallest temperature that could be reached (the Doppler limit T_D [36]) was determined by the competition between the cooling effect and the heating caused by spontaneous emission and was dependent on the natural linewidth of the transition:

$$k_B T_D \approx \frac{\hbar\gamma}{2}. \tag{1.20}$$

However, several years later, the temperature of Doppler-cooled sodium atoms was measured to be more than 5 times smaller than the Doppler limit [40]. This was

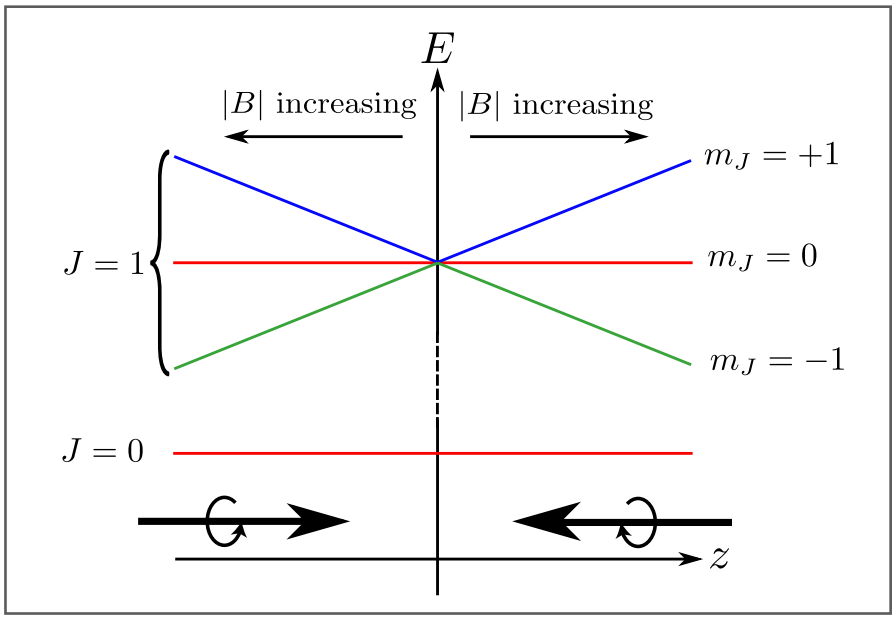

Figure 1.2: Position-dependent Zeeman shift of a simplified energy level structure with a linearly increasing B field.

discovered to be due to the interaction of the resultant polarisation of the MOT beams with the light-shifted energy levels of the atoms, which is not taken into account in the standard analysis of Doppler cooling. A full theoretical analysis of polarisation gradient cooling mechanisms may be found in [41] and here we simply give a qualitative description.

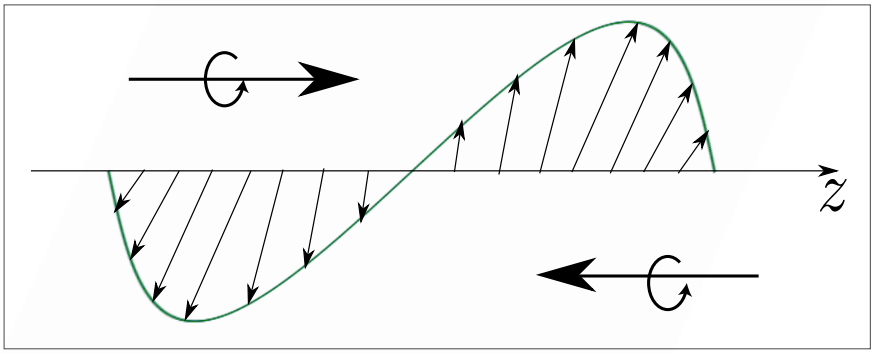

Figure 1.3: Two counter-propagating oppositely circular polarised field produce a resultant linear polarisation rotating in space. This is a view along an axis perpendicular to the wavevector of the light; the polarisation vector of the light follows the green path, rotating around the wavevector of the light.

Two counterpropagating beams with opposite circular polarisations set up a resultant polarisation which varies along their propagation axis. If the intensities of the two beams are equal, this polarisation is linear and rotates in space (Figure 1.3). The physics of an atom moving within this polarisation gradient is formally

equivalent to that of a stationary atom in a static magnetic field along z, interacting with a laser linearly polarised along y. The interaction between the magnetic sublevels of the atom and the magnetic field gives rise to a population difference within the sublevels which is sensitive to the atom's velocity; this means that the radiation pressures exerted by left and right circular polarisations are unbalanced. The resulting force is opposite to the atom's velocity and is thus a frictional force.

The coefficient of friction of this force is inversely proportional to the light shift produced by the electric fields and so the cooling mechanism is more efficient at low laser powers; to take advantage of this, in our experiment the power of the MOT lasers is ramped down after formation of the MOT.

Chapter 2

Two-photon transitions

In order to implement an atom interferometer, we require a method for the preparation, manipulation, and detection of atoms in particular quantum states. The two atomic states corresponding to the two interferometer paths must be stable on the timescale of the interferometer sequence (in our experiment, this would be of the order of microseconds), and so we choose two hyperfine sublevels of the ground state and implement two-photon (Raman) transitions between them.

In this chapter the theory of two-photon transitions is introduced and the equations describing the evolution of the amplitudes of the two interferometer states are derived. The Bloch sphere picture is outlined and we consider how the particular conditions of our experiment influence the paths taken by the state vectors during a pulse sequence.

2.1 Introduction

A quantum state is described by a list of amplitudes for a defined property of the system. In the Dirac notation we may represent this list of amplitudes as a state vector $|\psi\rangle$. The time evolution of a quantum state is determined by the Time-Dependent Schrödinger Equation (TDSE)

$$i\hbar\frac{\partial|\psi\rangle}{\partial t} = \hat{H}|\psi\rangle \tag{2.1}$$

where $\hat{H}$ is the Hamiltonian of the quantum system: the operator describing the total energy of the system. The method of determining the time evolution of a quantum system begins with writing down the Hamiltonian.

2.2 Three level atom in an electromagnetic field

2.2.1 Hamiltonian

We consider a 3-level atom interacting with two lasers [42]. The splitting between levels $|g\rangle$ and $|e\rangle$ is much smaller than that between $|g\rangle$ or $|e\rangle$ and $|i\rangle$. Two lasers with frequencies $\omega_1 \simeq \omega_i - \omega_g$ and $\omega_2 \simeq \omega_i - \omega_e$ interact with the atom (Figure 2.1).

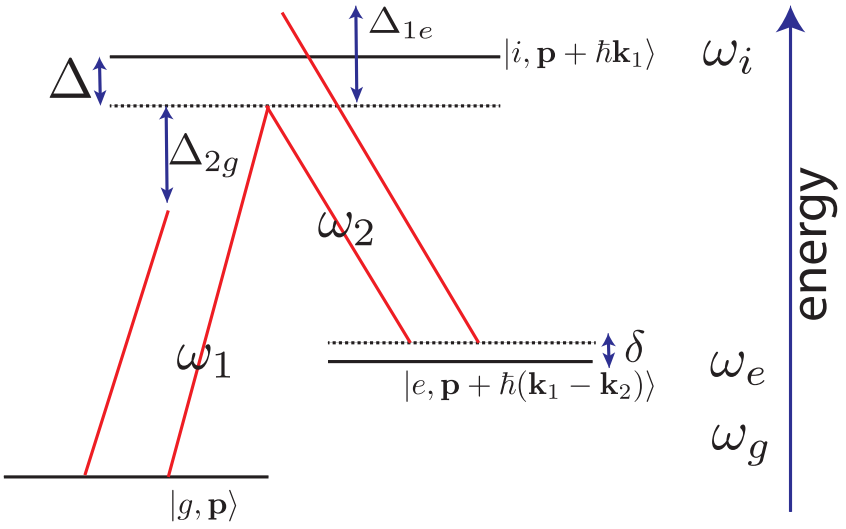

Figure 2.1: Atom-laser system

The Hamiltonian describing the atom-laser system is

$$\hat{H} = \frac{\hat{\mathbf{p}}^2}{2m} + \hbar\omega_e |e\rangle\langle e| + \hbar\omega_g |g\rangle\langle g| + \hbar\omega_i |i\rangle\langle i| - \mathbf{d}\cdot\mathbf{E}, \tag{2.2}$$

where the driving electric field has two frequency components

$$\mathbf{E} = \mathbf{E}_1 \cos(\mathbf{k}_1\cdot\mathbf{x} - \omega_1 t + \phi_1) + \mathbf{E}_2 \cos(\mathbf{k}_2\cdot\mathbf{x} - \omega_2 t + \phi_2) \tag{2.3}$$

and $\mathbf{d} = -e\mathbf{r}$ is the dipole moment of the atom.

The basis states have a momentum part which is operated on by the spatial part of the E-field $e^{\pm i\mathbf{k}_{1,2}\cdot\mathbf{x}}$. Considering the eigenstates of momentum in the position representation $\psi_{\mathbf{k}}(\mathbf{x}) = \frac{1}{\sqrt{2\pi}^3}e^{i\mathbf{k}\cdot\mathbf{x}}$, it is clear that the spatial part of the E-field operator acts to change the momentum of the atom:

$$\begin{aligned} e^{\pm i\mathbf{k}_{1,2}\cdot\mathbf{x}}\psi_{\mathbf{k}}(\mathbf{x}) &= \frac{1}{\sqrt{2\pi}^3}e^{i(\mathbf{k}\pm\mathbf{k}_{1,2})\cdot\mathbf{x}} \\ &= \psi_{\mathbf{k}\pm\mathbf{k}_{1,2}}(\mathbf{x}). \end{aligned} \tag{2.4}$$

Thus the absorption or emission of a photon of momentum $\hbar\mathbf{k}$ changes the momentum of the atom by the same amount, and there is a one-to-one correspondance between the internal (electronic state) and external (momentum) degrees of freedom of the atom: an atom beginning in a state $|g, \mathbf{p}\rangle$ will end in the state $|i, \mathbf{p} + \hbar\mathbf{k}\rangle$ upon absorption of a photon with momentum $\hbar\mathbf{k}$ when on resonance with the single photon transitions. We can now proceed to solve the Hamiltonian for the internal degrees of freedom.

The electric dipole coupling term $-\mathbf{d}\cdot\mathbf{E}$ incorporates the dipole moment of the single outer shell electron $\mathbf{d} = |e|\mathbf{r}$, where $\mathbf{r}$ is the vector drawn from the atomic nucleus to the position of the electron. We may write the element of this operator involving levels a and b in terms of the intensity-dependent Rabi frequency $\Omega_{ab}^{(n)}$ of the transition being coupled, summing over the $n = 2$ E-fields $\mathbf{E}_n$ present:

$$\begin{aligned} V_{ab} &= \sum_n -\langle a|\mathbf{d}\cdot\mathbf{E}_n|b\rangle \\ &= \sum_n \hbar\Omega_{ab}^{(n)}\frac{e^{i(\mathbf{k}_n\cdot\mathbf{z}+(-1)^n\omega_n t+\phi_n)} + e^{-i(\mathbf{k}_n\cdot\mathbf{z}+(-1)^n\omega_n t+\phi_n)}}{2}. \end{aligned} \tag{2.5}$$

When considering the energy level structure in Figure 2.1, we can see that the matrix elements V_{eg} and $V_{eg}^* = V_{ge}$ must be equal to zero as there is no dipole-allowed transition possible between $|g\rangle$ and $|e\rangle$, and so our Hamiltonian will only include terms coupling levels g and e with i.

The full Hamiltonian can be written as

$$\begin{aligned}\hat{H} = \frac{\hat{\mathbf{p}}^2}{2m} + \hbar\omega_e \left|e\right\rangle \left\langle e\right| + \hbar\omega_g \left|g\right\rangle \left\langle g\right| + \hbar\omega_i \left|i\right\rangle \left\langle i\right| \\ - \left(\left(V_{ig} \left|g\right\rangle + V_{ie} \left|e\right\rangle\right) \left\langle i\right| + V_{gi} \left|i\right\rangle \left\langle g\right| + V_{ei} \left|i\right\rangle \left\langle e\right|\right)\end{aligned} \tag{2.6}$$

and we use this Hamiltonian in the TDSE 2.1, with $\left|\psi\right\rangle = a_g \left|g\right\rangle + a_e \left|e\right\rangle + a_i \left|i\right\rangle$, to determine the time evolution of the state ampitudes $a_{g,e,i}$

$$\begin{aligned}i\hbar\dot{a_g}(t) &= \left(\frac{\mathbf{p}^2}{2m} + \hbar\omega_g\right) a_g(t) + V_{ig} a_i(t) \\ i\hbar\dot{a_e}(t) &= \left(\frac{\left(\mathbf{p} + \hbar(\mathbf{k}_1 - \mathbf{k}_2)\right)^2}{2m} + \hbar\omega_e\right) a_e(t) + V_{ie} a_i(t) \\ i\hbar\dot{a_i}(t) &= \left(\frac{\left(\mathbf{p} + \hbar\mathbf{k}_1\right)^2}{2m} + \hbar\omega_i\right) a_i(t) + V_{gi} a_g(t) + V_{ei} a_e(t)\end{aligned} \tag{2.7}$$

Now we consider the elements V_{ig} and V_{ie}, which include coupling terms from both lasers:

$$\begin{aligned}V_{ig} &= \frac{\hbar\Omega_{ig}^{(1)}}{2}\left(e^{i(\omega_1 t + \phi_1)} + e^{-i(\omega_1 t + \phi_1)}\right) + \frac{\hbar\Omega_{ig}^{(2)}}{2}\left(e^{i(\omega_2 t + \phi_2)} + e^{-i(\omega_2 t + \phi_2)}\right) \\ V_{ie} &= \frac{\hbar\Omega_{ie}^{(1)}}{2}\left(e^{i(\omega_1 t + \phi_1)} + e^{-i(\omega_1 t + \phi_1)}\right) + \frac{\hbar\Omega_{ie}^{(2)}}{2}\left(e^{i(\omega_2 t + \phi_2)} + e^{-i(\omega_2 t + \phi_2)}\right).\end{aligned} \tag{2.8}$$

Before continuing it is useful to define several detunings. Firstly, the single photon detuning Δ which is large enough that spontaneous emission from the intermediate state $\left|i\right\rangle$ may be neglected, and takes into account the Doppler shift due to the momentum of the atom and the change in this momentum effected by absorption of a photon:

$$\Delta = \left(\frac{\left(\mathbf{p} + \hbar\mathbf{k}_1\right)^2}{2m\hbar} + \omega_i\right) - \left(\frac{\mathbf{p}^2}{2m\hbar} + \omega_g\right) - \omega_1. \tag{2.9}$$

The detuning from the Raman resonance δ_{12} depends on the detuning of the laser from the hyperfine splitting and the momentum of the atom:

$$\delta = \left(\frac{\left(\mathbf{p}+\hbar(\mathbf{k}_1-\mathbf{k}_2)\right)^2}{2m\hbar}+\omega_e\right)-\left(\frac{\mathbf{p}^2}{2m\hbar}+\omega_g\right)-(\omega_2-\omega_1) \tag{2.10}$$

Additionally, we define the detunings of lasers 1 and 2 from the single photon transitions with which they are not resonant. The detuning of the laser with frequency ω_1 from the transition $|e\rangle$ to $|i\rangle$ is

$$\Delta_{1e} = \left(\frac{\left(\mathbf{p}+\hbar\mathbf{k}_1\right)^2}{2m\hbar}+\omega_i\right)-\left(\frac{\mathbf{p}^2}{2m\hbar}+\omega_g\right)-\omega_1 \tag{2.11}$$

and the detuning of the laser with frequency ω_2 from the transition $|g\rangle$ to $|i\rangle$ is

$$\Delta_{2g} = \left(\frac{\left(\mathbf{p}+\hbar\mathbf{k}_1\right)^2}{2m\hbar}+\omega_i\right)-\left(\frac{\mathbf{p}^2}{2m\hbar}+\omega_g\right)-\omega_2. \tag{2.12}$$

We now employ the rotating wave approximation. We may factor out the evolution at the atom frequency from each state amplitude by writing $a_m(t) = c_m(t)e^{i(\mathbf{p}_m/2m\hbar+\omega_m)t}$, where $m = g, e, i$ and $\mathbf{p}_m$ are the momentum terms in equations 2.7. We substitute these into 2.7.

We find the amplitude equations reduce to

$$i\hbar\dot{c_g}(t) = \hbar c_i\left[\frac{\Omega_{ig}^{(1)}}{2}\left(e^{-i((\Delta+2\omega_1)t-\phi_1)}+e^{-i(\Delta t+\phi_1)}\right)+\frac{\Omega_{ig}^{(2)}}{2}\left(e^{-i(\Delta_{2g}t-\phi_2)}+e^{-i((\Delta_{2g}+2\omega_2)t+\phi_2)}\right)\right] \tag{2.13}$$

$$\begin{aligned} i\hbar\dot{c_e}(t) = \hbar c_i\Bigg[&\frac{\Omega_{ie}^{(2)}}{2}\left(e^{-i(((\Delta+\delta_{12})+2\omega_2)t+\phi_2)}+e^{-i((\Delta+\delta_{12})t-\phi_2)}\right)\\ &+\frac{\Omega_{ie}^{(1)}}{2}\left(e^{-i(\Delta_{1e}t+\phi_1)}+e^{-i((\Delta_{1e}+2\omega_1)t-\phi_1)}\right)\Bigg] \end{aligned} \tag{2.14}$$

$$i\hbar\dot{c}_i(t) = \hbar c_g \left[\frac{\Omega_{gi}^{(1)}}{2} \left(e^{i(((\Delta+2\omega_1)t-\phi_1)} + e^{i(\Delta t+\phi_1)} \right) + \frac{\Omega_{gi}^{(2)}}{2} \left(e^{i(\Delta_{2g}t-\phi_2)} + e^{i((\Delta_{2g}+2\omega_2)t+\phi_2)} \right) \right] + \hbar c_e \left[\frac{\Omega_{ei}^{(2)}}{2} \left(e^{i((((\Delta+\delta_{12})+2\omega_2)t+\phi_2)} + e^{i((\Delta+\delta_{12})t-\phi_2)} \right) + \frac{\Omega_{ei}^{(1)}}{2} \left(e^{i(\Delta_{1e}t+\phi_1)} + e^{i((\Delta_{1e}+2\omega_1)t-\phi_1)} \right) \right] \tag{2.15}$$

If we consider the integral of the first of these equations, we see that the first and fourth terms under the integral oscillate at a much higher frequency than the second and third terms. The first and fourth terms will "average out" to zero and we can drop these terms from our analysis. Similarly we remove the higher frequency terms from the expressions for $\dot{c}_e$ and $\dot{c}_i$, leaving the amplitude equations as

$$i\hbar\dot{c}_g(t) = \hbar c_i \left[\frac{\Omega_{ig}^{(1)}}{2} e^{-i(\Delta t+\phi_1)} + \frac{\Omega_{ig}^{(2)}}{2} e^{-i(\Delta_{2g}t-\phi_2)} \right] \tag{2.16}$$

$$i\hbar\dot{c}_e(t) = \hbar c_i \left[\frac{\Omega_{ie}^{(2)}}{2} e^{-i((\Delta+\delta_{12})t-\phi_2)} + \frac{\Omega_{ie}^{(1)}}{2} e^{-i(\Delta_{1e}t+\phi_1)} \right] \tag{2.17}$$

$$i\hbar\dot{c}_i(t) = \hbar c_g \left[\frac{\Omega_{gi}^{(1)}}{2} e^{i(\Delta t+\phi_1)} + \frac{\Omega_{gi}^{(2)}}{2} e^{i(\Delta_{2g}t-\phi_2)} \right] + \hbar c_e \left[\frac{\Omega_{ei}^{(2)}}{2} e^{i((\Delta+\delta_{12})t-\phi_2)} + \frac{\Omega_{ei}^{(1)}}{2} e^{i(\Delta_{1e}t+\phi_1)} \right] \tag{2.18}$$

2.2.2 Elimination of intermediate state

The relative magnitudes of the various oscillation frequencies Δ, δ_{12} and $\Omega_{ig,e}^{(1,2)}$ are important in completing a reduction of the three level system to a pseudo-two level system. If the rates of change of c_e and c_g are small compared to the detuning from single photon resonance ($\Delta \gg \Omega_{ig,e}^{(1,2)}$), we can integrate the equation for $\dot{c}_i$ and consider $c_{e,g}$ to be effectively constant under the integral. We find the following expression for c_i:

$$c_i(t) = c_g \left(-\frac{\Omega_{gi}^{(1)}}{2\Delta} e^{i(\Delta t+\phi_1)} - \frac{\Omega_{gi}^{(2)}}{2\Delta_{2g}} e^{i(\Delta_{2g}t-\phi_2)} \right) + c_e \left(-\frac{\Omega_{ei}^{(2)}}{2(\Delta+\delta_{12})} e^{i((\Delta+\delta_{12})t-\phi_2)} - \frac{\Omega_{ei}^{(1)}}{2\Delta_{1e}} e^{i(\Delta_{1e}t+\phi_1)}) \right). \tag{2.19}$$

Substituting this relation into equations 2.16 and 2.17, we find terms oscillating at δ_{12}, at ω_{eg}, at $\delta_{12}+\omega_{eg}$ and at $\delta_{12}+2\omega_{eg}$. Since $\omega_{eg} \gg \delta_{12}$, we eliminate all terms with ω_{eg} in the oscillation frequency and find:

$$i\dot{c_g} = -\frac{|\Omega_{ig}^{(1)}|^2}{4\Delta} c_g - \frac{|\Omega_{ig}^{(2)}|^2}{4\Delta_{2g}} c_g - \frac{\Omega_{ig}^{(1)}\Omega_{ei}^{(2)}}{4\Delta} c_e e^{i(\delta_{12}t-\phi)} \tag{2.20}$$

$$i\dot{c_e} = -\frac{|\Omega_{ie}^{(2)}|^2}{4\Delta} c_e - \frac{|\Omega_{ie}^{(1)}|^2}{4\Delta_{1e}} c_e - \frac{\Omega_{ie}^{(2)}\Omega_{gi}^{(1)}}{4\Delta} c_g e^{-i(\delta_{12}t-\phi)} \tag{2.21}$$

where we have combined the phases of the two lasers as $\phi = \phi_2 - \phi_1$. We see that the state amplitude for $|i\rangle$ is oscillating at a much higher frequency than the amplitudes for $|g\rangle$ and $|e\rangle$. We can thus eliminate this state from our analysis on the basis that its fast oscillations will average to zero over the timescales we are considering.

To continue our analysis, it is helpful to contruct the matrix representation of this system. With the correspondence $|g\rangle$ to $(0,1)^T$ and $|e\rangle$ to $(1,0)^T$, we may use equations 2.20, 2.21 to write the Hamiltonian $\hat{H}$ in the form:

$$\hat{H} = \hbar \begin{pmatrix} \frac{|\Omega_{ig}^{(1)}|^2}{4\Delta} + \frac{|\Omega_{ig}^{(2)}|^2}{4\Delta_{2g}} & \frac{\Omega_{ie}^{(2)}\Omega_{gi}^{(1)}}{2\Delta} e^{-i(\delta_{12}t-\phi)} \\ \frac{\Omega_{ig}^{(1)}\Omega_{ei}^{(2)}}{\Delta} e^{i(\delta_{12}t-\phi)} & \frac{|\Omega_{ie}^{(2)}|^2}{4\Delta} + \frac{|\Omega_{ie}^{(1)}|^2}{4\Delta_{1e}} \end{pmatrix}. \tag{2.22}$$

The diagonal terms are the AC Stark shifts of the ground and excited states. We note here that the effect of lasers on transitions with which they are not resonant (ω_2 on the $|g\rangle$ to $|i\rangle$ transition and ω_1 on the $|e\rangle$ to $|i\rangle$ transition) is to add an extra term into the AC Stark shifts; the coupling (off-diagonal) terms are not affected. The off-diagonal terms may be written in terms of the effective Rabi frequency $\Omega_{eff} = \frac{\Omega_{ie}^{(2)}\Omega_{gi}^{(1)}}{2\Delta}$, giving

$$\hat{H} = \hbar \begin{pmatrix} \Omega_e^{AC} & \frac{\Omega_{eff}}{2} e^{-i(\delta_{12}t-\phi)} \\ \frac{\Omega_{eff}}{2} e^{i(\delta_{12}t-\phi)} & \Omega_g^{AC} \end{pmatrix}. \tag{2.23}$$

The relative AC Stark shift of the ground and excited states is $\delta_{AC} = \Omega_e^{AC} - \Omega_g^{AC}$. If we shift the energy scale uniformly by $-\hbar(\Omega_e^{AC} + \Omega_g^{AC})/2$, the two diagonal elements are antisymmetrised and the Hamiltonian becomes

$$\hat{H} = \frac{\hbar}{2} \begin{pmatrix} \delta_{AC} & \Omega_{eff} e^{-i(\delta_{12}t-\phi)} \\ \Omega_{eff} e^{i(\delta_{12}t-\phi)} & -\delta_{AC} \end{pmatrix}. \tag{2.24}$$

2.2.3 The rotating frame

This Hamiltonian is time-dependent and its solution is computationally intensive in all but the most trivial of cases. By transforming to a frame rotating about the z axis[1] at some frequency, it can be written in a time-independent form. This is useful because the time evolution of an arbitrary state $|\psi\rangle = a|\alpha\rangle + b|\beta\rangle$ under a time-independent Hamiltonian with eigenstates $|\alpha\rangle, |\beta\rangle$ and corresponding eigenvalues E_α, E_β is easily written down as $|\psi(t)\rangle = e^{-iE_\alpha t/\hbar}|\alpha\rangle + e^{-iE_\beta t/\hbar}|\beta\rangle$; hence, by transforming a state in the interaction picture to the rotating frame, evolving for a time τ, and transforming back to the interaction picture, the time evolution of that state can be easily calculated.

A rotation is a unitary operation, and the matrix U repesenting a rotation is unitary, which is to say $U^\dagger U = I$ (its inverse is equal to its Hermitian conjugate). This property is useful in constructing the form of the TDSE in a rotating frame. Define the operator representing a rotation of $-\delta t$ about the z axis as:

$$D(\mathbf{z}, -\delta t) = \exp(-i\sigma_z \delta t/2) = \begin{pmatrix} e^{-i\delta t/2} & 0 \\ 0 & e^{i\delta t/2} \end{pmatrix}, \tag{2.25}$$

and the transformation of states $|\psi\rangle$ from the reference frame of $\hat{H}$ to states $|\psi\rangle_R$ in the rotating frame as $|\psi\rangle_R = D|\psi\rangle$. Replacing $|\psi\rangle$ with $D^\dagger|\psi\rangle_R$ in equation 2.1, the Hamiltonian in the rotating frame $\hat{H}_R$ is found to be

[1] This is not the spatial z axis. The axis is along the state vector of the eigenstate on the Bloch sphere (see section 2.3) which is oriented in the positive z direction.

$$\hat{H}_R = D\hat{H}D^\dagger - i\hbar D\frac{\mathrm{d}D^\dagger}{\mathrm{d}t}. \tag{2.26}$$

Applying this to our Hamiltonian 2.24, we find the form of the Hamiltonian in the rotating frame:

$$\hat{H}_R = \frac{\hbar}{2}\begin{pmatrix} -(\delta_{12}-\delta_{AC}) & \Omega_{eff}e^{-i\phi} \\ \Omega_{eff}e^{i\phi} & \delta_{12}-\delta_{AC} \end{pmatrix}. \tag{2.27}$$

The next step is to find the eigenvalues $\lambda_\pm$ and eigenvectors $|\lambda_\pm\rangle = (a_\pm, b_\pm)$ of $\hat{H}_R$, defined by

$$\frac{\hbar}{2}\begin{pmatrix} -(\delta_{12}-\delta_{AC}) & \Omega_{eff}e^{-i\phi} \\ \Omega_{eff}e^{i\phi} & \delta_{12}-\delta_{AC} \end{pmatrix}\begin{pmatrix} a_\pm \\ b_\pm \end{pmatrix} = \lambda_\pm \begin{pmatrix} a_\pm \\ b_\pm \end{pmatrix}. \tag{2.28}$$

The two eigenvalues $\lambda_\pm$ are easily found by solving $\det(\lambda I - H_R) = 0$, giving $\lambda_\pm = \pm\hbar\Omega_r/2$, where we define the rotating frame Rabi frequency as $\Omega_r = \sqrt{{\Omega_{eff}}^2 + (\delta_{12}-\delta_{AC})^2}$. Considering the frequency terms to correspond to side lengths of a right-angled triangle as in Figure 2.2, we see that the trigonometric functions of the angle Θ may be written in terms of these frequency terms. We use these functions in the eigenvector expressions.

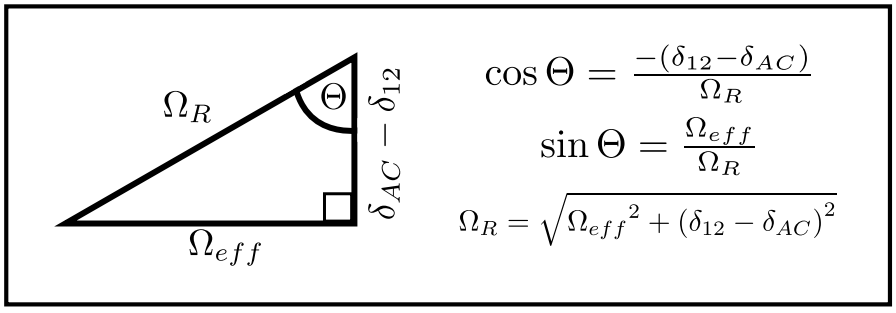

Figure 2.2: Relationship between frequencies and detunings

Following e.g. reference [33], by multiplying out 2.28 and using the expression found for the eigenvalues, we find the eigenvectors

$$\begin{aligned} |\lambda_+\rangle &= \sin\frac{\Theta}{2}e^{-i\phi/2}|g\rangle_R + \cos\frac{\Theta}{2}e^{i\phi/2}|e\rangle_R \\ |\lambda_-\rangle &= -\sin\frac{\Theta}{2}e^{i\phi/2}|e\rangle_R + \cos\frac{\Theta}{2}e^{-i\phi/2}|g\rangle_R. \end{aligned} \tag{2.29}$$

2.2.4 State amplitudes

We now have everything required to calculate the time evolution of the state amplitudes. If we begin at time t_0 with a state $|\psi(t_0)\rangle = c_e(t_0)|e\rangle + c_g(t_0)|g\rangle$ and wish to calculate the state amplitudes $c_g(t_0+\tau), c_e(t_0+\tau)$ after a period τ of evolution, we follow the procedure:

1. Transform to the rotating frame: $|\psi(t_0)\rangle_R = D(\mathbf{z}, -\delta_{12}t_0)|\psi(t_0)\rangle$

2. Project $|\psi(t_0)\rangle_R$ onto the basis of eigenstates of $\hat{H}_r$ and evolve for time τ: $|\psi(t_0+\tau)\rangle_R = \left(e^{-i\lambda_+\tau/\hbar}|\lambda_+\rangle\langle\lambda_+| + e^{-i\lambda_-\tau/\hbar}|\lambda_-\rangle\langle\lambda_-|\right)|\psi(t_0)\rangle_R$

3. Transform back to the interaction picture: $|\psi(t_0+\tau)\rangle = D^\dagger(\mathbf{z}, -\delta_{12}(t_0+\tau))|\psi(t_0+\tau)\rangle_R$

The solutions to the Hamiltonian are [42]

$$c_{e,\mathbf{p}+\hbar\mathbf{k}_{eff}}(t_0+\tau) = e^{-i(\Omega_e^{AC}+\Omega_g^{AC})\tau/2}e^{-i\delta_{12}\tau/2}\Bigg(c_{e,\mathbf{p}+\hbar\mathbf{k}_{eff}}(t_0)\left[\cos(\frac{\Omega_r\tau}{2}) - i\cos\Theta\sin(\frac{\Omega_r\tau}{2})\right] + c_{g,\mathbf{p}}(t_0)e^{-i(\delta_{12}t_0+\phi_{eff})}\left[-i\sin\Theta\sin(\frac{\Omega_r\tau}{2})\right]\Bigg) \tag{2.30}$$

$$c_{g,\mathbf{p}}(t_0+\tau) = e^{-i(\Omega_e^{AC}+\Omega_g^{AC})\tau/2}e^{i\delta_{12}\tau/2}\Bigg(c_{e,\mathbf{p}+\hbar\mathbf{k}_{eff}}(t_0)e^{i(\delta_{12}t_0+\phi_{eff})}\left[-i\sin\Theta\sin(\frac{\Omega_r\tau}{2})\right] + c_{g,\mathbf{p}}(t_0)\left[\cos(\frac{\Omega_r\tau}{2}) + i\cos\Theta\sin(\frac{\Omega_r\tau}{2})\right]\Bigg), \tag{2.31}$$

with

$$\sin\Theta = \Omega_{eff}/\Omega_r, \quad \cos\Theta = -(\delta_{12}-\delta_{AC})/\Omega_r, \qquad \Omega_r = \sqrt{{\Omega_{eff}}^2 + (\delta_{12}-\delta_{AC})^2}. \tag{2.32}$$

We can choose some typical experimental values and examine the behaviour of the calculated state amplitudes. We set the Raman beam intensity to $25\mathrm{kWm}^{-2}$ and the single-photon detuning to $\Delta = 2\pi \times 11\,\mathrm{GHz}$. The detuning from two-photon resonance δ is set such that an atom with zero momentum starting in the hyperfine state with $m_F = 0$ is on resonance[2]. The Raman beam polarisations are both right handed circular.

We first calculate some typical values for the frequencies involved. For an atom starting in $m_F = 0$, the single-photon Rabi frequency is approximately 850MHz. The Raman Rabi frequency is 2.6MHz, and the AC Stark shifts of the ground and excited states are 6.2MHz and 4.6MHz. If the intensity of the Raman lasers were to fluctuate by ±5%, the Stark shift would fluctuate by the same percentage. However, a ±5% fluctuation of the detuning Δ results in a fluctuation in the Stark shift of ±6.5%.

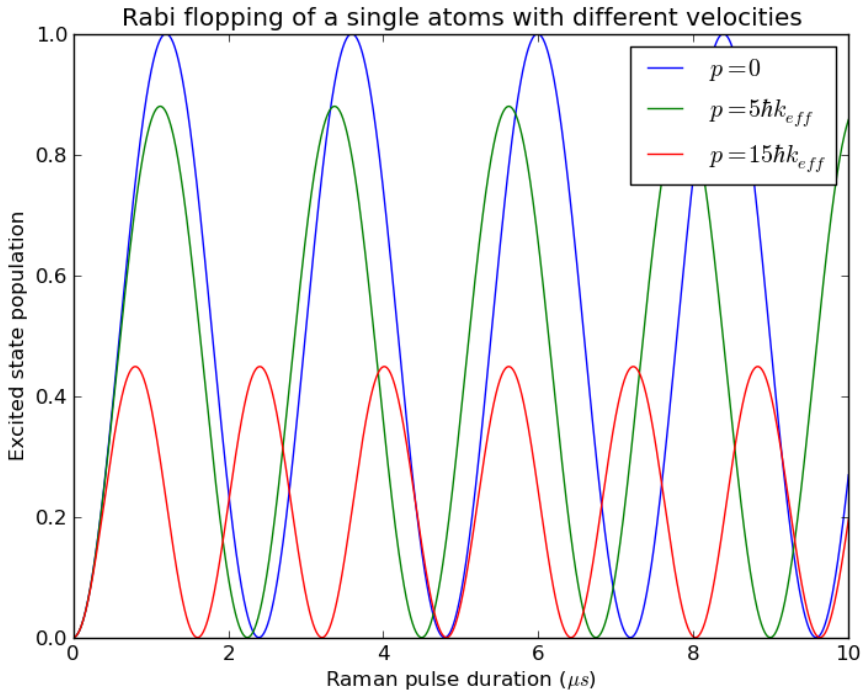

Figure 2.3: Excited state population as a function of Raman pulse duration for three atom momenta. Blue is on resonance, red and green are off resonance.

For an atom which is resonant with the driving field, $\delta_{12} - \delta_{AC} = 0$ and the terms proportional to $\cos\Theta$ vanish. We thus expect to see oscillation of population between the two states with a period equal to $2\pi/\Omega_{eff}$; this oscillation is Rabi flopping, and is shown in Figure 2.3 for atoms with several different velocities. In

[2]See Chapter 8 for a more detailed discussion of the simulations in this thesis.

the on-resonance case (the blue curve in Figure 2.3), all population is transferred coherently from the ground state to the excited state after a half period. For an atom which is off-resonance, achieved in the simulation by setting its momentum to a value such that $\delta_{12} - \delta_{AC} \neq 0$, the maximum population transfer is less than 1 (the green and red curves).

The off-resonance atoms, in addition to not achieving 100% population transfer, also have a different Rabi frequency to the on-resonance atom. This is because the Rabi frequency depends on the detuning from resonance (Equation 2.32). This means that the atoms do not remain in phase with each other. However, for all atoms, the amplitude of the Rabi oscillations remains constant in time; this is because there is no decoherence in the system[3].

Rabi flopping is simple to implement and analyse, but not particularly useful for coherent manipulation. Applying pulses of different lengths and phases allows the atom to be driven into different states. These pulse sequences will be introduced in subsequent chapters. To aid with their analysis we now outline a method of visualising quantum states and their evolution under the influence of a driving field.

2.3 Bloch sphere representation

The Bloch sphere provides a visual representation of a two-level quantum system under the influence of a perturbation, and is analogous to the Poincaré sphere used in optics to describe polarisation states. A detailed introduction can be found in [43], and is summarised briefly as follows.

Consider a two level system with Hamiltonian $H = H_0 + V$, where the background Hamiltonian H_0 is time-independent and V is some (possibly time-dependent) perturbation. Let the eigenstates of H_0 be $|a\rangle$ and $|b\rangle$, with corresponding energies $E_0 \pm \hbar\omega_0/2$; then an arbitrary state may be written as $|\psi\rangle = c_a(t)\,|a\rangle + c_b(t)\,|b\rangle$. The elements of V are $V_{mn} = \langle m|\,V\,|n\rangle$. The four parameters required to specify the real and imaginary parts of c_a, c_b are reduced to three when we note that the overall phase of the state is irrelevant. Feynman et al. [44] proceed by defining a vector $\mathbf{r} = (r_1, r_2, r_3)$ with the components

[3] We consider only dilute atomic gases, where collisions between atoms are rare enough to be neglected on the timescale of the experiments.

$$\begin{aligned} r_1 &\equiv ab^* + ba^* \\ r_2 &\equiv i(ab^* - ba^*) \\ r_3 &\equiv aa^* - bb^*. \end{aligned} \tag{2.33}$$

The time evolution of each component is found from the TDSE and the differential equation for $\mathbf{r}$ is thus

$$\frac{d\mathbf{r}}{dt} = \boldsymbol{\omega} \times \mathbf{r}, \tag{2.34}$$

where the driving field vector $\boldsymbol{\omega} = ((V_{ab} + V_{ba})/\hbar, i(V_{ab} - V_{ba})/\hbar, \omega_0)$. The state is represented by a vector of constant magnitude moving on the surface of a sphere (the Bloch sphere). The two eigenstates are located at the poles, and equally weighted superpositions on the equator of the sphere (Figure 2.4).

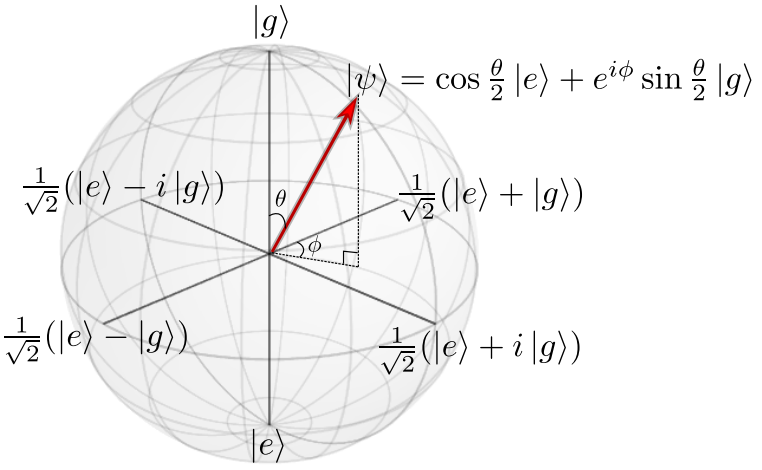

Figure 2.4: States on the Bloch sphere.

Figure 2.5 shows the trajectory of a state initially in $|g\rangle$ driven by an on-resonance field. The state vector moves in a direction perpendicular to itself and to the driving field vector.

To gain some insight into the representation on the Bloch sphere of coherent manipulations in our experiment, we revisit the rotating frame Hamiltonian.

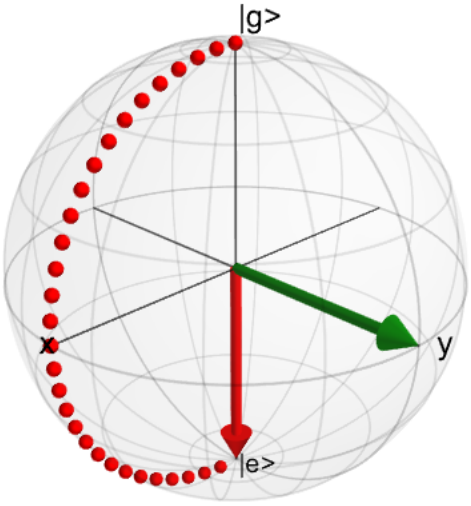

Figure 2.5: State vector trajectory (red points), final state vector (red arrow) and driving field vector (green arrow, constant) during a pi or inversion pulse.

2.3.1 The rotating frame Hamiltonian

The rotating frame Hamiltonian of our system is

$$\hat{H}_R = \frac{\hbar}{2}\begin{pmatrix} \delta^{AC} - \delta_{12} & \Omega_{eff}e^{-i\phi_{eff}} \\ \Omega_{eff}e^{i\phi_{eff}} & -(\delta^{AC} - \delta_{12}) \end{pmatrix} \tag{2.35}$$

and we see that the axis of rotation is $\boldsymbol{\omega} = (\Omega_{eff}\cos\phi_{eff}, -\Omega_{eff}\sin\phi_{eff}, \delta_{AC} - \delta_{12})$. There are two important features of $\boldsymbol{\omega}$ in our system.

Phase of the laser: the phase of the laser field ϕ_{eff} is a controllable parameter which can be used to influence the axis of rotation of the state vector; for example, a phase shift of π causes the rotation axis to be rotated through π about the z axis. This will be crucial in implementing error compensation schemes.

Detuning from resonance: the diagonal elements of $\hat{H}_R$ describe the detuning from resonance of a particular atom. This is dependent on the light shifts present, the velocity of the atom, and its hyperfine transition route(s). As noted in the previous section, the detuning affects the rotation rate $|\boldsymbol{\omega}| = \sqrt{{\Omega_{eff}}^2 + (\delta_{AC} - \delta_{12})^2}$ which is not the same for all the atoms. It also affects the axis about which the state vector rotates, causing it to be tilted out of the $x - y$ plane; this means that the path of the state vector driven by an off-resonant field is not a great circle as in the on-resonance case.

2.4 Polarisations and transition routes

Section 1.2 outlined the rules governing dipole transitions for a particular polarisation of the electric field. In this section we discuss the transition routes for two-photon transitions between the levels that concern us in Rubidium 85 and the requirements on the polarisations of the Raman beams.

Figure 2.6 shows the hyperfine structure of the Rubidium 85 D_2 transition. The two energy levels used as our ground and excited states are the $F = 2$ and $F = 3$ hyperfine sublevels of the $5S_{1/2}$ state. Each state has $2F + 1$ magnetic sublevels with m_F values in the range $-F \leq m_F \leq F$. The $F = 2, 3$ states couple to the $F' = 2, 3$ states of the $5P_{3/2}$ state, with the particular transitions taking place depending on the polarisation of the incident light.

There is no coupling to the $F' = 1$ or the $F' = 4$ states [45]. This is due to the selection rule $\Delta F = 0, \pm 1$. We can see this by writing out the single photon couplings due to each laser ($n = 1, 2$) as vectors, where the subscripts indicate the transition route $|F, m_F\rangle, |F', m'_F\rangle$:

$$\begin{aligned} \Omega_{gn} &= \left(\Omega^{(n)}_{|2,m_F\rangle,|1,m'_F\rangle}, \Omega^{(n)}_{|2,m_F\rangle,|2,m'_F\rangle}, \Omega^{(n)}_{|2,m_F\rangle,|3,m'_F\rangle}, \Omega^{(n)}_{|2,m_F\rangle,|4,m'_F\rangle} \right) \\ \Omega_{en} &= \left(\Omega^{(n)}_{|3,m_F\rangle,|1,m'_F\rangle}, \Omega^{(n)}_{|3,m_F\rangle,|2,m'_F\rangle}, \Omega^{(n)}_{|3,m_F\rangle,|3,m'_F\rangle}, \Omega^{(n)}_{|3,m_F\rangle,|4,m'_F\rangle} \right). \end{aligned} \tag{2.36}$$

The two-photon Rabi frequency is equal to the dot product of the two vectors Ω_{g1} and Ω_{e2}. As transitions from $F = 2$ to $F = 4$ and $F = 3$ to $F = 1$ are not allowed, the first element of Ω_{e2} and the fourth element of Ω_{g1} are zero, and so the two photon Rabi frequency only includes terms which couple the $F = 2$ and $F = 3$ levels.

A transition to a state with one more unit of angular momentum ($m_F \rightarrow m_F+1$) requires a photon with $+1$ unit of angular momentum - i.e., right-circularly polarised light (we denote the polarisation driving such transitions as σ^+). Decreasing angular momentum by one unit ($m_F \rightarrow m_F - 1$) requires left-circularly polarised light (we denote the polarisation driving such transitions as σ^-). A change in angular momentum of greater than one unit is not dipole allowed and such transitions are not considered here.

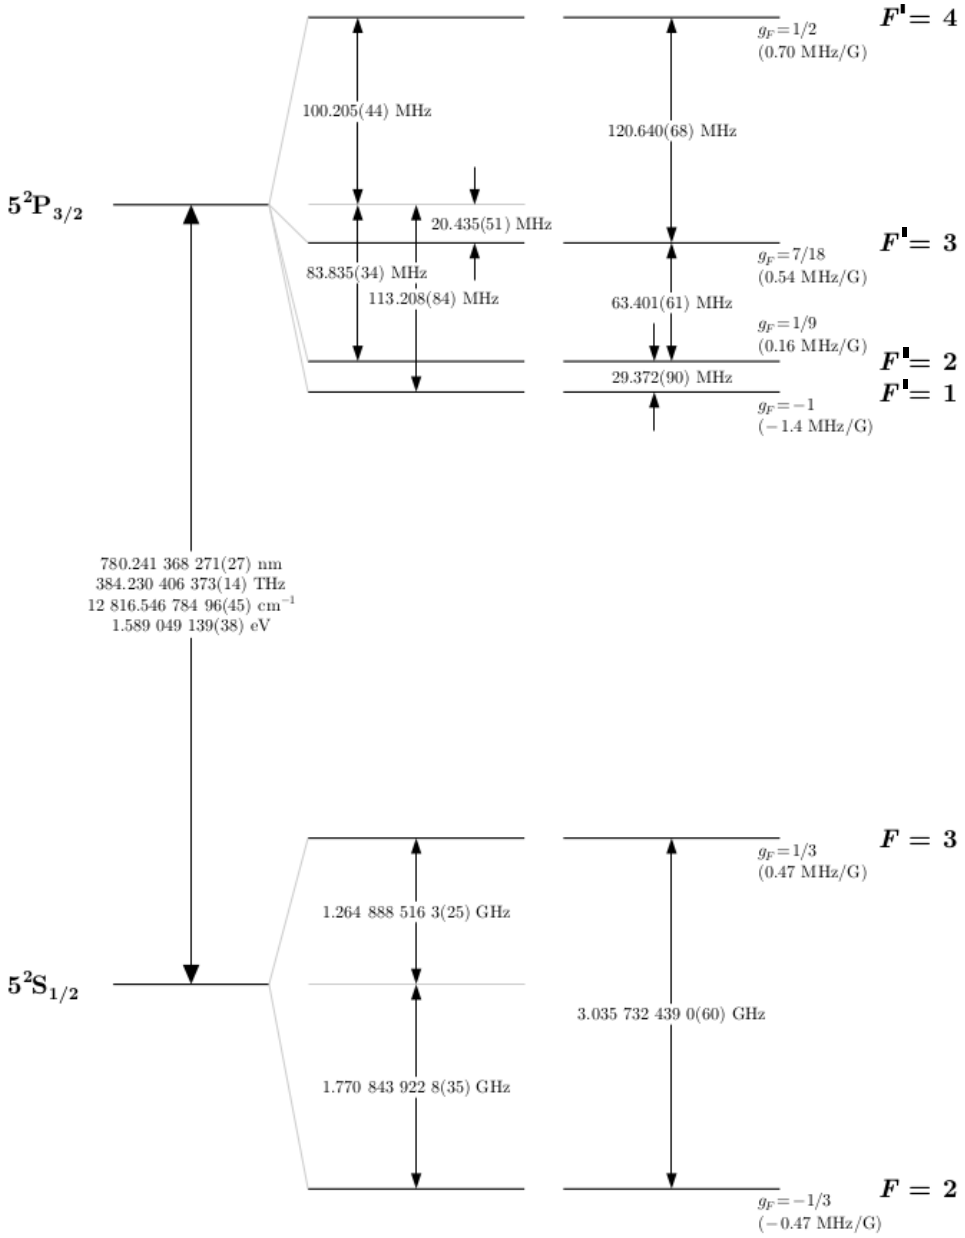

Figure 2.6: Hyperfine structure of the D_2 transition in Rubidium 85; diagram from [1].

Linear Raman beam polarisations are also possible, and are written as π^+, π^-. The $\pi^\pm$ may be written as a combination of circular polarisations: $\pi^+ = \frac{1}{\sqrt{2}}(\sigma^- - \sigma^+)$ and $\pi^- = \frac{i}{\sqrt{2}}(\sigma^- + \sigma^+)$, and correspond to the two orthogonal linear polarisations in the plane orthogonal to the propagation direction of the beam. The other linear polarisation which is perpendicular to $\pi^\pm$ is oriented along the direction of propagation of the beams and is not allowed as light is a transverse wave.

Several different polarisation configurations of the Raman beams are possible, but Raman transitions are only driven in particular situations. This can be shown by

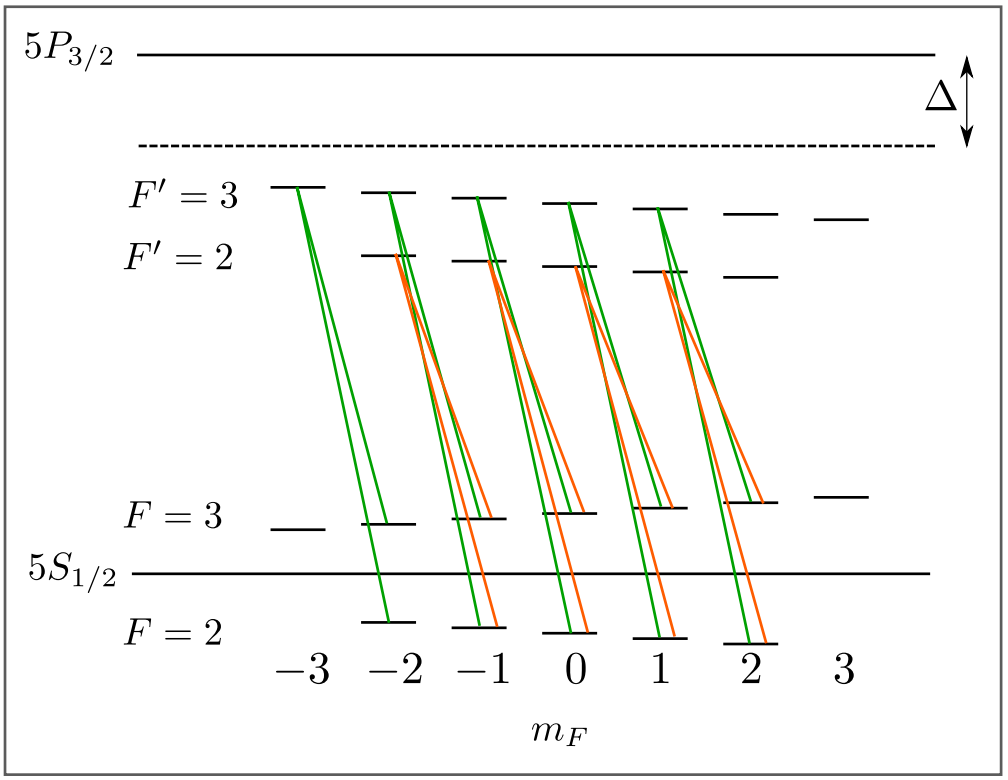

Figure 2.7: Two-photon transitions with $\sigma^- - \sigma^-$ Raman beam polarisations: via $F' = 2$ (orange) and $F' = 3$ (green). Transitions via $F' = 1$ and $F' = 4$ levels are not possible because of the selection rule $\Delta F = 0, \pm 1$. Not to scale.

calculating the dipole matrix element for these transitions, taking into account the two possible upper states ($F' = 2$ and $F' = 3$); where the dipole matrix elements via these states cancel each other out, Raman transitions are not possible. The allowed and disallowed polarisations can be found in Alex Dunning's thesis [45], along with a detailed description of the calculation of the dipole matrix elements. Here we note that the configurations of interest to us in this thesis are the $\sigma^+ - \sigma^+$, $\sigma^- - \sigma^-$, and $\pi^+ - \pi^-$. The possible routes for the $\sigma^- - \sigma^-$ configuration are illustrated in Figure 2.7. The dipole matrix elements for the $\pi^\pm$ case are found by combining the elements for the σ^+ and σ^- components of these polarisations.

Chapter 3

Atom interferometry

3.1 Introduction

An atom interferometer can be thought of as the matter-wave equivalent of an optical interferometer, with the roles of light and matter reversed: the matter has wavelike properties, and the splitting and deflection is achieved by interaction with light. The wavepacket of an atom is split into two components with different internal (electronic) and/or external (e.g., momentum) states; these components propagate and are redirected, propagate further and then are recombined. The population measured in each of the two states depends on the phase difference accrued along each path of the interferometer. There is no inter-atom interference; each atom forms its own interferometer and interferes with itself at the final beamsplitter.

The phase accumulated along each path is dependent on the evolution of the atomic wavefunction, and hence the potential experienced by each part of the wavepacket. The short de Broglie wavelength of ultracold atoms makes an atom interferometer ideal for precision measurements of inertial effects, for example gravity gradients [46], the rotation rate of the Earth [47]. Since the first demonstration of an atom interferometer inertial sensor over 20 years ago [48] (measurement of local gravitational acceleration g to 1 part in 3×10^{-6}), their precision has increased dramatically (measurement of g to 1 part in 6.7×10^{-12} [13]). Atom interferometers have also been used to make precision measurements of constants such as the Newtonian gravitational constant G [49] [50], and the fine structure constant α [51].

In this chapter we describe the components of an atom interferometer and their use in inertial sensing. We also discuss atom interferometry for velocity manipulation, leading to interferometric cooling. For a more detailed treatment of atom interferometry, we direct the reader to several comprehensive reviews ([42], [52], [53]).

3.2 Kasevich-Chu atom interferometer

Interferometry is a technique based on the coherent splitting and subsequent recombination of a wavelike object, in such a way as to gain information about the physics of the two paths taken through the interferometer. An interferometer usually has three essential components: a beamsplitter, to divide and recombine the wavelike object; a mirror, to redirect the two paths towards each other; and two or more periods of free propagation time, in which the wavelike object interacts with its environment only.

In this thesis we are concerned with atom interferometers which use light pulses to drive stimulated Raman transitions between two long-lived internal states of an atom. The exchange of momentum between atom and light field when the light field is used to drive a transition between internal states results in a correlation between the internal state of the atom and its momentum. Two-photon transitions such as those described in the previous chapter are ideal for atom interferometry because i) no spontaneous emission occurs from the two states, meaning they have a long coherence lifetime and ii) counter-propagating transitions are highly sensitive to the Doppler shift of the atoms.

The first atom interferometer using stimulated Raman transitions was demonstrated in 1991 [48] and is called the Kasevich-Chu (KC) interferometer; it is illustrated in Figure 3.2 a). Beginning with all atoms in the state $|g\rangle$, the first beamsplitter creates an equally weighted superposition of the states $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_{eff}\rangle$, where $\hbar\mathbf{k}_{eff}$ is the recoil momentum due to the simulated Raman transition. As the two parts of the superposition now have different momenta, they begin to separate spatially. Their relative velocity is reversed by the mirror pulse, and they are recombined with a final beamsplitter. The interference is subsequently detected by measuring the population in either of the two internal states as a fraction of the total number of atoms.

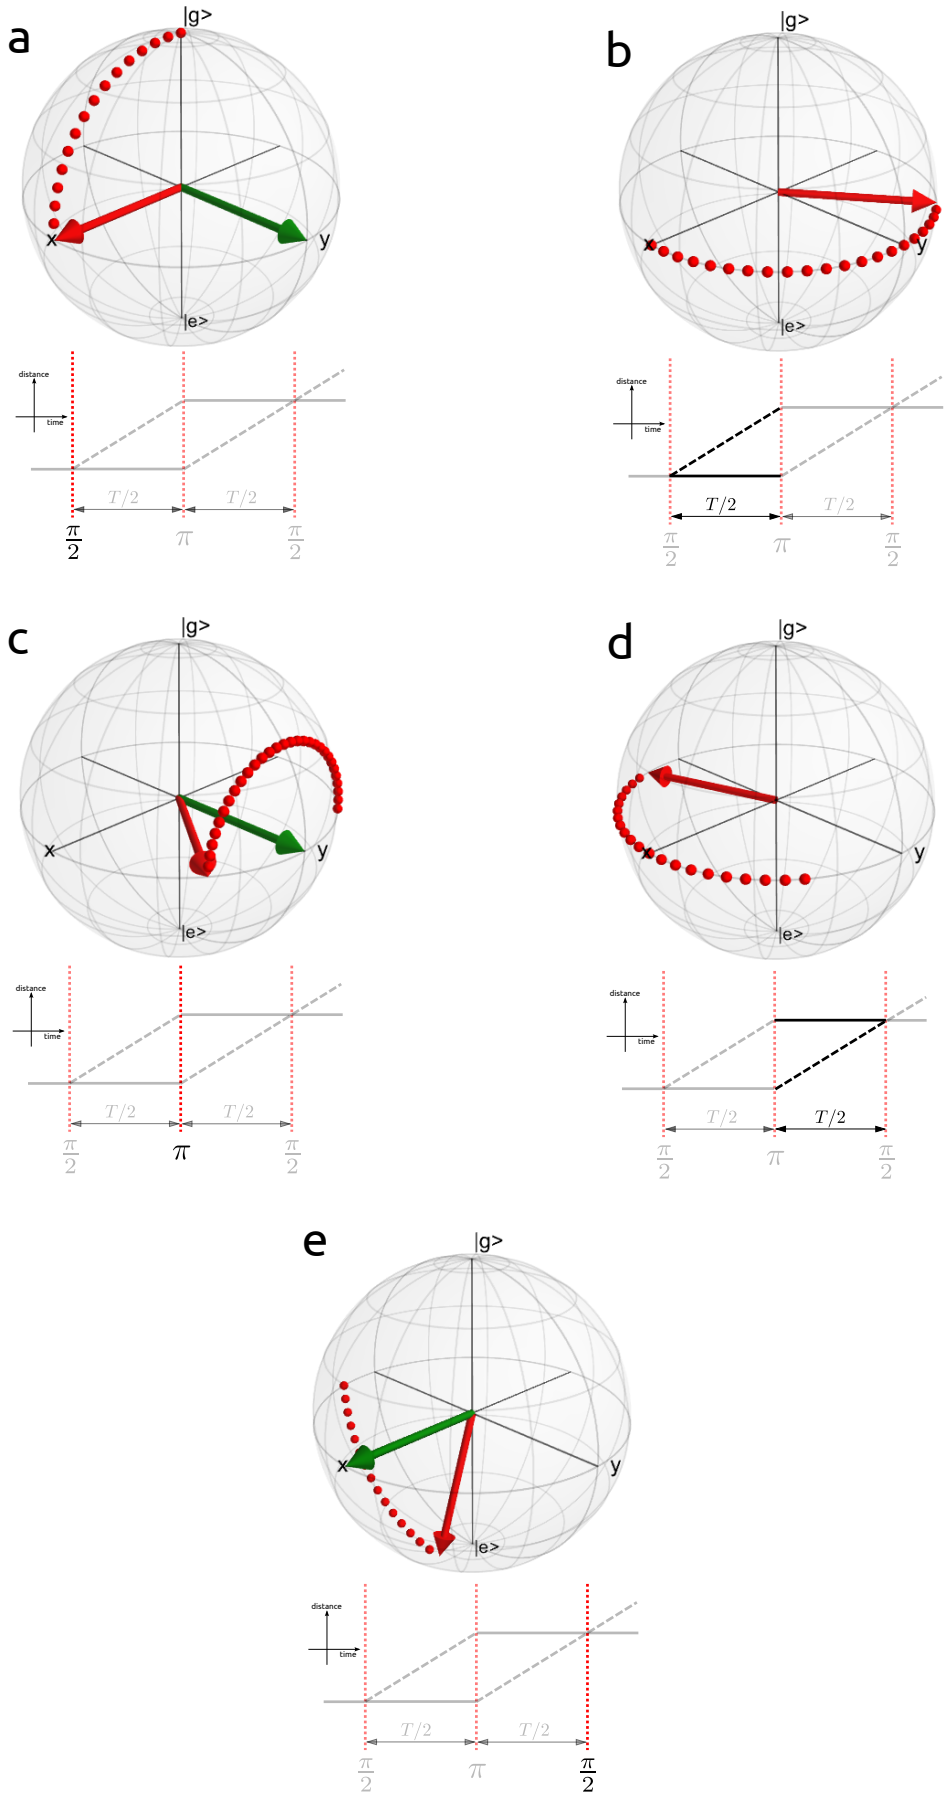

Figure 3.1: State vector trajectories during an interferometer sequence. a) first beamsplitter; b) first period of free evolution; c) mirror; d) second period of free evolution; e) final beamsplitter.

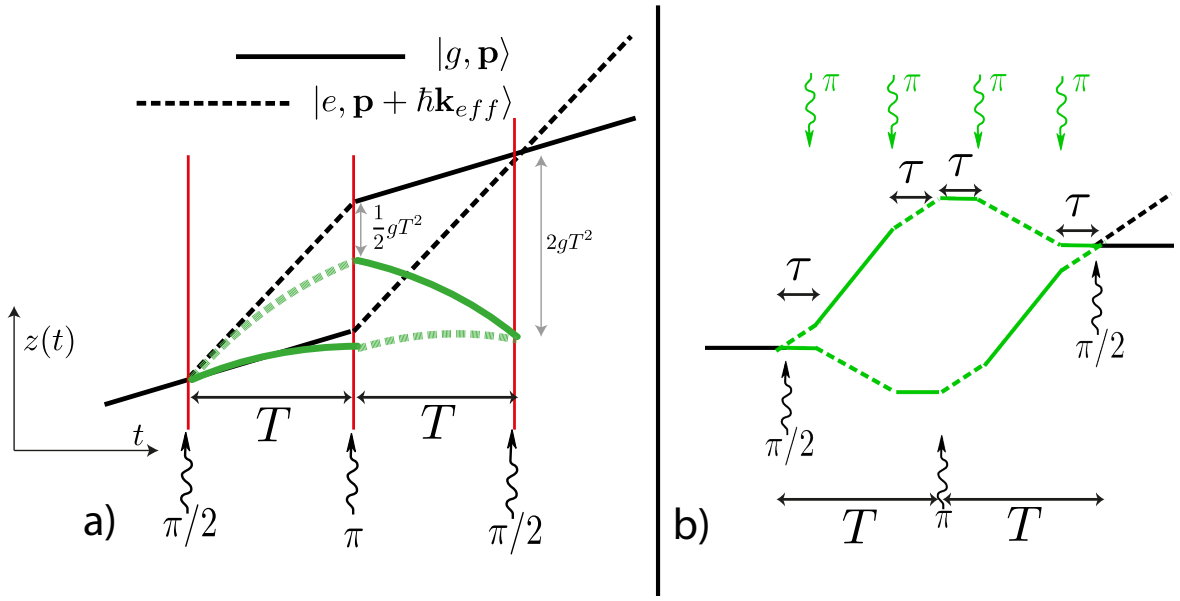

Figure 3.2: Trajectories of atomic wavepackets during an atom interferometer. a) Kasevich-Chu interferometer; the solid and dashed lines show the trajectories of atoms in, respectively, states $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_{eff}\rangle$. Green lines show the trajectories in a uniform gravitational field; black lines are without gravity. b) Interferometer with 4 augmentation pulses to increase the momentum transfer. Solid lines show trajectories of atoms with internal state $|g\rangle$, dashed lines show trajectories of atoms with internal state $|e\rangle$.

The state vector on the Bloch sphere during the various stages of an atom interferometer is shown in Figure 3.1. In this Figure, the phase of the atomic state at the end of the second period of free evolution is not equal to the phase of the final beamsplitter pulse, and so the state at the end of the beamsplitter is not aligned along the z axis.

3.2.1 Interferometer phase shift

The information about differences in propagation paths through an interferometer is contained in the phase difference between the two states at the interferometer output. The phase difference accumulated during a KC interferometer has three contributions: propagation phase $\Delta\phi_{prop}$, separation phase $\Delta\phi_{sep}$ and laser phase $\Delta\phi_{laser}$. The propagation phase is equal to the difference in the classical action evaluated along the upper and lower trajectories (divided by $\hbar$); it can be shown that, for the KC interferometer, this phase is equal to zero [54]. The separation phase accounts for an incomplete spatial overlap in the wavefunctions at the start of the interferometer, and is zero for the KC interferometer.

The laser phase *is* non zero and contains information on the inertial forces acting on the two paths. For example, the KC interferometer is commonly used to measure the local gravitational acceleration; under the Earth's gravity, the paths taken by the wavepackets are not straight lines (black trajectories in Figure 3.2 a)), but are in fact parabolas (green trajectories). At each point of interaction with the driving field, the phase of the field is imprinted on the component of the wavefunction which changes momentum due to the interaction. The points of interaction on the parabolic paths shown in green are shifted relative to the points of interaction along the straight paths, and so the phase imprinted is different by an amount depending on the time between pulses, the local gravitational acceleration, and the wavevector of the laser beam. For this interferometer, if the phases of the first beamsplitter, mirror, and second beamsplitter are respectively ϕ_1, ϕ_2 and ϕ_3, $\Delta\phi_{laser}$ in the absence of gravity is $\phi_1 - 2\phi_2 + \phi_3$ and in the presence of gravity is $k_{eff}gT^2(\phi_1 - 2\phi_2 + \phi_3)$. We see that the effect of inertial forces in the interferometer is to change the phase of the matter waves relative to the driving light field.

We note that the phase shift scales with the size of the momentum transfer. By increasing the momentum transfer per pulse, the sensitivity of the interferometer can be increased by the same factor. This was first demonstrated in [23]: in the experiment reported, a π pulse occurs immediately after the first beamsplitter, before and after the mirror, and before the final beamsplitter, as shown in Figure 3.2 b). The time separation between the interferometer pulses and the augmentation pulses is τ, where τ is small compared to the free evolution time $T - 2\tau$ (in [23], the temperature of the spin-polarised atoms used in the interferometer is 2μK and the unaugmented free evolution time T is 30ms). The momentum difference between the wavepackets during the long periods of free evolution is increased to $3\hbar\mathbf{k}_{eff}$, and the phase shift is increased by a factor of 3. The area enclosed by the interferometer is also increased by the same factor. The current record for such large momentum transfer (LMT) atom interferometers is 102 times the single-photon recoil momentum [55], realised through multiphoton Bragg pulses in a BEC; the authors show a wavepacket separation of 2mm whilst maintaining coherence.

The fidelity of the augmentation pulses is important in obtaining an acceptable signal-to-noise ratio. An augmentation pulse with a fidelity less than 1 will result in a fraction of the atoms not following the closed interferometer path; these atoms would instead follow paths which do not overlap spatially at the time of the final beamsplitter pulse. Their contribution to the measurement is then not coherent. For the interferometer shown in Figure 3.2 b) with 4 augmentation pulses, if each pulse has a fidelity of 70% then the fraction of atoms which follow

the closed interferometer path is $0.7^4 = 0.24$. With a fidelity of 90% this fraction increases to 0.65, and so there is a strong motivation for implementing high fidelity augmentation pulses.

3.3 Interferometric velocity manipulation

Atom interferometers for precision measurement are designed so that the phase difference between the two paths is dependent on the inertial forces experienced, and the measurement of the fraction of atoms in each state at the output of the interferometer gives information on the paths taken. An alternative interferometer design, proposed in 2000 [29], relies on a velocity-dependent phase difference to manipulate atomic velocities; the purpose is not to make a measurement of the fraction of atoms in each state, but to engineer a velocity-dependent impulse which, when combined with a spontaneous emission event, gives rise to cooling.

The interferometer consists of two beamsplitters separated by a period of free evolution (Figure 3.3 a)). We can use equations 2.30 and 2.31 to calculate the excited state probability $|c_e|^2$ after such an interferometer; with on-resonance beamsplitter pulses (with phases ϕ_1 and ϕ_2), and a free evolution time T, we find $|c_e|^2 = \frac{1}{2}(1 - \cos(\phi_2 - \phi_1 + \delta_{12}T))$. We recall from Section 2.2 that $\delta_{12} = \delta_L + \mathbf{p}\cdot\mathbf{k}_{eff}/m + \hbar k_{eff}^2/2m$; thus we see that the interferometer produces a velocity-selective excitation. The impulse as a function of atom velocity is plotted in Figure 3.3 b), for positive k_{eff} (red) and negative k_{eff} (blue); the phases of the beamsplitters are chosen so that the average impulse (dashed line) is zero when the velocity of the atom is zero, and the direction of impulse is most likely to occur in the direction opposite to the atom's velocity.

To cause cooling of the ensemble, the interferometer is followed by a spontaneous emission event. If the direction of the Raman beams is alternated, then the average impulse given to the atoms is that of the dashed curve in Figure 3.3 b). As the spontaneously emitted photon is emitted in a random direction, its impulse averages to zero, and the net effect of the interferometer-spontaneous emission cycle is to push atoms towards $v = 0$. The limiting temperature is the Raman recoil temperature, which for the transition we use in Rb 85 is $\hbar^2 k_{eff}^2/mk_B = 1.5\mu$K. The dwell time T of the interferometer sets the velocity capture range Δv: from Figure 3.3 b) we can see that $\Delta v = \pi/k_{eff}T$.

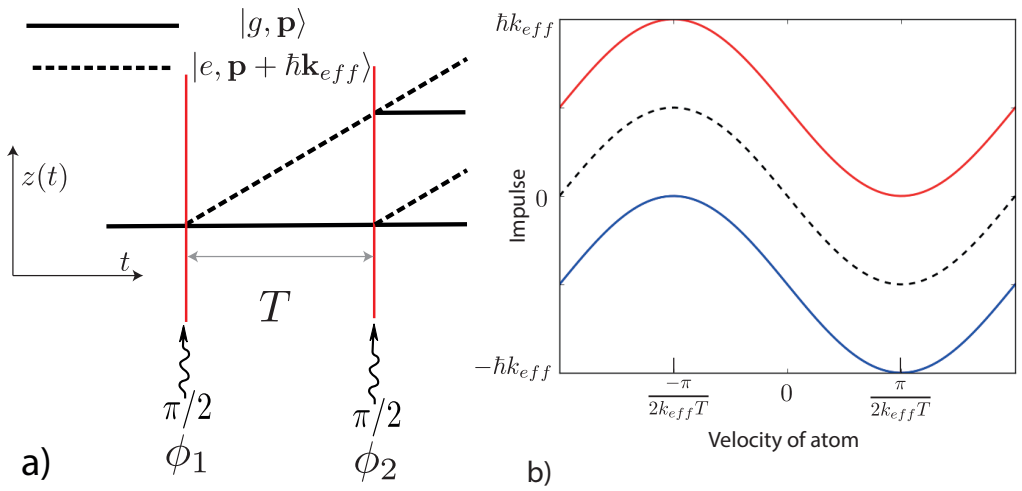

Figure 3.3: Interferometric cooling: a) shows wavepacket trajectories during the interferometer sequence of two $\pi/2$ pulses separated by time T. b) shows the impulse given to an atom at the end of the interferometer as a function of its velocity: red and blue plots are for opposing driving laser k vectors, and the dashed line shows the average impulse when the k vector is inverted for each sequence.

The advantage of this cooling method over conventional Doppler cooling is not immediately obvious - it requires high power lasers to reach the Rabi frequencies which allow broadband beamsplitter pulses (required to address atoms with a wide range of velocities i.e. at high temperatures 1K), and does not have a much smaller limiting temperature than the Doppler limit (in Rb 85, this is $\hbar\gamma/2k_B = 145\mu$K) or sub-Doppler mechanisms (which approach the single-photon recoil temperature $\hbar^2 k_{eff}^2/mk_B < 1\mu$K [41]). Indeed the interferometer as depicted in Figure 3.3 a) can only operate on one atomic transition. However, the original proposal in [29] considers an interferometer with two π pulses between the beamsplitters, which serve to eliminate the dependence of the phase shift on the laser detuning from resonance - thus allowing cooling on several transitions simultaneously. The authors suggest applications in the laser cooling of species which normally require many repumping lasers, and cooling of the motional degrees of freedom of diatomic molecules.

Chapter 4

Composite pulses

In systems where inhomogeneities in the environment or the control mechanisms can affect the useful signal output, methods of compensating for these inhomogeneities are invaluable. Composite pulses offer one such method.

This chapter introduces composite pulses and describes the two main types of correction sequence. The conditions of our experiment are considered and differences from the standard analysis are outlined.

4.1 Introduction

In coherent manipulation experiments, which may involve manipulating the state of trapped atoms [56] [57], ions [58], [59], or electron spins in NV centres [60], compensation methods are usually concerned with optimising the phase, intensity, and duration of the driving field. Optimal control approaches often aim to replace a unitary operation such as a π pulse, which would normally be a square pulse of fixed duration and phase, with a longer pulse with time-varying intensity and phase. Although such methods have been shown to be very effective [61], here we consider a simpler approach - one which is not computationally intensive and can easily be implemented using equipment common to almost all atomic physics experiments.

Composite pulses [62] have long been used in the field of NMR spectroscopy and offer a method of compensating for such uncontrollable systematic errors. The errors are considered to be systematic insofar as they do not vary over the duration of the correction sequence. In the uncompensated case, a square pulse implementing

a rotation of θ on the Bloch sphere has a duration τ set by the Rabi frequency Ω: $\Omega\tau = \theta$, and the phase of the driving field is constant over this time. A composite pulse replaces this square pulse with a continuous series of pulses of varying lengths and phases, designed to compensate for particular types of systematic error.

In our experiment, we are concerned with correcting two types of error. Pulse length errors occur when the strength of the control field is different from the optimal value; this causes the angle through which the state vector rotates on the Bloch sphere to be different from the desired angle. Off-resonance errors arise when the transition is not resonant with the control field, and this causes the axis of rotation to be tilted out of the $x - y$ plane. In our system, pulse length errors are due to the spatially non-uniform intensity of the Raman beams and the range of coupling strengths between the different m_F sublevels of the hyperfine states, and off-resonance errors are due to the range of velocities in the atom cloud.

4.2 Propagators and fidelities

Two tools are useful in evaluating the efficacy of composite pulse sequences: the propagator of the system, and the fidelity of a given sequence.

It is useful to consider the form of the propagator of the system both in the error–free case and in the situation where one or both types of error is present. The evolution of a state $|\psi\rangle$ under a time-independent Hamiltonian $\hat{H}$ is given by $|\psi(t)\rangle = U(t)\,|\psi(0)\rangle$, where $U(t)$ is the propagator of the system, defined by $U(t) = \exp(-i\hat{H}t/\hbar)$. We represent the propagator describing a pulse intended to rotate the state vector through an angle θ about an axis $(\cos(\phi)\cos(\alpha), \sin(\phi)\cos(\alpha), \sin(\alpha))$ by

$$\begin{aligned} U(\theta,\phi,\alpha) = &\cos\left(\frac{\theta}{2}\right)\sigma_0 - \\ &i\sin\left(\frac{\theta}{2}\right)\left[\sigma_x\cos(\phi)\cos(\alpha) + \sigma_y\sin(\phi)\cos(\alpha) + \sigma_z\sin(\alpha)\right], \end{aligned} \tag{4.1}$$

where $\sigma_{x,y,z}$ are the Pauli matrices and σ_0 is the identity.

We are interested in resonant control fields, where α is zero (for example, an ideal beamsplitter pulse in an atom interferometer); in this case we can also represent

the pulse in the θ_ϕ notation. For example, a $\pi/2$ pulse about the x–axis would be written as $\left(\frac{\pi}{2}\right)_0$ or 90_0. It is worth noting that in the case of a cloud of atoms at non-zero temperature, a range of atomic velocities is present and thus a pulse with α set to zero will be resonant only with a certain velocity class; hence it is only possible to implement a rotation about an axis in the $x-y$ plane for a proportion of the atoms at any one time.

Composite pulses are usually classified as one of two types: point-to-point and general rotor. A point-to-point (P2P) pulse is only effective for a particular initial and final state, and may not correct errors well in other cases (and may indeed have a detrimental effect). Conversely, a general rotor (GR) is designed to implement a rotation through a particular angle and is error-tolerant for any initial state.

The fidelity of a given sequence is a measure of how closely it matches the desired operation. Several ways of defining a fidelity exist and it is important to choose the most appropriate method for a particular situation. For a general rotor, the fidelity of a sequence described by the propagator V where the propagator in the absence of errors is U may be defined as the modulus of the Hilbert–Schmidt inner product:

$$\mathcal{F}_{GR}(U,V) = \frac{\left|\mathrm{tr}(U^\dagger V)\right|}{2}. \tag{4.2}$$

The fidelity of a pulse which perfectly reproduces the desired rotation will be 1, and a pulse operating under conditions where errors ϵ of order n are present will have a fidelity with corresponding terms ϵ^{2n}.

The fidelity of a P2P operation (with target state $|\psi\rangle$) may be evaluated by the point-to-point fidelity

$$\mathcal{F}_{PP}(U,V,|\psi\rangle) = \left|\langle\psi| U^\dagger V |\psi\rangle\right|^2 \tag{4.3}$$

4.3 Pulse length errors

A pulse length error (PLE) arises when the strength of the control field differs from the optimum value, changing the angle through which the state vector rotates. Figure 4.1 shows the paths of three vectors with different pulse length errors: the state represented by the green vector (no pulse length error) arrives exactly at

the south pole of the Bloch sphere, whereas the red and blue vectors under- and over-shoot the target. If the desired angle to rotate through is θ, and there exists an error ϵ in the strength of the control field such that the actual angle rotated through is $\theta(1+\epsilon)$, the propagator becomes

$$\begin{aligned} V(\theta,\phi,\alpha) &= U(\theta(1+\epsilon),\phi,\alpha) \\ &= \cos\left(\frac{\theta(1+\epsilon)}{2}\right)\sigma_0 - i\sin\left(\frac{\theta(1+\epsilon)}{2}\right)\times \\ &\quad [\sigma_x\cos(\phi)\cos(\alpha) + \sigma_y\sin(\phi)\cos(\alpha) + \sigma_z\sin(\alpha)]\,. \end{aligned} \tag{4.4}$$

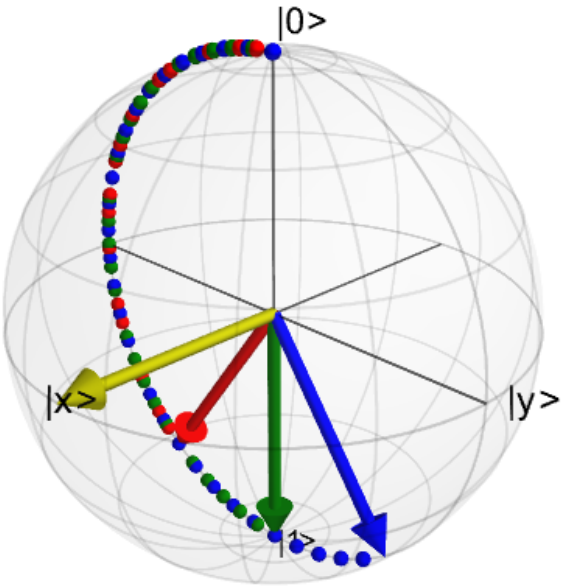

Figure 4.1: Bloch sphere diagram showing state vector paths during a π pulse. The yellow vector represents the control field. Green: no pulse length error, red and blue: fractional pulse length error of ± 0.2.

We initially consider a pulse length error for an on-resonance pulse, where $\alpha = 0$. In this case, the propagator is

$$V(\theta,\phi) = \cos(\frac{\theta(1+\epsilon)}{2})\sigma_0 - i\sin(\frac{\theta(1+\epsilon)}{2})\,[\sigma_x\cos(\phi) + \sigma_y\sin(\phi))]\,. \tag{4.5}$$

If we now assume the error ϵ is small compared to θ and expand the above equation to first order in ϵ, the propagator becomes

$$V(\theta,\phi) \approx U(\theta,\phi) - \frac{\epsilon\theta}{2}\left(\sigma_0 \sin(\frac{\theta}{2}) + i\cos(\frac{\theta}{2})(\sigma_x\cos(\phi) + \sigma_y\sin(\phi))\right) + \mathcal{O}(\epsilon^2). \tag{4.6}$$

It can be seen that a pulse-length error in isolation results in the desired rotation being performed plus an additional rotation about the same axis (Figure 4.1). The fidelity of this operation is

$$\mathcal{F}(U(\theta,\phi), V(\theta,\phi)) = 1 - \epsilon^2\theta^2/8 + \mathcal{O}(\epsilon^4). \tag{4.7}$$

Therefore, any composite pulse sequence aiming to correct for pulse length errors should have a fidelity which contains no terms of order $\mathcal{O}(\epsilon^2)$. Note that the correction of errors of order n requires a fidelity which removes terms of order $2n$.

The collinearity of the rotation axes for both the error-free propagation part and the additional propagation part makes the design of pulse length error correcting composite pulses somewhat simpler than the design of off-resonance error correcting composite pulses, as it enables the reliable implementation of inverse pulses (a pulse which reverses the evolution of the original pulse, even in the presence of errors). Since a shift of π in the phase of the control field reverses the direction of rotation on the Bloch sphere, and under the influence of systematic pulse length errors the distance moved on the Bloch sphere will be equal and along the same great circle during both the pulse and its inverse, an inverse pulse can be easily implemented, allowing the reliable return of a state vector to its original position on the Bloch sphere. This arises from the equivalence of $U(\theta,\phi+\pi)$ and $U(-\theta,\phi)$, which remains true in the presence of pulse length errors; the sequence $U(\theta,\phi+\pi)U(\theta,\phi)$ is equal to the identity (do-nothing) operation.

One example of a composite pulse designed to remove first order PLEs (and additionally, although not by design, the second order errors) is the BB1 pulse [63]. The sequence to correct pulse length errors for an arbitrary desired rotation θ_ϕ is $W(\theta,\phi) = \theta_\phi 180_{\phi_1} 360_{3\phi_1} 180_{\phi_1}$, where $\phi_1 = \arccos(-\theta/4\pi)$. The BB1 is an interesting case as it is not designed to replace the uncorrected pulse; rather, it follows the uncorrected pulse with an operation which compensates first order pulse length errors when they are present, and reduces to a do nothing operator in their absence. The fidelity of a 180_0 BB1 sequence is in fact 6th order in ϵ - although not by design, the sequence eliminates the 4th order term in addition to the 2nd order term.

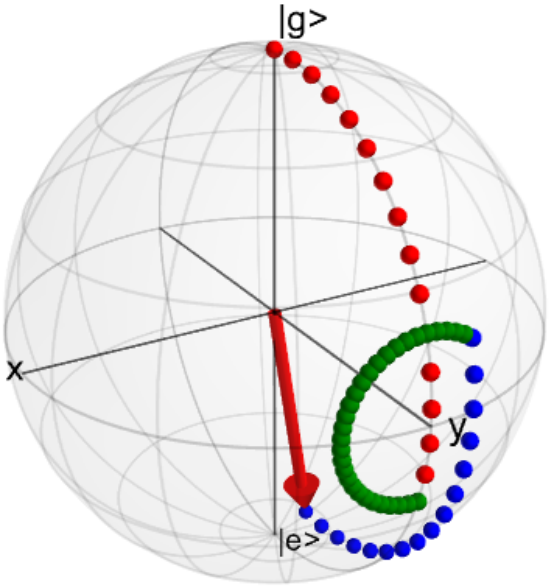

Figure 4.2: Bloch sphere diagram showing state vector path during a Levitt pulse sequence: 90_x (red points) 180_y (green points) 90_x (blue points). A fractional pulse length error of +0.2 is present. The red vector shows the final state.

BB1 is a general rotor and thus compensates for PLEs for any given initial state and desired rotation angle; however, it is not easily visualised on the Bloch sphere. One sequence which is not particularly versatile but does illustrate well the error-correction potential of controlling the driving field phase is the Levitt pulse [64]. Originally conceived of as a "sideways spin-echo", it is an inversion pulse sequence which replaces a 180_x rotation with $90_x180_y90_x$. As expected, the propagator fidelity (equation 4.2) of the Levitt pulse is no better than that of the uncorrected pulse, and the point-to-point fidelity shows an improvement over the uncorrected pulse ($\mathcal{F}_{PP} = 1 - \pi^4\epsilon^4/16 + \mathcal{O}(\epsilon^6)$). Figure 4.2 shows this sequence on the Bloch sphere for a fractional PLE of +0.2. The 180_y rotation approximately swaps the position of undershooting and overshooting vectors with fractional PLEs of opposite sign, so that their position at the start of the final 90_x rotation allows their under/overshooting to bring them closer to the target state than in the non-correcting case.

4.4 Off-resonance errors

An off-resonance error arises when the control field frequency differs from the transition frequency, changing the angle (relative to the z axis) of the vector about which the state vector rotates on the Bloch sphere. Figure 4.3 shows the paths of three vectors with different off-resonance errors: the state represented by the green vector (no off resonance error) rotates about an axis in the $x-y$ plane and thus arrives exactly at the south pole of the Bloch sphere, whereas the red and blue vectors rotate about axes tilted out of the $x-y$ plane. Their paths are not great circles and thus cannot pass through the south pole when starting at the north pole.

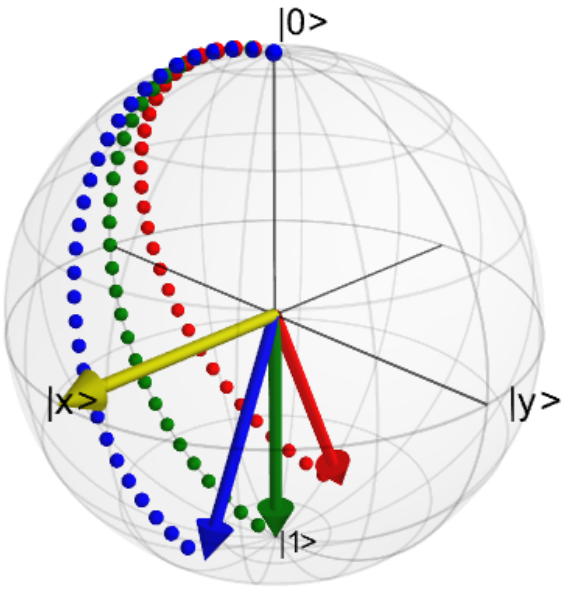

Figure 4.3: Bloch sphere diagram showing state vector paths during a π pulse. The yellow vector represents the control field for the resonant case. Green: on resonance, red and blue: fractional off-resonance error of ± 0.2.

If the control field frequency is ω and the deviation is $\delta\omega$, we can write the propagator in terms of the off-resonant fraction $f = \delta\omega/\omega$:

$$R(\theta, \phi) = \exp\left(-i\frac{\theta}{2}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)\right). \tag{4.8}$$

In Appendix D this exponential is expanded and it is shown that the propagator can also be written as

$$R(\theta,\phi) = \sigma_0 \cos\left(\frac{\theta}{2}\sqrt{1+f^2}\right) - i(1+f^2)^{-1/2}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)\sin\left(\frac{\theta}{2}\sqrt{1+f^2}\right). \tag{4.9}$$

Expanding this for small f,

$$R(\theta,\phi) = U(\theta,\phi) - if\sin\left(\frac{\theta}{2}\right)\sigma_z + \mathcal{O}(f^2). \tag{4.10}$$

This suggests that the effect of an off-resonance error (for small f) is to add a z rotation onto the error-free operation. However, for larger off-resonance errors, the pulse length is also affected and unlike in the case of pulse length errors, the angle rotated through is increased regardless of the sign of the error ($\sqrt{1+f^2}$ is always positive).

The fidelity of this operation is

$$\mathcal{F}(U(\theta,\phi), R(\theta,\phi)) = 1 - \frac{f^2\sin^2(\theta/2)}{2} + \mathcal{O}(f^4), \tag{4.11}$$

and so any composite pulse sequence aiming to correct for off-resonance errors should have a fidelity which does not contain second order terms in f.

We can see that generation of an inverse pulse is not as trivial as in the pulse length error case: the phase ϕ does not feature in the term describing the off-resonant part of the propagator, and so it is not possible to design a sequence which will exactly return all state vectors to their starting positions. In contrast to the pulse length error case, $R(\theta,\phi+\pi) \neq R(-\theta,\phi)$ and indeed the off-resonance errors accumulate:

$$\begin{aligned} R(\theta,\phi+\pi)R(\theta,\phi) = \sigma_0(1 - f^2\sin^2\frac{\theta}{2}) - 2if\sigma_z\sin\frac{\theta}{2}\cos\frac{\theta}{2} \\ - 2if(\sigma_x\sin\phi - \sigma_y\cos\phi) + \mathcal{O}(f^2). \end{aligned} \tag{4.12}$$

An example of an off-resonance error correcting composite pulse is the WALTZ sequence $90_0180_{180}270_0$ [65]. This is a point-to-point pulse designed to transfer population between the poles of the Bloch sphere. Its behaviour is interesting as its fidelity is

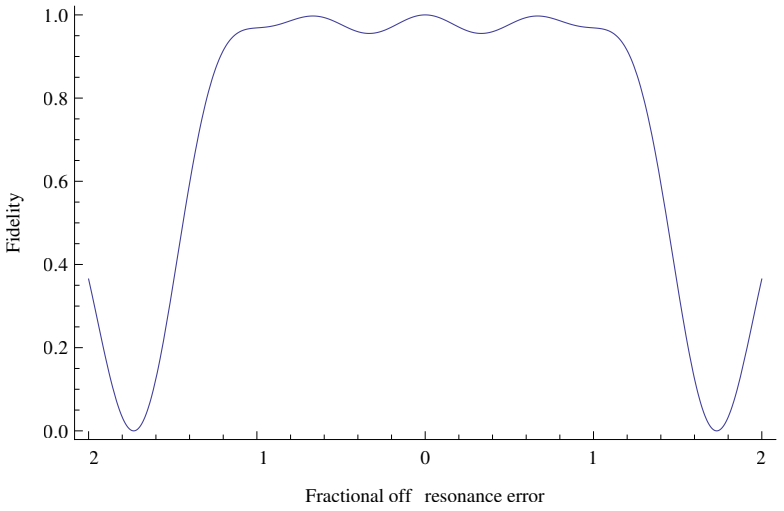

Figure 4.4: Fidelity of a WALTZ sequence as a function of fractional off-resonance error

$$\mathcal{F}_{WALTZ} = 1 - f^2 + 1/16(80 + 16\pi - \pi^2)f^4 + \mathcal{O}(f^6), \tag{4.13}$$

which at first glance suggests that it would not compensate for first order errors due to the presence of the f^2 term. However, there is partial cancellation between the second and fourth order terms for a wide range of values of f: Figure 4.4 plots the fidelity for $-2 < f < 2$, and it can be calculated that $\mathcal{F}_{WALTZ}$ does not drop below 0.95 until $f > 1.14$.

4.5 Composite pulses in our system

We now analyse the propagator derived from the Hamiltonian of our system and its behaviour in the presence of errors.

In the experiment, several different transition routes are possible, with their own coupling strengths and light shifts; however, we initially consider just one of these transition routes. Equation 2.27 gives the form of the Hamiltonian in the rotating frame. The rotating frame propagator $U(t) = \exp(-i\hat{H}t/\hbar)$ is[1]

[1]A diagonalisable matrix A is exponentiated by writing it as $A = PDP^{-1}$, where D is a diagonal matrix whose elements are the eigenvalues of A, and P is the matrix of eigenvectors of A. Since PP^{-1} is equal to the identity, the exponential of A is $e^A = Pe^DP^{-1}$. These eigenvalues and eigenvectors are calculated in Section 2.2.3.

$$U(t) = \begin{pmatrix} (\cos\frac{\Omega_r t}{2} + i\cos\Theta\sin\frac{\Omega_r t}{2}) & e^{-i\phi}(-i\sin\Theta\sin\frac{\Omega_r t}{2}) \\ e^{i\phi}(-i\sin\Theta\sin\frac{\Omega_r t}{2}) & (\cos\frac{\Omega_r t}{2} - i\cos\Theta\sin\frac{\Omega_r t}{2}) \end{pmatrix}, \tag{4.14}$$

which may also be written in terms of the Pauli matrices:

$$U(t) = \cos\frac{\Omega_r t}{2}\sigma_0 + i\cos\Theta\sin\frac{\Omega_r t}{2}\sigma_z - i\sin\Theta\sin\frac{\Omega_r t}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi). \tag{4.15}$$

$\cos\Theta$ and $\sin\Theta$ were defined previously and it is clear that when the laser is exactly on resonance ($\delta_{12} = \delta_{AC}$), Equation 4.15 reduces to the ideal propagator and $\Omega_r = \Omega_{eff}$.

This suggests, as shown in the previous section, that an off-resonance error additionally causes a pulse length error. To see this in more detail, we define the off-resonant fraction $f = -(\delta_{12} - \delta_{AC})/\Omega_r$ and expand the propagator for small f. Retaining terms up to $\mathcal{O}(f^2)$ we find that U is approximately

$$\begin{aligned} U(t) \approx{}& \sigma_0\cos\frac{\Omega_{eff}t}{2} - i\sin\frac{\Omega_{eff}t}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi) \\ & - if\sigma_z\sin\frac{\Omega_{eff}t}{2} \\ & - \frac{f^2}{2}\frac{\Omega_{eff}t}{2}\left(\sigma_0\sin\frac{\Omega_{eff}t}{2} + i\cos\frac{\Omega_{eff}t}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi)\right) \\ & + \frac{if^2}{2}\sin\frac{\Omega_{eff}t}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi) + \mathcal{O}(f^3). \end{aligned} \tag{4.16}$$

This propagator is of the same form as the expansion of $R(\theta, \phi)$ (Appendix D); this is promising as suggests that standard composite pulses designed to correct for off-resonance errors may be of some use. The first line of this equation is the error-free propagator and the subsequent lines are perturbations to this ideal. The second line is an off-resonance error of the same form as Equation 4.10; the third line is a pulse length error, proportional to the off-resonance error and of the same form as Equation 4.6. The final line is an additional pulse length error.

This propagator describes one possible transition route only. Composite pulse sequences are made up of a series of rotations through specific angles, and the duration of the laser on-time t required to produce such a rotation is calculated

according to $\theta = \Omega t$. When only one transition route is present, Ω can be chosen to be Ω_{eff}: the Rabi frequency when $\delta_{12} = \delta_{AC}$. However, when several transition routes are possible (e.g. where the atoms are initially distributed across more than one m_F sublevel and/or the Raman beam polarisation configuration is $\pi^+ - \pi^-$), there is no single on-resonance Rabi frequency. Each transition route has a slightly different coupling strength: even if there were no off-resonance errors there would still exist pulse length errors due to the range of Rabi frequencies. For example, in the $\sigma^+ - \sigma^+$ configuration with the atoms initially equally distributed across all 5 m_F sublevels of the $F = 2$ state, there are 3 possible Rabi frequencies: the $m_F = 0$, the $m_F = \pm 1$ (94% of the $m_F = 0$ Rabi frequency), and the $m_F = \pm 2$ (74%).

Clearly the choice of Rabi frequency used to define pulse durations is important as deviations from it determine the pulse length errors; this is explored in further detail in Chapter 8.

Part II

Experiment

Chapter 5

Magneto-optical trapping of Rb 85

Before coherent manipulation experiments may take place, the atoms to be manipulated must be cooled and trapped. In this chapter the experimental setup of the MOT lasers and the MOT chamber is described, along with the methods for state preparation and detection. The MOT setup was initially constructed by James Bateman [66] and Matt Himsworth [67] and details regarding the setup may be found in their theses. The setup was further optimised and characterised by Alex Dunning [45].

5.1 Lasers

As described in 1.3, laser cooling ideally requires a closed two-level transition. We use the $5S_{1/2}\,(F = 3) \rightarrow 5P_{3/2}\,(F' = 4)$ transition within the D_2 line, which may be addressed by one laser, which we call the cooling laser. Transitions between levels of high F are used because the hyperfine splitting between levels $F' = 3$ and $F' = 4$ is approximately 4 times larger than the splitting between $F' = 1$ and $F' = 2$, and so the overlap with neighbouring states for high F transitions is much smaller (Figure 5.1).

However, there is still a small probability to excite to the $5P_{3/2}\,(F' = 3)$ state, which may decay to the dark state $5S_{1/2}\,(F = 2)$; atoms undergoing this transition then would not "see" the cooling laser and are removed from the absorption-spontaneous emission cycle. To bring these atoms back into the cooling cycle, we

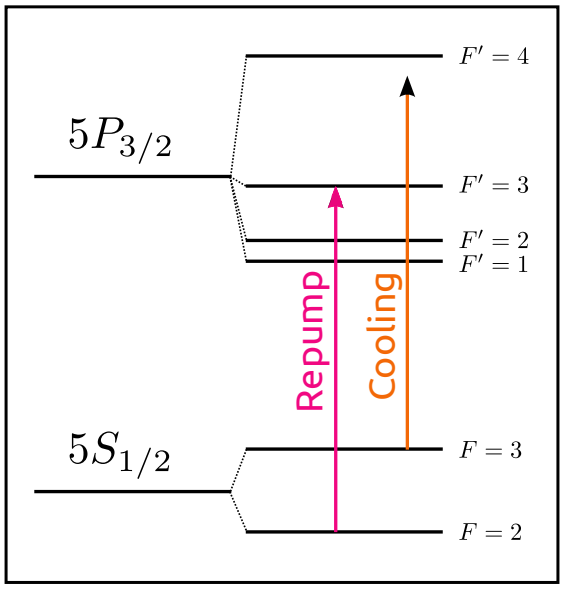

Figure 5.1: Rb 85 energy level structure used for laser cooling (relative hyperfine splittings approximately to scale, other splittings not to scale. For accurate values see [1]).

apply also a repump laser to excite these atoms to the $5P_{3/2}\,(F' = 3)$ state, from which they may decay back to $5S_{1/2}\,(F = 3)$.

5.1.1 External Cavity Diode Laser (ECDL)

Rubidium 85 is commonly used in atomic physics experiments as the frequencies required for cooling and manipulation are easily accessible through diode lasers. One requirement for laser cooling is that the linewidth of the laser be smaller than the linewidth of the atomic transition [37]; for the D_2 line in Rubidium 85 this is approximately $\gamma \approx 2\pi \times 6\,\mathrm{MHz}$. The linewidth of a free running diode laser is typically tens of MHz [68] [69], and the frequency of such a laser is not reliably controlled; however, both of these problems may be addressed by creating a cavity with a laser diode with one end facet having an anti-reflection coating and a reflection grating (an External Cavity Diode Laser or ECDL). The grating introduces frequency tunability (by changing its angle relative to the axis of the cavity) and the longer length of the cavity results in a smaller linewidth.

The cooling and repump lasers in this experiment are Littrow configuration ECDLs: changing the angle of the grating changes the angle of the output beam and (in the case of large angles) requires realignment of the following optics. A PZT stack controls the angle of the grating and allows frequency tuning, as described in the next section.

5.1.2 Frequency stabilisation

The second restriction on a laser used for Doppler cooling is that the frequency of the laser be stable in time. We achieve this by locking the laser frequency to the particular transition required using Doppler-free spectroscopy and a PID (proportional-integral-differential) feedback circuit; the error signal from this circuit is fed back to the piezo-electric transducer stack. Details of the setup can be found in Matt Himsworth's thesis [67] and are not reproduced here; we simply note that the frequency stabilisation is sufficient to achieve cooling and trapping. The detuning of the cooling laser from the cooling transition is typically 1 to 3 linewidths.

5.1.3 Beam switching

The MOT beams are only 'on' for the time it takes to cool and trap the atoms; during manipulation sequences the atoms are in free fall. We therefore require a method of switching the MOT beams off quickly and without disturbing their frequency stability. We use acousto-optical modulators (AOMs) for this purpose.

A sound wave travelling through an acousto-optic crystal causes a periodic modulation of the refractive index of the crystal. This forms a diffraction grating and light is scattered from the acoustic wavefronts; an incident beam of light may absorb phonon(s) from or emit phonon(s) to the crystal lattice, causing a shift in the frequency (by conservation of energy) and direction (by conservation of momentum) of the beam. The condition for constructive interference of the scattered beam results in diffracted orders m with frequency shifts mf_{phonon} and wavevectors $\mathbf{k}'_{\text{light}} = \mathbf{k}_{\text{light}} + m\mathbf{k}_{\text{phonon}}$, and experimentally the alignment is optimised to obtain maximum power in the 1st (or -1st, depending on the frequency shift required) order.

An AOM is driven by a radio frequency signal and has at the opposing end of the crystal an absorber to prevent reflections of the sound wave. It has a finite rise time (the time between applying the RF signal to the AOM and maximum power being diffracted into the 1st order) due to the time it takes for the sound wave to travel across the incident beam. This is not an issue for the MOT beams but is important when considering short pulses; see Section 8.5.3 for further discussion.

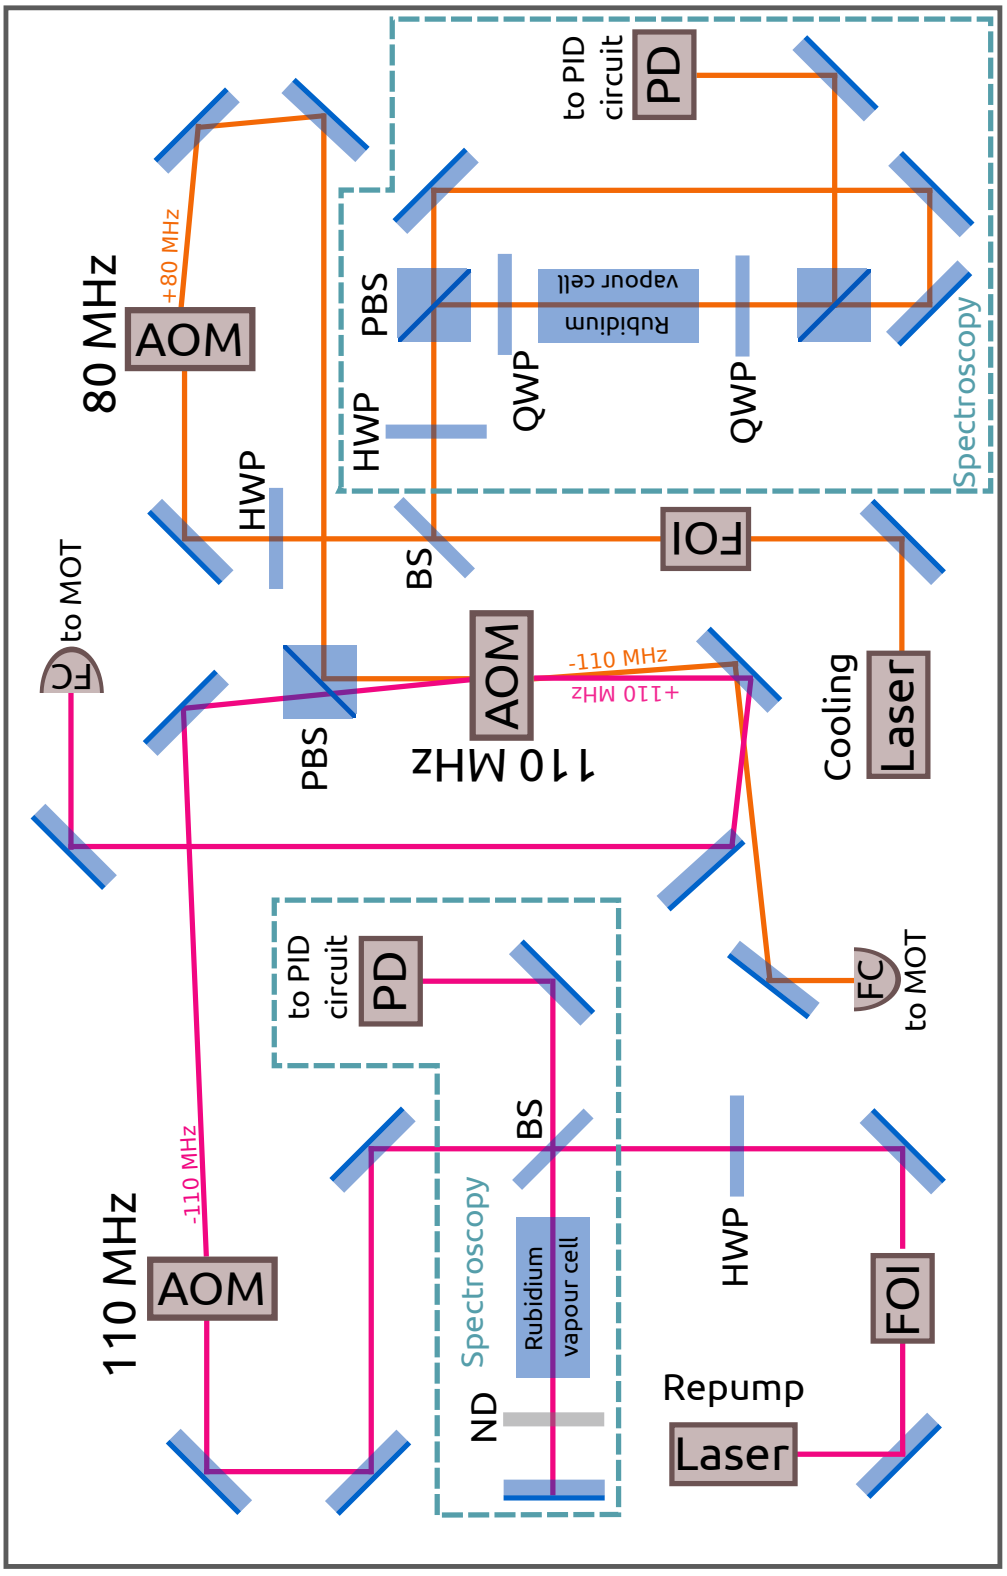

Figure 5.2: Experimental setup for generation of cooling and repump beams. HWP: half wave plate, QWP: quarter wave plate, BS: beam sampler, PD: photodiode, FC: fibre coupler, AOM: acousto-optical modulator, FOI: Faraday optical isolator

5.1.4 Experimental setup

Figure 5.2 shows the layout of the MOT laser system. Each laser passes through an optical isolator which provides greater than 30dB of attenuation of any back reflections from components further along the beam path. This is important because optical feedback can disrupt the stability of the laser. After the isolator, a small amount of each beam is diverted to the spectroscopy optics.

Each laser passes through two AOMs, aligned such that the sign of the frequency shift is opposite through successive AOMs. The cooling laser is shifted up by 80MHz when passing through the cooling AOM (AA Optoelectronics M080-2B/F-GH2), and down by 110MHz (M110-2B/F-GH2) when passing through the common AOM[1]. The repump is shifted down by 110MHz by the repump AOM (M110-2B/F-GH2), and up by 110MHz by the common AOM. The two beams are combined using a polarising beamsplitter cube before the common AOM. The opposite sign of the frequency shift of each beam through the common AOM means that the first orders are either side of the zeroth order beam and so the cooling and repump beams can be separated and directed towards their own fibres. Using fibres to bring the beams to the MOT chamber means that any change in alignment of the lasers or spectroscopy will only affect coupling into the fibres and does not require realignment of the beams in the MOT chamber.

The AOMs are powered by RF drivers which have a DC power supply (always on) and a modulation input, which can take a signal up to $5V$ and controls the intensity in the first order. The driver for the common AOM is controlled by a 5V TTL signal from a SpinCore PulseBlaster (Section 6.4), whereas the cooling and repump AOMs are individually controlled by waveforms output from a Tabor 8024 arbitrary waveform generator. This means that the intensity of the cooling and repump beams can be independently controlled, which is essential for the state preparation, readout, and molasses cooling stages.

5.2 MOT chamber

The experiments take place within the custom-built, stainless steel MOT chamber. The chamber has windows and is kept at a pressure of approximately 10^{-9}mbar by

[1]The net frequency shift of the cooling beam is -30 MHz from the cooling transition, which is more than the ideal 1-3 linewidths; we compensate for this by shifting the lock point electronically.

vacuum pump. The windows/apertures are variously used for the MOT beams, Raman beams, the getters for dispensing Rubidium, the photomultiplier tube (PMT), the vacuum pump, and the MOT camera (5.3).

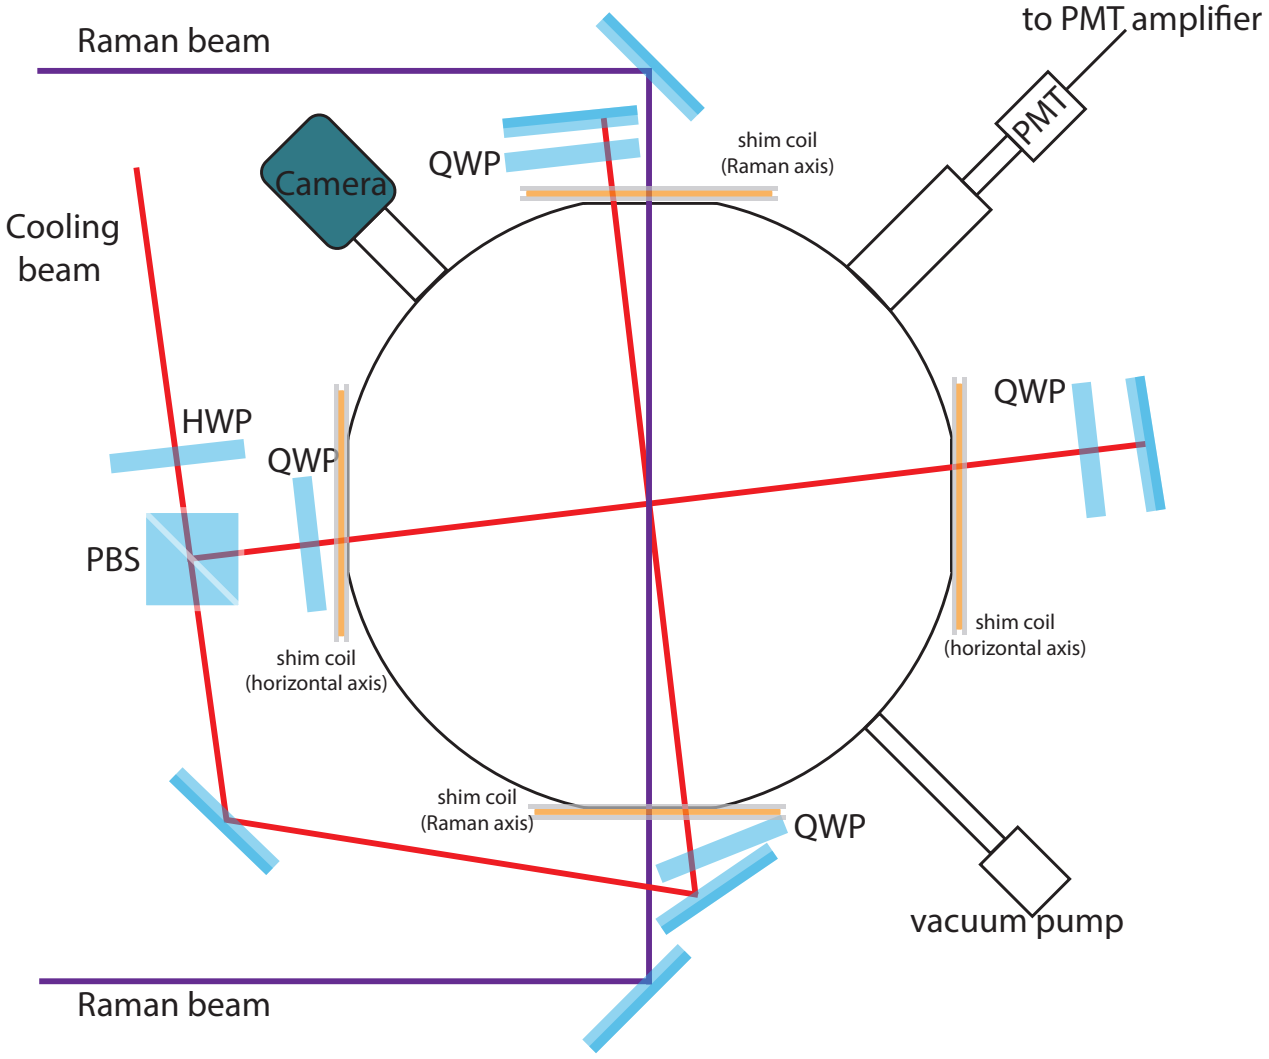

Figure 5.3: Top view of MOT chamber. HWP: half wave plate, QWP: quarter wave plate, PBS: polarising beamsplitter, PMT: photomultiplier tube.

5.2.1 MOT and Raman beam alignment

For cooling to occur in all directions, three perpendicular retroreflected beams are required. These are generated by directing the cooling beam after its fibre output coupler into a polarising beamsplitter cube to deflect one-third of the power upwards, the remainder of the beam being directed to the horizontal plane of the chamber and split again to generate the two horizontal beams. The power in each

perpendicular component is controlled by two half waveplates, and each component is back reflected after passing through two quarter wave plates.

The simplest configuration of the MOT beams would be for the beams to pass through the centre of the windows located on the 3 perpendicular axes. However, to allow easier alignment of the Raman beams onto the MOT, the Raman beams are directed through the centres of two of the windows (Figure 5.3; additionally, the shim coils used to apply a non-zero magnetic field along the Raman beam axis are fixed to the chamber such that the field is produced along an axis passing through the centre of the windows, and so offsetting the Raman beams from this axis would complicate the determination of the currents to the shim coils). The MOT beams in the horizontal plane are thus rotated about the vertical axis such that they each pass through a point slightly offset from the centre of each window. Alignment is achieved by placing a plastic cover containing a hole at the desired offset position over each window in turn and adjusting the steering and back reflecting mirrors. Fine adjustments can be made by viewing the back reflected beams at the fibre output coupler and adjusting mirrors so that they overlap with the beam coming out of the fibre.

The repump beam does not provide a velocity or position dependent force and so it is not required to propagate in all three dimensions. The repump beam in our setup is directed diagonally upwards through the vertical windows of the chamber and back reflected (Figure 5.4). Alignment of the Raman beams onto the MOT is achieved by imaging the MOT and the Raman beams along the axis of the Raman beams and is described in more detail in Section 6.3.

5.2.2 Rubidium dispensers

The MOT is formed when atoms are captured from a background gas of Rubidium. This gas is created by dispensing Rubidium from resistively heated getters, which are run at approximately 3.5A for a few minutes until the MOT is visible on the MOT camera. The chamber was last opened in 2008 and the inside walls are coated with condensed Rubidium, which may be released in quantities sufficient to achieve a MOT if the MOT coils are run without the water cooling, but this is not advisable.

If the MOT beams need to be realigned, for example if a mirror has been moved or replaced, it is useful to view the fluorescence of the atoms within the MOT

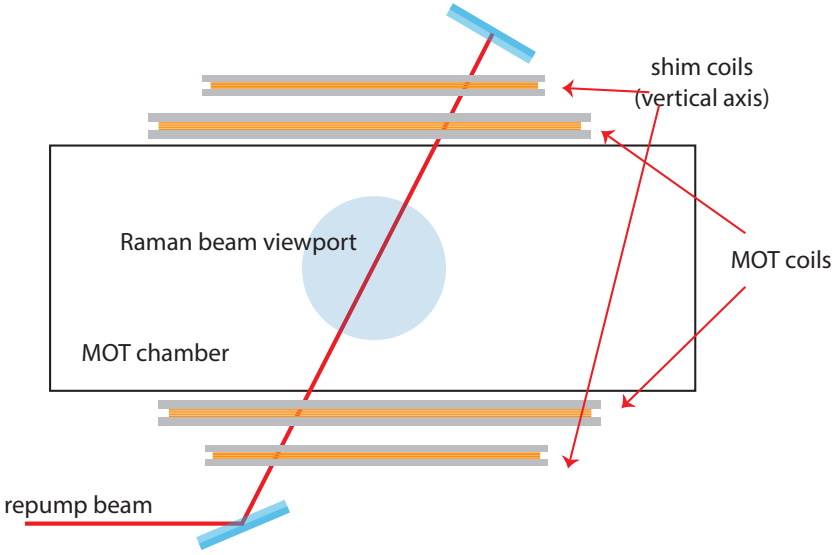

Figure 5.4: Side view of MOT chamber

beams as part of the alignment process. In this case the getters are run at a higher current (up to 5A) for a short period of time.

5.2.3 MOT camera

The MOT is observed through one of the viewports using a Watec 902H3 CCD camera and a zoom lens. This camera is used throughout the experiment to check the presence of the MOT as the locking of the cooling laser is very sensitive to vibrations on the optical table. The camera is also used when aligning the MOT; once the alignment of the beams is optimised as described in 5.2.1, the camera output is monitored and the mirrors and coils tweaked until the MOT is seen.

5.3 Magnetic fields

5.3.1 MOT coils

As described in Section 1.3.2, a magnetic field with a constant gradient in all three directions can be produced in a small region around the midpoint of the axis of two coils in the anti-Helmholtz configuration. The coils are water cooled to allow high currents (typically 2.2A in each coil, although the individual values are adjusted to maximise the brightness of the MOT).

Various magnetic components surrounding the MOT may sustain eddy currents after the MOT coils are switched off. To counteract this effect, the holders for the MOT coil wires are cut in one place so that a current loop cannot form. It was measured that the switchoff time (time for the field measured close to the coils to decay to 10% of its starting value) decreased from 4.8ms to 1.4ms after the holders were cut. Details of this measurement and further information regarding the switching of the MOT coils is found in Alex Dunning's thesis [45].

5.3.2 Shim coils

There are also three pairs of shim coils whose function is to cancel out any residual magnetic field at the MOT when the MOT coils are switched off. The shim coil currents are determined by switching off the MOT coils and observing the fluorescence of the MOT in the beams as it ballistically expands and falls under gravity. The shim currents are roughly set by adjusting until the atoms are seen to fall under gravity and are not "thrown" by any residual magnetic field. Fine adjustment of the shim currents is achieved by performing a velocity-insensitive Raman scan. If there is a non-zero magnetic field along the Raman beam axis, the m_F sublevels will not be Zeeman degenerate and five peaks in the Raman scan are observed at different frequencies. The shim coil currents are tweaked until the observed Zeeman splitting is minimised. Again, more information on the nulling of the magnetic fields along with the corresponding Raman scans may be found in Alex Dunning's thesis.

5.4 Preparation and detection

Before a coherent manipulation experiment takes place, it is necessary to initialise the state of the atoms; i.e., to prepare them all in the same internal state. We prepare the atoms in the lower hyperfine state $5S_{1/2}, F = 2$ by applying the cooling laser without the repump laser for a period of time sufficient to optically pump the atoms into $F = 2$. This is possible because, as described in Section 5.1, the cooling laser has a small probability to excite atoms to the $5P_{3/2}, F = 3$ state, from which they can decay to $5S_{1/2}, F = 2$, and without the repump laser present during the state preparation, these atoms will now be in a "dark" state and will not be affected by the cooling laser.

We also require a method of determining the proportion of the atoms which are transferred into the excited state $5S_{1/2}, F = 3$ by the Raman pulses. This is achieved by applying the cooling laser to pump these atoms into $F = 2$ and observing the fluorescence signal with a photomultiplier tube (PMT). The PMT (Hamamatsu H7422-50) is preceded by a telescope whose first lens has a large diameter to maximise the amount of light collected. At the focus of the telescope is an iris, adjusted so that any scatter from the inside of the chamber is cut out. The PMT output is amplified by a Stanford Research Systems SR570 low-noise current amplifier.

The fluorescence signal is an exponential decay because the number of atoms remaining in $F = 3$ after each absorption-spontaneous emission cycle is equal to P multiplied by the number of atoms in $F = 3$ before that absorption-emission cycle, where P is the probability to be excited into the $F' = 4$ state [2]. After measuring this fluorescence signal, all atoms are pumped into $F = 3$ by applying the repump laser, then the cooling laser is applied again and the fluorescence decay is measured. An exponential function is fitted to the two fluorescence signals, and the ratio of the amplitudes of the fits gives the proportion of atoms transferred to the $F = 3$ state by the Raman pulses.

Figure 5.5 shows one such series of fluorescence decays. There is a background signal with a magnitude of approximately 0.25V, which is present even when the MOT lasers are off; this is most likely due to scatter of stray light off the inside of the chamber. The blue lines show the PMT signal and the red lines are the exponential fits; in this case the fraction of atoms transferred is 0.1.

[2] For a cooling laser detuned by 3 linewidths, P is approximately 0.985 [67]

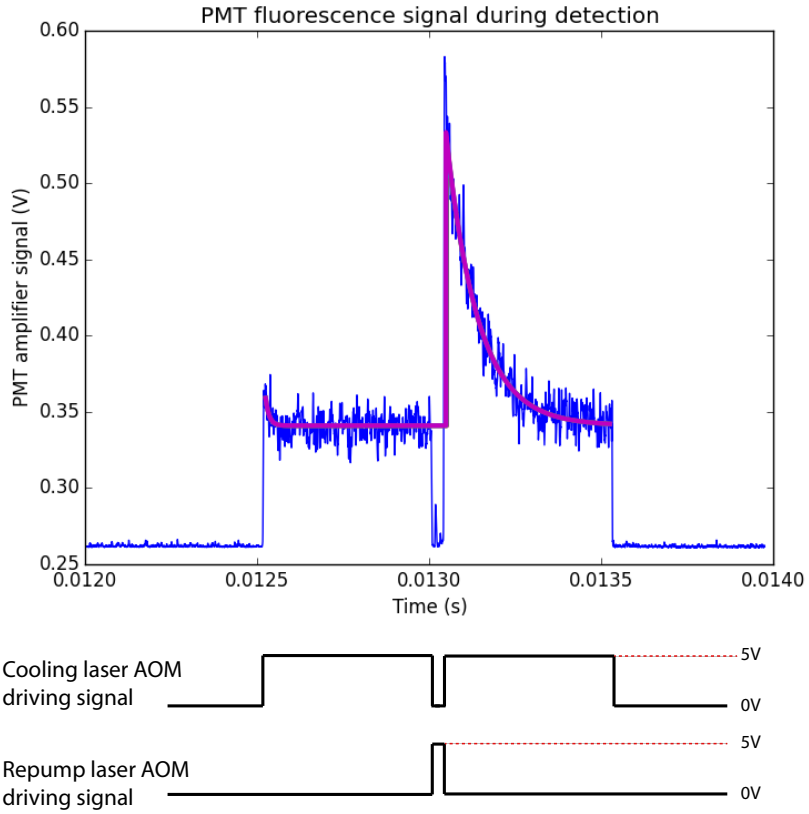

Figure 5.5: Top: Fluorescence signal during detection measured by PMT. Bottom: Cooling and repump AOM driving signals during detection.

Chapter 6

Manipulation lasers

In order to induce two-photon transitions of the kind described in Chapter 2, we require two phase-coherent laser beams with a defined frequency difference equal to the hyperfine splitting. This may be achieved by phase-locking two separate lasers; however, we choose to derive both components from one 'master' laser and generate the frequency difference using a microwave frequency source. This is because the linewidth of the Raman transition is then defined only by the uncertainty of the frequency synthesizer, and high accuracy microwave frequency sources are readily available.

In this chapter we describe the preparation of the two frequency components required for Raman transitions. Various stages of spectral filtering and amplification are necessary; these are outlined along with the procedure for switching the beams during experiments. Finally, we describe the spatial beamshaping required to produce a uniform intensity profile.

6.1 Generating two frequency components

A schematic of the Raman laser setup is shown in Figure 6.1. The master laser beam is split spectrally and spatially by an AOM. The first order is amplified and sent on to the pulsing AOM; the zeroth order is frequency shifted further by an EOM. This is followed by spectral filtering and amplifying stages before proceeding to the pulsing AOM. Both beams are then coupled into fibres for transport to the experimental chamber, each of which has beamshaping optics at the output. The zeroth order of the pulsing AOM is sent to an optical spectrum analyser and a

wavemeter for analysis. The various stages are described in more detail in the following subsections.

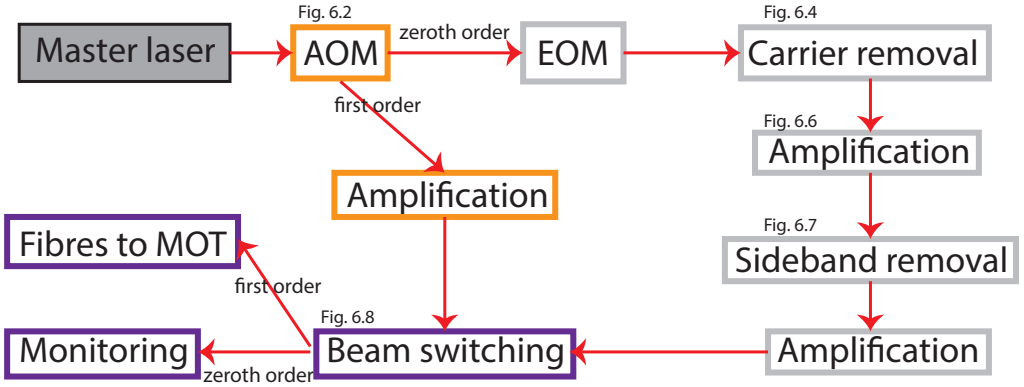

Figure 6.1: Schematic of stages required to produce two spatially and spectrally separated, phase-coherent laser beams

6.1.1 Master laser

Our master laser from which the two frequency components are derived is the Sacher Tiger, an external cavity diode laser. Unfortunately, although the specifications suggest a typical output power of between 500mW and 1000mW, we find the actual maximum to be around 150mW. Additionally, the 'rugged' construction of the cavity and grating make adjustments to the tuning angles difficult and regular tuning is required. The lasing wavelength is set by adjusting the grating angle (with subsequent optimisation of the cavity and collimation screws necessary to maximise output power), and measured by an Advantest Q8326 wavelength meter. The wavelength is typically set to 780.270 nm, which provides a sufficient detuning from the single-photon frequency to suppress these transitions.

6.1.2 AOM beam

We refer to the two frequency components as the AOM beam (the first order beam deflected by the AOM) and the EOM beam (the zeroth order which goes on to be modulated by the EOM). The AOM is driven by an RF signal at 310MHz from a Marconi 2022 signal generator. The amount of power deflected into the first order depends on the power of the RF signal applied, and we set this at -6dBm, which

results in an efficiency of approximately 5%. This beam is then passed through a tapered amplifier (TA), the active element of which is an m2k diode (TAL-0780-2000-DHP, centre wavelength 780 nm, maximum output power 2 W). The TA may be supplied with up to 3A of current (this current is only safe to supply when the TA is seeded), and the maximum seed power is 30 mW. We choose the power of the AOM input signal as a balance between seeding the TA with enough power to obtain high powers at the output, and supplying enough power to the EOM branch, which suffers a large loss of power in the spectral filtering process. The output of this TA is typically around 350mW.

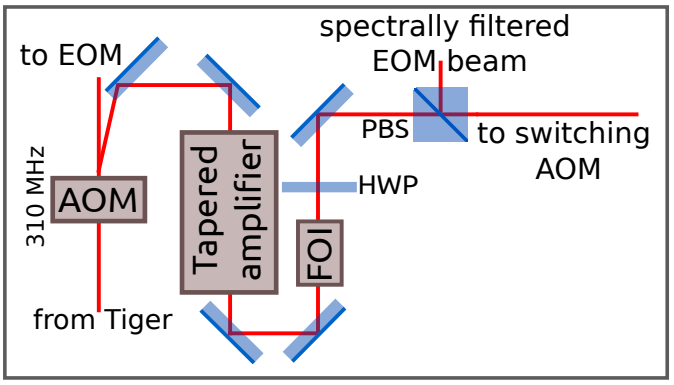

Figure 6.2: Spectral and spatial splitting of master laser: AOM beam deflection and amplification. Abbreviations as previously.

The TA output is passed through an isolator to block back reflections which could damage the diode, and subsequently a half wave plate. This is to ensure the polarisation of the AOM beam is orthogonal to that of the spectrally filtered EOM beam before they are combined using a polarising beamsplitter.

6.1.3 EOM beam

The AOM provides a frequency difference of 310 MHz between the two beams; to increase this to the hyperfine splitting of 3.030 GHz, we require a microwave frequency shift. The advantage of an AOM is that the modulated beam contains only the frequency at which the AOM is driven, and so no subsequent spectral preparation is required; however, AOMs can only operate at frequencies at which the acoustic crystal can respond. AOMs are available up to approximately 2 GHz but these devices are very costly and we choose instead to generate the remainder of the frequency difference with an electro-optic modulator (EOM).

Whereas an AOM relies on the changing response of a crystal to a sound wave, an EOM is a crystal whose properties (in our case, the refractive index) vary when an

electric field is applied. A changing refractive index results in phase modulation at the applied frequency. An incident laser beam $A_0 e^{i\omega_0 t}$ at frequency ω_0 passing through the EOM, which is driven at ω_m with an amplitude A_m, is modulated and the output field is $Ae^{i(\omega_0 t + A_m \cos\omega_m t)}$. The spectrum of this field consists of the carrier and an infinite number of sidebands with amplitudes determined by Bessel functions:

$$A_0 e^{i(\omega_0 t + A_m \cos\omega_m t)} = A_0 \left(J_0(A_m) e^{i\omega_0 t} + e^{i\omega_0 t} \sum_{n\neq 0} i^n J_n(A_m) e^{in\omega_m t} \right). \tag{6.1}$$

The EOM used in our experiment is a New Focus 4431M driven by a microwave signal (via an amplifier) from an Agilent 8648C frequency source. The frequency of the microwave signal is 2.725732440GHz, i.e. the hyperfine splitting minus the frequency shift introduced by the AOM, and this is remotely programmable to enable a varying offset during experiments (for example, to investigate the response of the atoms to a range of frequencies). Due possibly to internal damage to the crystal, the modulation depth we observe is well below that stated in the specification [66], and approximately 40% of the input power is modulated. However, this is not detrimental as we have several subsequent amplification stages, and indeed it may be helpful for our experiment as the power in the higher-order sidebands is low enough that we do not need to worry about removing them from the spectrum.

6.1.3.1 Phase shifting

The signal from the Agilent signal generator is passed through an IQ modulator (Miteq SDM0104LC1CDQ) in order to implement phase (and/or frequency) shifts. Before the signal proceeds to the amplifier, the IQ (in phase and in quadrature) modulator applies a phase shift ϕ to the signal $\cos\omega_0 t$ which depends on the values of the inputs I and Q (Figure 6.3), resulting in an output of $\cos(\omega_0 t + \phi)$. The inputs may be DC voltages for a pure phase shift, or sinusoidal signals for a frequency shift. This allows us to create Raman pulses with periods of different phases (e.g. composite pulses); additionally, we find this method of changing the frequency difference of the AOM and EOM beams to be more convenient for our experiments, compared to programming the Agilent frequency generator.

An in-depth characterisation of the response of the IQ modulator can be found in Alex Dunning's thesis [45]. Here we note two important points: 1) the intensity of the output signal varies depending on the phase shift, and 2) the measured

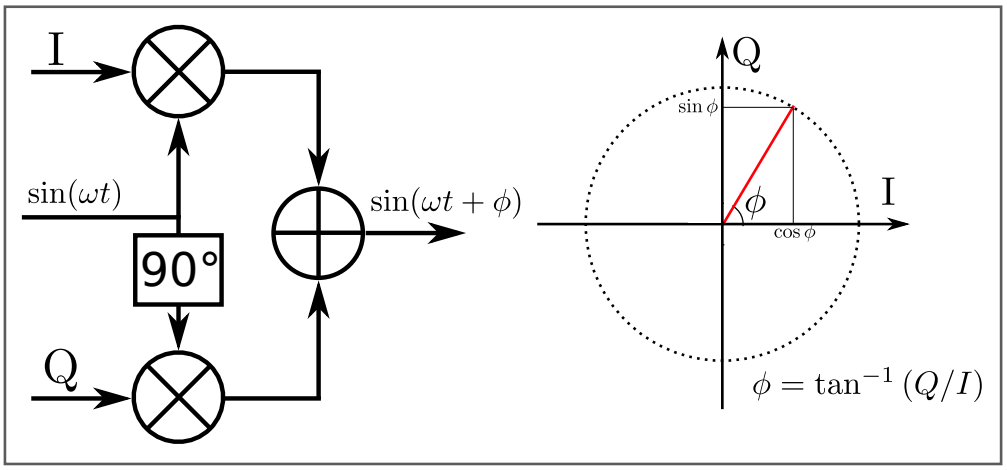

Figure 6.3: Applying a phase shift to a sinusoidal signal with an IQ modulator

phase shift is not a linear function of the applied phase shift. These effects will be discussed in further detail in Chapter 8.

6.1.3.2 Carrier removal

The beam at the output of the EOM contains many frequency components (Equation 6.1) in addition to the one required for the two-photon transition (the first order lower sideband). These other electric fields cause disruptive light shifts and it is advantageous to remove them. However, as the frequency components are only spectrally separated and not spatially separated (as in the AOM case), spectral filtering is a little more challenging. In our experiment we employ two methods, developed by the Quantum Control group, to remove the carrier and the off-resonant first order sideband.

The carrier removal technique (published as [70], with detailed analyses in previous group theses [45] and [71]) exploits the fact that the EOM only modulates one polarisation component. If the polarisation of the input beam is slightly offset from that required for maximum modulation, then a half wave plate and polarising beam splitter placed after the EOM will reject the carrier when the HWP is correctly aligned. In practice, the EOM exhibits some temperature-dependent birefringence; the carrier becomes elliptically polarised and is not completely rejected by the polarising beamsplitter. We actively correct for this birefringence by applying an AC signal at 1 kHz with a varying amplitude to a liquid crystal cell.

Figure 6.4 shows the setup for active feedback carrier removal. A half wave plate is placed before the liquid crystal (LC) to rotate the polarisation of the incoming

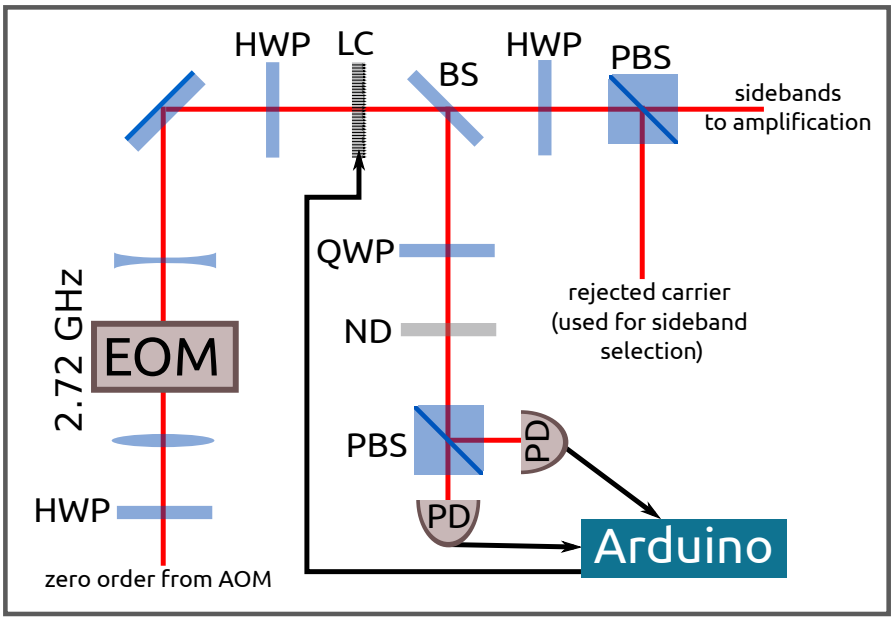

Figure 6.4: Carrier removal from EOM beam using a liquid crystal for active feedback. EOM: electro optic modulator, LC: liquid crystal, PD: photodiode

light, to ensure the maximum possible range of phase modulation. After the LC, a beam sampler picks off a small part of the beam which passes through a quarter wave plate and a polarising beamsplitter cube (PBSC). The two photodiodes at the outputs of the PBSC measure the intensity of the left and right circular polarisations present in the beam. When these two components are equal in intensity, the polarisation after the LC is linear. When these intensities are unbalanced, the polarisation is elliptical, and the difference of the two values determines the correction required. A PID function controls the amplitude of the AC signal applied to the LC based on this difference.

The spectrum is monitored by directing both the EOM and the AOM beams into a scanning Fabry-Perot spectrum analyser (FP-SA). The FP-SA is an optical cavity whose length is scanned using a piezo-electric transducer (PZT). The cavity transmits only modes of a specific frequency and the intensity of the transmitted light is monitored using a photodiode whilst the length is scanned. The PZT voltage (a sawtooth which repeatedly sweeps through one free spectral range of the cavity by scanning its length) and photodiode amplification are provided by a Spectral Products SAC-2 Optical Spectrum Analyser controller. The amplified signal from the SAC-2 is monitored on an oscilloscope. The free spectral range of the FP-SA is 2 GHz.

The procedure for carrier elimination is as follows. The half waveplate before the polarising beamsplitter is rotated until the intensity of the carrier (as observed on the scope) is at a minimum. The amplitude of the AC signal applied to the LC is varied (by commands sent to the Arduino) until a value is found where the carrier is minimised. At this AC signal amplitude, the quarter waveplate is rotated until the difference between the photodiode values (intensities of the oppositely circular polarised components of the EOM beam) is minimised. The AC signal amplitude is fed into the PID loop as the setpoint and carrier is then suppressed.

The degree of suppression is determined by monitoring the amplified photodiode signal on the scope. However, at carrier suppressions greater than approximately 18 dB, it is difficult to see the carrier component within the noise, and indeed the difference between 20 and 25 dB of suppression is almost undetectable, as shown in Figure 6.5. To distinguish between these cases, the spectrum was acquired by a Thorlabs OSA 202 which has a sensitivity of -60 dBm/nm (top plot in Figure 6.5). This OSA is not normally used in this experiment and was employed to determine the level of carrier suppression possible; in practice, carrier suppression to the extent that the carrier is not distinguishable from the noise appears to be sufficient.

From this we can conclude that during the normal running of the experiment, the carrier is suppressed by at least 18 dB. The power in the sidebands is also reduced by an amount depending on the orientation of the half wave plate just before the EOM as a side effect of the carrier removal, typically resulting in a combined sideband power of approximately 12mW and amplification is necessary before proceeding to the next stage of spectral filtering.

6.1.3.3 First stage amplification: injection-locked laser

Losses occur due to the carrier removal (as described in the previous section) and coupling into the fibre Mach-Zehnder interferometer for sideband removal (see following section), so the beam is amplified by an injection-locked laser (ILL) before sideband removal. The ILL is a laser diode contained in the same type of mount as is used for the MOT lasers, with the grating and mirror at the output removed. The anti-reflection coated diode reproduces all the spectral components of the input beam and amplifies the power - typically, running at 70mA with an input power of 12mW, the output is 35mW. Figure 6.6 shows the setup of the ILL; the anamorphic prism pair reshapes the input beam to match the diode output

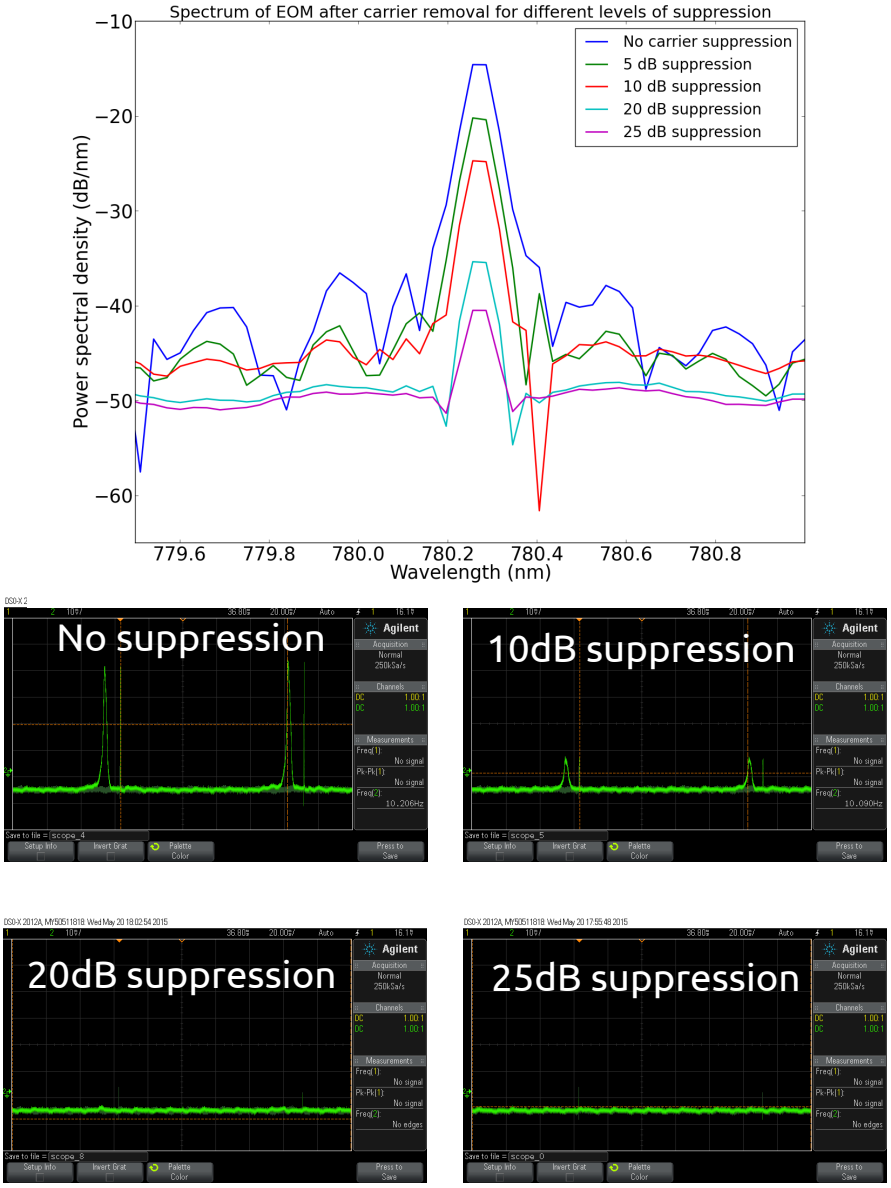

Figure 6.5: Scope trace images from the scanning Fabry-Perot spectrum analyser (Spectral Products) and spectrum measurements from the Michelson Interferometer OSA (Thorlabs) for various levels of carrier suppression. The gain of the Fabry-Perot OSA amplifier is at its maximum value. The images taken from the oscilloscope have a horizontal axis scale such that the carrier is visible twice; the frequency separation is the free spectral range of 2 GHz.

profile, and the isolator prevents back reflections from the fibre disrupting the stability of the laser.

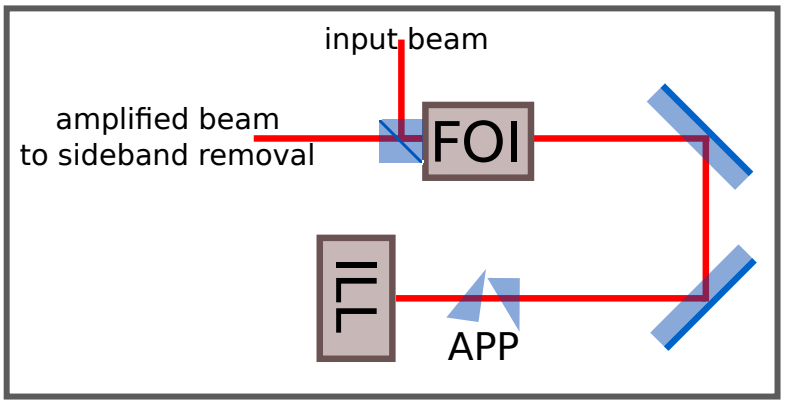

Figure 6.6: Injection-locked laser for amplification of sidebands before Mach Zehnder spectral filtering. FOI: Faraday optical isolator, APP: anamorphic prism pair, ILL: injection locked laser.

6.1.3.4 Sideband removal

The off-resonant sideband is removed for two reasons. Firstly, its presence at the MOT causes a light shift of the atoms' energy levels. In some polarisation configurations, this light shift is intensity-dependent, and this is problematic when carrying out experiments where the Raman beam intensity is varied. Secondly, the EOM beam requires further amplification at this stage in the experiment, and sending two frequency components through a tapered amplifier may result in the generation of additional sidebands [72].

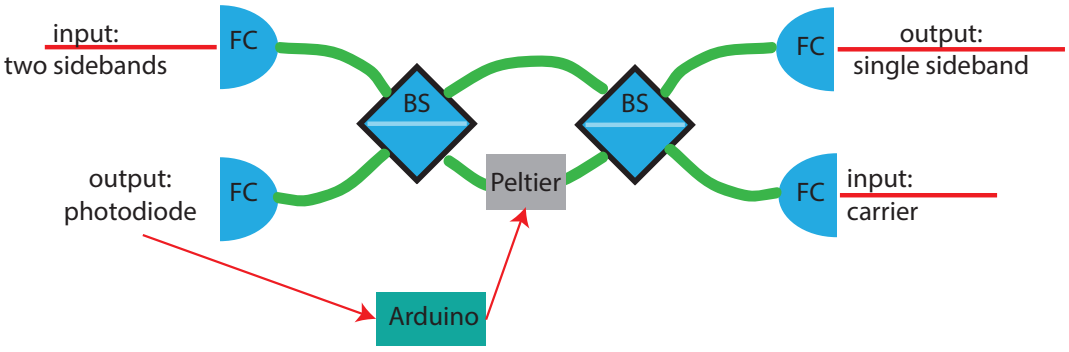

Figure 6.7: Fibre Mach-Zehnder interferometer for the removal of the off-resonant sideband. FC: fibre coupler, BS: non-polarising beamsplitter.

The sideband removal technique is described in detail in [73]. The carrier-removed beam is passed through a fibre Mach Zehnder interferometer with refractive index $\eta(\omega)$. The phase difference between the two arms is $\Delta\phi(\omega) = \omega\eta(\omega)\Delta l/c$ where Δl is the difference in the lengths of the two arms (not the optical path difference). Since $\Delta\phi$ is frequency dependent, when two frequency components are sent

through the interferometer they will each have a different phase difference between the two paths. The two components are separable at the final beamsplitter if their difference in phase difference is equal to π.

The experimental setup is show in Figure 6.7. The fibre Mach Zehnder is a custom made Oz Optics interferometer whose arms differ in length by an amount sufficient to separate the two sideband frequencies at the output of the interferometer. However, the optical path lengths are strongly temperature dependent and so active stabilisation is necessary. The photodiode monitors the intensity of the backwards transmitted carrier and the resulting error signal is fed into a PID loop. The PID output determines the current supplied to the Peltier cooler. The setpoint is found by varying this voltage without the PID loop running and monitoring the spectrum on the Fabry-Perot spectrum analyser.

6.1.3.5 Second stage amplification

After removal of the off-resonant sideband, the EOM beam requires a further stage of amplification. The power in the EOM beam at the output of the fibre Mach Zehnder is approximately 3 mW and it is passed through a Toptica BoosTA. The output power is typically 250mW, of which approximately 20mW is amplified spontaneous emission (ASE), which can be removed by the heated Rubidium vapour cell (Figure 6.8).

6.2 Beam switching

After their final stages of amplification, the AOM and the EOM beam are combined using a polarising beamsplitter. The output of the beamsplitter is directed into an 80MHz AOM (Gooch and Housego M080-2B/F) for beam switching: the first order of the AOM output continues to the fibres used to take the Raman beams to the MOT (Figure 6.8). This AOM controls the duration of the Raman pulses and also their intensity (by varying the voltage level on the modulation input of the AOM driver). The signal applied to the modulation input of the driver is generated by a Tabor 8024 arbitrary waveform generator.

The zero order of the switching AOM (containing a fraction of both the EOM and the AOM beam, the exact proportion depending on the efficiency of coupling into the switching AOM) is directed to the Fabry-Perot spectrum analyser and

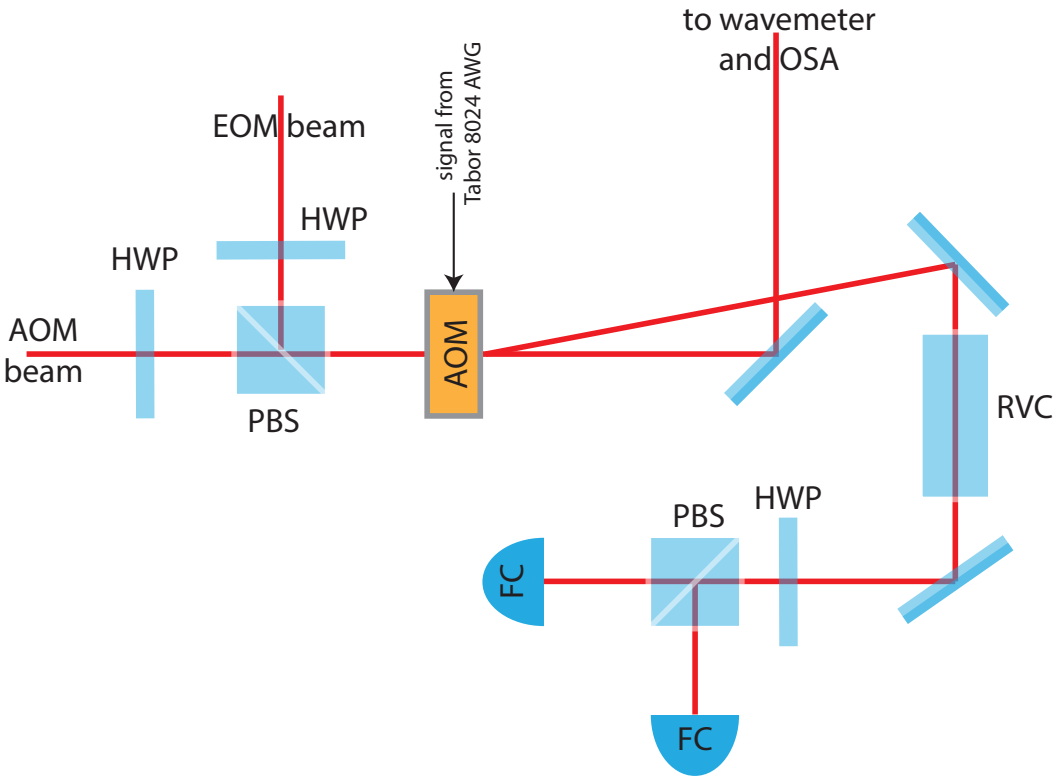

Figure 6.8: Raman beam paths after their final amplification stages. PBS: polarising beamsplitter, HWP: half wave plate, FC: fibre coupler, RVC: Rubidium vapour cell.

the Advantest wavemeter for analysis and monitoring. The Advantest wavemeter measures the wavelength of the combined Raman beams to a precision of 0.001nm and is used to determine the detuning Δ from the two-photon transition; typically, the Sacher Tiger laser is wavelength tuned such that the wavemeter reads around 780.270nm.

6.3 Beam shaping and alignment

After passing through the switching AOM and vapour cell, the two Raman beams are separated by a polarising beamsplitter and coupled into separate fibres. The coupling efficiency is at maximum around 50% due to the non-Gaussian beam shape of the outputs of the tapered amplifiers. There are separate fibres for the two Raman beams so that both co- and counter- propagating experiments may be performed. At the output of each fibre is a series of beam shaping optics to shape the Gaussian fibre output into a square flat top beam. This is to ensure that the

intensity of the Raman beams is as uniform as possible across the atom cloud and to maximise the available power; one alternative would be to expand the Gaussian output so that the beams are approximately constant in intensity across the atom cloud but this would use only a small fraction of the available power in the Raman beams and limit the achievable Rabi frequencies. A detailed characterisation of the beam shaping optics is found in Alex Dunning's thesis [45].

Alignment of the Raman beams onto the MOT is achieved by using a camera to look down the path of the beam to be aligned. The Raman beam is blocked and a zoom lens is used to focus on the MOT. The camera is placed to that the MOT is in the centre of the image. Then the Raman beam is unblocked and mirrors used to overlap the top hat with the position of the MOT. This alignment can also be checked by placing the repump output fibre into the fibre coupler for the Raman beam shaping and adjusting the Raman beam mirrors until the MOT is seen.

6.4 Experimental sequence

Once the Raman beam preparation has been adjusted by tuning the master laser to the appropriate wavelength, optimising the spatial filtering, maximising the powers of the Raman beams and aligning them into the MOT, our coherent manipulation experiments may take place. The waveforms controlling the pulse duration and intensity of the Raman beams and the MOT beams are generated and uploaded to the Tabor waveform generators which control the relevant AOMs. A SpinCore PulseBlaster is used to trigger the Tabor waveform generators to output a waveform, and additionally to switch the MOT coils and the common MOT AOM, and to trigger the oscilloscope reading the PMT signal.

The sequence of events occurring during a typical coherent manipulation experiment is as follows.

Initialise MOT: The MOT coils and both MOT beams are switched on to cool and trap the atoms.

Molasses cooling: The MOT coils and the repump beam are switched off; the cooling beam is ramped down in intensity.

State preparation: The cooling beam remains on at a constant intensity to pump atoms into the $F = 2$ state.

Raman pulses: The Raman beams are applied for a period of time depending on the experiment being carried out.

State detection: First the cooling beam is applied to pump any atoms which were transferred to the $F = 3$ state by the Raman pulse into the $F = 2$ state, and the fluorescence signal is detected by the PMT. Then the atoms are all returned to $F = 3$ by the repump beam, and the cooling beam is applied again to detect the fluorescence as a result of pumping all atoms from $F = 3$ to $F = 2$.

MOT recovery: The MOT beams and MOT coils are switched on and the atoms are cooled and trapped again in preparation for the next sequence.

This sequence would be repeated several times for each Raman beam configuration (pulse duration, pulse phase, single-photon detuning) to be investigated - for example, during a Raman scan, the duration of the pulse is kept constant and the single-photon detuning is varied; for each detuning value, the sequence might be repeated 8 times and the fluorescence signals averaged.

Chapter 7

Gaussian to tophat beamshaping

The Gaussian to tophat shaping of our Raman beams is achieved by a refractive beamshaper in combination with a spherical lens, which allows control over the side length of the square profile of the shaped beam. An alternative beamshaping method was also investigated and is described in this chapter: the use of a spatial light modulator (SLM), which applies an arbitrary, controllable phase profile to a laser beam, allowing control over the intensity in the Fourier plane.

The operation of a spatial light modulator is described. We also describe and implement an algorithm for the calculation of the phase profile applied to the SLM, and investigate a feedback correction method.

7.1 Spatial Light Modulators

A spatial light modulator is a device which may change the phase, amplitude, or polarisation of a beam in one plane, usually with the aim of controlling the properties of the beam in the Fourier plane. The ability of an SLM to dynamically change the potential landscape created by light finds uses in many areas of physics; in atomic physics, amongst other applications, SLMs have been used to create arrays of dipole traps [74], to create a 'stirrer' for inducing rotation in a BEC [75], and for transport and addressing of single atoms in dipole traps [76]; in semiconductor physics, they find uses in creating potentials for polariton condensates [77].

For the purposes of beamshaping we are interested in amplitude or phase modulation. A beam with electric field $E_0(x, y) = A_0(x, y)e^{i\phi_0(x,y)}$ in the input plane will

have the Fourier plane field $E_1(x', y') = \mathcal{F}(E_0(x, y)) = A_1(x', y')e^{i\phi_1(x', y')}$; varying the phase or amplitude of the beam in the input plane will affect the phase and amplitude in the Fourier plane. As described in 6.3, an atom does not move a large enough distance within the Raman beams during a pulse to be affected by amplitude or phase variations, and so we focus on control over the intensity only: shaping a beam from a Gaussian profile to a tophat (a square region of uniform intensity).

SLMs exist to modulate phase and/or amplitude and in reflective or transmissive modes. Here we investigate the performance of the HEO 1080P, a phase-only reflective LCOS (liquid crystal on silicon) SLM produced by Holoeye. The active area consists of a layer of liquid crystal between two transparent electrodes, on a silicon substrate. When no electric field is applied to the electrodes, the liquid crystal molecules align parallel to the substrate. An applied voltage will rotate the molecules so that their long axis is no longer in the plane of the substrate, changing the effective refractive index of the liquid crystal, and causing a phase shift on reflection. This phase shift is maximised for light which is linearly polarised along the long axis of the liquid crystals in their 'off' state; phase modulation without polarisation modulation will only occur for light polarised in this way. The microdisplay of the HEO 1080P has a resolution of 1920×1080 individually addressable pixels, each of area $8\mu\text{m} \times 8\mu\text{m}$. It operates as a 'second screen' - in order to apply a phase pattern to the SLM, it is connected to a computer as a second monitor and the phase image is displayed on this monitor.

7.2 Iterative Fourier Transform Algorithms

The electric field in the Fourier plane after reflection off the SLM is the Fourier transform of the field incident on the SLM with an additional phase modulation due to the SLM. The problem is thus to find the appropriate phase modulation to apply to a given incident field which will recreate the desired intensity distribution in the Fourier plane. There is not one unique solution to this problem as the phase in the Fourier plane is not constrained; iterative Fourier transform algorithms take advantage of this phase freedom to approximate the desired pattern as closely as possible.

7.2.1 Gerchberg-Saxton Algorithm

The Gerchberg-Saxton (GS) algorithm [78] was the first of these to be proposed. A flow diagram is shown in Figure 7.1. It takes as input: the amplitude of the incident beam A_0, a guessed initial phase distribution e^{iK_0}, and the desired intensity distribution at the output (Fourier) plane I_{out}.

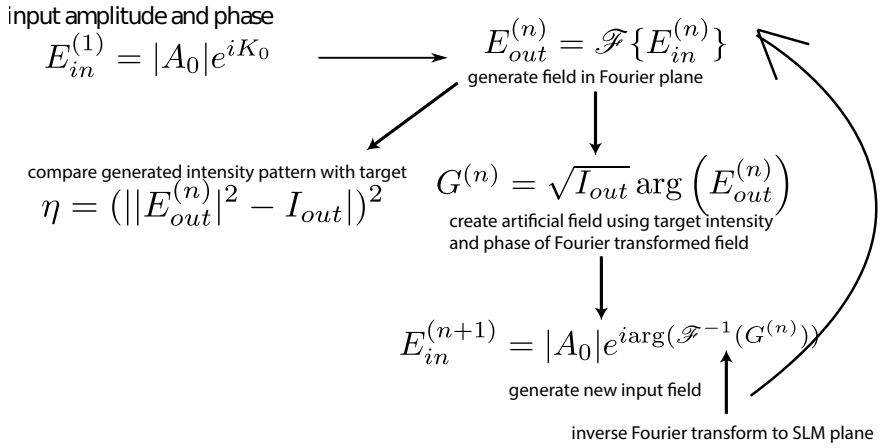

Figure 7.1: Gerchberg-Saxton algorithm flow diagram

The input field $E_{in}^{(n)}$ is propagated to the Fourier plane. The amplitude of this Fourier transform is discarded and the phase is combined with the desired output distribution. This function is propagated back to the input plane by an inverse Fourier transform and the phase is combined with the input amplitude. When a sufficient number of iterations has taken place (judged by comparing the algorithm output intensity $\left|E_{out}^{(n)}\right|^2$ distribution to the desired output), the phase distribution to be applied to the SLM is then given by $\arg\left[E_{in}^{N+1}\right]$. A proof of the convergence of the GS algorithm is outlined in Appendix B.

7.2.2 Mixed-Region Amplitude Freedom modification

The GS algorithm is not optimal in situations where the target intensity distribution is required to be smooth over a large area; one alternative is the Mixed-Region Amplitude Freedom (MRAF) modification to the GS algorithm. The MRAF method [79] trades a drop in the efficiency of the beam shaping for a smoother intensity profile, by allowing some amplitude freedom in the output plane outside of the boundary of the region of interest. In contrast, the GS algorithm has no freedom of amplitude and constrains the entire region of the output plane outside of the region of interest to zero intensity.

The MRAF algorithm divides the Fourier plane into two regions: the signal region (SR), which in our case defines the region which is close to the atoms and includes the non-zero part of the target intensity profile and an area of zero intensity around it, and the noise region (NR), which is the rest of the Fourier plane. In the algorithm, step 2 of Figure 7.1 is replaced with $G = (m\sqrt{I_{out}}SR + (1 - m)NR\,|E_{out}|) \times \exp(i\arg(E_{out}))$. The phase of the Fourier transformed input field is kept over all regions of the Fourier plane. In the signal region the amplitude of the Fourier transformed field is replaced with the target amplitude, as in the GS algorithm, but in the noise region the amplitude of the Fourier transformed field is retained - i.e. it is not constrained to be zero. The amount of amplitude freedom allowed is controlled by the value of the parameter m, which is generally chosen to be a compromise between the smoothness of the pattern (better for small m) and the amount of light lost to the noise region (smaller for large m). We also define a measure region (MR) over which the flatness of the calculated tophat is evaluated. We use the RMS error to evaluate the deviation from a flat tophat

$$\eta = \sqrt{\frac{1}{N_{pixels}} \sum_{MR} (I_{out} - I_{av})^2}, \tag{7.1}$$

where I_{av} is the average of all pixels in the MR.

7.2.3 Target profile

The target intensity profile is a region of uniform intensity. As the power available for the Raman beams is limited and we wish to produce as high an intensity as possible at the MOT cloud, the size of the target is chosen accordingly. A square tophat (as opposed to a circular tophat) is used in these tests for better comparison with the refractive beamshaper.

Although the SLM has a resolution of 1920×1080 pixels, we require a square array for the Fourier transform, and so we use an area of 1000×1000 pixels in the centre of the SLM as the region over which the phase may be controlled. This array is padded with zero pixels on all sides, giving an array of 2000×2000 pixels. The coordinates x and y in the following sections represent distances in μm, with the origin being at the top left corner of the SLM plane.

Recreating a shape with sharp edges would require high spatial frequencies and so we use instead a super-Gaussian function:

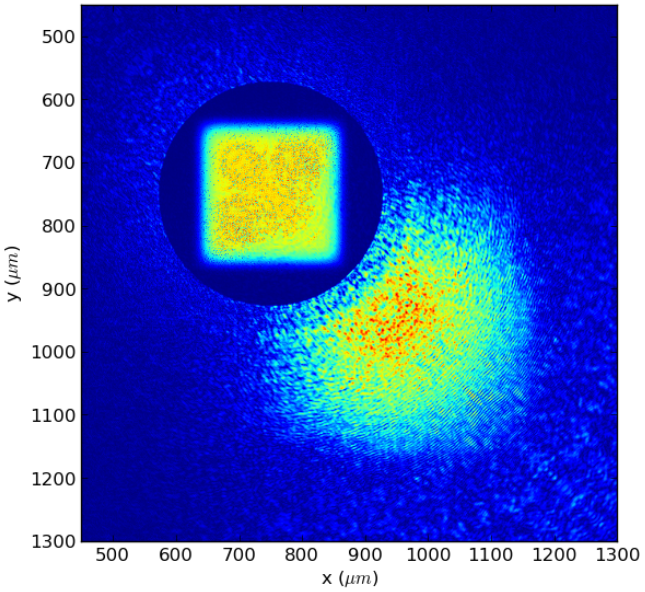

Figure 7.2: Calculated MRAF algorithm output intensity distribution when the starting phase is not set so that the first propagated field overlaps with the target distribution in the Fourier plane. Optical vortices have been created in the tophat.

$$I(x,y) = \exp(-\alpha((x - x_{offset})^m + (y - y_{offset})^m)). \tag{7.2}$$

Here x and y are the distances in μm from the top left corner of the array of pixels over which the phase is controlled. The offset coordinates (x_{offset}, y_{offset}) determine the position of the centre of the tophat. We include this offset in order to locate the shaped beam away from the centre of the Fourier plane due to the presence of the zero order (unmodulated) reflection from the SLM at this position.

7.2.4 Choice of initial phase

The choice of initial phase e^{iK_0} is crucial to the later accuracy of the algorithm; it is most useful to choose it such that the initial propagation of the field to the output plane causes $\left|E_{out}^{(1)}\right|^2$ to be a patch of light coinciding with the position of the desired intensity distribution. This can be achieved by quadratic (to focus/defocus the light and spread it out over the extent of the target) and linear (to deflect the position of the intensity maximum) phase functions.

The quadratic phase function is of the form $K_q = R(ax^2 + (1-a)y^2)$, where R determines the amount of focussing or defocussing of the beam and a allows ellipticity in the case of targets which are asymmetric in x and y (a is set to 0.5 in all work presented here). The linear phase function $K_{lin} = B(x\cos u + y\sin u)$ allows positioning of the beam in the Fourier plane: in this work u is set to $\pi/4$ as the offset coordinates of the target are equal, and varying B controls the position along the line $x = y$.

Choosing an initial phase where the first propagation of the input field to the Fourier plane does not overlap with the target results in inefficiency (light is lost to the NR) and creation of optical vortices (regions of zero intensity); Figure 7.2 shows the algorithm output for a misaligned initial propagation. These two effects are established in the early iterations of the algorithm [80] and so it is advantageous to ensure the initial phase is optimal. In the case of the MRAF algorithm, R is a compromise between allowing enough light initially in the NR to take advantage of the amplitude freedom and having enough light in the SR to give an acceptable efficiency; the choice of R is also influenced by the value of m.

7.3 Experimental setup

Figure 7.3 shows the setup for investigating the effects of applying phase patterns to a laser beam. We use a diode laser which is expanded so as to address a large area on the SLM. A half wave plate before the SLM ensures that the incident beam polarisation is aligned for maximum phase modulation. After the SLM, the zero order (not phase modulated) beam is blocked and the modulated beam is focussed onto a camera, which is connected to the computer for image processing and analysis in the case of the feedback correction method outlined in later sections. The phase patterns are calculated by the computer using an Iterative Fourier Transform Algorithm (IFTA) and then displayed on the SLM.

The angle of incidence on the SLM is ideally as small as possible so that the area of a pixel projected onto the surface perpendicular to the propagation direction of the beam is maximised; this means that a greater percentage of the incident light can be phase modulated.

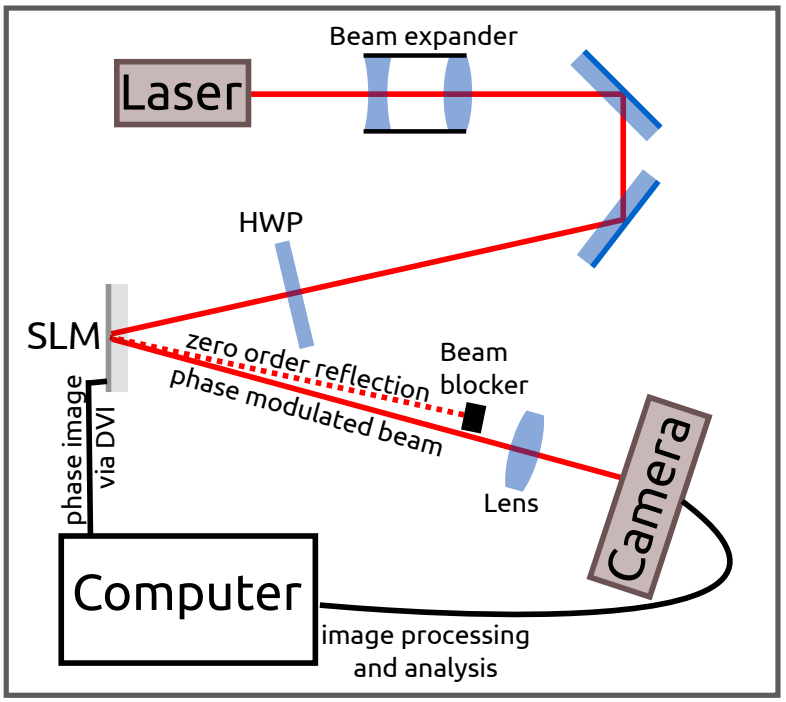

Figure 7.3: Experimental setup for SLM

7.4 Results

The MRAF algorithm was used to calculate a phase pattern to produce a square tophat with side length 1.5mm. The value of the parameter m was chosen to be 0.39 as this resulted in a tophat with intensity variations of approximately 10%, which is comparable to the performance of the refractive beamshaper. The algorithm was run for 100 iterations. The output intensity distribution in the entire Fourier plane as calculated by the algorithm is shown in Figure 7.4(a): the accumulation of light in the noise region is clear, and intensities at some points in this region are larger than the intensities in the signal region. The inset shows the intensity distribution in the signal region, rescaled to the maximum value in that region. The calculated phase distribution is shown in Figure 7.4(b). We measure the RMS deviation over a region defined as a square centred on the tophat with dimensions 150×150 pixels (side length 1.2mm); we find the RMS deviation to be 0.3% of the average intensity value in this region.

The output intensity distribution when the phase pattern in Figure 7.4 b) is displayed on the SLM is shown in Figure 7.5, along with horizontal and vertical profiles through the centre of the top hat. As in the algorithm output, the accumulation of noise in the noise region is clear. The sharp transition from signal region to noise region results in areas of high intensity at the border between the

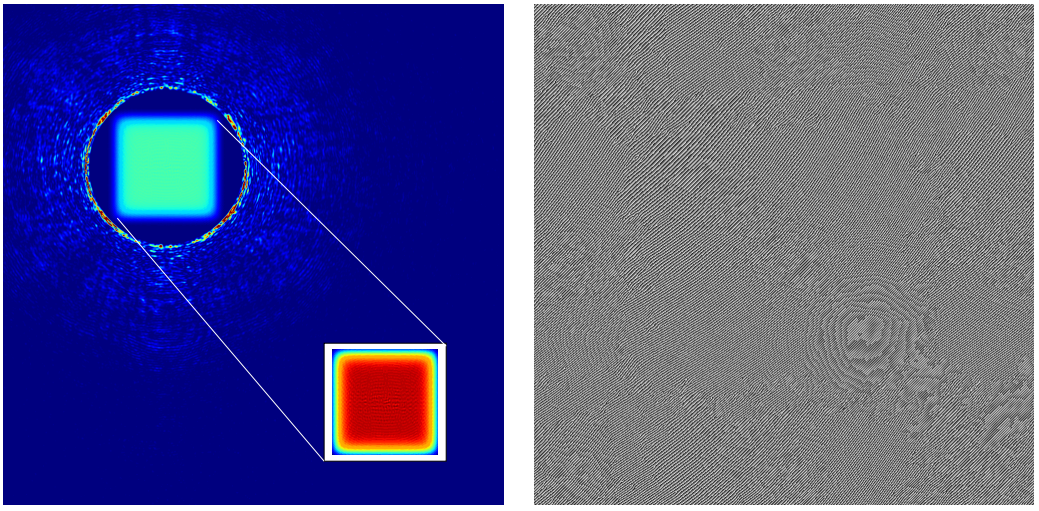

(a) (b)

Figure 7.4: (a) Tophat intensity distribution as calculated by MRAF algorithm and (inset) rescaled distribution of tophat area only; (b) calculated phase distribution to produce (a) in the Fourier plane.

two regions. This could potentially be smoothed out by making the transition between signal and noise regions smooth, although preliminary investigations suggest that the smoothness of the top hat profile may be affected. The RMS deviation over a measure region defined as a square centred on the top hat with side length 1.2mm is 6.6% of the average intensity value in this region. This is not equal to that predicted by the MRAF algorithm. This could be due to mismatch between the algorithm and the experiment of the input amplitude at the SLM. It could also be due to a mismatch between the assumed transformation between SLM and image plane (the Fourier transform) and the transformation which occurs in the experiment, which is likely to be distorted from a Fourier transform due to, for example, slight misalignment of the lens. One possible way to improve on the smoothness of the experimentally generated top hat is to implement a feedback correction algorithm, for example as described in [75] and [81]; in the latter paper, an improvement in RMS error from 22% to 7% was achieved.

The quoted RMS deviation as a percentage of average intensity of the top hat of the beamshaper used in our experiment (Topag Lasertechnik GTH- 4.2.2) is 2.5%; this beamshaper has an efficiency of 95%. Due to the unavoidable loss of light to the noise region, we achieve an efficiency of 30% with the SLM.

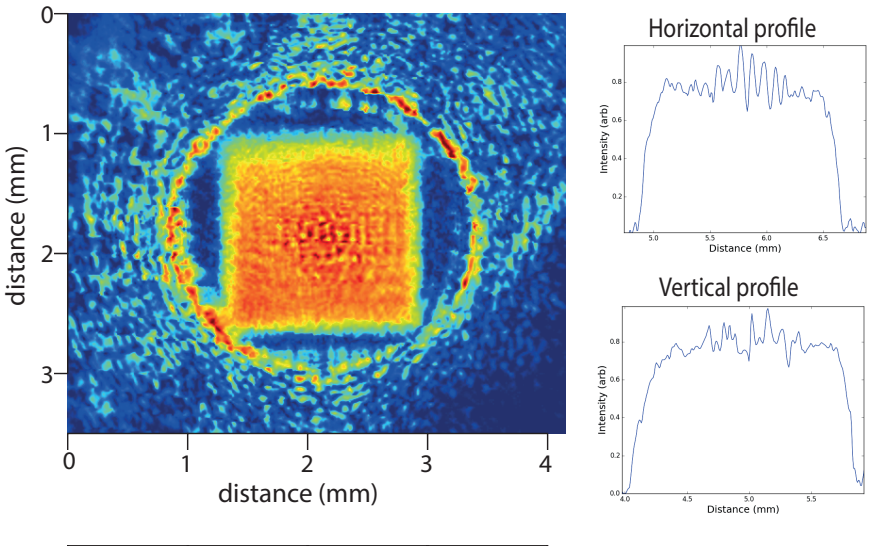

Figure 7.5: Top hat intensity observed in the Fourier plane when the phase distribution from Figure 7.4 b) is displayed on the SLM.

7.5 Conclusions

Although the achievable performance of the SLM is comparable to that of the refractive beam shaper in terms of flatness of the intensity profile, the SLM's low efficiency renders it impractical for use in our experiment, where maximising the power in the Raman beams is crucial. However, if more power were available, the ability of the SLM to reshape the Raman beams could have potential uses in our experiment. For example, the velocity distribution in our MOT is a double Gaussian function because molasses cooling is more effective in the region of higher intensity MOT beams at the centre of the MOT; so if atoms towards the edge of the MOT are on average faster, the spatial intensity distribution of the Raman beams could be adjusted to produce a different Rabi frequency for these atoms.

Part III

Results

Chapter 8

Modelling and characterisation of composite pulses

This chapter presents an experimental and theoretical study of composite pulses in ultracold Rubidium. We demonstrate that composite pulses can compensate for errors present in our system and show an improved fidelity. We establish a good agreement between the model and the experiment and use our model to investigate further composite pulse sequences, including one sequence which has not been previously proposed. We also investigate the performance of composite pulses in atom interferometers and compare this to the performance as inversion pulses.

Parts of this chapter are published in [2]. All atomic physics experimental data in this chapter were taken jointly with Alex Dunning.

8.1 Motivation

As described in Chapter 4, composite pulses can be used to increase the fidelity of inversion pulses in a system which experiences off resonance and pulse length errors. As a precursor to the results presented in the rest of this chapter, we show in Figure 8.1 the simulated performance of several composite inversion pulses as the range of off-resonance errors (quantified by the atomic momentum distribution width) increases. It is clear that as the temperature of the atoms decreases, the peak fidelity of the composite pulses approaches 1 whereas the fidelity of the π pulse is limited to ≈ 0.93.

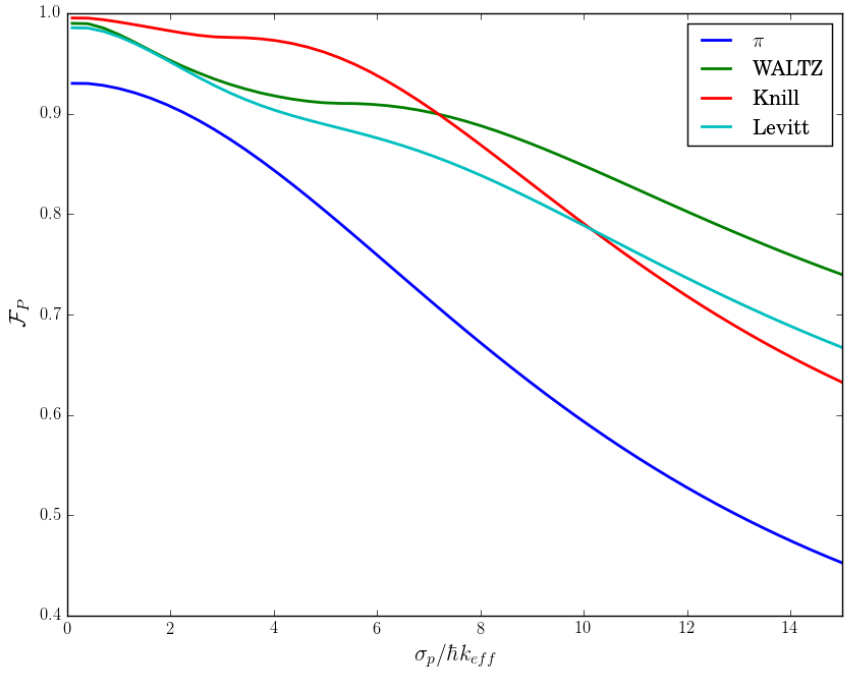

Figure 8.1: Peak fidelity of various inversion pulses as a function of momentum distribution width. Blue: π, green: WALTZ, red: Knill, light blue: Levitt.

As discussed in Section 3.2.1, the fraction of atoms following the closed interferometer path in a large-area atom interferometer is related to the fidelity of the augmentation pulses. An improvement in the fidelity from 0.7 to 0.93, as is shown in Figure 8.1 for the WALTZ pulse at a momentum distribution width of $7.5\hbar k_{eff}$, results in 75% of the atoms following the closed path as opposed to 25%.

This Figure serves as motivation for further study of composite pulses in atomic systems with a range of off-resonance and pulse length errors: it shows that fidelities approaching 1 are achievable, and the π pulse can always be improved upon; furthermore, high fidelity state preparation without optical pumping can be achieved.

8.2 Experiments to be simulated: Temporal and spectral scans

Two types of experiment are of interest in evaluating pulse sequences applied to atoms: temporal scans and spectral scans. In a spectral scan, a pulse sequence is applied in its entirety and the final excited state population is measured as a

function of the detuning δ from two photon resonance. This is of interest when investigating the performance of a pulse sequence in a system with off-resonance errors; it is also used as a diagnostic tool, for example to determine the required magnetic fields for Zeeman degeneracy (Section 5.3.2).

In a temporal scan, the evolution of the state amplitude during the pulse sequence is of interest. The two-photon detuning is constant and is set to the frequency which gives maximum transfer to the excited state, as determined by the results of a spectral scan. The pulse sequence is applied for varying durations and the excited state population is measured as a function of pulse length. A Rabi flopping experiment is one example of a temporal scan, and is performed prior to composite pulse experiments, to determine the Rabi frequency which then sets the π pulse duration.

8.3 Design of simulations

In all atomic physics experiments described in this work, the measurement we make is of the relative populations of two states; in order to simulate these experiments, we need to be able to calculate these populations given various experimental parameters. Chapter 2 outlined a derivation of the time evolution of the probability amplitudes for our pseudo-two-level system, and the populations are simply the modulus-squared of the amplitudes. The simulations use equations 2.30 and 2.31 to calculate these amplitudes for each velocity class and Raman route.

From Raman velocimetry measurements [45] we find that the 1-D momentum distribution of the MOT $n(p)$ can be well modelled by a double Gaussian profile:

$$n(p) = n_1 \exp\left(-\frac{p^2}{2\sigma_1^2}\right) + n_2 \exp\left(-\frac{p^2}{2\sigma_2^2}\right), \tag{8.1}$$

where $n_{1,2}$ are the weights of the two Gaussians and $\sigma_{1,2}$ are the widths. The double Gaussian distribution is most likely due to the increased efficiency of molasses cooling at the centre of the MOT, resulting in cooler atoms in the centre and hotter atoms towards the edge of the MOT [82].

The following is an outline of the stages of a simulation of a Rabi flopping experiment.

Define the time evolution function: this is a function which takes as inputs:

Symbol	Parameter	Value
$I_{1,2}$	Raman beam intensities	$2\,\mathrm{kWm}^{-2}$
Δ	Detuning from single-photon resonance	$10\,\mathrm{GHz}$
m	Rubidium-85 atomic mass	$85 \times 1.66 \times 10^{-27}\,\mathrm{kg}$
$\omega_{1,2}$	Raman laser frequencies	$2\pi(384.230406373\,\mathrm{THz} - 1.264888\,\mathrm{GHz})$, $2\pi(384.230406373\,\mathrm{THz} + 1.770843\,\mathrm{GHz})$
$p_{width1,2}$	Width of momentum distributions	$4\,\hbar k_{eff}$, $16\,\hbar k_{eff}$
t_π	π pulse duration	$1.2\,\mu\mathrm{s}$

Table 8.1: Simulation parameters and typical values.

- the initial ground and excited state amplitudes
- the start and end times of the period of evolution
- the phase of the Raman beams
- the detuning from two photon resonance,

and calculates the resulting ground and excited state amplitudes.

Define the parameters of the simulations: this includes the single-photon detuning, the intensities of the Raman beams, and the momentum distribution (see below and Table 8.1).

Calculate the state amplitudes: this is done by the time evolution function, for each combination of pulse length, momentum, and m_F level.

Sum the state amplitudes: for each pulse length, the amplitudes are averaged and normalised first across the momentum values, and then across the m_F values. For each momentum value, the population is weighted by the momentum distribution function.

Table 8.1 shows the parameters of the simulations. The detuning, intensities, and momentum distribution widths are free parameters[1], the Rb mass and Raman laser frequencies are fixed, and the π pulse duration is determined from a Rabi flopping experiment or simulation. All other values used in the calculation of the probability amplitudes are derived from these parameters.

[1] In simulations to fit experimental results, these parameters are 'free' in the sense that they are allowed to be adjusted within a small range of the measured (I, Δ) or estimated (p_{width}) values.

The simulation of pulse sequences which can be described by a single propagator is relatively fast on a standard desktop computer. This is because the state amplitudes at the end of the sequence may be calculated in one time evolution step. The total time to run the simulation is broken up into initialisation time (defining functions, creating the arrays to hold the state amplitudes) and calculation time (calling the function and writing data to the arrays). Typically there are 100 momentum values[2] and 5 m_F values; on the computer used for these simulations (Intel Core i7 CPU @ 2.2 GHz) the initialisation time is $0.54\text{ms} \pm 0.03\text{ms}$ and the calculation time is $0.13\text{ms} \pm 0.02\text{ms}$. The total time to simulate one propagator and average over the amplitudes is $37\text{ms} \pm 2\text{ms}$.

8.3.1 Simulating sequences with many propagators: IRIDIS compute cluster

In situations where the sequence propagator is the product of a series of different propagators, for example where the intensity and/or phase of the pulse varies over the sequence, the amplitudes must be calculated for each time period for which the propagator describes the system, and the amplitudes at the end of such a period are the starting amplitudes for the subsequent period. The computing time scales approximately linearly with the number of different propagators and this quickly adds up when, for example, the intensity is ramping up or down to account for the finite rise time of the AOM.

Consider the simulation of the temporal evolution of the state amplitudes during composite pulse sequence when the AOM rise time is accounted for. There may be typically 50 pulse lengths, for each of which the pulse sequence is simulated over the momentum and m_F values (although momentum is a continuous variable, to numerically simulate the momentum distribution we must discretise it). In this case the total simulation time is approximately 40 minutes. This is not equal to the time per propagator multiplied by the number of propagators because simulating the intensity during switch on/off of the AOM introduces some additional calculations which are discussed in more detail in Section 8.5.3. This makes fitting to experimental data a laborious process as the entire simulation must be run many times with different parameters.

These computations are significantly sped up when the calculations for each velocity class can be distributed among many computing cores, which compute in

[2] In the interferometric cooling simulations (Chapter 9), this number is much greater.

parallel, and the results combined at the end. This is possible because interactions between atoms are rare enough to be negligible and so their state amplitudes may be calculated independently. For simulations of sequences with many propagators, the IRIDIS Compute Cluster is used. IRIDIS is the University of Southampton's high performance computing facility[3] and was designed with parallel computing tasks in mind.

IRIDIS has 750 available compute nodes, each having 16 CPUs (2.6 GHz Intel Sandybridge processors[4]). Jobs are submitted by requesting an appropriate number of compute nodes and processors per node. When simulating the time evolution of pulses with many propagators, one CPU is requested per momentum value; the amplitudes are calculated and averaged for each m_F value for this momentum value. When simulating spectral scans, the calculations are divided up by their detuning value, and each CPU calculates the amplitude after the whole pulse sequence, averaged over all momentum and m_F values.

8.4 Simulations to fit experimental results

As we wish to use simulations to test the efficacy of pulse sequences in our experiment, it is important to first establish the extent to which the simulations agree with experimental data.

8.4.1 Rabi flopping

Composite pulse sequences are defined in terms of angles rotated through on the Bloch sphere, which depend on the Rabi frequency Ω and the duration of the pulse t: $\theta = \Omega t$. The Rabi frequency must be measured before doing any composite pulses so that the correct pulse durations may be used. This is achieved by applying pulses of varying duration and measuring the excited state population.

Considering equation 2.30, when on resonance ($\delta_{12} = \delta^{AC}$, $\cos\Theta = 0$, $\sin\Theta = 1$) with the initial conditions $t_0 = 0$, $c_e(t_0) = 0$, $c_g(t_0) = 1$ and $\phi = 0$, the excited state population varies as

[3] IRIDIS was ranked 261 in the TOP500 list of supercomputers [83] in November 2014.

[4] As the IRIDIS CPUs have a higher clock speed, the single propagator calculation time is decreased; however, jobs submitted to IRIDIS are queued depending on the resources requested and available, and so the speed increase for simple simulations such as those involving a single propagator only is not particularly useful. We exploit the parallel processing functionality, not the greater clock speed.

$$|c_e(\tau)c_e^*(\tau)| = \sin^2\left(\frac{\Omega_{eff}\tau}{2}\right). \tag{8.2}$$

However, our MOT contains atoms with a distribution of momenta and so not all atoms will be on resonance with the applied laser beams. When off-resonance, the population oscillates as

$$|c_e(\tau)c_e^*(\tau)| = \left(\frac{\Omega_{eff}}{\Omega_r}\right)^2 \sin^2\left(\frac{\Omega_r\tau}{2}\right). \tag{8.3}$$

As $\Omega_r > \Omega_{eff}$, the amplitude of the off-resonance oscillations is always less than one, and the frequency of the oscillations $\Omega_r = \sqrt{{\Omega_{eff}}^2 + (\delta_{12} - \delta_{AC})^2}$ is always greater than the on-resonance Rabi frequency Ω_{eff}. The detuning from resonance depends on the momentum of the atoms and so, as the MOT in our experiment has a distribution of momenta, we expect that the contrast of the Rabi flopping will decay with time.

This is indeed what happens, as shown in Figure 8.2, where the contrast decreases from 0.56 to 0.25 after one period of oscillation. The simulated fit to the data uses parameters which are within the error range of the measured values: the Raman beam intensities are $I_1 = 12\text{kW/m}^2$, $I_2 = 17\text{kW/m}^2$, the single photon detuning is 15GHz, and the two-photon detuning is 70kHz. Fitting was done by eye, by varying the values of the parameters within reasonable limits and comparing the resulting simulated data to the experimental data. The momentum distribution widths are $3\hbar k_{eff}$ and $10\hbar k_{eff}$, with weights 0.75 and 0.25.

In the following section, we quantify the sources of the dephasing by investigating through simulation the behaviour of various subgroups of atoms.

8.4.2 Inhomogeneities contributing to dephasing

As introduced in Chapter 4, we can identify errors as one of two kinds: off-resonance errors and pulse length errors. In our experiment, the source of off-resonance errors is a combination of the distribution of momenta and the AC Stark shift: the detuning from two-photon resonance is $\delta_{12} - \delta_{AC}$, where $\delta_{12} = \delta_L - pk_{eff}/m - \hbar{k_{eff}}^2/2m$ and $\delta_{AC} = \Omega_e^{AC} - \Omega_g^{AC}$. δ_{12} depends solely on the momentum of the atoms, whereas δ_{AC} depends on the intensity of the Raman beams (which may vary across the atom cloud) and the strength of the transition

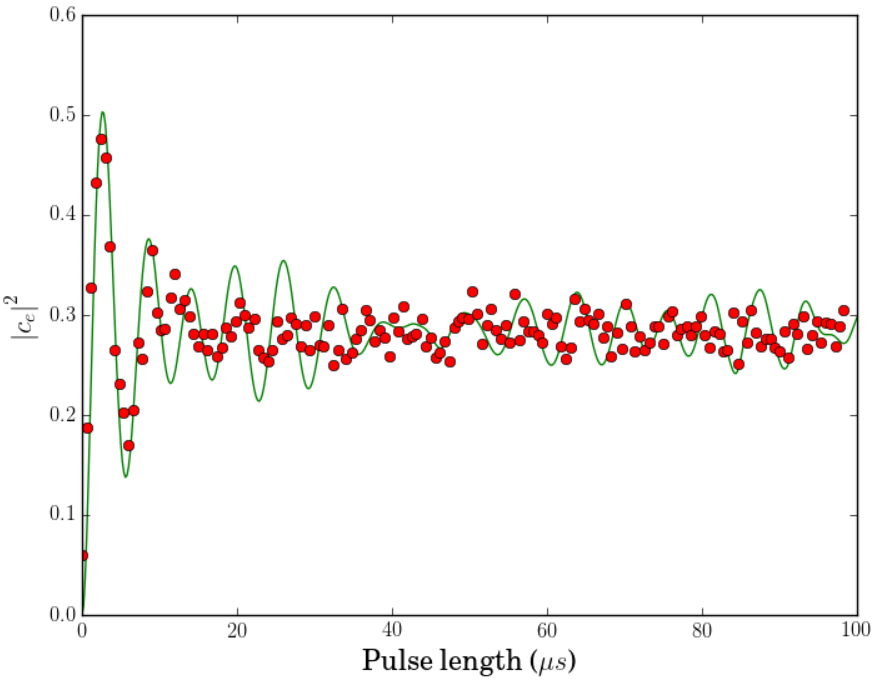

Figure 8.2: Time evolution of the excited state population during a Rabi flopping experiment. The Raman beams are counter propagating with σ^+ polarisations. Red circles: experimental data, green line: simulation. Experimental data taken by Alex Dunning.

concerned (see Section 4.5). Pulse length errors arise from the differing Rabi frequencies of the transition routes and spatial inhomogeneity of the Raman beam intensity.

8.4.2.1 Off-resonance errors

Although the shim coil currents are chosen such that the 5 m_F levels are Zeeman degenerate, they still each have different transition strengths through the Clebsch-Gordan coefficients for the particular transition. This manifests itself in a different Rabi frequency and a different AC Stark shift for each level. δ_{AC} is a constant shift across all momentum values for a particular m_F value, whereas δ_{12} varies with momentum. The detuning from resonance as a function of momentum for each m_F level is plotted in Figure 8.3, for experimental parameters identical to those used in the Rabi flopping simulation; the plot is weighted by the momentum distribution $n(p)$, which is why the detuning tends to zero for large values of p. We multiply by $n(p)$ because although the momentum-dependent detuning is large at

outlying p values, very few atoms occupy these momentum states and so they do not contribute as much to the dephasing as the atoms closer to zero momentum.

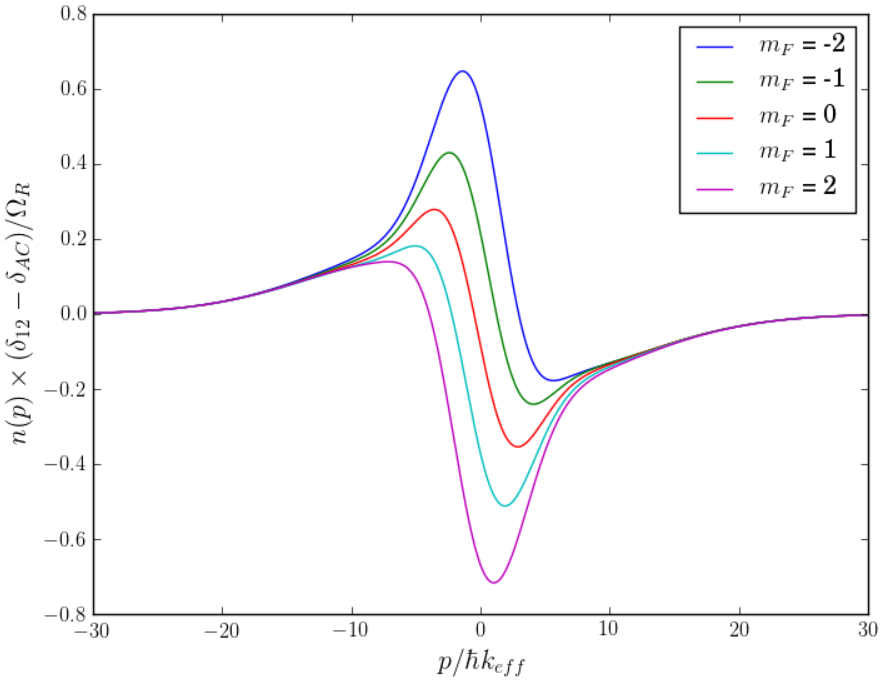

Figure 8.3: Fractional off-resonance error for each m_F level as a function of atomic momentum, for counterpropagating σ^+ polarised Raman beams. Parameters are as in Figure 8.2.

Figure 8.3 shows that for a particular value of momentum, whether an atom is on resonance or not is dependent on its m_F state. For these parameters the $m_F = 0$ route has the smallest range of detuning values. This is because the value of the laser detuning δ_L, which features in δ_{12}, is set to the frequency where the transfer of population to the excited state is maximal (as determined from a Raman scan) and, because the differences in consecutive Stark shifts are equal and opposite on either side of $m_F = 0$,[5] the detuning where the transfer is maximal is always going to be one where the Stark shift for $m_F = 0$ is minimised.

8.4.2.2 Intensity dependence of off-resonance errors

δ_{12} is independent of intensity, but δ_{AC} is not. As the intensity of the Raman beams increases, δ_{AC} begins to dominate the detuning and the fractional off-resonance error begins to look like a constant multiplied by the momentum distribution, as shown in Figure 8.4.

[5] i.e. $\delta_{AC}(m_F = 2) - \delta_{AC}(m_F = 0) = -(\delta_{AC}(m_F = -2) - \delta_{AC}(m_F = 0))$

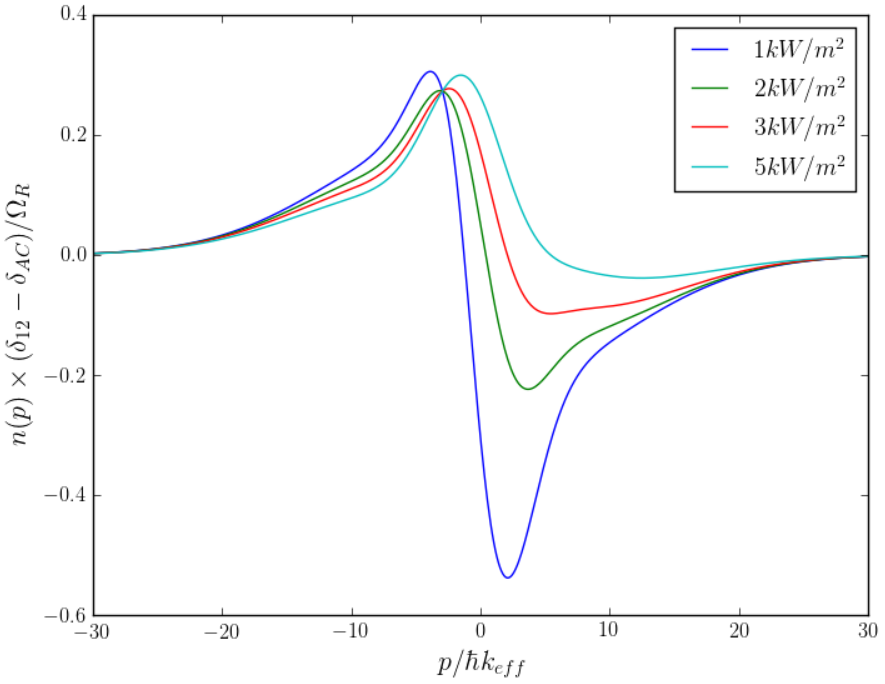

Figure 8.4: Fractional off-resonance error for $m_F = 0$, for several values of Raman beam intensities. Raman beams are counterpropagating with polarisations σ^+.

We can calculate the average off-resonance error f_{av} over the distribution of momentum and transition routes (where N_p is the number of momentum values and N_m is the number of transition routes):

$$f_{av} = \frac{1}{N_p N_m} \sum_p \sum_{m_f} n(p) \times \frac{|\delta_{12}(p) - \delta_{AC}(m_F)|}{\Omega_R}. \tag{8.4}$$

As the intensity increases, the fractional off-resonance error for each transition route changes (the average decreasing for $m_F = 0$ and increasing for other m_F values) but the average over the transition routes stays essentially constant.

8.4.2.3 Pulse length errors

There are two primary sources of pulse length error: the spatially varying intensity of the Raman beams, and the different Rabi frequencies for each transition route. Pulse length errors contribute to the dephasing because the pulse length required to transfer all atoms to the excited state depends on the Rabi frequency of a particular atom and thus on the intensity which it experiences.

First we address the intensity variations. To isolate this effect from the other inhomogeneities we consider a single atom which is on resonance with the Raman beams ($\delta_{12} - \delta_{AC} = 0$). The resonant two-photon Rabi frequency is proportional to the Raman beam intensities $I_{1,2}$:

$$\Omega_{res} \propto \sqrt{I_1 I_2}. \tag{8.5}$$

The fractional intensity deviations $\epsilon_{1,2}$ in each beam are not equal at each point in space and so, neglecting error terms higher than first order, the modified Rabi frequency is

$$\Omega'_{res} \propto \sqrt{I_1 I_2 (1 + \epsilon_1 + \epsilon_2)} \tag{8.6}$$

The maximum deviation of the Raman beam intensity from its average value was measured to be approximately $\pm 5\%$ [45]; according to the above equation, the maximum Rabi frequency would be $\Omega_{res} \times \sqrt{1.1}$ i.e. approximately $1.05\Omega_{res}$, or a fractional error of 0.05. Comparing this to Figure 8.3 where the fractional errors exceed 0.6, we see that pulse length errors due to spatially nonuniform intensity are much smaller.

Pulse length errors due to the different transition routes are slightly easier to quantify. The relative transition strengths, as a fraction of the $m_F = 0 - m'_F = 0$ strength, are 0.75 for $m_F = \pm 2$ and 0.94 for $m_F = \pm 1$ for all Raman beam polarisations considered in this thesis (both σ^+, both σ^-, or $\pi^+ - \pi^-$). This translates directly into fractional pulse length errors of 0.25 and 0.06; the maximum fractional pulse length error here is smaller than the maximum fractional off-resonance error.

We finally consider the second-order pulse length errors arising from off-resonance errors, as described in Section 4.4. According to equation 4.9, a fractional off-resonance error of f causes the desired angle of rotation θ to be replaced with $\theta\sqrt{1+f^2}$. This results in (for small f) a fractional pulse length error of

$$\begin{aligned} \epsilon &= \sqrt{1+f^2} - 1 \\ &= \frac{f^2}{2} - \frac{f^4}{8} + \mathcal{O}(f^6). \end{aligned} \tag{8.7}$$

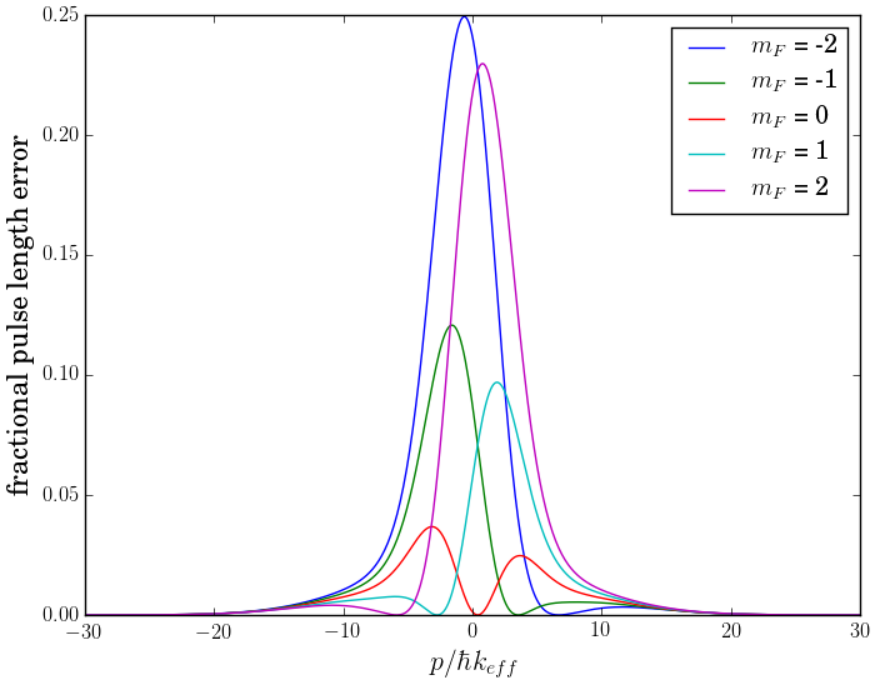

Figure 8.5: Fractional pulse length error arising from off-resonance error. Parameters are as in Figure 8.2.

This is plotted as a function of momentum, for the 5 transition routes, in Figure 8.5. We can see that for the $m_F = 0$ transition the maximum fractional pulse length error is approximately 0.04, which is of the order of the errors due to the intensity variations, but the non-zero m_F routes experience larger errors which are comparable in size with the errors due to the differing coupling strengths of the transitions.

We see that no source of pulse length error results in a maximum fractional error which is greater than that of the off-resonance errors. We note also that the pulse length errors arising from Doppler broadening are larger than those arising from the different transition routes or spatially nonuniform intensity of the Raman beams. We conclude that in our experiment the primary source of both off-resonance and pulse length errors is the Doppler broadening arising from the range of atomic momenta.

8.4.3 Rotary echoes

In the previous section, we considered the errors present in our system and how they contribute to dephasing. However, the Rabi flopping experiment offers no

insight into whether the rapid decay in contrast is due to dephasing (due to the range of pulse lengths and resonance frequencies), which is a reversible process, or decoherence (due to, for example, randomisation of phase upon collision between atoms), which is an irreversible process. One way to distinguish between the two mechanisms is to apply a pulse sequence to bring the atoms back into phase with each other. If dephasing is occurring without decoherence, the phase information is still present; the decay in the contrast is simply due to the addition of coherent oscillations in state populations with different frequencies, and it may be possible to recover the contrast. However, if decoherence is present, the atoms' phase information is lost as soon as the spontaneous emission takes place, and we expect no recovery of contrast.

One sequence which may be used to test this is the rotary echo: repeated application of the pulse 360_0360_π. First proposed and investigated in 1959 [84], it has been demonstrated in many systems including neutral atoms trapped in an optical lattice [85] and neutral atom clouds [24]. Assuming uncorrelated errors, the rotary echo compensates pulse length errors perfectly (the point to point fidelity $\mathcal{F}_{PP}$, equation 4.3, contains no terms depending on the pulse length error), and off-resonance errors to 10th order. The state evolution during a rotary echo sequence is shown in Figure 8.6, where the blue states/vectors are error-free and the red and green states/vectors exhibit pulse length and off-resonance errors. The deviation from the target state $|0\rangle$ after the second period of evolution (8.6(b)) is much smaller than after the first period of evolution (8.6(b)).

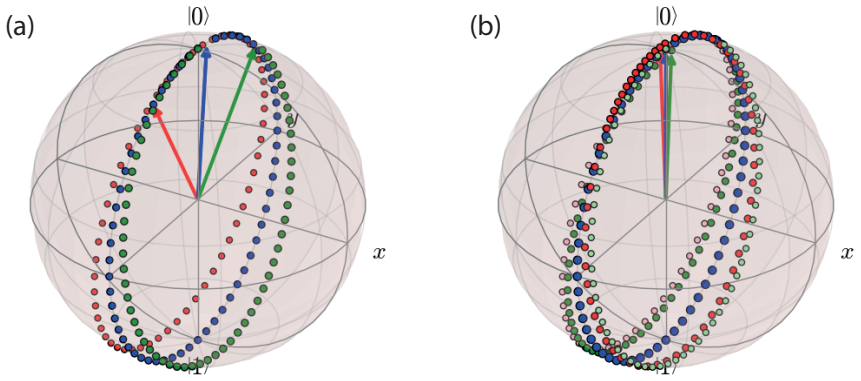

Figure 8.6: Bloch sphere trajectories during a rotary echo sequence in the presence of pulse length and off-resonance errors. Blue points and blue vectors are error free; green and red points and vectors have pulse length and off-resonance errors. (a): After the 360_0 pulse. (b): after the 360_0360_π pulse. States during the first period of evolution are shown in light red/green points, states during the second period of evolution are shown in dark red/green points. The blue points during the first period of evolution are not visible because the error-free state vector exactly retraces its previous trajectory during the second period of evolution.

Figure 8.7 shows rotary echoes in our experiment. The data were taken at the same time as the Rabi flopping data in Figure 8.2, and the simulations used to fit to the data have identical parameters. The data points from the experiment are plotted in red, with the simulation in green. There is a small decay in the contrast of the rotary echoes, which may be explained by incorporating the effects of phase noise in the IQ modulator which applies phase shifts to the Raman beams; we find a good fit with a maximum phase noise of $\pm\pi/6$ (the value for each pulse being a randomly generated value within this range). However, this value is greater than that specified by the manufacturer ($\pi/25$). It is possible that there are other sources of phase noise in the system which add to this specified value, for example variations in the path lengths of the Raman beams. Another possible cause of this decay in contrast is decoherence, which we investigate in the following section.

8.4.3.1 Decoherence

Decoherence is the process by which a superposition becomes a statistical mixture. We recall the density matrix description of a two level system where, with $|\psi\rangle = a\,|a\rangle + b\,|b\rangle$, the density matrix is

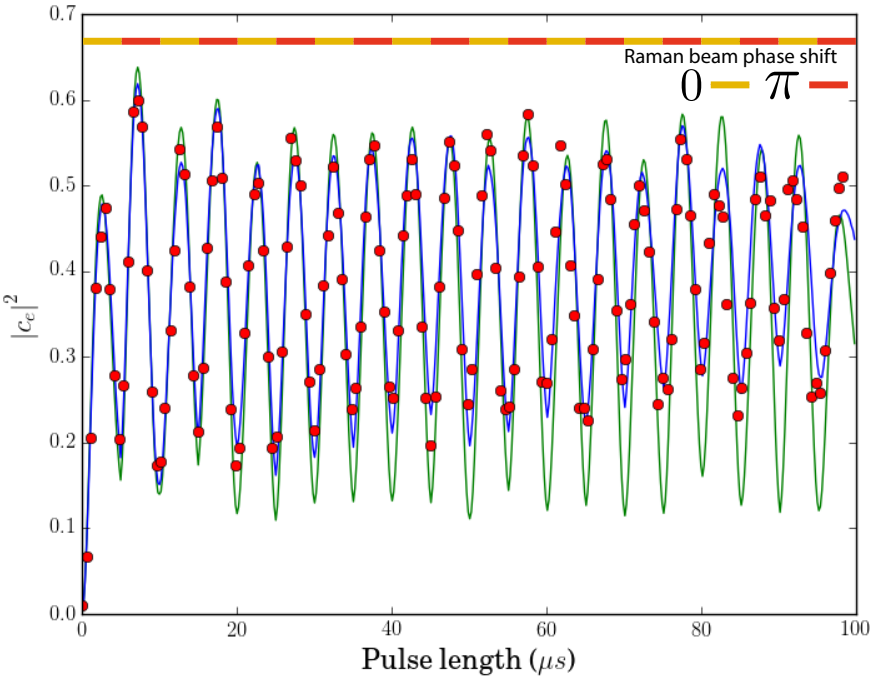

Figure 8.7: Time evolution of the excited state population during a rotary echo experiment. Red circles: experimental data, green line: simulation without phase noise, blue line: simulation with phase noise randomly generated for each pulse in the range $\pm\pi/6$. The phase of the Raman beams during the experiment is shown at the top of the plot.

$$\rho = |\psi\rangle\langle\psi| = \begin{pmatrix} aa^* & ab^* \\ ba^* & bb^* \end{pmatrix}. \tag{8.8}$$

The diagonal elements are populations and the off-diagonal elements are coherences. The density matrix describing a statistical mixture has non-zero elements on the diagonal only.

To demonstrate the potential effects of decoherence in our system, we can show the loss of contrast in Rabi flopping by simulating the evolution of an ensemble of atoms with the same parameters (same momentum value and m_F), close to resonance with the Raman beams. This is not a physically realistic situation as it would not be possible to produce such a momentum distribution in our experiment, but it is sufficient to demonstrate decoherence in isolation from dephasing.

We assume for simplicity that following a decoherence event, the phase of the atom's wavefunction is randomised but the excited and ground state populations are unchanged. We make this assumption because the most likely decoherence

event is a collision between atoms, and (on average) the atoms at their post-molasses cooling stage do not have enough kinetic energy to cause a change in internal state. The probability to experience a decoherence event after time t is modelled as an exponential: $p_{dec}(t) = 1 - \exp(-t/\tau)$, with $\tau = 2.5\mu$s (chosen to give a similar decay of contrast as in the experimental data). Figure 8.8 (a) shows the results of a Monte-Carlo simulation of Rabi flopping with and without decoherence mechanisms; the decoherence-free oscillations (blue) continue undamped indefinitely, and when decoherence is present the contrast of the Rabi flopping decays rapidly and the oscillations are indistinguishable from noise after two periods.

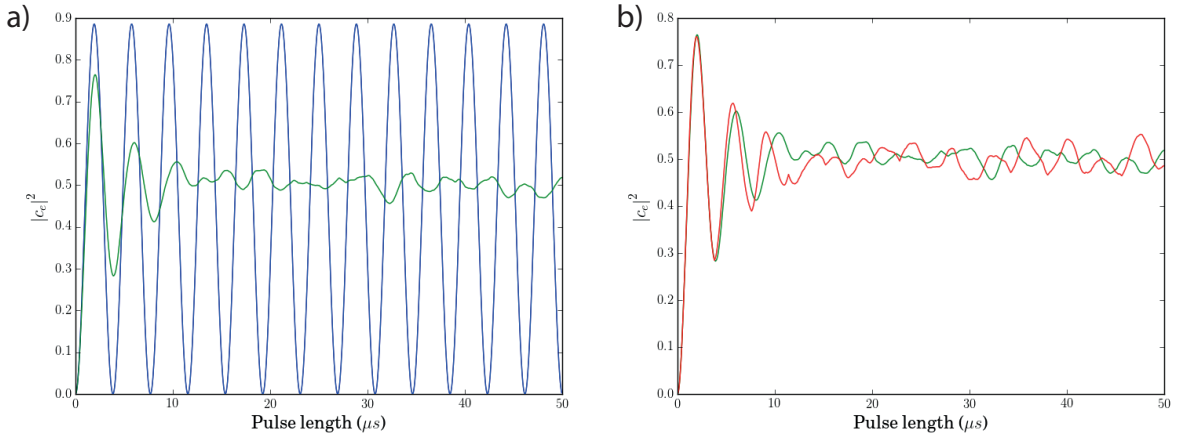

Figure 8.8: The effect of decoherence mechanisms on Rabi flopping and rotary echoes. a) Rabi flopping without (blue) and with (green) decoherence effects. b) Rabi flopping with decoherence (green, as in (a)) and rotary echoes with decoherence (red).

We can test the behaviour of rotary echoes in a system with a short coherence lifetime by modifying the Rabi flopping simulation. Figure 8.8 b) shows normal Rabi flopping (green), where the phase of the Raman beams is unchanged throughout the experiment, and rotary echoes (red), where the phase is switched by π every 3.8μs. The contrast under rotary echoes is not recovered and decays at approximately the same rate as the Rabi flopping.

If we now consider the decay of contrast in the rotary echoes, we can incorporate decoherence into our simulation to fit the data in Figure 8.7. If we assume the level of phase noise present in the Raman beams is equal to that specified by the manufacturer of the IQ modulator, we find a good fit to the data by assuming that decoherence events may happen with a lifetime of 200μs (much longer than that used in the simulations in Figure 8.8 b)). This is plotted in Figure 8.9.

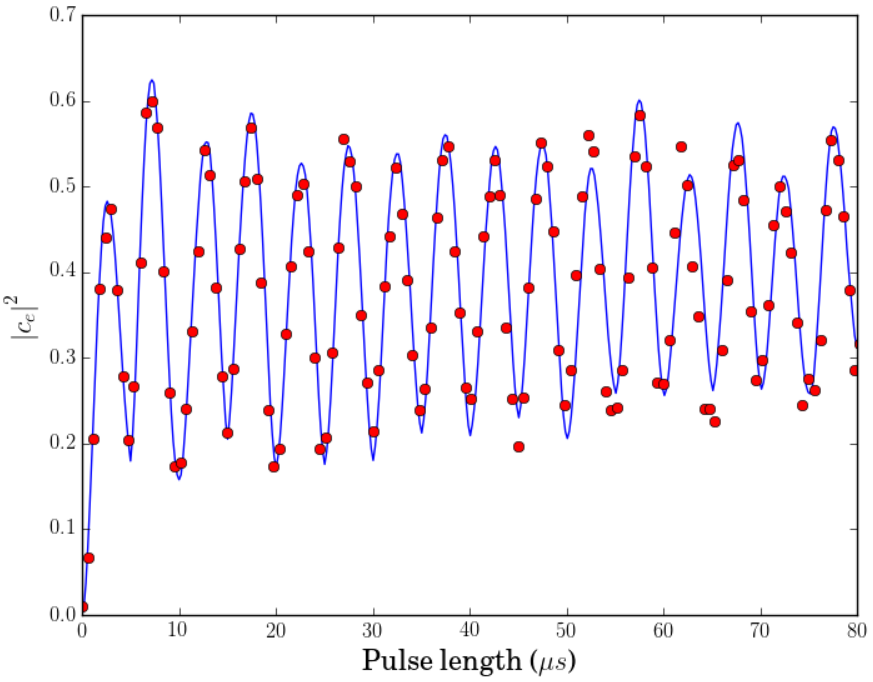

Figure 8.9: Time evolution of the excited state population during a rotary echo experiment. Red circles: experimental data, blue line: simulation. The simulation assumes phase noise is present in the Raman beams and decoherence events occur with a lifetime of 200μs.

Although the simulation of Figure 8.9 is remarkably consistent with the experimental data, the coherence lifetime which we must assume in order to achieve it is unrealistic experimentally. The single photon scattering rate, given the value of detuning Δ used to fit to the data, is approximately 0.15Hz, giving a coherence lifetime of $\sim$ 7s. Given that the detuning and intensity of the Raman beams are measured before performing experiments and that the data are well fitted assuming these values, we can be confident in this calculation of the coherence lifetime, and thus assume that the decay in the contrast of the rotary echo oscillations is not due to decoherence caused by single photon scattering. Previous experiments on this MOT have shown [67] that collisions between the atoms are negligible on the timescale of the experiment; [86] predicts a coherence lifetime (given the parameters of our experiment) of $\sim$ 100s, which is longer than that used for the fit in 8.9), and thus collisional dephasing is not contributing to the decay in contrast.

It is also worth noting that an assumption of random intensity fluctuations of $\pm$5%, which affect the single photon Rabi frequencies and thus the detuning Δ and the Raman Rabi frequency, is not sufficient alone to account for the decay in contrast of the rotary echo oscillations. Also, a similar magnitude of fluctuations

in the detuning Δ (which is reasonable to assume given that the laser from which the Raman beams are derived is not stabilised) does not account for the decay in contrast.

We conclude that in our system, decoherence is not the primary cause of the decay of contrast in Rabi flopping, and that the coherence time of the atoms is more than sufficient to carry out atom interferometry experiments. Although some irreversible dephasing is present, it does not cause a significant decay in the contrast of the rotary echoes; it is most likely due to phase noise due to path length changes between the Raman beams over the timescale of the experiment. It is possible to fit a model incorporating collisional dephasing to the data to explain the decay in contrast of rotary echoes, but the coherence lifetime that we are required to assume is unreasonable.

8.4.4 Composite pulses: spectral scans

The rotary echo experiment confirmed that dephasing of Rabi flopping is reversible and the atoms can be brought back into phase with each other by applying pulses with varying phase. Although it is an insightful diagnostic tool, this particular sequence does not perform a rotation on the Bloch sphere that is useful in atom interferometry. For interferometry, we require high fidelity beamsplitter and mirror pulses, and for this purpose we look to composite pulses (as introduced in Chapter 4). In this Section we investigate the experimental performance of selected composite pulses and show that the data can be well modelled by our simulations.

8.4.4.1 Assessing the performance of composite pulses

We can quantify the performance of a composite pulse relative to a π pulse in several ways. One is to divide the peak fidelity of the composite pulse by that of the π pulse, to find the fractional increase in fidelity. As this is at zero detuning, we expect that this gives a good measure of the tolerance to pulse length errors. A pulse having a spectral scan with a high peak fidelity and small width is not effective over a large range of off resonance errors, and so we also measure the FWHM of the fitted curve to determine the tolerance to off-resonance errors.

We can also measure the area under the spectral scan and divide by the area under a spectral scan where all of the atoms are in the excited state; this gives a measure of performance over the entire ensemble.

Spectral scans of the experimentally investigated composite pulses, along with the parameters used in the simulations, are shown in Figure 8.10. Simulations were used to find a fit to the π pulse using the parameters listed in Figure 8.10 f), and these parameters were subsequently used to plot the spectral scans for the composite pulses. As all of the data were taken in the same experimental run, they are fitted with a common set of parameters. The polarisation of the Raman beams in these experiments was $\sigma^+ - \sigma^+$. In this experiment, the m_F sublevels were not fully degenerate, as this would require compensation magnetic fields which shift the magnetic field zero point of the MOT out of the paths of the MOT beams. As a consequence, the off-resonance errors are larger than they would be in the optimal configuration.

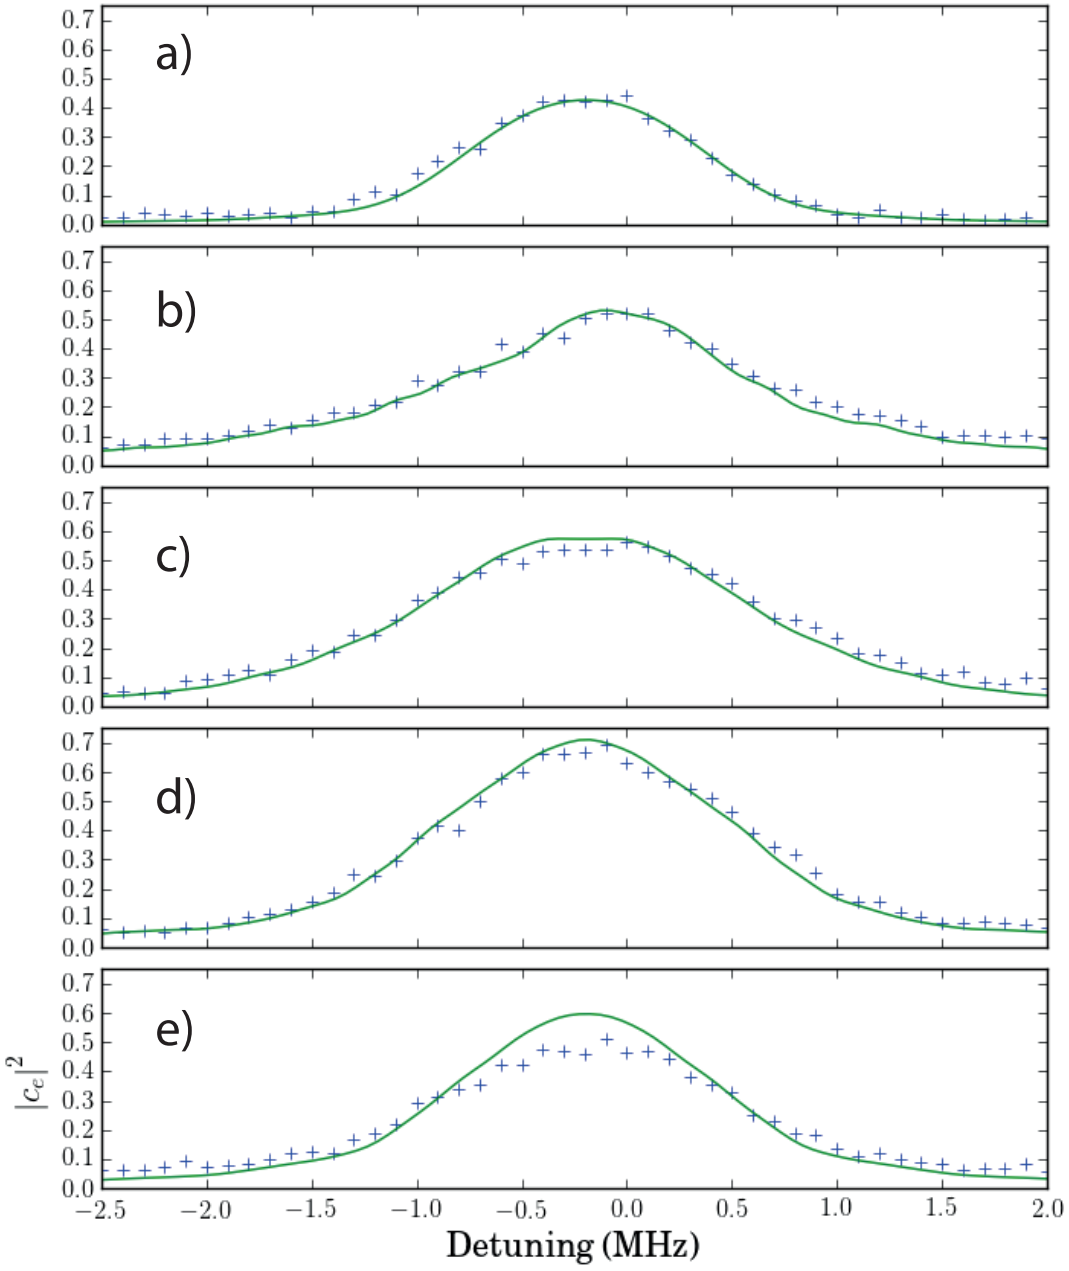

f)

I_1	14kW/m^2	n_1	2	σ_1	$3\hbar k_{eff}$
I_2	21kW/m^2	n_2	1	σ_2	$9\hbar k_{eff}$
Δ	$2\pi \times 9.2$GHz	t_π	1.4μs		

Figure 8.10: Spectral scans of composite pulses. Blue cross: experiment, green lines: simulation. a) π pulse; b) BB1; c) CORPSE; d): WALTZ; e): Knill; f): parameters used in simulations.

Table 8.2 lists the pulse sequences along with their fidelities and widths. In the

following sections we discuss features of each pulse and its performance in our system.

Pulse name	Pulse duration	$\mathcal{F}_P$	$\mathcal{F}_P/\mathcal{F}_\pi$	FWHM ratio	Area ratio
Pi pulse	τ_π	0.42	1	1	0.13
BB1	$5\tau_\pi$	0.53	1.26	1.26	0.22
CORPSE	$13\tau_\pi/3$	0.57	1.35	1.44	0.25
WALTZ	$3\tau_\pi$	0.71	1.69	1.27	0.27
Knill	$5\tau_\pi$	0.59	1.40	1.14	0.21

Table 8.2: Characterisation of simulations to fit to composite pulse data as shown in Figure 8.10. For each pulse, we list the duration in terms of the π pulse duration τ_π, peak fidelity $\mathcal{F}_P$, the ratio of the peak fidelity of the pulse to the peak fidelity of a π pulse $\mathcal{F}_\pi$, the ratio of the full width at half maximum (FWHM) to that of the π pulse, and the ratio of the area under the spectral scan of the pulse to the area under the spectral scan of a perfect inversion pulse.

8.4.4.2 BB1

The BB1 pulse sequence is $180_{104.5}360_{313.4}180_{104.5}180_0$ and it is is a general rotor designed to compensate for pulse length errors[63] ; its point to point fidelity is 6th order in ϵ, and its tolerance to off-resonance errors is not decreased from that of a π pulse. Given that we established previously that off-resonance errors are dominant in our experiment, we do not expect this pulse to perform as well as those designed to correct off-resonance errors. However, we do see an increase in both the peak fidelity and the width compared to the π pulse.

8.4.4.3 CORPSE

The CORPSE (COmpensation for Off-Resonance errors with a Pulse SEquence) pulse [87], [88] is a general rotor; it is a generalisation of a previous pulse designed to perform a $\pi/2$ rotation [89]. It consists of three pulses of duration $\pi/3$, $5\pi/3$ and $7\pi/3$, with phases set by the desired angle of rotation. For an inversion pulse, the sequence is $60_0300_{180}420_0$. Its tolerance to pulse length errors is identical to that of a plain π pulse, and it eliminates off-resonance errors up to the 4th order, its point to point fidelity being $1-6.5\times10^{-3}f^4+\mathcal{O}(f^5)$. CORPSE outperforms BB1 in both peak fidelity and FWHM relative to the π pulse. Being a general rotor, we might expect that it would perform better than point-to-point off-resonance

correcting pulses in an atom interferometer; this is discussed further in Section 8.6.

8.4.4.4 WALTZ

The WALTZ pulse is named for its "1-2-3" structure: $90_0180_{180}270_0$. It was first investigated in 1983 [65] and is a point-to-point pulse which aims to correct off-resonance errors. As described in Section 4.4, the fidelity of this pulse contains 2nd order terms in both pulse length ϵ and off-resonance error f; its error correction depends on cancellation of higher order terms in f. WALTZ exhibits the greatest increase in peak fidelity of all the experimentally tested pulses; it also shows the greatest increase in total fraction of atoms transferred (final column in Table 8.2).

8.4.4.5 Knill

The Knill pulse [90] is a time-symmetric sequence of 5 π pulses: $180_{240}180_{210}180_{300}180_{210}180_{240}$. It is designed to correct both types of error: its point-to-point fidelities are, for pulse length errors $1-(21+12\sqrt{3})\pi^6\epsilon^6/64+\mathcal{O}(\epsilon^7)$ and for off-resonance errors $1-(21-12\sqrt{3})f^6+\mathcal{O}(f^7)$. It has the worst increase in FWHM and in total fraction of atoms transferred; additionally, the fit to the data is worse than that of the other pulses around the peak of the scan. This could be due to the IQ modulator having difficulty with synthesising the phases required for this sequence; previous characterisation has shown that it reproduces 0 and π phase shifts more accurately than angles in between. However, if we incorporate the measured deviation from the desired phase shifts into the simulation, there is no great improvement in the fit to the data. Given that the fits to the other pulse sequences are good, we are not overly concerned by the discrepancy in this case, and (as shown in the following section) we can fit well to the temporal scans of these pulses.

8.4.5 Composite pulses: temporal scans

The temporal scan of a composite pulse is built up by terminating the composite pulse at various times during the sequence and measuring the excited state population. Before presenting the experimental data, we describe a modification to the simulations which takes into account the finite rise time of the AOM which switches the Raman beams and its effect on the intensity-time dependence of the Raman beams.

The Raman beams which perform the coherent manipulation are passed thorough an AOM to switch them on and off (Section 6.2). The rise time of the AOM (the time between receiving the switching signal and achieving maximum diffraction into the first order) is non-zero, because it takes a finite time for the sound wave within the AOM crystal to traverse the input beam. During this time, the amount of light diffracted into the first order is submaximal. This means that the pulses we are applying to the atoms do not have an intensity which is constant in time, and so the evolution of the state amplitudes is different to the case where they experience a constant intensity.

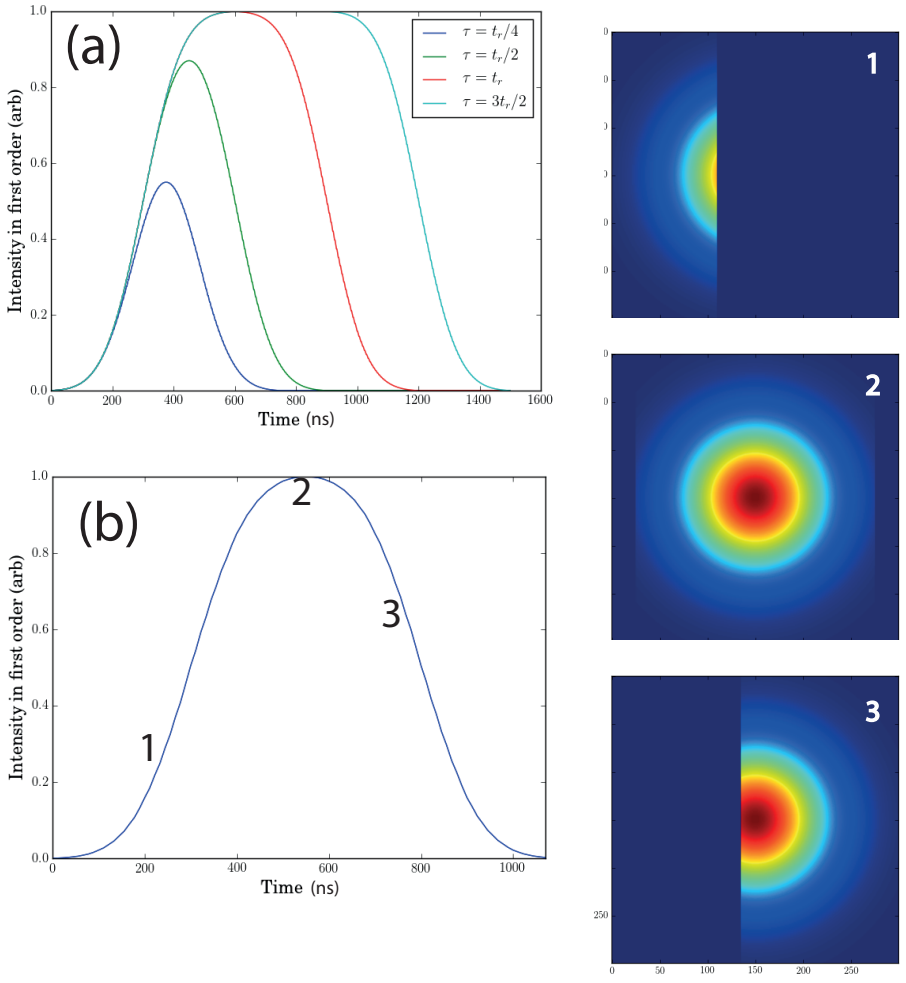

Figure 8.11: AOM diffracted intensity versus time. a) Intensity-time functions plotted for several different pulse lengths. b) Intensity-time function for a pulse length of 500 ns, with numbered images showing the portion of the Gaussian beam covered by the sound wave at different points during the pulse.

We can model the AOM sound wave as a moving window function and the Raman beams as Gaussians (both in two dimensions). The intensity in the first order

at any time is equal to the convolution of the two functions, as shown in Figure 8.11 b); the numbered images show the amount of the beam which is covered by the sound wave at several points during the pulse. The duration of the rise time depends on the speed of the sound wave and the width of the Raman beam, and in this figure the rise time is 600ns. Figure 8.11 a) shows the intensity as a function of time for several pulse lengths. If a pulse is applied which is shorter than the rise time of the device, the intensity in the first order never reaches the maximum value, because the width of the travelling wave is less than the width of the Gaussian.

When including the effect of the finite rise time of the AOM in the simulations, the time to simulate a single pulse is greatly increased compared to the time to simulate a pulse with constant intensity in time. Instead of a single propagator which determines the evolution for the duration of the pulse, we instead evolve under a piecewise propagator where the pulse is divided up into time steps during which the intensity may be assumed to be constant.

This is particularly computationally intensive when simulating temporal scans. When simulating a pulse with intensity constant in time, the state amplitudes at any point in time may be found by evolving the state amplitudes at a previous point in time, under one propagator. However, when including the effects of the AOM rise time, the time-varying intensity of a pulse depends on its duration, and thus the amplitude at a point in time (when the pulse length is shorter than the rise time) is not equal to the amplitude at a previous point in time evolved under a propagator applied for the duration of the difference in the two times. This means that the amplitudes at any point in time which is less than the AOM rise time must be calculated by evolving the amplitudes from the start time of the entire pulse sequence. At times greater than the rise time of the AOM, the intensity can be assumed to be constant after the rise time; however, there is also a time-varying intensity after the switch-off signal is sent to the AOM, and this must also be simulated. These simulations are performed on IRIDIS (see Section 8.3.1).

The temporal scans for the composite pulses investigated in the previous section are shown in Figure 8.12. For each plot, the experimental data is shown in blue and the fits to the data are shown with (red) and without (green, dashed) the effect of the AOM rise times included. The rise time of the AOM is 380ns, found by fitting the model described above to the 1st order intensity as a function of time, as measured by a fast photodiode during switching. The parameters used in

these simulations are listed in Table 8.3; these values are all within experimental error of the measured values.

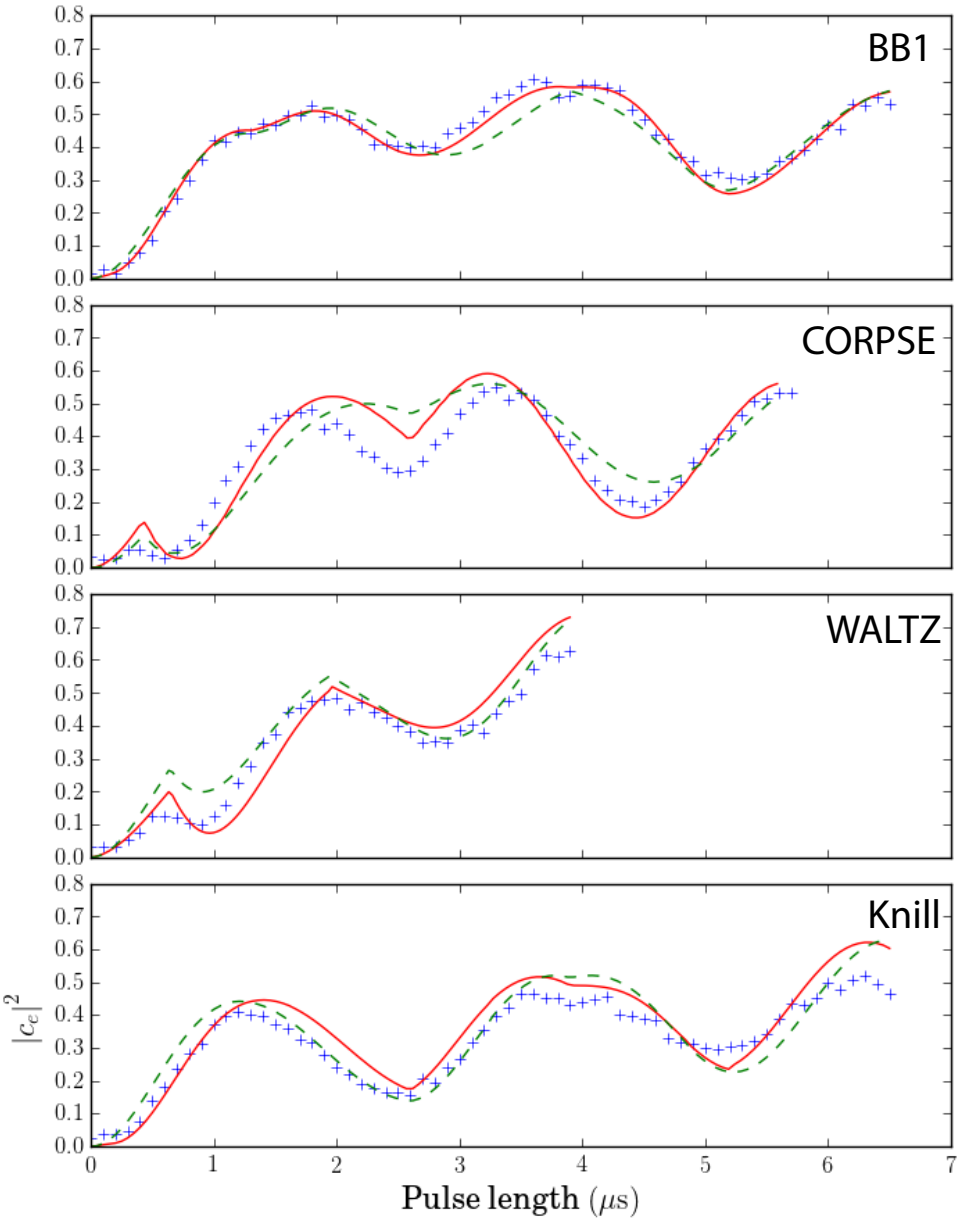

Figure 8.12: Temporal scans of composite pulses, with experimental data shown in blue, fits to the data with (red) and without (green, dashed) the effect of the AOM rise times included. a) BB1; b) CORPSE; c) WALTZ; d) Knill.

We see that the simulated fits to the data vary in accuracy depending on the pulse sequence. For example, for the BB1 temporal scan, the simulation incorporating

Symbol	Parameter	Value
I_1	Raman beam 1 intensity	$14\,\mathrm{kWm^{-2}}$
I_2	Raman beam 2 intensity	$21\,\mathrm{kWm^{-2}}$
Δ	Detuning from single photon resonance	$2\pi \times 8.5\mathrm{GHz}$
$\sigma_{1,2}$	Momentum distribution widths	$3\hbar k_{eff}$, $11\hbar k_{eff}$
t_π	Pi pulse duration	$1.3\,\mu\mathrm{s}$

Table 8.3: Simulation parameters used for the temporal scans plotted in Figure 8.12.

the finite rise time of the AOM is a better fit to the data than the simulation which assumes instantaneous switching; however, the situation is reversed for the WALTZ pulse, and for the Knill pulse neither seems to have the advantage.

These differences in fit could be due to various experimental fluctuations, for example: Raman beam intensity fluctuations caused by drift in the spectral filtering processes; misalignment of the Raman beams on the MOT; phase or intensity shifts caused by the IQ modulator. However, we are satisfied with the existing model of the experiment as it produces good fits to the data with parameter values which are within the range of experimentally measured values, and are confident in using our model to investigate further coherent manipulation experiments through simulation only.

8.5 Investigation through simulation

Having established that our simulations fit well with the experimental data, we investigate the performance of other composite pulses through simulation only. We first simulate these composite pulses with the same parameters as were used to fit the spectral scans in Figure 8.10, to enable direct comparison. We then simulate all composite pulses in an experimentally ideal system, where the m_F sublevels are completely Zeeman degenerate, to investigate the optimal performance of the composite pulses.

8.5.1 Composite pulses in a Zeeman non-degenerate system

8.5.1.1 180 - 360 composite pulse

Before investigating previously established composite pulses, we return to the rotary echo plot of Figure 8.7 and note that the second maximum, at $t \approx 5.7\mu$s, is greater than the first. This suggests that the pulse $180_\pi 360_0$ could offer improved performance over an uncorrected π pulse (180_0). To our knowledge this pulse has not been previously reported. This pulse is a point-to-point pulse, because its operator fidelity (Equation 4.2) is identical to that of a π pulse.

The point-to-point fidelity of this pulse in the presence of pulse length errors only is identical to that of the π pulse; this is as expected because the 180_π stage of the pulse cancels out the pulse length errors accumulated during half of the 360_0 stage. The fidelities for off-resonance errors (f) are, for the π pulse

$$\mathcal{F}_{PP}(\pi_0) = 1 - f^2 + \left(1 - \frac{\pi^2}{16}\right) f^4 + \left(-1 + \frac{3\pi^2}{32}\right) f^6 + \mathcal{O}(f^8) \tag{8.9}$$

and for the 180-360 composite pulse

$$\mathcal{F}_{PP}(\pi_0) = 1 - f^2 + \left(1 - \frac{\pi^2}{16}\right) f^4 + \left(-1 + \frac{35\pi^2}{32}\right) f^6 + \mathcal{O}(f^8). \tag{8.10}$$

At first glance this would suggest that the composite pulse offers no improvement over a π pulse, as there is no difference in their fidelties until the f^6 term; however, if we plot these two fidelities, we find that the fidelity of the 360 - 180 pulse is greater than that of the π pulse for all values of off-resonance error, as shown in Figure 8.13. This is possibly due to fortuitous cancellation of high order terms in f, as the shape of the fidelity as plotted in this figure can only be approximated by a series expansion of the expression for the fidelity when terms of order f^{10} are included.

We can also simulate the behaviour of this composite pulse in our system. This is plotted in Figure 8.14 in green, where it can be seen that the transferred fraction of atoms is greater than that of the π pulse at all frequencies. Its peak fidelity is 0.58, which is just above that of the BB1 pulse. However, the FWHM is 1.13 times that of the π pulse, which is worse than any of the experimentally tested

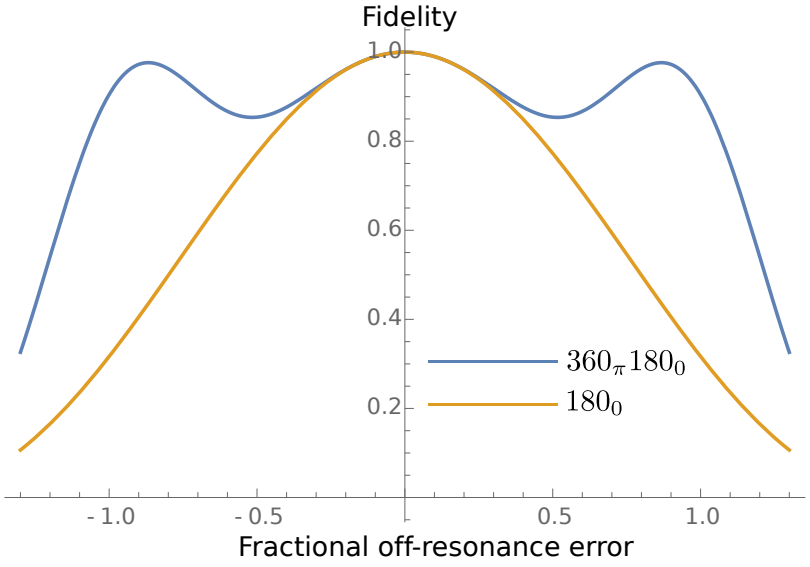

Figure 8.13: Point-to-point fidelities as a function of fractional off-resonance error, for a plain π pulse (orange) and for the 180 - 360 pulse(blue)

pulses; most surprisingly, the BB1 (a pulse which corrects pulse length errors) has a greater FWHM than the $180\pi360_0$ pulse which, from the analysis above, appears to correct off-resonance errors.

It is interesting to note that this pulse is similar to the WALTZ pulse: for both pulses, the total angle rotated through with phase π is 180° and the total angle rotated through with phase 0 is 360°. Augmentation pulses in atom interferometers appear in pairs, and two consecutive WALTZ pulses can be written as $90_0180_\pi360_0180_\pi270_0$, i.e. a WALTZ sequence with a 180 - 360 pulse after the first period of evolution. The performance of pairs of composite pulses is investigated further in Section 8.6.

8.5.1.2 The Levitt pulse

The Levitt pulse was introduced in Section 4.3 (see Figure 4.2 for a representation on the Bloch sphere). It is a point-to-point, pulse length error compensating pulse with the sequence $90_0180_{\frac{\pi}{2}}90_0$. Its behaviour in the Zeeman non-degenerate system is show in Figure 8.14, in red; it has the best performance of the pulses presented in this plot. Its peak fidelity is 0.64, which is 1.52 times that of the π pulse; this peak fidelity is better than any experimentally tested composite pulse except the WALTZ pulse. Its FWHM is 1.16 times that of the π pulse, which is to be expected as it offers no compensation for off-resonance errors.

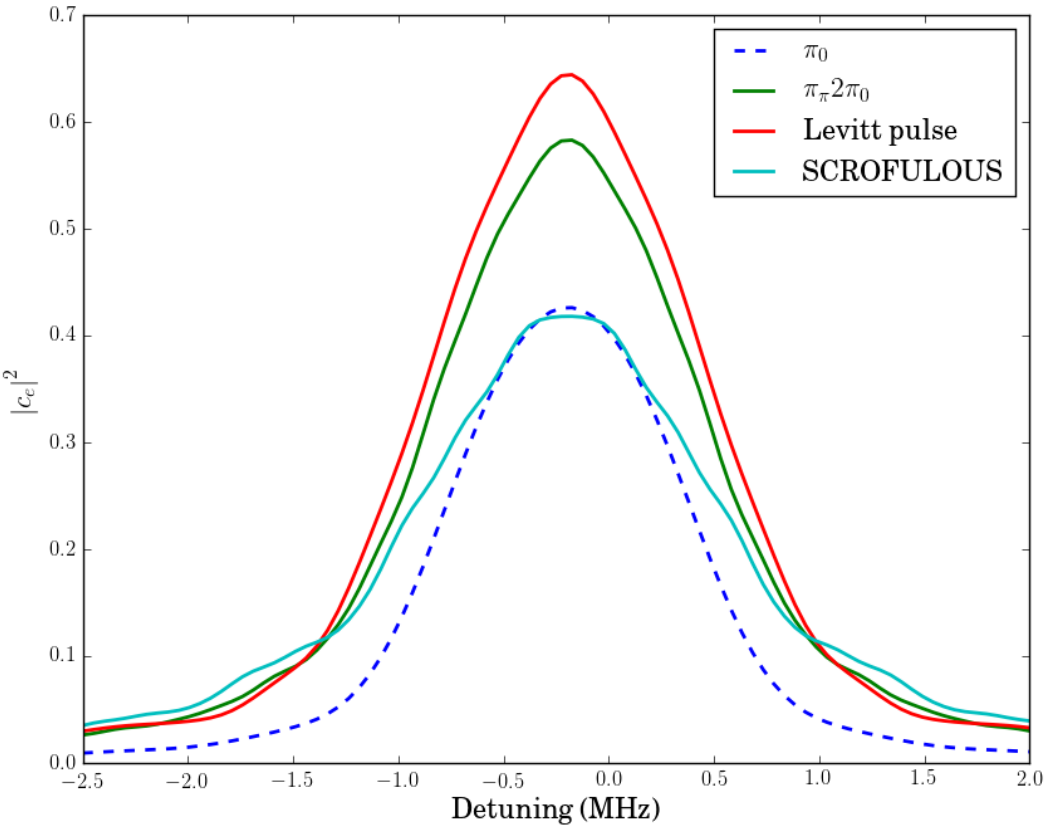

Figure 8.14: Spectral scans of composite pulses in a Zeeman non-degenerate system. The parameters of the simulations are identical to those of Figure 8.10. Blue dashed line: π pulse, green: $360\pi180_0$, red: Levitt, light blue: SCROFULOUS.

8.5.1.3 SCROFULOUS

SCROFULOUS [88] is a composite pulse which has been previously used in cold atoms experiments as part of a spin echo sequence [91]. It is a general rotor and its inversion form is $180_{\frac{\pi}{3}}180_{\frac{5\pi}{3}}180_{\frac{\pi}{3}}$. In theory, it corrects pulse length errors, with a fidelity which is $\mathcal{O}(\epsilon^6)$; however, we observe a decrease in peak fidelity compared to the π pulse. It does show an increase in FWHM, being 1.31 times that of the π pulse, which is second only to the CORPSE.

In the Zeeman non-degenerate system, none of the pulses presented in this section offer an improvement over the WALTZ or the CORPSE in peak fidelity or FWHM respectively. In the following section we investigate the performance of composite pulses in a Zeeman degenerate system to determine whether the behaviour is different.

8.5.2 Composite pulses in a Zeeman degenerate system

The simulations in this section use parameters which are experimentally realisable and represent an optimised setup: i.e., the intensities of the Raman beams are equal and the m_F levels are degenerate. We use Raman beam intensities of $I_{1,2} = 20\text{kW/m}^2$ and a detuning of $\Delta = 10\text{GHz}$; the momentum distribution is identical to that used in the simulations of the spectral scans in Section 8.4.4. The spectral scans are plotted in Figure 8.15 and the fidelities and widths are listed in Table 8.4.

Pulse name	$\mathcal{F}_P$	$\mathcal{F}_P/\mathcal{F}_\pi$	FWHM ratio	Area ratio
Pi pulse	0.61	1	1	0.13
BB1	0.67	1.10	1.21	0.23
CORPSE	0.74	1.21	1.83	0.27
WALTZ	0.80	1.31	1.61	0.29
Knill	0.77	1.26	1.17	0.22
180 - 360	0.72	1.18	1.61	0.22
Levitt	0.76	1.25	1.41	0.24
SCROFULOUS	0.59	0.97	1.09	0.18

Table 8.4: Characterisation of simulations of composite pulses in a Zeeman-degenerate system. For each pulse, we list the peak fidelity $\mathcal{F}_P$, the ratio of the peak fidelity of the pulse to the peak fidelity of a π pulse $\mathcal{F}_\pi$, the ratio of the full width at half maximum (FWHM) to that of the π pulse, and the ratio of the area under the spectral scan of the pulse to the area under the spectral scan of a perfect inversion pulse.

We find, once again, that the WALTZ pulse offers the greatest increase in peak fidelity and the CORPSE the greatest increase in FWHM. The performance of SCROFULOUS is not improved compared to the Zeeman non-degenerate case, and still provides a peak fidelity less than that of the π pulse. The Knill pulse, which is designed to correct for both pulse length and off-resonance errors, does not perform well by either the peak fidelity or the FWHM measurement; it is outperformed on both counts by the WALTZ.

In these plots, the FWHM is a less reliable measure of the range of off-resonance errors which a particular sequence can address, as most scans exhibit some sort of structure away from resonance. We might consider ranking pulses by the area under the spectral scan as a fraction of the area under the spectral scan of a perfect inversion pulse. By this metric we find that the WALTZ again offers the best

performance (more than double that of the π pulse), with little to choose between the other pulses and the SCROFULOUS again showing the worst performance by far.

These results suggest that either WALTZ or CORPSE would offer good performance as composite inversion pulses, with the choice depending on the priority given to the correction of each type of error.

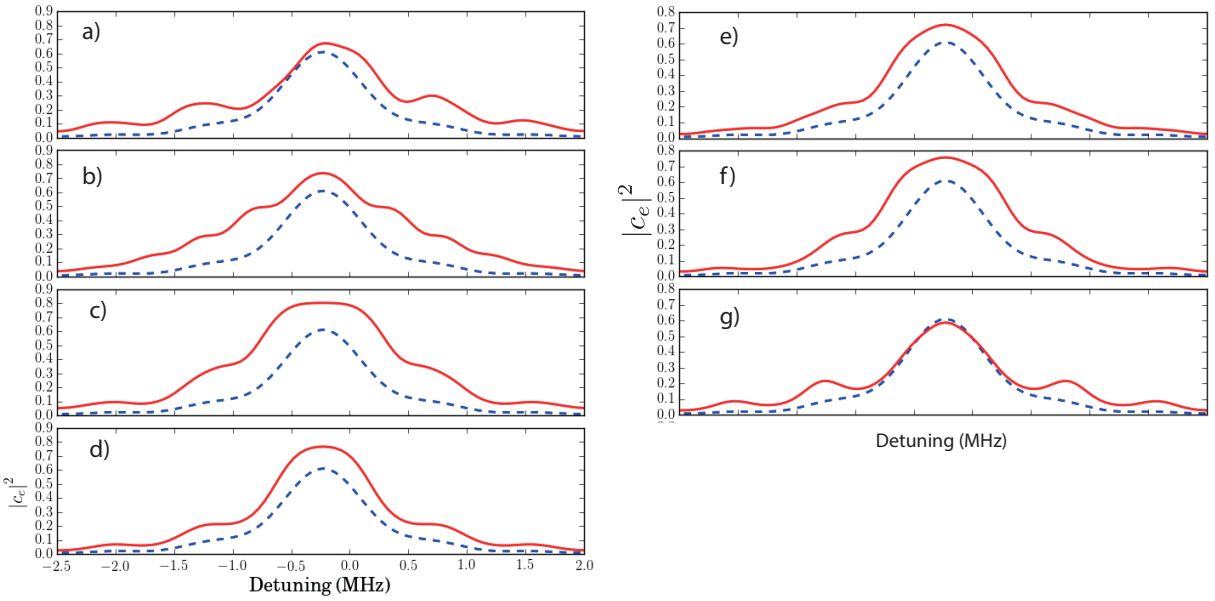

Figure 8.15: Spectral scans of composite pulses in a Zeeman degenerate system. For each plot, the spectral scan of the π pulse is shown in dashed blue for reference. a) BB1; b) CORPSE; c) WALTZ; d) Knill; e) 180-360; f) Levitt; g) SCROFULOUS.

An alternative way to analyse the performance of composite inversion pulses is to calculate their fidelity for a range of momentum distribution widths. This gives some insight into the best possible performance under achievable experimental conditions (i.e., temperatures that can be reached in a standard MOT). To investigate this, we plot the peak fidelity $\mathcal{F}_P$ of various composite pulses as a function of σ_p, the width of the momentum distribution, in Figure 8.16. We plot the three pulses with the largest increase in peak fidelity according to Table 8.4: WALTZ, Knill, and Levitt, and the π pulse for reference.

We see that the Knill pulse has the best performance when the width of the momentum distribution is small, and the WALTZ takes over as the width increases. BB1 also outperforms WALTZ at $\sigma_p < 2$ but shows a significant reduction in peak fidelities at higher values of momentum width (much more so than Knill) and is

not plotted here. The Levitt pulse is included here because it offers improved performance over the π and is also the shortest composite inversion pulse that we test here; this can be an advantage in atom interferometry where the phase sensitivity increases with the time allowed for free evolution and it is advantageous to minimise the duration of interactions with the lasers to maximise the free evolution time.

We note again, as described at the start of this chapter, that fidelities approaching 1 are achievable: as the temperature of the atoms decreases, the pulse length errors due to differing Rabi frequencies begin to dominate, and these are compensated well by all composite pulses plotted here. The π pulse can always be improved upon.

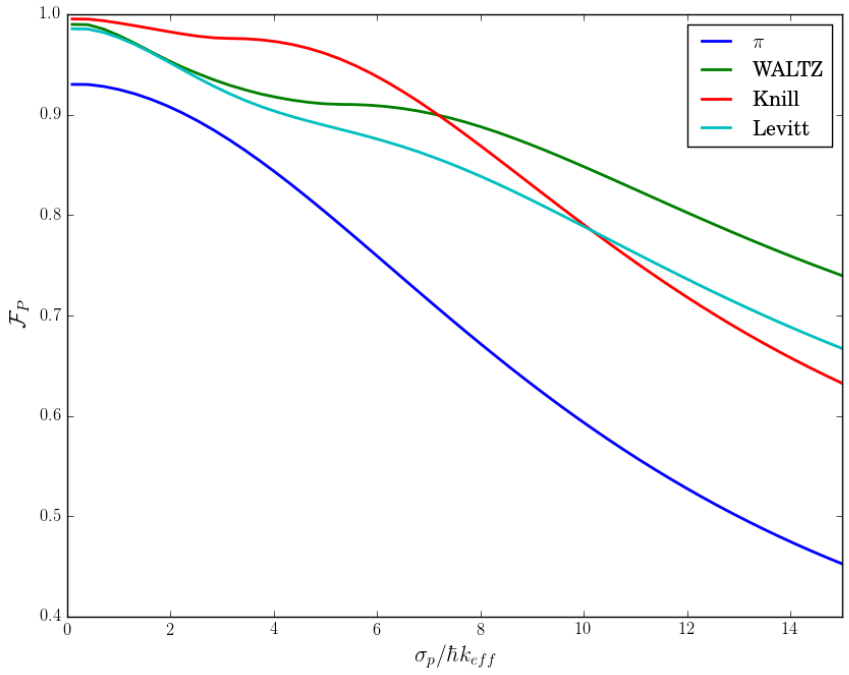

Figure 8.16: Peak fidelity of various inversion pulses as a function of momentum distribution width. Blue: π, green: WALTZ, red: Knill, light blue: Levitt.

The results presented in this section suggest that, on the basis of peak fidelity achievable, WALTZ or Knill are the best choices for an inversion pulse. Section 8.6 will consider an alternative method of assessing the performance, by simulating the interferometers themselves.

8.5.3 Compensating for AOM rise times

As discussed in Section 8.4.5, the Raman beams are not switched instantaneously and the finite rise time of the AOM results in a time-varying intensity during the pulse. In situations where this is unavoidable, it is possible to adjust the timings of the phase shifts within the pulse sequence and achieve a fidelity which is comparable to the case where the beams are switched instantaneously.

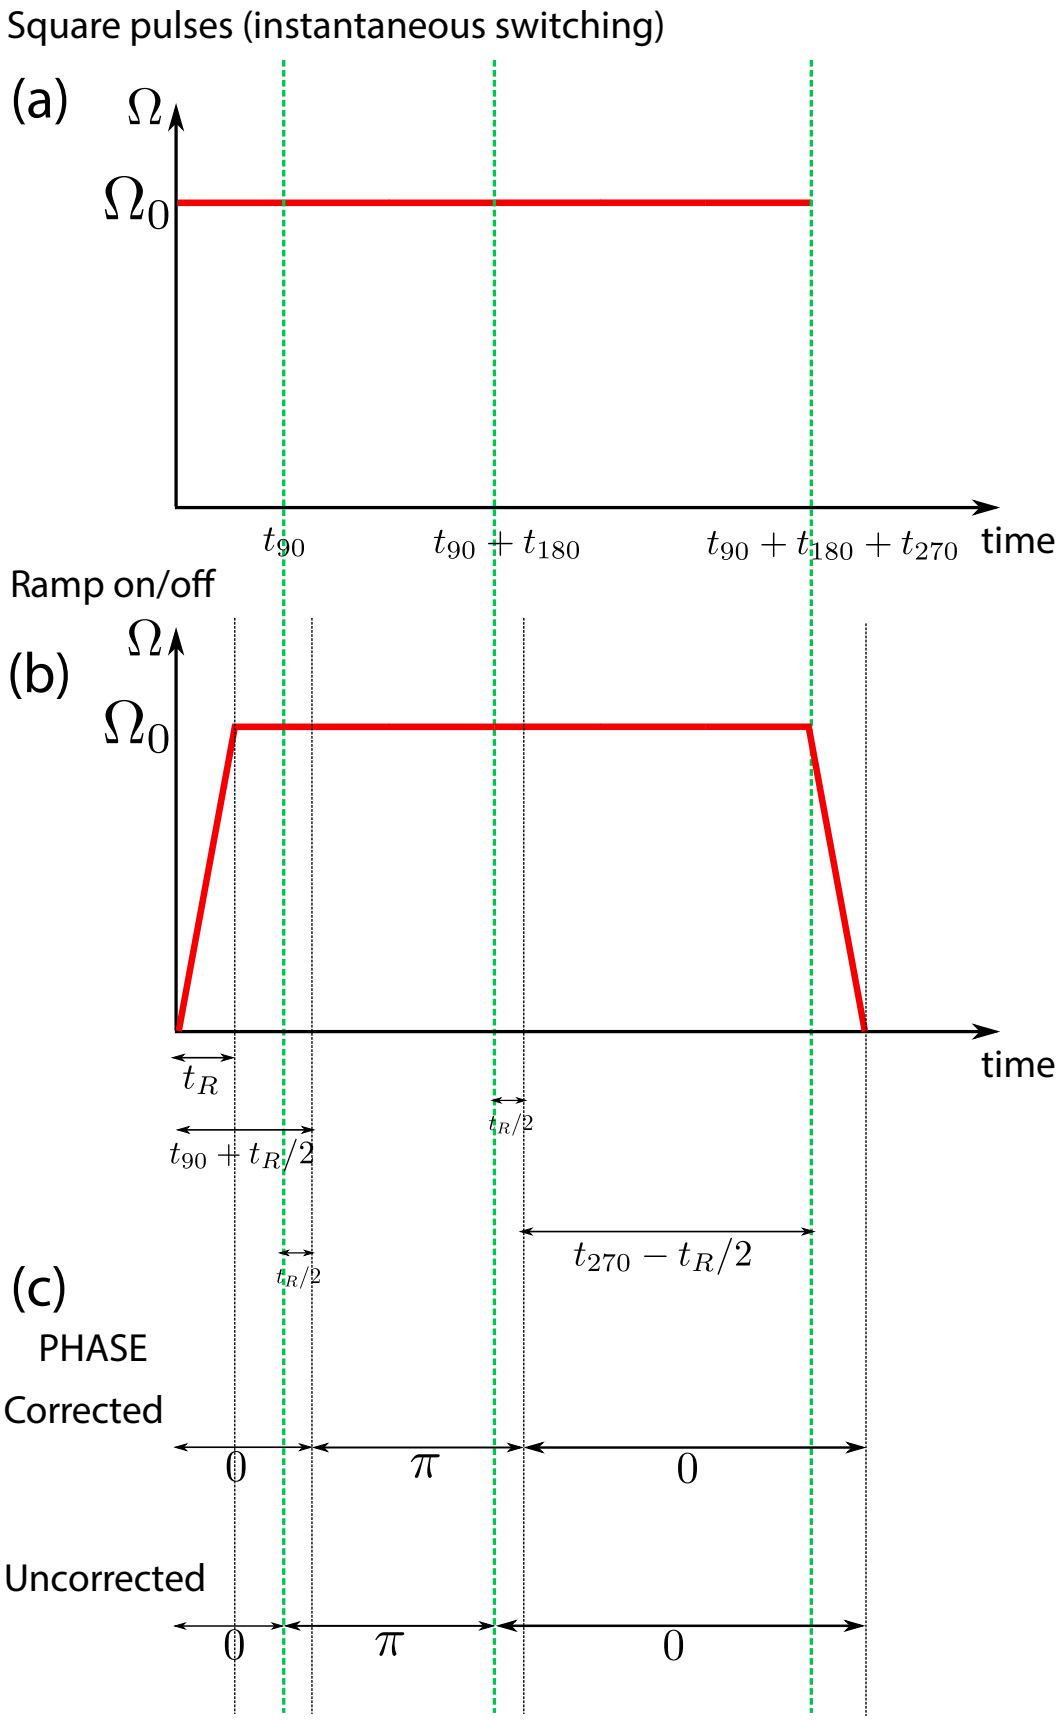

Figure 8.17: Rabi frequency-time and phase-time plots during the WALTZ pulse. a) Rabi frequency - time plot for instantaneous switching of the Raman beams. b) Rabi frequency - time plot when the AOM rise time is finite and the ramping is linear. c) Timings of the phase shifts during the pulse.

Figure 8.17 a) shows the Rabi frequency Ω and phase during a WALTZ pulse when the beam switching is instantaneous. Here we assume equal intensities I in the two beams so that $\Omega \propto I$. The area under the $\Omega - t$ curve is equal to the angle rotated through on the Bloch sphere and we require the angles for the three stages of the pulse to be 90, 180, and 270 degrees; i.e., $\Omega_0 t_{90} = \pi$ etc..

The Rabi frequency - time plot for finite AOM rise times is shown in Figure 8.17 b). Here for simplicity we assume a linear ramp on/off of intensity. If the phase shift timings are kept as in the instantaneous switching case, the angles rotated through during the three stages of the pulse are $\Omega_0 t_{90} - \Omega_0 t_R/2$, $\Omega_0 t_{180}$, and $\Omega_0 t_{270} + \Omega_0 t_R/2$. The total area under the curve is equal to that under the constant intensity curve, but the evolution time under each phase is not the same.

We can adjust the timings of the phase shifts so that the area under the curve during each period of the pulse sequence is equal to that in Figure 8.17 a); this is achieved by adding $t_R/2$ to the time at which the phase is shifted. Figure 8.17 c) shows the phase during the pulse sequence for uncorrected and corrected cases. To determine the effect on the fidelity of the composite pulse, we simulate a spectral scan of the WALTZ pulse assuming i) instantaneous switching of the Raman beams, ii) a finite ramp time with uncorrected pulse timings, and iii) a finite ramp time with corrected timings. This is plotted in Figure 8.18. We can see that the pulse sequence with corrected phase timings has an almost identical behaviour to the sequence with instantaneous switching, and both perform better than the sequence with uncorrected phase timings.

This shows that the detrimental effect of the finite rise time of the switching AOM can be removed very simply, by considering the angles rotated through during each part of the pulse and adjusting the timings of the phase shifts accordingly.

8.6 Composite pulses in atom interferometers

Spectral scans of composite pulses give an indication of the behaviour of the pulse as a point-to-point inversion operation over a range of detunings from two-photon resonance. However, we would like to know if the performance of a composite pulse as an inversion pulse is correlated with its performance as an augmentation pulse in an atom interferometer. It is important to have high fidelity augmentation pulses because each pulse with fidelity $\mathcal{F}$ will only implement a π pulse on a fraction of

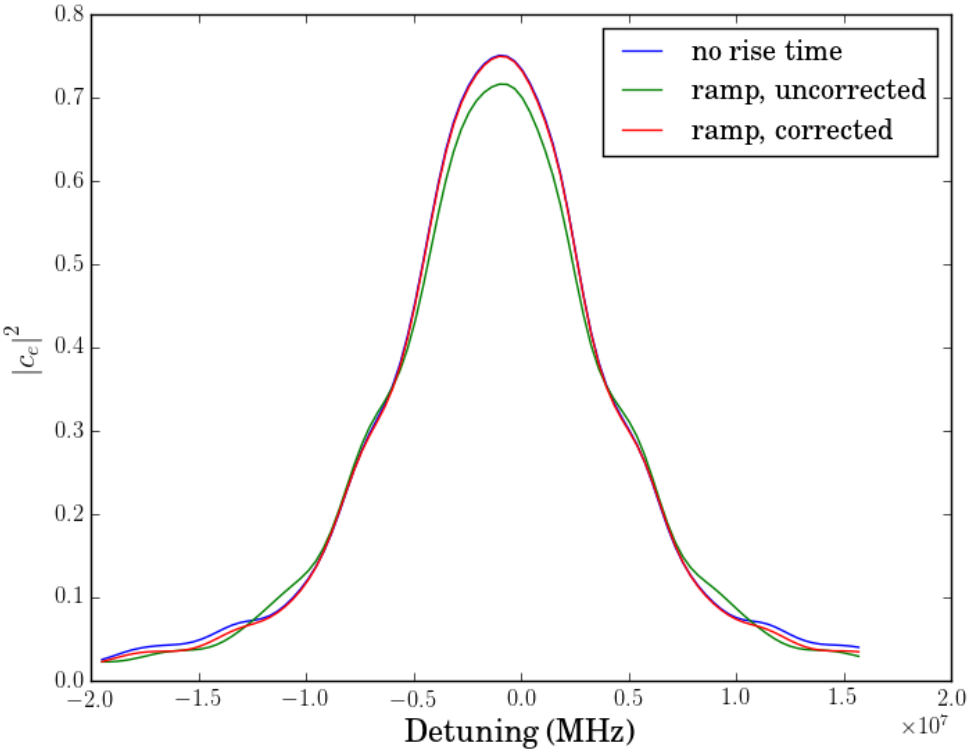

Figure 8.18: Spectral scans of the WALTZ pulse, where the Raman beam switching is instantaneous (blue), has a finite rise time and the phase timings are uncorrected (green), and has a finite rise time with phase timings corrected (red).

atoms $\mathcal{F}$ - the remainder of the atoms will be lost from the interferometer, and cause the output contrast to be reduced by $\mathcal{F}^n$ for n augmentation pulses.

We simulate an atom interferometer with augmentation pulses as shown in Figure 3.2 b): one at $t = \tau$ after the first beamsplitter, one either side of the mirror pulse at $t = T \pm \tau$, and one before the final beamsplitter at $t = 2T - \tau$. The parameters of this simulation are identical to those used in Section 8.5.2, and the timings are $T = 100\mu$s and $\tau = 10\mu$s. Figure 8.19 a) shows the excited state population as a function of the phase of the final beamsplitter pulse when the momentum distribution is a single Gaussian with width $\sigma_p = 4\hbar k_{eff}$. The fringes are plotted for π, WALTZ, Knill and Levitt augmentation pulses; it can be seen that all composite pulses offer an improvement in contrast compared to the π pulse.

In Figure 8.19 b), we plot the interferometer contrast as a function of the momentum distribution width. We see that all composite pulses outperform the π pulse at all points. It is interesting to compare this plot to Figure 8.16, where

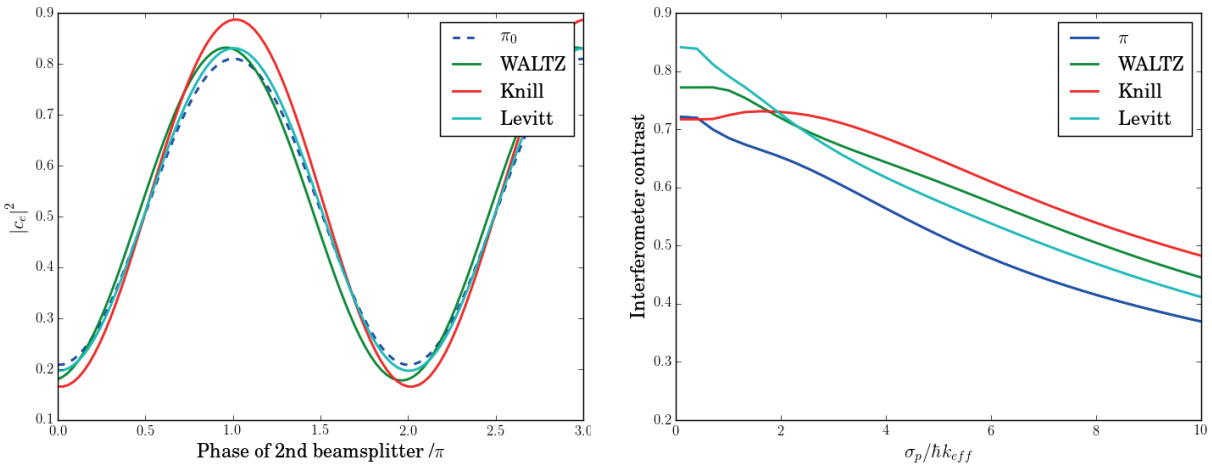

Figure 8.19: Augmentation pulses in atom interferometers. a) Excited state population after final beamsplitter as a function of final beamsplitter phase, plotted for π, WALTZ, Knill and Levitt augmentation pulses. b) Interferometer contrast as a function of momentum distribution width, for the same composite pulses as in a).

the Knill pulse offered the best performance as an inversion pulse up to approximately $7\hbar k_{eff}$, after which point the WALTZ gave the highest fidelity. We find that when used as augmentation pulses, Knill does still offer improvement over WALTZ at small σ_p, but Levitt is better than both Knill and WALTZ in this range ($\sigma_p < 2\hbar k_{eff}$). These results agree with the experiment reported in [24], where WALTZ is found to improve the interferometer contrast both when 4 and 8 augmentation pulses are applied.

8.7 Conclusion

In this chapter, the ability to coherently manipulate an ensemble of ultracold Rubidium atoms was established. We showed that although the wide range of resonance frequencies present in the ensemble resulted in rapid dephasing of Rabi flopping, the contrast can be regained simply by changing the phase of the driving laser. We showed that coherence times necessary for atom interferometry are achievable.

Composite pulses were investigated experimentally and through simulations. All tested composite inversion pulses offered an improvement over a basic π pulse, although the extent of the improvement varied between pulses. Simulations were used to fit to experimental spectral and temporal scans of composite inversion

pulses; the good agreement between experiment and theory justified further investigation through simulation.

Further composite pulses, not tested in our experiment, were simulated as inversion pulses. This included one pulse which we believe to be novel. We then compared the performance of composite pulses as inversion pulses with their performance as augmentation pulses in large area atom interferometers. We showed that good performance as an inversion pulse is not necessarily indicative of good performance as an augmentation pulse, and that the optimal choice of pulse depends on dominant type of error present.

Further work in this area might involve testing composite pulses experimentally in interferometers, and development of new pulse sequences which are optimised for the requirements of atom interferometry.

Chapter 9

Modelling and characterisation of interferometric cooling

Although it was proposed 15 years ago, no experimental implementation of the interferometric cooling method described in Section 3.3 was reported in the intervening years. In this chapter, the first experimental demonstration of interferometric cooling is described and we show that it can be modelled well by our theoretical framework. We also investigate extensions to the experiment. The experiment was reported in [3]. All atomic physics data in this chapter was taken by Alex Dunning.

9.1 Experimental details

As a first step towards studies of interferometric cooling in complex structures, the method was tested with already ultracold atoms on a single atomic transition. We showed in the previous chapter that coherent manipulation of internal states is readily achievable in our experiment, and so to test the interferometric cooling the $F = 2$ and $F = 3$ levels of the ground state $5S_{1/2}$ are used for the ground and excited states of the interferometer. As we are only using one atomic transition, the interferometer consists of only the two beamsplitter pulses, as described in Section 3.3. Additionally, we do not alternate the directions of the Raman beams between cycles, and so the atoms are both cooled and accelerated in the direction of k_{eff}.

As spontaneous emission is necessary to cause cooling, and there is no dipole-allowed transition between the $F = 2$ and $F = 3$ states, we force spontaneous emission by applying a 'depump' laser which is tuned to the $5S_{1/2}, F = 3$ to $5P_{3/2}, F' = 3$ transition in order to only address those atoms left in the excited state after the interferometer. The depump laser is applied after each cycle of the interferometer.

The velocity distribution of the atoms is measured by Raman velocimetry. This is essentially velocity selection by applying a low-intensity Raman pulse and measuring the population in the excited state; by varying the detuning of the Raman beams and repeating this measurement, the velocity distribution of the atoms can be constructed. Recall that the detuning of the Raman beams from the two photon transition is $\delta = \delta_L + \mathbf{p} \cdot \mathbf{k}_{eff}/m + \hbar k_{eff}^2/2m$, which is velocity-dependent. At various values of δ_L, different velocity classes will be on resonance and hence transferred to the excited state. The resolution of this method is limited by the spectral width of the pulse: the smallest velocity width Δv that can be resolved is related to the range of frequencies present in the pulse used to excite the atoms. According to Fourier transform theory the bandwidth Δf of a pulse of duration T is limited to $\Delta f \geq 1/T$; similarly the resolvable velocity range of a pulse with wavelength λ and duration T is $\Delta v \approx \lambda/T$ [30]. In the Raman velocimetry used to measure the velocity distributions after the interferometric cooling, the pulse duration used is 200μs.

9.2 Modelling the interferometric cooling experiment

Appendix C contains Python code to simulate interferometric cooling. The following is an outline of the stages of the simulation to fit to the experimental data. A Monte Carlo method is used to account for the random nature of the spontaneous emission - in particular, the direction of the impulse given to the atom by a spontaneous emission event. Monte Carlo methods have been used to simulate laser cooling in many situations (see e.g. [92], [93], [94] for more details).

Apply the interferometer sequence to the atoms: this is done using the same model of the evolution of state amplitudes as developed in Chapter 8. We calculate the amplitude after the pulse-free evolution-pulse sequence for each combination of momentum and transition route.

Modify the momentum distribution: we store the momentum distribution of atoms in the ground state and atoms in the excited state as separate arrays. If we consider an atom with a momentum p which, after the interferometer sequence, is in the state $c_e |e\rangle + c_g |g\rangle$, we modify the ground state momentum distribution at p by $|c_g|^2$ and the excited state momentum distribution at $p + \hbar k_{eff}$ by $|c_e|^2$ (with a weighting dependent on the momentum distribution at p before the interferometer sequence). The momentum transfer occurs along the k vector of the Raman beams and so the momentum distribution of the atoms is modulated in this dimension only.

Spontaneous emission: after the interferometer, we are left with a momentum distribution for atoms in the ground and in the excited state. The depump laser only affects atoms in the excited state, which are excited to the $F' = 3$ state and then spontaneously decay with a branching ratio of 47% and 53% to $F = 3$ and $F = 2$ respectively; at this branching ratio, 98% of atoms are transferred to the ground ($F = 2$) state after 5 absorption-emission cycles, and so we simulate this number. The spontaneous emission is simulated as follows: for a particular momentum p, we simulate n atoms with this momentum, where n is proportional to the height of the excited state momentum distribution at p. We generate a random number in the range $(0, 1)$; if this number is less than 0.47, the spontaneous decay returns the atom to the excited state (F=3), and otherwise the atom is transferred to the ground state (F=2). A random number θ is generated in the range $\pm\pi$, and the component along the Raman beam axis of the change in momentum due to spontaneous emission is calculated as $\cos\theta$. The momentum distribution of the ground or excited state (depending on the first random number generated) at the momentum value $p + (\hbar k_{eff} \cos\theta)/2$ is modified by an amount proportional to the height of the momentum distribution at p. This is repeated for each momentum value and the ground and excited state momentum distributions at the end of one spontaneous emission cycle are used as the starting distributions for the next spontaneous emission cycle.

Velocity probe pulse: after 5 spontaneous emission cycles, the velocity probe pulse is applied. This is simulated using the same model for the Raman pulses as in the interferometer sequence, with the momentum distribution being that achieved at the end of the spontaneous emission stage of the simulation. The probe is applied at varying values of detuning δ from two photon resonance.

We use this framework to simulate the interferometric cooling experiment, with parameters matching those measured experimentally: Raman beam intensity of

25kW/m^2, π pulse duration of 1.0μs, single-photon detuning of 10GHz. The velocity distribution is a double Gaussian with widths $\sigma_1, \sigma_2 = 3\hbar k_{eff}, 13\hbar k_{eff}$ having a ratio 3:1. The two beamsplitter pulses have phases $\pm 3\pi/16$; this is slightly different to the values set in the experiment, which we may attribute to the IQ modulator's difficulty in reproducing phase shifts accurately. A good fit is made to the velocity scan before any interferometry has taken place, and subsequently the parameters remain unchanged for simulation of the velocity distributions after the cooling cycles. Approximately 220,000 atoms are simulated, and the continuous range of momentum values is discretised into 1600 possible values. The results are shown in Figure 9.1. The cooling cycle is applied n times with $n = 0, 4, 8, 12$.

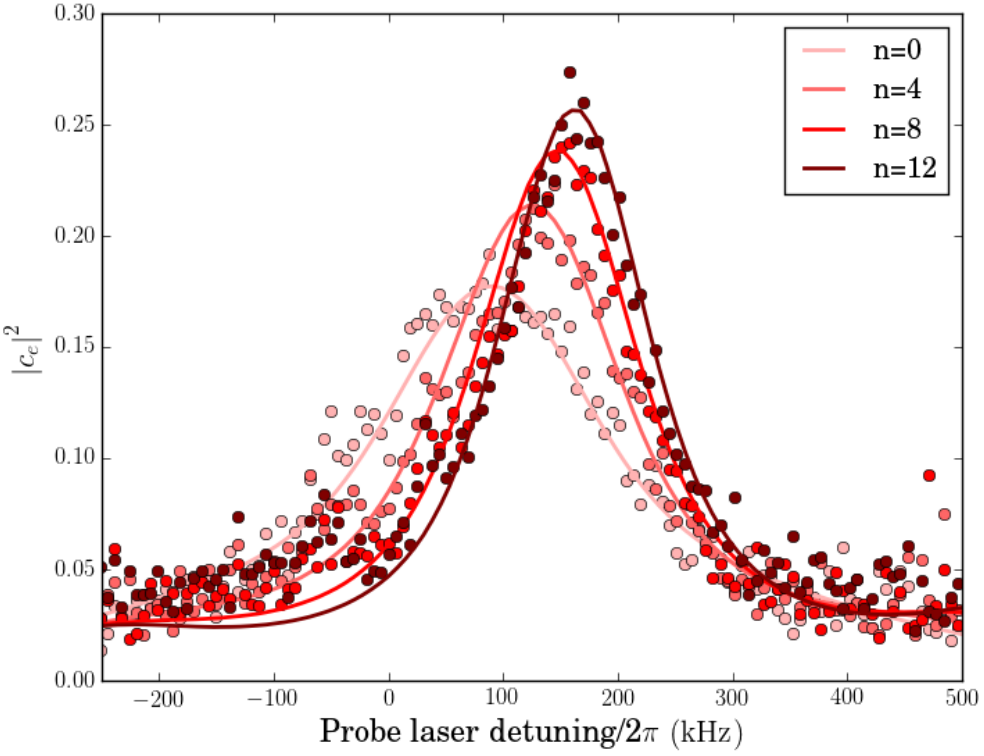

Figure 9.1: Raman velocimetry scans after n cycles of interferometer-spontaneous emission, where $n = 0, 4, 8, 12$. Circles: experimental data, lines: simulations to fit to data.

The velocity distribution becomes narrower as n increases, and the area under the simulated curves remains essentially the same (a decrease of 1.2% from $n = 0$ to $n = 12$) - i.e., the number of atoms that are being measured does not change. This allows us to say that cooling has occurred, as opposed to simply selection of a narrower velocity class. To quantify the change in temperature, we can investigate the velocity distributions as calculated by the simulations after n repetitions of the interferometer. This is plotted in Figure 9.2. The velocity distribution for $n = 0$

is a smooth curve whereas for $n = 4, 8, 12$ it is not smooth - this is because of the random direction of the change in momentum after a spontaneous emission event.

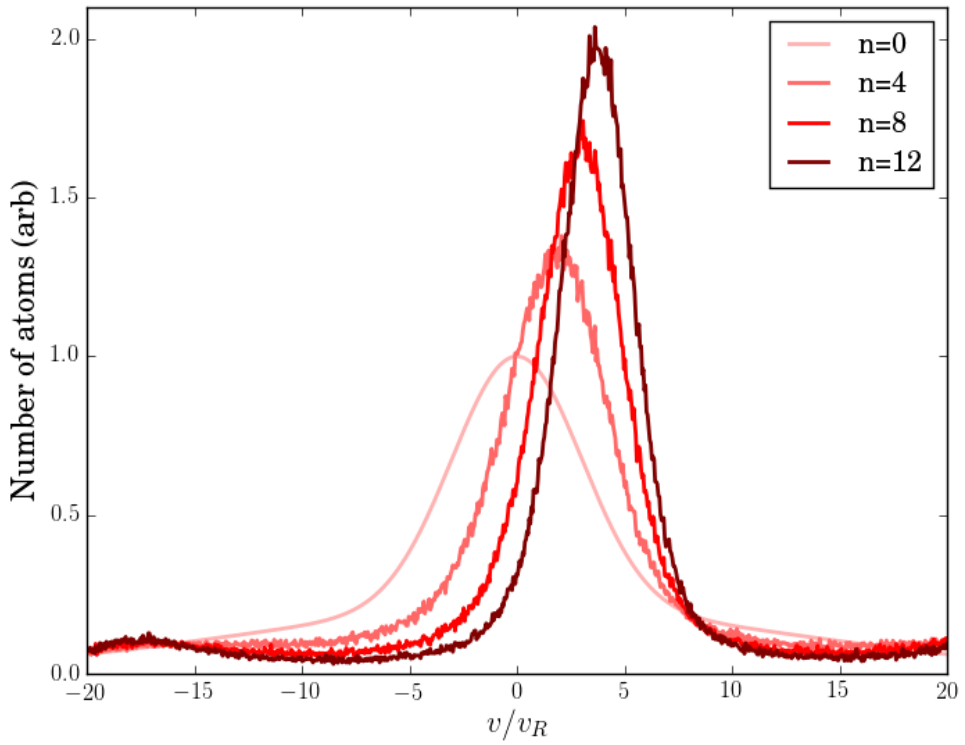

Figure 9.2: Velocity distributions after n cycles of interferometer-spontaneous emission, where $n = 0, 4, 8, 12$. These distributions are from the simulations plotted in Figure 9.1.

We fit double Gaussian distributions with weights equal to those used in the simulation to these velocity distributions. The peak of the distribution shifts by $1.9v_R$, $3.0v_R$, and $3.7v_R$ for $n = 4, 8, 12$. We assume that the atoms belonging to the Gaussian with the wider width remain at the same temperature throughout the experiment, and the width of the narrower Gaussian decreases with the number of interferometer cycles. The temperature of the atoms within the narrower Gaussian distribution (width $\sigma_v = 3v_R$) before the interferometer is $T_0 = m\sigma_v^2/k_B = 19\mu$K. After 4 cycles ($\sigma_v = 2.6v_R$) the temperature is 10μK, after 8 cycles ($\sigma_v = 2.0v_R$) the temperature is 5.9μK, and after 12 cycles ($\sigma_v = 1.7v_R$) the temperature is 4.3μK. As well as being cooled, the atoms are accelerated: the peak of the distribution is at $1.9v_R$ for $n = 4$, $3.0v_R$ for $n = 8$, and $3.7v_R$ for $n = 12$.

9.3 Interferometric cooling: investigation through simulation

The previous section showed a good agreement between experiment and our model for interferometric cooling, and we continue in this section to investigate cooling through simulations.

9.3.1 Cooling without acceleration

Figure 9.1 shows both cooling and acceleration of the atoms. If the impulse given to the atoms alternates in sign in successive interferometers, then cooling should be possible without acceleration. Experimentally, there are two possible ways of achieving this: the direction of the Raman beams can be alternated each cycle (for example, by switching with a Pockels cell), or the atoms can be prepared alternately in the $F = 2$ and $F = 3$ sublevels of $5S_{1/2}$.

We implement direction switching of the Raman beams in the simulation by changing the sign of k_{eff} on each cycle of the interferometer. The resulting Raman velocimetry scan after 12 cycles of the interferometer, along with the scans for $n = 0$ and for $n = 12$ without direction switching, is plotted in Figure 9.3. We see that after 12 cycles of the interferometer with switched Raman beam directions, there is no acceleration of the atoms, and the same amount of cooling occurs as when the beams are not switched.

9.3.2 Limits of cooling in our experiment

The theoretical limiting temperature of interferometric cooling in our experiment is the Raman recoil temperature $\hbar^2 k_{eff}^2/mk_B = 1.5\mu$K. We would like to investigate how the temperature of the atoms evolves with the number of repetitions of the interferometer.

As described in Section 3.3, the dwell time T sets the velocity capture range Δv of the interferometer according to $\Delta v = \pi/k_{eff}T$. We first simulate the resulting velocity distributions after n repetitions of the interferometer, with T set such that Δv encompasses 90% of the atoms. We choose the phases $\phi_{1,2}$ so that the velocity at which the gradient of the excitation profile is a maximum is equal to the most probable velocity (the value at which the velocity distribution is a

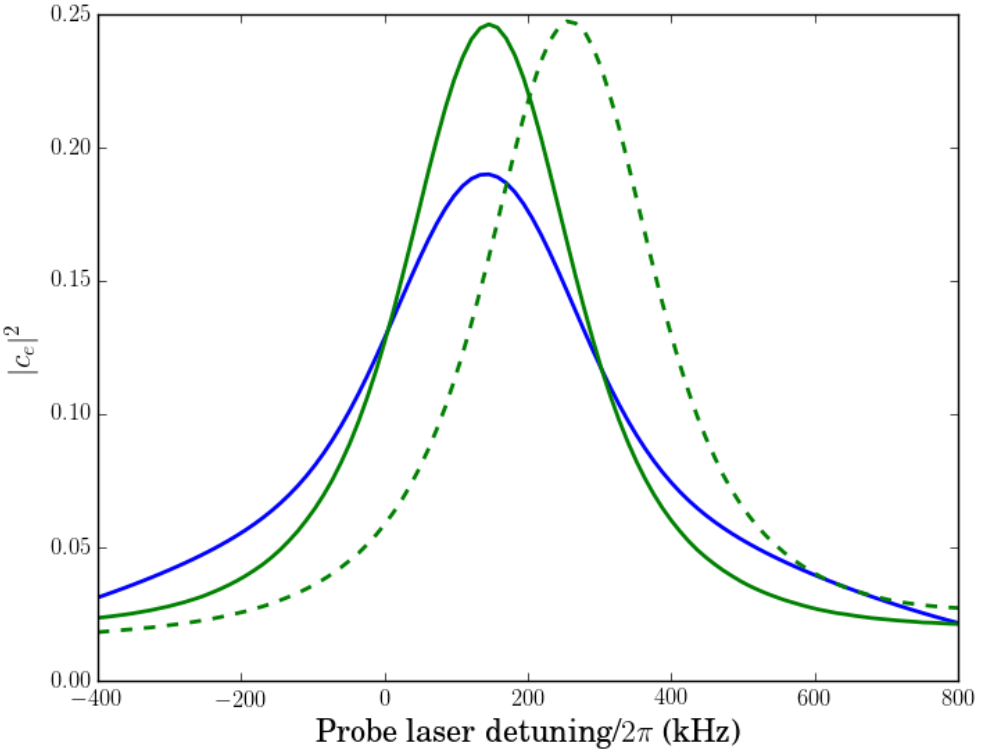

Figure 9.3: Simulated Raman velocimetry scans of the velocity distribution of the atoms. Blue: before interferometry. Dashed green: after 12 cycles of the interferometer, without switching the direction of the Raman beams. Solid green: after 12 cycles of the interferometer, where the direction of the Raman beams is alternated each cycle.

maximum). Figure 9.4 shows the velocity distribution before the interferometer (blue), and the excitation profile - the fraction of atoms in the excited state after one interferometer sequence - in green; the dwell time and phases are $T = 0.75\mu$s and $\phi_1 = \pi/4$, $\phi_2 = -\pi/4$. The blue vertical lines are placed so that 95% of the atoms fall between them; they also mark the limits of the velocity capture range. In this figure the width of the initial velocity distribution is $3v_R$. Also plotted in Figure 9.4 is the velocity distribution after 50 repetitions of the interferometer. The distribution is plotted in red and the vertical red lines show the range within which 95% of atoms are found - this range is now smaller than Δv.

Consider an atom with initial velocity $-5v_R$. Being at the edge of the velocity capture range, this atom can receive the maximum impulse of $\hbar k_{eff}$ upon excitation. After several repetitions of the interferometer, the velocity of the atom will move closer to $v = 0$, and if T remains unchanged so that the excitation profile remains the same, the impulse it receives on excitation will be less than $\hbar k_{eff}$. Thus as

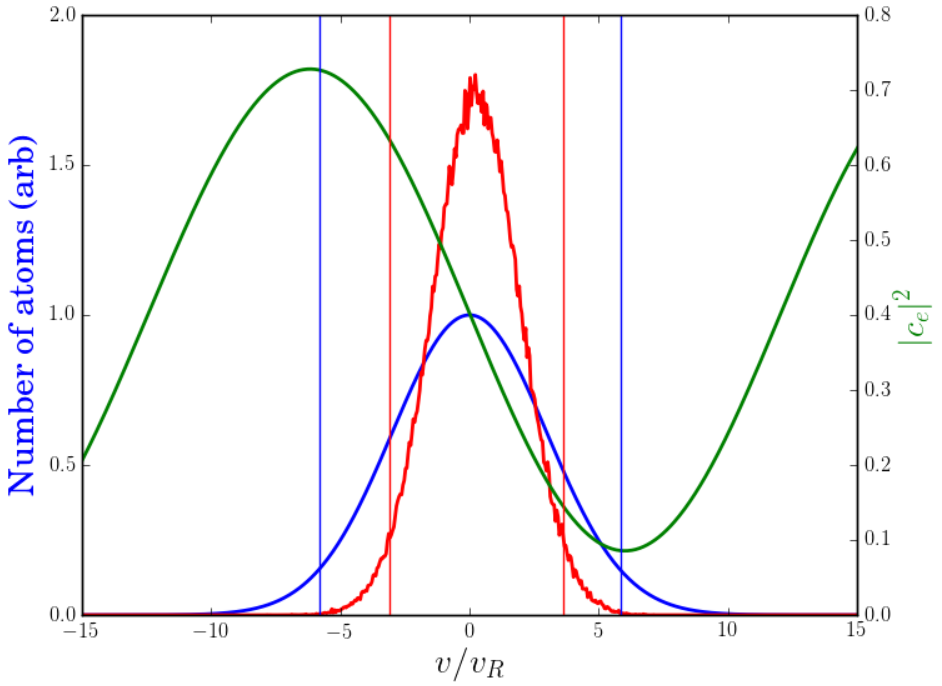

Figure 9.4: Excitation profiles and velocity distributions for interferometric cooling. Green: excitation profile at start of cooling, with T and $\phi_{1,2}$ chosen such that the velocity capture range encompasses 95% of the atoms. Blue: initial velocity distribution (width $3v_R$), with blue lines showing the range within which 95% of atoms are found. Red: velocity distribution after 50 repetitions of the interferometer, with T and $\phi_{1,2}$ unchanged; red vertical lines show the range within which 95% of the atoms are found.

the atom cloud cools and its velocity distribution narrows, if the dwell time T is constant, we expect that the cooling becomes less efficient with increasing n.

We can investigate this by changing T and $\phi_{1,2}$ with n in order to to maintain a velocity capture range which is only wide enough to include 90% of the atoms. After each cycle of the interferometer the required velocity capture range is recalculated, which sets the value of T, and then $\phi_{1,2}$ are set to maintain the position of the maximum gradient of the excitation profile at the centre of the momentum distribution.

Figure 9.5 shows the temperature of the atoms as a function of the number of repetitions of the interferometer, both with and without varying T with n. Here the initial velocity distribution width is $30v_R$, corresponding to a temperature of approximately 1.4mK. The temperature, found by fitting a Gaussian to the simulated velocity distribution after each repetition of the interferometer, is plotted in blue (circles). The blue line is a fit to the data of the form $T(n) = A_1 + B_1 \exp(-n/C_1)$,

where $T(n)$ is the temperature in mK after the nth repetition of the interferometer. We find a good fit with $A_1 = 0.098$mK, $B_1 = 1.27$mK, and $C_1 = 42.5$. The green circles and solid line show the temperatures, and their fit, when the dwell time and phases are recalculated after each interferometer-spontaneous emission event. The fit is again exponential, with parameters $A_2 = 0.035$mK, $B_2 = 1.39$mK, and $C_2 = 20.5$. Clearly the rate at which the temperature decreases is greater when the dwell time and relative phase are recalculated. The decay constant in the case of corrected times and phases is less than half that when the times and phases are constant.

The fits suggest that in both cases the limiting temperature is greater than the recoil temperature; for the blue plot the limiting temperature is 98mK, and for the green it is 35mK. There are several possible reasons for this. First, the fits are exponential and appear to be good fits, but the data does show a trend which is not quite exponential; especially below $n = 40$, both sets of data are consistently above and then below the fit. It is possible that as n increases there is deviation from the exponential fit. Second, only the dwell time and the phases are changed with n; the intensity and π time remain constant, and is it possible that more effective and efficient cooling could be achieved by also recalculating their optimal values at each step. Finally, the recalculated phase and dwell time are chosen to include 90% of the atoms; a small number of high velocity atoms, outside the region of efficient cooling, can have a large influence on the RMS velocity. This could be countered by only including atoms within a particular velocity range in the temperature calculation, and/or increasing the region of efficient cooling to include greater than 90% of the atoms. This is the subject of continuing investigation.

9.4 Conclusion

In this chapter we have presented data from the first experimental demonstration of interferometric cooling and showed that very good fits are achievable using our model of the experiment, with parameters in good agreement with those measured experimentally. Cooling of ultracold Rubidium from 19μK to 4.3μK was inferred from the data.

With confidence in the ability of the model to predict experimental outcomes, we investigated through simulation some modifications to the cooling procedure. It was shown that cooling without acceleration is possible with only minor experimental adjustments to switch the directions of the Raman beams. It was also

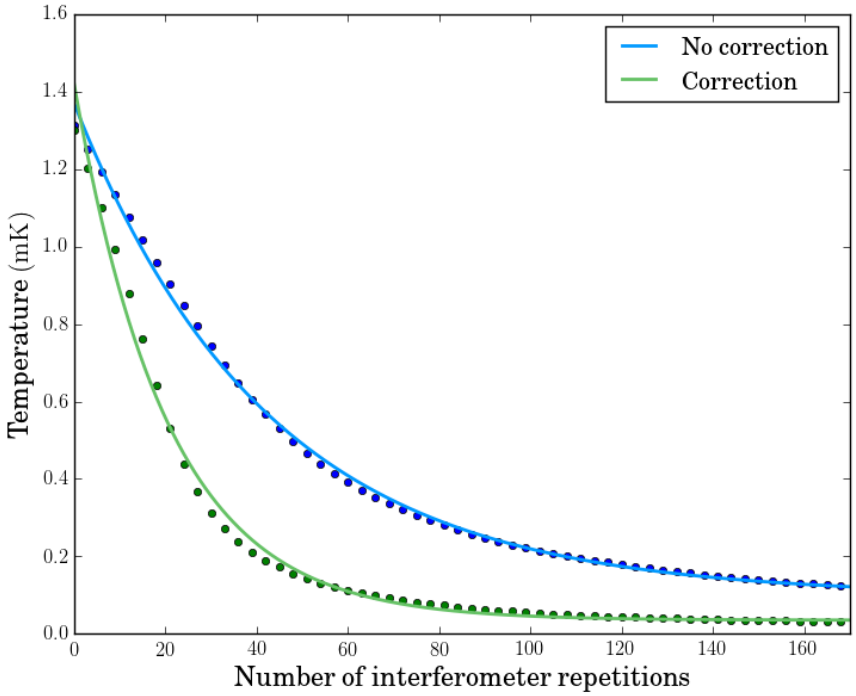

Figure 9.5: Temperature as a function of number of repetitions of interferometer, for atoms beginning at 1.4mK. Points are values calculated from fits to simulated velocity distributions at each value of n; lines are exponential fits to the points. Blue: dwell time and phases remain unchanged with n; green: dwell time and phases recalculated at each step.

shown that the efficiency of the cooling can be doubled by recalculating interferometer parameters after each step, resulting in half as many spontaneous emissions required to reach the same temperature.

Although the simulations do not agree with the theory in terms of the achievable temperature, this is most likely due to limitations of the simulations in their current form. Future investigations will focus on optimisation of controllable experimental parameters such as detuning, Raman beam intensity, interferometer dwell time and interferometer phases in order to reduce the number of spontaneous emissions required to reach the limiting temperature.

Chapter 10

Conclusions and future work

This thesis presented a study of the characteristics and limitations of coherent manipulation in an ensemble of ultracold Rubidium, with a particular focus on atom interferometry applications.

We first investigated the achievable coherence times in the experiment. Initial Rabi flopping experiments showed a rapid decay in contrast. However, the contrast could be recovered by applying rotary echoes, and Rabi flopping was observed to persist with minimal decay in contrast for hundreds of microseconds. We were able to use a model derived from the equations governing the time evolution of state amplitudes in a pseudo-two-level system to produce simulated fits with very good agreement with the data.

These simulations were used to investigate the possible causes of the small decay in contrast that is observed in the rotary echo experiment. This could have been due to phase noise at a level higher than that predicted by the specifications of the IQ modulator (for example, due to small variations in the path length difference of the Raman beams), and a good fit to the data was found in this case. A good fit was also found under the assumption of no additional phase noise when randomisation of the phase of the quantum states occurs upon collision between atoms; however, the fit assumed an unreasonably short coherence lifetime and was discounted for this reason.

The model was also found to give very good fits to experimental data regarding composite inversion pulses. It was found that several established composite pulses could implement a higher fidelity broadband inversion pulse than a simple π pulse. Again we built on the good agreement between theory and experiment to further

investigate composite pulses through simulation, including one pulse which we believe to be novel. We also showed that the decrease in fidelity due to the finite rise time of the switching AOM can be reversed by appropriate shifts in the timings of the stages of the composite pulse. The performance of the composite pulses offering the highest fidelity as inversion pulses was compared to their performance as augmentation pulses in atom interferometers. We found that the best choice of augmentation pulse is dependent on the temperature of the atoms, and that a high fidelity inversion pulse is not necessarily also a high fidelity augmentation pulse.

The first experimental demonstration of interferometric cooling was modelled by simulations. We showed that a good fit to the data is achievable, and investigated possible improvements to the experiment. This included switching the direction of the Raman beams every other cycle in order to produce cooling without acceleration, and changing the dwell time and phases of the interferometer as the temperature of the atoms decreases in order to maintain an efficient cooling effect.

Future work

The experiment in its current form is operating at the limits of the capability of the equipment, as regards the power in the Raman beams, the stability of the lasers, and the efficiency of the spectral filtering. The experiment will soon be upgraded to provide more power in the Raman beams; this will be achieved by putting an extra tapered amplifier in the EOM beam. The replacement master laser is a DBR laser which offers much greater stability and tuneability than the current master laser. The spectral filtering (liquid crystal and Mach Zehnder) stages will be replaced with fibre circulator components for greater stability. The higher powers will bring interferometric cooling with higher temperature atoms into the reach of the experiment. This will necessitate investigation through simulation and experiment of the optimal evolution of the dwell time and relative phase of the interferometer during the cooling process. It would also be interesting to investigate whether the cooling efficiency can be improved by optimising the intensity of the Raman beams after each cycle.

Appendix A

Optical Bloch equations

In 2.2 we derived the time evolution of the state amplitudes for a pseudo-two level system under coherent excitation. Here, we are also concerned with the incoherent process of spontaneous emission which is required to cause cooling. A convenient way of describing such a two level system interacting with radiation is the density matrix approach. The density matrix for a two-level system in the state $|\Psi\rangle = c_g |g\rangle + c_e |e\rangle$, with ground and excited state amplitudes c_g and c_e, is

$$\rho = |\Psi\rangle \langle\Psi| = \begin{pmatrix} c_e c_e^* & c_e c_g^* \\ c_g c_e^* & c_g c_g^* \end{pmatrix}. \tag{A.1}$$

The diagonal elements ρ_{ee} and ρ_{gg} are clearly the populations of the states, and the off-diagonal elements ρ_{eg} and ρ_{ge} are called the coherences. Since we are considering a closed two-level system, the population is conserved: $\rho_{ee} + \rho_{gg} = 1$. To calculate the time dependence of the elements of the density matrix, we first solve the Schrödinger equation for a two level atom interacting with a laser, where the Hamiltonian is

$$\hat{H} = \hbar\omega_e |e\rangle \langle e| + \hbar\omega_g |g\rangle \langle g| - V_{eg} |g\rangle \langle e| - V_{ge} |e\rangle \langle g|). \tag{A.2}$$

Making the rotating wave approximation and defining $\delta = \omega - (\omega_e - \omega_g)$, we find the amplitudes

$$\begin{aligned} \frac{dc_g}{dt} &= i\frac{\Omega_{eg}}{2}c_e e^{i\delta t} \\ \frac{dc_e}{dt} &= i\frac{\Omega_{ge}}{2}c_g e^{-i\delta t}, \end{aligned} \tag{A.3}$$

and thus

$$\frac{d\rho_{ee}}{dt} = \gamma\rho_{ee} + i\frac{\Omega_{ge}}{2}\tilde{\rho}_{ge} - i\frac{\Omega_{eg}}{2}\tilde{\rho}_{eg} \tag{A.4a}$$

$$\frac{d\rho_{gg}}{dt} = -\gamma\rho_{ee} + i\frac{\Omega_{eg}}{2}\tilde{\rho}_{eg} - i\frac{\Omega_{ge}}{2}\tilde{\rho}_{ge} \tag{A.4b}$$

$$\frac{d\tilde{\rho}_{eg}}{dt} = -(\frac{\gamma}{2} - i\delta)\tilde{\rho}_{eg} + i\frac{\Omega_g e}{2}(\rho_{gg} - \rho_{ee}) \tag{A.4c}$$

$$\frac{d\tilde{\rho}_{ge}}{dt} = -(\frac{\gamma}{2} + i\delta)\tilde{\rho}_{ge} + i\frac{\Omega_e g}{2}(\rho_{ee} - \rho_{gg}) \tag{A.4d}$$

where the rotating frame coherences are $\tilde{\rho}_{eg} = \rho_{eg}e^{i\delta t}$ and $\tilde{\rho}_{ge} = \rho_{ge}e^{-i\delta t}$. The effects of spontaneous decay have been included through the relation $\left(\frac{d\rho_{eg}}{dt}\right)_{spon} = -\gamma\rho_{eg}/2$, where γ is the spontaneous decay rate of the excited state.

In the steady state the elements of the density matrix do not change in time. When considering the scattering force, we are interested in the population of the excited state and so setting equation A.4a to zero, we find

$$\rho_{ee} = \frac{\Omega^2}{\gamma^2 + 4\delta^2 + 2\Omega^2}. \tag{A.5}$$

Appendix B

Convergence of the Gerchberg-Saxton algorithm

We consider an input plane (containing the SLM) and an output (or Fourier) plane. The field in the output plane is the Fourier transform of the field in the input plane. The input amplitude and the output target intensity are fixed; the output phase is free and the input phase is controllable. The Gerchberg–Saxton algorithm (Section 7.2.1) refines the input phase pattern each iteration in an attempt to produce an intensity distribution in the output plane which is closer to the target than in the previous step. It is not immediately clear that the algorithm should converge; here we outline one proof of convergence from [95].

Define the following quantities at the k-th step of the algorithm:

1. $G_k(u) = |G_k(u)|\, e^{i\phi_k(u)} = \mathcal{F}\{g_k(x)\}$ is the field in the Fourier plane as propagated from the input plane;
2. $G'_k(u) = |F(u)|\, e^{i\phi_k(u)}$ is the modified field in the Fourier plane: the phase of $G_k(u)$ is retained and the amplitude is replaced in order to satisfy the Fourier plane constraint;
3. $g'_k(u) = |g'_k(u)|\, e^{i\theta'_k(x)} = \mathcal{F}^{-1}\{G'_k(u)\}$ is the backwards-propagated field in the input plane, calculated as the inverse Fourier transform of the modified field in the Fourier plane;
4. $g_{k+1}(x) = |f(x)|\, e^{i\theta_{k+1}(x)} = |f(x)|\, e^{i\theta'_k(x)}$ is the backwards-propagated field modified to satisfy the input plane constraint.

At the k-th interation, the squared error in the Fourier plane is

$$\begin{aligned} E_{F,k}^2 &= \sum_u |G_k(u) - G'_k(u)|^2 \\ &= \sum_x |g_k(x) - g'_k(x)|^2, \end{aligned} \tag{B.1}$$

where the second line follows from Parseval's theorem. In the input plane, the squared error is

$$E_{I,k}^2 = \sum_x |g_{k+1}(x) - g'_k(x)|^2. \tag{B.2}$$

The objects g_k and g_{k+1} satisfy the input plane constraint that the amplitude is constant and as defined at the start of the algorithm. At any point x, $g_{k+1}(x)$ is the closest value to $g'_k(x)$ that satisfies the input plane constraints, and so $|g_{k+1}(x) - g'_k(x)| \leq |g_k(x) - g'_k(x)|$, i.e.

$$E_{I,k}^2 \leq E_{F,k}^2. \tag{B.3}$$

This shows that the squared error in the input plane at the k-th step (i.e. after propagating the modified field back to the input plane) is less than or equal to the squared error in the Fourier plane (of the modified field) at the k-th step; for convergence, the squared error in the Fourier plane at the k+1-th step must be less than or equal to both of these.

Applying Parseval's theorem to Equation B.2,

$$E_{Ik}^2 = \sum_u |G_{k+1}(u) - G'_k(u)|^2. \tag{B.4}$$

As before, since G'_k and G'_{k+1} both satisfy the Fourier plane constraint and at any point u, $G'_{k+1}(u)$ is the closest value to $G_{k+1}(u)$ that satisfies that constraint, $\left|G_{k+1}(u) - G'_{k+1}(u)\right| \leq |G_{k+1}(u) - G'_k(x)|$. Therefore $E_{F,k+1}^2 \leq E_{I,k}^2$ and, combining this with Equation B.3:

$$E_{F,k+1}^2 \leq E_{I,k}^2 \leq E_{F,k}^2 \tag{B.5}$$

i.e. the squared error can only decrease or stay the same at each iteration.

Appendix C

Simulation code

This Appendix details code used to simulate interferometric cooling. Due to image size, the code begins on the following page.

```
##########################################################################################################
########       PYTHON CODE FOR SIMULATING INTERFEROMETRIC COOLING FOLLOWED BY RAMAN VELOCIMETRY     ########
########                                  Rachel Gregory, June 2015                                   ########
##########################################################################################################



from numpy import *
from pylab import *
import sys
from joblib import Parallel, delayed


# Define the function which calculates the state amplitudes after a pulse, for particular momentum and m_F
sublevel.
# Dipole matrix element, Rabi frequency, Stark shift calculations from code by Alex Dunning


def pulse(cg0,ce0,t0,tau,z,p,phi,laserson,dirn,delta):
    #sigma- dipole matrix elements for Zeeman sublevel n
    DMe1M=DMeValsMinus[0,z]
    DMe2M=DMeValsMinus[1,z]
    DMe3M=DMeValsMinus[2,z]
    DMe4M=DMeValsMinus[3,z]
    DMe5M=DMeValsMinus[4,z]
    DMe6M=DMeValsMinus[5,z]
    #sigma+ dipole matrix elements for Zeeman sublevel n
    DMe1P=DMeValsPlus[0,z]
    DMe2P=DMeValsPlus[1,z]
    DMe3P=DMeValsPlus[2,z]
    DMe4P=DMeValsPlus[3,z]
    DMe5P=DMeValsPlus[4,z]
    DMe6P=DMeValsPlus[5,z]
    #----------------------SINGLE PHOTON RABI FREQUENCIES-----------------------
    #M/P indicates sigma-/+ and 1/2 indicates beam 1/2)
    Omega_eM2 = -(sqrt(DMe1M**2+DMe3M**2+DMe5M**2)/hbar)*E(laserson*I_2)
    Omega_gM1 = -(sqrt(DMe2M**2+DMe4M**2+DMe6M**2)/hbar)*E(laserson*I_1)
    Omega_eM1 = -(sqrt(DMe1M**2+DMe3M**2+DMe5M**2)/hbar)*E(laserson*I_1)
    Omega_gM2 = -(sqrt(DMe2M**2+DMe4M**2+DMe6M**2)/hbar)*E(laserson*I_2)
    Omega_eP2 = -(sqrt(DMe1P**2+DMe3P**2+DMe5P**2)/hbar)*E(laserson*I_2)
    Omega_gP1 = -(sqrt(DMe2P**2+DMe4P**2+DMe6P**2)/hbar)*E(laserson*I_1)
    Omega_eP1 = -(sqrt(DMe1P**2+DMe3P**2+DMe5P**2)/hbar)*E(laserson*I_1)
    Omega_gP2 = -(sqrt(DMe2P**2+DMe4P**2+DMe6P**2)/hbar)*E(laserson*I_2)
    #----------------------RESONANT RAMAN RABI FREQUENCY (Omega_R)----------------------
    Omega_R =
(-(abs((DMe1M*DMe2M)+(DMe3M*DMe4M))/hbar)*E(laserson*I_1))*(-1./hbar)*E(laserson*I_2)/(2*Delta)*cos(THETA)**2
+
(-(abs((DMe1P*DMe2P)+(DMe3P*DMe4P))/hbar)*E(laserson*I_1))*(-1./hbar)*E(laserson*I_2)/(2*Delta)*sin(THETA)**2
    #----------------------LIGHT SHIFTS-----------------------
    #light shifts due to sigma- components (of both beams)



    Omega_e_ACM = Omega_eM2**2/(4*Delta) + Omega_eM1**2/(4*(Delta-omega_eg))
    Omega_g_ACM = Omega_gM1**2/(4*Delta) + Omega_gM2**2/(4*(Delta+omega_eg))
    #light shifts due to sigma+ components (of both beams)
    Omega_e_ACP = Omega_eP2**2/(4*Delta) + Omega_eP1**2/(4*(Delta-omega_eg))
    Omega_g_ACP = Omega_gP1**2/(4*Delta) + Omega_gP2**2/(4*(Delta+omega_eg))
    #total light shifts of each hyperfine level
    Omega_e_AC = Omega_e_ACP*sin(THETA)**2 + Omega_e_ACM*cos(THETA)**2
    Omega_g_AC = Omega_g_ACP*sin(THETA)**2 + Omega_g_ACM*cos(THETA)**2
    #Relative light shift \delta^{AC}
    delta_AC = Omega_g_AC - Omega_e_AC
    #Raman detuning \delta
    delta_12 = laserson*omega_2 - laserson*omega_1 + delta - (omega_eg + dirn*p*k_eff/m +
hbar*k_eff**2/(2*m)) + (z-2.0)*zshift

    Omega_r = sqrt(Omega_R**2 + (delta_12 - delta_AC)**2)
    Omega_eff = Omega_R

    stheta = Omega_eff/Omega_r
    ctheta = -(delta_12 - delta_AC)/Omega_r

    # Calculate the state amplitudes
```

```
    # Calculate the state amplitudes

    ce1=exp(-i*(Omega_e_AC + Omega_g_AC)*tau*0.5)*exp(-i*delta_12*tau*0.5)*(ce0*(cos(Omega_r*tau*0.5) -
i*ctheta*sin(Omega_r*tau*0.5)) + cg0*exp(-i*(delta_12*t0 + phi))*(-i*stheta*sin(Omega_r*tau*0.5)))

    cg1=exp(-i*(Omega_e_AC + Omega_g_AC)*tau*0.5)*exp(i*delta_12*tau*0.5)*(cg0*(cos(Omega_r*tau*0.5) +
i*ctheta*sin(Omega_r*tau*0.5)) + ce0*exp(i*(delta_12*t0+phi))*(-i*stheta*sin(Omega_r*tau*0.5)))

    return cg1, ce1


# Experimental parameters

# Single photon detuning
Delta = 2*pi*15.0e9
# Rubidium 85 atomic mass
m = 85*1.66e-27
h = 6.63e-34
hbar = h/(2*pi)


# Frequencies of Raman beams
omega_1 = 2*pi*(384.230406373e12 - 1.264888e9)
omega_2 = 2*pi*(384.230406373e12 + 1.770843e9)
i=1j
# splitting of hyperfine levels of 5S1/2
omega_eg =  2*pi*3.035732439e9

# k vector for counterpropagating beams


k_eff = (omega_1+omega_2)/(3e8)

# reduced dipole matrix element
RME = 3.58e-29
#set THETA to 0 for sigma+ transitions, pi for sigma- or pi/2 for lin-perp-lin
THETA = 0

e0 = 8.85e-12
# Raman beam intensities
I_1 = 2.5e4
I_2 = 2.5e4
c = 2.99e8

zshift=-500.e3
#sigma+ transition dipole matrix elements (see Steck)
# (DMeValsAB indicates transition FROM F=A TO F'=B)
DMeVals33 = array([sqrt(25./216),sqrt(5./36),sqrt(5./36),sqrt(25./216),sqrt(5./72)])
DMeVals23 = array([sqrt(2./135),sqrt(2./45),sqrt(4./45),sqrt(4./27),sqrt(2./9)])
DMeVals32 = array([sqrt(10./189),sqrt(2./63),sqrt(1./63),sqrt(1./189),0.])
DMeVals22 = array([sqrt(7./54),sqrt(7./36),sqrt(7./36),sqrt(7./54),0.])
DMeVals21 = array([sqrt(3./10),sqrt(3./20),sqrt(1./20),0.,0.])
DMeVals34 = array([sqrt(3./56),sqrt(3./28),sqrt(5./28),sqrt(15./56),sqrt(3./8)])
DMeValsPlus = RME*array(((DMeVals33),(DMeVals23),(DMeVals32),(DMeVals22),(DMeVals34),(DMeVals21)))

#sigma- transition dipole matrix elements (see Steck)
DMeVals33 = array([-sqrt(5./72),-sqrt(25./216),-sqrt(5./36),-sqrt(5./36),-sqrt(25./216)])
DMeVals23 = array([sqrt(2./9),sqrt(4./27),sqrt(4./45),sqrt(2./45),sqrt(2./135)])
DMeVals32 = array([0.,sqrt(1./189),sqrt(1./63),sqrt(2./63),sqrt(10./189)])
DMeVals22 = array([0.,-sqrt(7./54),-sqrt(7./36),-sqrt(7./36),-sqrt(7./54)])
DMeVals21 = array([0.,0.,sqrt(1./20),sqrt(3./20),sqrt(3./10)])
DMeVals34 = array([sqrt(3./8),sqrt(15./56),sqrt(5./28),sqrt(3./28),sqrt(3./56)])
DMeValsMinus = RME*array(((DMeVals33),(DMeVals23),(DMeVals32),(DMeVals22),(DMeVals34),(DMeVals21)))

def E(I): return sqrt(2*I/(c*e0))

# Pi pulse duration: found by running a Rabi flop simulation with the same parameters as used in this code

pi_time = 1.02e-6

#momentum distribution
p_width1 = 3.0*hbar*k_eff
p_width2 = 13*hbar*k_eff
p_max = 40*hbar*k_eff
```

```
p_min = -40*hbar*k_eff
Nps = 1600

# number of places to move within array of momentum values upon two-photon excitation
p_step_raman = abs(int(Nps/ceil((p_max-p_min)/(hbar*k_eff))))



# limits of number of places to move within array of momentum values upon spontaneous decay
# this is half that of the two-photon transition
p_range_spontem=p_step_raman/2

# define the arrays for momentum values and momentum distribution
p_vals = linspace(p_min,p_max,Nps)
p_init = 0.77*e**(-0.5*(p_vals/p_width1)**2) + 0.23*e**(-0.5*(p_vals/p_width2)**2)
p_init /= max(p_init)
Nz=5


# Set the duration of beamsplitter and mirror pulses
BStime=0.5*pi_time
Mtime=pi_time

# Set the number of repetitions of the interferometer
n_reps = 12

# Array to hold the momentum distributions at each step
p_array=zeros((n_reps,Nps))

# Initialise counter for number of atoms lost if their momentum becomes greater than the range we are
simulating
lost=0

# Switch beam direction after each interferometer 1=yes 0=no
switchbeams=1

if switchbeams==1:
    direcs=[1.,-1.]
else:
    direcs=[1.,1,]


# Set beamsplitter phases
phases=[pi/4,-pi/4]
dphi=+0.0
dT=-0.0e-6
d_offset=2*pi*180.e3#960


###  SIMULATE THE INTERFEROMETER
for n_r in range(n_reps):
    # call up the required momentum distribution
    p_dist=p_init if n_r==0 else p_array[n_r-1]

    c_e0=zeros(Nps)
    c_g0=ones(Nps)
    # arrays for excited and ground state populations




    gp=zeros((Nps,Nz))
    ep=zeros((Nps,Nz))
    # arrays for momentum distributions at various stages
    pd2=zeros(Nps)
    pd2_g=zeros(Nps)
    pd2_e=zeros(Nps)
    # set direction of two-photon impulse
    k_eff = direcs[n_r%2]*(omega_1+omega_2)/(3e8)
    # set dwell times
```

```
    T=0.6e-6
    d_int=-2*pi*120e3

    # do the interferometer for all momentum and m_F values

    for n in range(Nps-p_step_raman):
        for z in range(Nz):
            #beamsplitter
            c_g1,c_e1 = pulse(c_g0[n],c_e0[n],0.,BStime,z,p_vals[n],phases[0],1.,1.,d_int)
            #free evolution
            c_g2,c_e2 = pulse(c_g1,c_e1,BStime,T,z,p_vals[n],0.,0.,1.,0.)
            #beamsplitter pulse
            c_g3,c_e3 = pulse(c_g2,c_e2,BStime+T,BStime,z,p_vals[n],phases[1],1.,1.,d_int)

            # add the population to the population arrays
            gp[n,z]=p_dist[n]*(c_g3.real**2+c_g3.imag**2)
            ep[n,z]=p_dist[n]*(c_e3.real**2+c_e3.imag**2)
        # store the excitation profile (for monitoring only)
        exc_prof[n]=p_dist[n]*sum(ep[n,:])/(sum(ep[n,:])+sum(gp[n,:]))
        # amend the momentum distribution for atoms in the ground state
        pd2_g[n]+=p_dist[n]*sum(gp[n,:])/(sum(ep[n,:])+sum(gp[n,:]))
        # and the same for the excited state, but move atoms to their new place in the momentum distribution
        pd2_e[n+direcs[n_r%2]*p_step_raman]+=p_dist[n]*sum(ep[n,:])/(sum(ep[n,:])+sum(gp[n,:]))


    #spontaneous emission

    #number of spontaneous emissions (most atoms [99%] are in |g> after 5 spont ems)
    n_se=5
    # simulate a number of atoms with each momentum value which is proportional to the momentum dsitribution
    n_atoms=(1e3)*pd2_e


    for ns in range(n_se):
        pd2_e_se=zeros(Nps)     #array holding the atoms that decay back into the |e> state
        for n in range(Nps):
            n_a=int(n_atoms[n]) # how many atoms to simulate at this momentum value
            for a in range(int(n_a)):
                decay_prob = 1 if uniform(0,1)>0.53 else 0      #probability to decay to |g>
                if decay_prob==1:
```

```
                    #randomly generate theta for z component of momentum
                    theta_SE = uniform(-pi, pi)
                    p_add  = int(p_range_spontem*cos(theta_SE))
                    #if it falls outside the current momentum range, that atom is lost
                    if (n+p_add<0 or n+p_add>(Nps-1)):
                        lost+=1
                    #else, put an atom in the bin corresponding to its new momentum in the |g> state
                    else:
                        pd2_g[n+p_add]+=(1./n_a)*pd2_e[n]
                        n_atoms[n]-=1       #number of atoms in |e> state decreases by 1
                else:
                    #randomly generate the fraction of hbar*k that is lost (from -1 to 1)

                    #randomly generate theta for z component of momentum
                    theta_SE = uniform(-pi, pi)
                    p_add  = int(p_range_spontem*cos(theta_SE))
                    #if it falls outside the current momentum range, that atom is lost
                    if (n+p_add<0 or n+p_add>(Nps-1)):
                        lost+=1
                    #else, put an atom in the bin corresponding to its new momentum in the |e> state
                    else:
                        pd2_e_se[n+p_add]+=(1./n_a)*pd2_e[n]
                        n_atoms[n+p_add]+=1
                        n_atoms[n]-=1

    # create the new momentum distribution and store in the array
    pd2=pd2_g+pd2_e_se
    p_array[n_r,:]=pd2


# After the interferometric cooling has taken place, we now need to perform the Raman velocimetry
```

```
#Raman scans

# Number of frequencies to scan
ndeltas=100
deltas=linspace(-0.4,0.8,ndeltas)*2*pi*1e6

# scan the momentum distribution after the last repetition of the interferometer
p_dist=p_array[-1,:]

# set the Raman beam intensities
I_1 = 0.45e4
I_2 = 0.45e4


# probe pulse duration
```

```
time_on=200.e-6

# define Raman scan function
def ramanscan(detuning):
    zsum_after=0
    ramanpops=zeros((2,Nps,Nz))
    ramanstates=zeros((2,Nps))
    for n in range(Nps):
        for z in range(Nz):
            c_g,c_e = pulse(1.,0.,0.,time_on,z,p_vals[n],0.,1.,1.,detuning-d_offset)#-d_offset)
            ramanpops[0,n,z] = abs(c_g)**2*p_dist[n]
            ramanpops[1,n,z] = abs(c_e)**2*p_dist[n]
    for z in range(Nz):
        ramanstates[0,z]=sum(ramanpops[0,:,z])/(sum(ramanpops[0,:,z]) + sum(ramanpops[1,:,z]))
        ramanstates[1,z]=sum(ramanpops[1,:,z])/(sum(ramanpops[0,:,z]) + sum(ramanpops[1,:,z]))

    return sum(ramanstates[1,:])/(sum(ramanstates[0,:]) + sum(ramanstates[1,:]))

# Raman scan of modified momentum distribution
after_scan=Parallel(n_jobs=-1, verbose=3)(delayed(ramanscan)(deltas[nd]) for nd in range(ndeltas))

# Raman scan of initial momentum distribution
p_dist=p_init

init_scan=Parallel(n_jobs=-1, verbose=3)(delayed(ramanscan)(deltas[nd]) for nd in range(ndeltas))
```

Appendix D

The off-resonance error propagator

The simplest form of the propagator in the presence of off resonance errors is given in Equation 4.8:

$$R(\theta,\phi) = \exp\left(-i\frac{\theta}{2}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)\right); \tag{D.1}$$

whilst compact, it is not easy to visualise the operations taking place. Here we expand the exponential and derive a form which better suits analysis of the behaviour of errors.

We first expand the exponential as a Taylor series and divide into sums over odd and even integers:

$$\exp\left(-i\frac{\theta}{2}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)\right) = \sum_{\text{odd}\, n} \frac{\left(\frac{-i\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi + f\sigma_z)\right)^n}{n!} + \sum_{\text{even}\, n} \frac{\left(\frac{-i\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi + f\sigma_z)\right)^n}{n!}. \tag{D.2}$$

The Pauli matrix terms are simplified by using the relations $\sigma_i\sigma_j = \epsilon_{ijk}\sigma_k$ and $\sigma_i\sigma_i = \sigma_0$. The even powers simplify to

$$\begin{aligned}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)^{2n} &= \sigma_0(1+f^2)^n \\ &= \sigma_0((1+f^2)^{1/2})^{2n}\end{aligned} \tag{D.3}$$

and the odd powers to

$$\begin{aligned}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)^{2n+1} &= (\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)(1+f^2)^n \\ &= (\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)(1+f^2)^{-1/2}((1+f^2)^{1/2})^{2n+1}.\end{aligned} \tag{D.4}$$

Putting these into Equation D.2,

$$\begin{aligned}R(\theta,\phi) = &\sum_{n=0}^{\infty} \frac{(-1)^n(\frac{\theta}{2}\sqrt{1+f^2})^{2n}}{2n!} \\ &- i(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)(1+f^2)^{-1/2}\sum_{n=0}^{\infty}\frac{(-1)^n(\frac{\theta}{2}\sqrt{1+f^2})^{2n+1}}{(2n+1)!}.\end{aligned} \tag{D.5}$$

The sums in these two terms are the Taylor expansions of $\cos(\frac{\theta}{2}\sqrt{1+f^2})$ and $\sin(\frac{\theta}{2}\sqrt{1+f^2})$; hence the propagator may be written as

$$R(\theta,\phi) = \sigma_0 \cos\left(\frac{\theta}{2}\sqrt{1+f^2}\right) - i(1+f^2)^{-1/2}(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y + f\sigma_z)\sin\left(\frac{\theta}{2}\sqrt{1+f^2}\right). \tag{D.6}$$

Evaluating this propagator for small f, we find

$$\begin{aligned}R(\theta,\phi) \approx &\sigma_0\cos\frac{\theta}{2} - i\sin\frac{\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi) \\ &- if\sigma_z\sin\frac{\theta}{2} \\ &- \frac{f^2\theta}{4}(\sigma_0\sin\frac{\theta}{2} + i\cos\frac{\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi)) \\ &+ i\frac{f^2}{2}\sin\frac{\theta}{2}(\sigma_x\cos\phi + \sigma_y\sin\phi).\end{aligned} \tag{D.7}$$

References

[1] D. A. Steck. Rubidium 85 D Line Data, 2013. *http://steck.us/alkalidata/*.

[2] A.J. Dunning, R.L. Gregory, J. Bateman, N. Cooper, M. Himsworth, T. G. M. Freegarde, and J. A. Jones. Composite pulses for interferometry in a thermal atom cloud. *Phys. Rev. A*, 90(033608), 2014.

[3] A. J. Dunning, R. L. Gregory, J. Bateman, M. Himsworth, and T. G. M. Freegarde. Interferometric Laser Cooling of Atomic Rubidium. *Phys. Rev. Lett.*, 115:073004, 2015.

[4] T.W. Hansch and A.L. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13(1):68 – 69, 1975.

[5] W. D. Phillips and H. Metcalf. Laser deceleration of an atomic beam. *Phys. Rev. Lett.*, 48:596–599, 1982.

[6] 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.

[7] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of bose-einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, 1995.

[8] 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.*, 75:3969–3973, 1995.

[9] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, 80:885–964, 2008.

[10] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.

[11] C.S. Adams, M. Sigel, and J. Mlynek. Atom Optics. *Physics Reports*, (240):143, 1994.

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

[13] S. M. Dickerson, J. M. Hogan, A. Sugarbaker, D. M. S. Johnson, and M. A. Kasevich. Multiaxis Inertial Sensing with Long-Time Point Source Atom Interferometry. *Phys. Rev. Lett.*, 111:083001, 2013.

[14] T. L. Gustavson, P. Bouyer, and M. A. Kasevich. Precision Rotation Measurements with an Atom Interferometer Gyroscope. *Phys. Rev. Lett.*, 78:2046–2049, 1997.

[15] N. Gaaloul, H. Ahlers, T.A. Schulze, Y. Singh, S.T. Seidel, W. Herr, W. Ertmer, and E. Rasel. Quantum tests of the equivalence principle with atom interferometry. *Acta Astronautica*, 67(9 - 10):1059 – 1062, 2010.

[16] P. Hamilton, M. Jaffe, J. M. Brown, L. Maisenbacher, B. Estey, and H. Müller. Atom interferometry in an optical cavity. *Phys. Rev. Lett.*, 114:100405, 2015.

[17] R. Bouchendira, P. Cladé, S. Guellati-Khélifa, F. Nez, and F. Biraben. New determination of the fine structure constant and test of the quantum electrodynamics. *Phys. Rev. Lett.*, 106:080801, 2011.

[18] G. Rosi, F. Sorrentino, L. Cacciapuoti, M. Prevedelli, and G. M. Tino. Precision measurement of the Newtonian gravitational constant using cold atoms. *Nature*, 510:518, 2013.

[19] P. Hamilton, A. Zhmoginov, F. Robicheaux, J. Fajans, J. S. Wurtele, and H. Müller. Antimatter Interferometry for Gravity Measurements. *Phys. Rev. Lett.*, 112:121102, 2014.

[20] W. F. Holmgren, R. Trubko, I. Hromada, and A. D. Cronin. Measurement of a Wavelength of Light for Which the Energy Shift for an Atom Vanishes. *Phys. Rev. Lett.*, 109:243004, 2012.

[21] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krger. Matter-wave interferometry in a double well on an atom chip. *Nature Physics*, 1:57, 2005.

[22] Ch.J. Borde. Atomic interferometry with internal state labelling. *Physics Letters A*, 140(1):10 – 12, 1989.

[23] J. M. McGuirk, M. J. Snadden, and M. A. Kasevich. Large Area Light-Pulse Atom Interferometry. *Phys. Rev. Lett.*, 85:4498–4501, 2000.

[24] D. L. Butts, K. Kotru, J. M. Kinast, A. M. Radojevic, B. P. Timmons, and R. E. Stoner. Efficient broadband Raman pulses for large-area atom interferometry. *J. Opt. Soc. Am. B*, 30(4), 2013.

[25] J. F. Barry, D. J. McCarron, E. B. Norrgard, M. H. Steinecker, and D. DeMille. Magneto-optical trapping of a diatomic molecule. *Nature*, 512:286, 2014.

[26] E. S. Shuman andJ. F. Barry and D. DeMille. Laser cooling of a diatomic molecule. *Nature*, 467:820, 2010.

[27] I. Manai, R. Horchani, H. Lignier, P. Pillet, D. Comparat, A. Fioretti, and M. Allegrini. Rovibrational Cooling of Molecules by Optical Pumping. *Phys. Rev. Lett.*, 109:183001, 2012.

[28] M. Viteau, A. Chotia, M. Allegrini, N. Bouloufa, O. Dulieu, D. Comparat, and P. Pillet. Optical Pumping and Vibrational Cooling of Molecules. *Science*.

[29] M. Weitz and T. W. Hänsch. Frequency-independent laser cooling based on interferometry. *Europhys. Lett.*, 49:302, 2000.

[30] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005.

[31] M. Bander and C. Itzykson. Group Theory and the Hydrogen Atom. *Rev Mod Phys*, 38(2):330, 1966.

[32] J. D. Jackson. *Classical Electrodynamics*. Wiley, New York, 1962.

[33] C. Cohen-Tannoudji, B. Diu, and F Laloë. *Quantum Mechanics*. Wiley, New York, 1977.

[34] P. A. M. Dirac. The Quantum Theory of the Emission and Absorption of Radiation. *Proc. Roy. Soc. A*, 114:243, 1927.

[35] J. J. Sakurai and J. Napolitano. *Modern Quantum Mechanics*. Addison-Wesley, 2010.

[36] S. Stenholm. The semiclassical theory of laser cooling. *Rev. Mod. Phys.*, 58(3):699, 1986.

[37] J. T. Mendona and Hugo Teras. *Physics of Ultra-Cold Matter: Atomic Clouds, Bose-Einstein Condensates and Rydberg Plasmas.* Springer, 2013.

[38] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping.* Springer, 2002.

[39] T. Bergeman, G. Erez, and H. J. Metcalf. Magnetostatic trapping fields for neutral atoms. *Phys Rev A*, 35(4):1535, 1987.

[40] P D.. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf. Observation of Atoms Laser Cooled below the Doppler Limit. *Phys. Rev. Lett*, 61:169, 1988.

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

[42] P. Berman, editor. *Atom Interferometry.* Academic Press, 1997.

[43] R. P. Feynman, F. Vernon, and R. Hellwarth. Geometrical Representation of the Schrödinger Equation for Solving Maser Problems. *J App Phys*, 28(1):2–5, 1957.

[44] P. A. M. Dirac. Richard P. Feynman and Frank L. Vernon and Robert W. Hellwarth. *J. Appl. Phys*, 28:49, 1957.

[45] A. J. Dunning. *Coherent atomic manipulation and cooling using composite optical pulse sequences.* PhD thesis, University of Southampton, UK, 2014.

[46] J. M. McGuirk, G. T. Foster, J. B. Fixler, M. J. Snadden, and M. A. Kasevich. Sensitive absolute-gravity gradiometry using atom interferometry. *Phys. Rev. A*, 65:033608, 2002.

[47] T. L. Gustavson, P. Bouyer, and M. A. Kasevich. Precision Rotation Measurements with an Atom Interferometer Gyroscope. *Phys. Rev. Lett.*, 78:2046–2049, 1997.

[48] M. Kasevich and S. Chu. Atomic interferometry using stimulated Raman transitions. *Phys. Rev. Lett.*, 67:181–184, 1991.

[49] F. Sorrentino, Y-H Lien, G. Rosi, L. Cacciapuoti, M. Prevedelli, and G. M. Tino. Sensitive gravity-gradiometry with atom interferometry: progress towards an improved determination of the gravitational constant. *New J. Phys.*, 12:095009, 2010.

[50] J. B. Fixler, G. T. Foster, J. M. McGuirk, and M. A. Kasevich. Atom Interferometer Measurement of the Newtonian Constant of Gravity. *Science*, 315:74, 2007.

[51] R. Bouchendira, P. Cladé, S. Guellati-Khélifa, F. Nez, and F. Biraben. New Determination of the Fine Structure Constant and Test of the Quantum Electrodynamics. *Phys. Rev. Lett.*, 106:080801, 2011.

[52] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.*, 81:1051–1129, 2009.

[53] J. M. Hogan, D. M. S. Johnson, and M. A. Kasevich. Light-pulse atom interferometry. In E. Arimondo, W. Ertmer, W. P. Schleich, and E. M. Rasel, editors, *Atom Optics and Space Physics*, page 411, 2009.

[54] B. Barrett, P.-A. Gominet, E. Cantin, L. Antoni-Micollier, A. Bertoldi, B. Battelier, P. Bouyer, J. Lautier, and A. Landragin. Mobile and remote inertial sensing with atom interferometers. *Proceedings of the International School of Physics "Enrico Fermi*, pages 493–555, 2014.

[55] Chiow, S. and Kovachy, T. and Chien, H-C and Kasevich, M. A. $102\hbar k$ large area atom interferometers. *Phys. Rev. Lett.*, 107:130403, 2011.

[56] M. G. Bason, M. Viteau, N. Malossi, P. Huillery, E. Arimondo, D. Ciampini, R. Fazio, V. Giovannetti, R. Mannella, and O. Morsch. High fidelity quantum driving. *Nature Physics*, 8:147, 2012.

[57] D. Meshulach and Y. Silberberg. Coherent quantum control of two-photon transitions by a femtosecond laser pulse. *Nature*, 396:239, 1998.

[58] S. Gulde, M. Riebe, G. P. T. Lancaster, C. Becher, J. Eschner, H. Häffner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt. Implementation of the DeutschJozsa algorithm on an ion-trap quantum computer. *Nature*, 421(48), 2002.

[59] V. Nebendahl, H. Häffner, and C. F. Roos. Optimal control of entangling operations for trapped-ion quantum computing. *Phys. Rev. A*, 79:012312, Jan 2009.

[60] C. D. Aiello, M. Hirose, and P. Cappellaro. Composite-pulse magnetometry with a solid-state quantum sensor. *Nature Communications*, 4(1419), 2013.

[61] C. M. Tesch and R. de Vivie-Riedle. Quantum Computation with Vibrationally Excited Molecules. *Phys. Rev. Lett.*, 89:157901, Sep 2002.

[62] M. H. Levitt. Composite pulses. *Prog. Nucl. Magn. Reson. Spectrosc.*, 18(61), 1986.

[63] S. Wimperis. Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments. *J. Mag. Reson.*, Series A 109(221), 1994.

[64] M. H. Levitt and R. Freeman. NMR population inversion using a composite pulse. *Journal of Magnetic Resonance (1969)*, 33(2):473 – 476, 1979.

[65] A. J. Shaka, J. Keeler, T. Frenkiel, and R. Freeman. An improved sequence for broadband decoupling: WALTZ16. *J. Mag. Reson.*, 52:335, 1983.

[66] J. Bateman. *Novel schemes for the optical manipulation of atoms and molecules*. PhD thesis, University of Southampton, UK, 2009.

[67] M. Himsworth. *Coherent manipulation of ultracold Rubidium*. PhD thesis, University of Southampton, UK, 2009.

[68] A. L. Schawlow and C. H. Townes. Infrared and optical masers. *Phys. Rev.*, 112:1940–1949, 1958.

[69] M. W. Fleming and A. Mooradian. Fundamental line broadening of singlemode (gaal)as diode lasers. *Applied Physics Letters*, 38(7):511–513, 1981.

[70] N. Cooper, J. Bateman, A. J. Dunning, and T. G. M. Freegarde. Actively stabilized wavelength-insensitive carrier elimination from an electro-optically modulated laser beam. *J. Opt. Soc. Am. B*, 29(4):646–649, Apr 2012.

[71] N. Cooper. *Novel techniques for the trapping and manipulation of ultracold atoms* . PhD thesis, University of Southampton, UK, 2014.

[72] H. Luo, K. Li, D. Zhang, T. Gao, and K. Jiang. Multiple side-band generation for two-frequency components injected into a tapered amplifier. *Opt. Lett*, 38:1161–1163, 2013.

[73] N. Cooper, J. Woods, J. Bateman, A. J. Dunning, and T. G. M. Freegarde. Stabilized fiber-optic machzehnder interferometer for carrier-frequency rejection. *Applied Optics*, 52(23):5713–5717, 2013.

[74] S. Bergamini, B. Darquié, M. Jones, L. Jacubowiez, A. Browaeys, and P. Grangier. Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator. *J. Opt. Soc. Am. B*, 21(11):1889–1894, 2004.

[75] G. D. Bruce, J. Mayoh, G. Smirne, L. Torralbo-Campo, and D. Cassettari. A smooth, holographically generated ring trap for the investigation of superfluidity in ultracold atoms. *Phys. Scr*, 2011:014008, 2010.

[76] C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsworth, and A. Kuhn. Control and manipulation of cold atoms in optical tweezers. *New J. Phys.*, 14:073051, 2012.

[77] H. Ohadi, A. Dreismann, Y. G. Rubo, F. Pinsker, Y. del Valle-Inclan Redondo, S. I. Tsintzos, Z. Hatzopoulos, P. G. Savvidis, and J. J. Baumberg. Spontaneous Spin Bifurcations and Ferromagnetic Phase Transitions in a Spinor Exciton-Polariton Condensate. *Phys. Rev. X*, 5:031002, 2015.

[78] R W Gerchberg and W O Saxton. A Practical Algorithm for the Determination of Phase from Image and Diffraction Plane Pictures. *Optik*, 35(2):237, 1972.

[79] M. Pasienski and B. DeMarco. A high-accuracy algorithm for designing arbitrary holographic atom traps. *Opt. Express*, 16(3):2176–2190, 2008.

[80] Vortex stagnation problem in iterative fourier transform algorithms.

[81] G. D. Bruce, M. Y. H. Johnson, E. Cormack, D. A. W. Richards, J. Mayoh, and D. Cassettari. Feedback-enhanced algorithm for aberration correction of holographic atom traps. *J. Phys. B: At. Mol. Opt. Phys.*, 48:115303, 2015.

[82] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard. Phase-space density in the magneto-optical trap. *Phys. Rev. A*, 52:1423–1440, 1995.

[83] IRIDIS page at TOP500 ranking website. http://www.top500.org/system/178203. Accessed: 14/5/2015.

[84] I. Solomon. Rotary Spin Echoes. *Phys. Rev. Lett.*, 2:301–302, Apr 1959.

[85] W. Rakreungdet, J. H. Lee, K. F. Lee, B. E. Mischuck, E. Montano, and P. S. Jessen. Accurate microwave control and real-time diagnostics of neutral-atom qubits. *Phys. Rev. A*, 79:022316, Feb 2009.

[86] S. J. J. M. F. Kokkelmans, B. J. Verhaar, K. Gibble, and D. J. Heinzen. Predictions for laser-cooled rb clocks. *Phys. Rev. A*, 56:R4389–R4392, 1997.

[87] H K Cummins and J A Jones. Use of composite rotations to correct systematic errors in NMR quantum computation. *New J. Phys.*, 2:6, 2000.

[88] H. K. Cummins, G. Llewellyn, and J. A. Jones. Tackling systematic errors in quantum logic gates with composite rotations. *Phys. Rev. A*, 67:042308, 2003.

[89] R. Tycko. Broadband Population Inversion. *Phys. Rev. Lett.*, 51:775–777, 1983.

[90] Ryan, C. A. and Hodges, J. S. and Cory, D. G. Robust decoupling techniques to extend quantum coherence in diamond. *Phys. Rev. Lett.*, 105:200402, 2010.

[91] I. D. Leroux, M. H. Schleier-Smith, and V. Vuletic. Implementation of Cavity Squeezing of a Collective Atomic Spin. *Phys. Rev. Lett*, 104(073602), 2010.

[92] *Levy Statistics and Laser Cooling.*

[93] Monte Carlo simulation of the atomic master equation for spontaneous emission.

[94] Klaus Mølmer, Yvan Castin, and Jean Dalibard. Monte Carlo wave-function method in quantum optics. *J. Opt. Soc. Am. B*, 10(3):524–538, 1993.

[95] J. R. Fienup. Phase–retrieval algorithms for a complicated optical system. *Applied Optics*, 32(10), 1993.