University of Southampton

Velocimetry, trapping and optimal coherent manipulation of atomic rubidium

Author:
Max S. Carey

Supervisors:
Dr. Tim Freegarde &
Dr. Andrea Bertoldi

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

in the

Faculty of Engineering and Physical Sciences
School of Physics & Astronomy

26 June, 2020

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
School of Physics & Astronomy

Doctor of Philosophy

Velocimetry, Trapping And Optimal Coherent Manipulation Of Atomic Rubidium

by Max S. Carey

Ultracold samples of laser-cooled atoms are quantum systems over which modern atomic physicists can exert exquisite control. Largely decoupled from their environment, they can act as near-ideal test masses for inertial sensors based on atom interferometry and are well suited to experiments in coherent control of quantum systems that probe the fundamental nature of quantum mechanics and pave the way for practical quantum simulation and computation.

This thesis details a series of experimental results that arise from the coherent control of rubidium atoms with laser light, focusing on the interplay between these interactions and atomic velocities; the laser frequency an atom 'sees' is Doppler shifted according to its velocity, while conservation of momentum dictates that, in exchanging photons with a laser, an atom's velocity is altered.

Coherent light–atom interactions can thus be tailored either to measure or to narrow the spread of velocities in an ultracold atomic gas. Alternatively, it can be desirable to design interactions that are homogeneous across a large spread of atomic velocities. All of these aspects are explored in this thesis.

The velocity-sensitive interactions that lie at the heart of atom-interferometric inertial sensors are reexamined in a manner that yields considerable insight into the underlying processes and culminates in a novel, precise and elegant technique for measuring the velocity of ultracold atoms that is used to reveal the Gaussian nature of the velocity spread in a cloud with an effective temperature of 18.7(6) µK, undistorted by artefacts that plague other methods.

Furthermore, optimal control techniques are applied to the problem of coherently and uniformly manipulating the quantum states of atoms in an ensemble with a large spread of velocities, even subject to variations in laser intensity. A broadband inversion pulse is demonstrated to change the internal state of 99.8(3) % of atoms in a ∼35 µK ensemble and — for the first time — this technique is used to optimise an entire atom interferometry sequence, yielding a threefold enhancement in the measurement contrast.

Finally, a version of *grey molasses* cooling — in which atoms accumulate in velocity-dependent 'dark' states, narrowing their momentum spread and increasing their phase space density — is demonstrated with phase-coherent cooling beams; dark states that exist in this system prove to be particularly resilient to the spatially varying light shifts that are present in an optical dipole trap, and this is used both to enhance the number of atoms loaded into such a trap — by a factor of 7 compared to loading from a conventional optical molasses — and to further cool them once they are loaded in a technique that has promising prospects for the rapid production of ultracold, trapped, atoms.

Declaration of authorship

I, Max Sebastian Carey, declare that the thesis entitled *Velocimetry, trapping and optimal coherent manipulation of atomic rubidium* and the work presented in the thesis are 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 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: [1], [2] and [3].

Signed: _____

Date: _____

List of publications

- **M. Carey**, M. Belal, M. Himsworth, J. Bateman, and T. Freegarde, "Matterwave interferometric velocimetry of cold Rb atoms," Journal of Modern Optics **65**, 657 (2018).

- J. C. Saywell, I. Kuprov, D. Goodwin, **M. Carey**, and T. Freegarde, "Optimal control of mirror pulses for cold-atom interferometry," Physical Review A **98**, 023625 (2018).

- **M. Carey**, J. Saywell, D. Elcock, M. Belal, and T. Freegarde, "Velocimetry of cold atoms by matter-wave interferometry," Physical Review A **99**, 023631 (2019).

- J. C. Saywell*, **M. Carey***, M. Belal, I. Kuprov, and T. Freegarde, "Optimal control of Raman pulse sequences for atom interferometry," Journal of Physics B **53**, 085006 (2020). **equal contribution*

- D. S. Naik, H. Eneriz-Imaz, **M. Carey**, T. Freegarde, F. Minardi, B. Battelier, P. Bouyer, and A. Bertoldi, "Loading and cooling in an optical trap via hyperfine dark states," Physical Review Research **2**, 013212 (2020).

- A. Bertoldi, C.-H. Feng, H. Eneriz, **M. Carey**, D. S. Naik, J. Junca, X. Zou, D. O. Sabulsky, B. Canuel, P. Bouyer, and M. Prevedelli, "A control hardware based on a field programmable gate array for experiments in atomic physics," Review of Scientific Instruments **91**, 033203 (2020).

Acknowledgements

To single out just a few people for particular gratitude feels almost cruel as, in retrospect, the omissions are bound to stand out far more prominently than those names that make it onto the list. Of course, I am immensely indebted to all of the people below, but likely no more so than to many others whom my thesis-addled mind will forget.

Nonetheless, I will start by thanking my supervisor in Southampton, Tim Freegarde, from whom I have not just received immense help and support throughout my PhD, but have also inherited a way of thinking about and approaching questions in physics that is both eminently practical and extremely rewarding.

It seems a cop out to say that exactly the same can be said of my supervisor in Bordeaux, Andrea Bertoldi. Nonetheless, this is true, and I have been incredibly lucky to be supervised by two physicists from whom I have learned so much. Thank you both.

I am similarly grateful to Matt Himsworth and Devang Naik who, on opposite sides of the English channel, offered immense day-to-day support in the laboratory; what competence I now possess as an experimental atomic physicist I owe largely to them.

Mohammad Belal — or just Belal — I feel that we embarked on a journey into the unknown wilderness of atomic physics together. We both started with much to learn, and it has been a pleasure to learn together. Thank you for sharing this journey with me, and also for all that you had to teach me about lasers (of course), but also about how to navigate a career in academia.

To all the PhD students I have worked alongside, past and present, I am grateful not just for your support and un-bolted-down optical components but also for your continued friendship. There are far too many to list but, in particular, I count myself lucky to have shared office and lab space with Jonathan Woods, Matt Aldous, Andrei Dragomir, David Elcock, Jack Saywell, Chester Camm, Grigor Kuyumjyan, Grégoire Lefèvre and Hodei Eneriz. I would like nothing more than for these friendships to persist once I emerge from this thesis bubble. To those whom I have left out, be it due to space constraints or absent mindedness: please forgive me, and I hope that our friendships persist beyond this omission!

Two PhD students who should be singled out are Alex Dunning and Rachel Gregory, my direct predecessors. We did not have much overlap or time for direct knowledge transfer, but the fact that I was able to reconstruct the experiment and achieve the results presented here in the time I had in Southampton is a direct testament to the quality of both of your theses.

Thanks also to all of the Physics and Astronomy technical support staff in Southampton, in particular Justin, Gary and Gareth — for your technical support, and for a common sense attitude to health and safety that kept me honest but never impeded research progress — and the entire staff of the mechanical workshop — not just for the components you produced during my PhD, but also for your guidance.

That I had the opportunity to embark upon this incredible journey at all is thanks to the support and funding of the Dstl UK-France PhD scheme. In particular, I would like to thank Susannah Jones and Duleep Wickramasinghe for their support that has always enabled and never hindered.

Outside of the academic community, I must of course thank Hannah, she of a thousand freckles; for diligently reproducing the increasingly random designs that I would ask you to tattoo on my body as this thesis drove me slowly mad, I thank you. Jason Isaacs must also be mentioned, in addition to Stephen Fry and those members of Fairport Convention who are still kicking around.

My sincerest gratitude, though, goes to my parents — Mark and Sue — without whose emotional and financial support I would never have made it this far. I am more indebted to you both than I could ever hope to express. Thank you.

Contents

List of figures

List of abbreviations

AOM	acousto-optic modulator. 73–78, 90, 93, 106, 107, 147, 148, 163–165, 174
AWG	arbitrary waveform generator. 90, 147, 148
BAT	broadband atomic transfer. 146–152, 157
BEC	Bose-Einstein condensate. 3, 141, 144, 187
BIARO	condensation de Bose-Einstein et interférométrie atomique dans un résonateur optique de grande finesse. 4, 161, 162, 176
BSC	beamsplitter cube. 76–79
CtS	course to steer. 206–209
DBR	distributed Bragg reflector. 74, 77
DFB	distributed feedback. 76, 77, 165, 172, 173
EDFA	erbium-doped fibre amplifier. 163, 164, 174
EOM	electro-optic modulator. 73, 76–81, 93, 143, 147, 164, 172–175
FFT	fast Fourier transform. 104–106, 111, 113, 116–118
FM	frequency modulation. 74, 165
FOI	Faraday optical isolator. 76, 77
FORT	far off-resonant optical dipole trap. xiii, 44, 161–163, 169–171, 174–180, 183, 187
GRAPE	gradient ascent pulse engineering. 143, 146, 148, 154–158, 185
HDS	hyperfine dark state. 171–174, 176–181, 187
HWP	half waveplate. 75, 76, 78
LMT	large momentum transfer. 3, 142, 144, 145, 149, 150, 156–158, 185, 186
MOT	magneto-optical trap. xv, 47–49, 73–75, 78, 81–84, 86–93, 115, 161, 165–167, 169–174, 176, 178, 185
NMR	nuclear magnetic resonance. 3, 10, 16, 17, 141–143, 147, 152
OBE	optical Bloch equation. 16, 17, 19, 23, 38–40, 45, 109, 191, 192, 202
PBSC	polarising beamsplitter cube. 74–78, 85, 164
PDH	Pound-Drever-Hall. 163, 164, 174
PI	proportional-integral. 81
PID	proportional-integral-differential. 76, 164
PM	polarisation maintaining. 73, 76, 78, 85, 165
PMT	photo-multiplier tube. 73, 82–84
PP	point-to-point. 144, 151, 152, 157
QED	quantum electrodynamics. 2, 130, 132, 162

RF radio frequency. 76–78, 81, 143, 147, 164, 172, 174, 175, 180
RWA rotating wave approximation. 15, 16, 32
RYA Royal Yachting Association. 206, 207, 209

TA tapered amplifier. 76, 77, 165
TOF time of flight. 119, 120, 163, 173, 174, 180

UHV ultra-high vacuum. 161, 164, 198
UR universal rotation. 144, 152, 153

VSCPT velocity-selective coherent population trapping. 51–53

ZDS Zeeman dark state. 171, 176–178

To "Brexit" Britain, for making the process of writing a thesis just that little bit more stressful.

1 Introduction

Begin at the beginning and go on till you come to the end: then stop.

— Lewis Carroll, *Alice in Wonderland*

That matter should be constructed from innumerable indivisible atoms has been hypothesised and thought upon since the times of ancient Greece, though no means was available to demonstrate their existence. Post-enlightenment, tantalising hints accumulated from the study of chemistry [4], and from the Brownian motion of particles visible under a microscope [5]. Further, indirect, evidence following the discovery of the electron [6, 7] and the scattering experiments of Rutherford *et al.* [8] allowed theorists to establish a full quantum model of the atoms that they now understood to be very much divisible, with rich and complex internal structures.

From a modern perspective, it is quite remarkable that so much was discovered about the structure and quantum nature of atoms from the experiments that followed during the early part of the 20^{th} century. It is a great demonstration of human ingenuity that, prior to the invention of the laser — which provides both a high-resolution tool for resolving spectral features [9, 10] and a means of cooling and trapping atoms in order to mitigate Doppler broadening and observe quantum coherence effects [11–13] — enough knowledge could be generated to facilitate the invention of the laser.

Nonetheless, modern atomic physics experiments owe much to molecular beam systems developed in this period. In 1950 Ramsey proposed, and later conducted, an experiment that had all the ingredients of what would now be recognised as a *matter-wave interferometer* [14, 15]. A molecular beam underwent two interactions with oscillating magnetic fields, one preparing a superposition between two internal states and one resolving the phase acquired by the superposition during a period of free evolution. This method was employed to measure the transition frequency, and is still employed in atomic clocks to this day [16–18].

The advent of the laser allowed for such experiments to be conducted on optical transitions, with pulses of laser light replacing

the oscillating magnetic fields. At optical frequencies, the absorption of a photon entails a non-negligible change in an atom's momentum. Atoms can thus be placed in a superposition of momentum states that propagate along discrete spatial paths [19–21], allowing for atom interferometers that enclose a finite area and offer the potential for precision measurement of external forces. Modern atom interferometers have since been used in the precise measurement of Earth's gravity [22–27], testing of fundamental physical laws such as Einstein's equivalence principle [28–32] or quantum electrodynamics (QED) [33–36], and, lately, highly sensitive inertial sensors [37–42].

These interferometers have evolved a lot since Ramsey's seminal experiment. As the atomic wavefunction is separated into discrete wavepackets by the optical fields, additional pulses are employed to overlap these wavepackets again and resolve coherence at the output which is then a measure of the differential phase acquired between paths that share common start and end points. The long interrogation times and ultracold atomic samples employed in state-of-the-art atom interferometers allow these differential measurements to be made exceedingly precisely.

When optical fields are employed in an interferometer like Ramsey's, the atomic wavepackets are allowed to separate for a period of free evolution before a final measurement. This makes the interferometer sensitive to the absolute atomic velocity but, as the wavepackets no longer overlap perfectly at the output, the coherence of the output signal rapidly deteriorates. It is perhaps for this reason that not much attention has been given to this sort of interferometer with velocity-sensitive pulses, despite the fact that it has considerable applications and can elicit much physical insight. Part of this thesis seeks to redress that balance.

Weitz and Hänsch identified that, in being sensitive to the absolute atomic velocity, the final measurement of such an interferometer also results in a velocity-dependent impulse that could be employed to cool atoms and molecules that are impervious to other laser cooling methods [43]. This was only recently experimentally verified by Dunning *et al.* [44].

This thesis builds upon that work, employing a Ramsey-type interferometer to measure the velocity distribution of laser cooled atoms [1, 3]. In doing so, we explore the insights that such velocity-sensitive interferometers have to offer; giving them some well-deserved attention in a trend that would since appear to have

developed traction [45, 46].

All of these atom interferometers depend upon coherent control of atomic states with optical fields. The *fidelity* of these coherent operations is reduced by systematic inhomogeneities, and this ultimately affects the contrast of the interferometer output. A second strand of this thesis focuses on the application of optimal control techniques developed in the field of nuclear magnetic resonance (NMR) spectroscopy [47–50] to the problem of maximising the fidelity of these operations over the range of inhomogeneities present in our interferometry experiment in order to realise an enhancement in signal contrast [2, 51].

The optimised pulses that result are a natural extension of the *composite pulses* — also originating from NMR studies [52, 53] — that have previously been applied to atom interferometry for the same purpose [54–56]. Optimal control algorithms, however, offer many more exciting possibilities.

There are many problems in atomic physics experiments that lend themselves to such treatment, and it is a growing field of study. Notable successes include optimal production [57, 58] and manipulation [59, 60] of cold atom Bose-Einstein condensates (BECs), and the realisation of high-fidelity quantum logic gates [61, 62] and state preparation [63] for quantum computing. We demonstrate the potential for realising high-fidelity atom interferometry pulses that are well suited to *large momentum transfer (LMT)* atom interferometry [55, 64–67], and present a strategy for optimising an entire interferometry sequence.

An alternative approach is taken to the problem of optimising the production of cold atom BECs in a third, and final, line of investigation pursued at the Institut d'Optique, Bordeaux. The path to reaching the critical phase-space density required to produce a cold atom BEC is typically long, and it is not uncommon to lose two or more orders of magnitude between the number of atoms in the initial sample and the final condensate [68, 69].

There is room for optimising every step in this process, and here we focus on employing advanced laser cooling techniques to increase the number of atoms loaded into an optical dipole trap and subsequently cool them within the trapping potential. As this cooling is optical, the energy is carried away by spontaneously emitted photons not — as in conventional evaporative cooling methods — by

the loss of atoms from the trap [70, 71]. This method thus has the potential to cool large numbers of atoms to quantum degeneracy.

All of the experiments described in this thesis were performed on the 780 nm D_2 line of ^{85}Rb and ^{87}Rb atoms. Rather than dedicating any one section to describing these systems in full, the relevant energy levels and optical properties are introduced in sufficient detail as and when they are pertinent. For a more complete picture, the reader is referred to the works of Daniel Steck [72, 73], who has catalogued both structures in exquisite detail.

1.1 Thesis outline

This thesis ties these strands together in three parts:

Part I establishes a theoretical background, and encompasses a lot of the physical insights gleaned from the experimental work presented later. Chapter 2 relates to the coherent control of atoms with laser light that lies at the heart of much of the following work. Chapter 3 explores what happens when this coherence is lost via spontaneous emission, the basis of laser cooling and trapping techniques. Chapter 4 covers the application of all these techniques to atom interferometry, while Chapter 5 then takes some time to consider nautical navigation in conjunction with highly precise mechanical clocks. Inspired by a paper currently in preparation, this latter chapter has more than just curiosity value, tying in closely to the former and proving relevant to the rest of the work.

Part II presents the experimental results acquired in a basement in Southampton. Chapter 6 details the experimental apparatus that is found there, before Chapter 7 describes the results of the interferometric velocimetry experiments. A further dive into the world of boats and clocks cannot be resisted in Chapter 8, in order to introduce the theory of optimal control that is used to obtain the results presented in Chapter 9.

Part III then forms something of an epilogue, with Chapter 10 detailing the BIARO experiment in Bordeaux on which the work to optimise the loading of an optical dipole trap, presented in Chapter 11, was conducted. Finally, Chapter 12 takes a moment to reflect, discussing the results of the previous two Parts in a broader context and considering their future development.

Part I

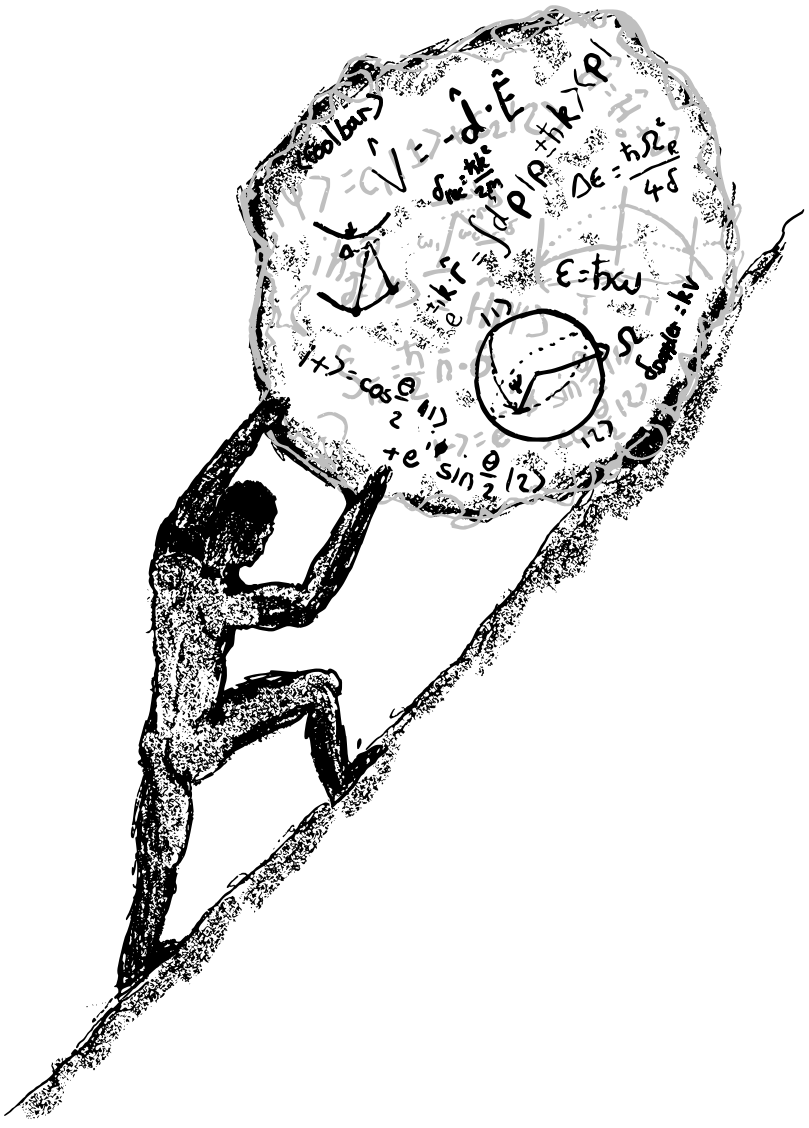

Theoretical background

Coherent manipulation of atoms with laser light

2

The most curious part of the thing was, that the trees and the other things round them never changed their places at all: however fast they went, they never seemed to pass anything. 'I wonder if all the things move along with us?' thought poor puzzled Alice.

— Lewis Carroll, *Through the Looking Glass*

(And into a rotating frame?)

We will start our attempt to ground the results of this thesis in some sort of theoretical footing by considering the interaction of light with a two-level atom. We will review textbook results to show that the three-level Raman system at the heart of all of the atom interferometry and coherent control experiments of later chapters can, in certain limits, be well described by the two-level picture that is developed here and employed throughout the rest of this work.

2.1 Spin precession in a magnetic field

We begin by considering the mathematics of a spin-$^1/_2$ particle in the presence of an external, time-independent magnetic field. This well-studied problem provides a nice 'sand-box' in which to explore the quantum dynamics of any 2-level Hamiltonian, and so it is worthwhile spending some time to develop an intuition for it.

Using $|\pm\rangle$, the eigenstates for the projection of the particle's spin along a unit vector $\boldsymbol{e_z}$, as basis states, we can define operators and eigenstates for spin along the orthogonal Cartesian axes:

$$\hat{S}_x = \frac{\hbar}{2}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \hat{S}_y = \frac{\hbar}{2}\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \hat{S}_z = \frac{\hbar}{2}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.1)$$

and

$$|x,\pm\rangle = \frac{1}{\sqrt{2}}\left(|+\rangle \pm |-\rangle\right) \ , \ |y,\pm\rangle = \frac{1}{\sqrt{2}}\left(|+\rangle \pm i\,|-\rangle\right). \tag{2.2}$$

It is convenient to introduce the Pauli spin matrices

$$\hat{\sigma}_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \ , \ \hat{\sigma}_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \ , \text{ and } \hat{\sigma}_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.3}$$

so that the Cartesian spin operators can be succinctly expressed as $\hat{S}_i = \frac{\hbar}{2}\hat{\sigma}_i$. A spin operator in the direction of any unit vector $\boldsymbol{n}$ can then be defined as

$$\hat{S}_{\boldsymbol{n}} = \frac{\hbar}{2}\boldsymbol{n}\cdot\hat{\boldsymbol{\sigma}} = \frac{\hbar}{2}\left(n_x\hat{\sigma}_x + n_y\hat{\sigma}_y + n_z\hat{\sigma}_z\right). \tag{2.4}$$

[1] Conventions here vary on phase factors. The salient point is that an azimuthal rotation induces a phase shift between $|+\rangle$ and $|-\rangle$, and their respective amplitudes are constrained by normalisation and the fact that $\boldsymbol{n} = \pm\boldsymbol{e_z}$ when $\theta = 0, \pi$.

It is trivial to verify that this has eigenvectors[1] which can be conveniently expressed in terms of standard polar θ and azimuthal ϕ angles as

$$|n,+\rangle = \cos\frac{\theta}{2}\,|+\rangle + e^{i\phi}\sin\frac{\theta}{2}\,|-\rangle\,, \tag{2.5a}$$

$$|n,-\rangle = e^{-i\phi}\sin\frac{\theta}{2}\,|+\rangle - \cos\frac{\theta}{2}\,|-\rangle\,. \tag{2.5b}$$

Now consider how an arbitrary spin state behaves in a constant magnetic field $\boldsymbol{B} = B\boldsymbol{e_z}$. The Hamiltonian will be of the form

$$\hat{H}_S = -\gamma B\hat{S}_z, \tag{2.6}$$

where γ quantifies how the particle's magnetic dipole moment scales in proportion to its spin angular momentum. The Hamiltonian commutes with $\hat{S}_z$, so they share a common eigenbasis in which the time evolution of a general spin state can be evaluated. The Hamiltonian is time-independent, with the associated time evolution operator

$$\mathcal{U}(t,0) = \exp\left(-i\frac{\hat{H}_S t}{\hbar}\right) \ = \ \exp\left(-i\frac{(-\gamma B t)\hat{S}_z}{\hbar}\right). \tag{2.7}$$

A spin state aligned along the direction defined by angles θ_0 and ϕ_0 at $t = 0$ then evolves as

$$\begin{aligned}
|\Psi(t)\rangle &= \mathcal{U}(t,0)\,|\Psi(0)\rangle \\
&= e^{-i(-\gamma Bt)\hat{S}_z/\hbar}\left(\cos\frac{\theta_0}{2}\,|+\rangle + e^{i\phi_0}\sin\frac{\theta_0}{2}\,|-\rangle\right) \\
&= \cos\frac{\theta_0}{2}e^{i\gamma Bt/2}\,|+\rangle + e^{i\phi_0}\sin\frac{\theta_0}{2}e^{-i\gamma Bt/2}\,|-\rangle \\
&= e^{i\gamma Bt/2}\left(\cos\frac{\theta_0}{2}\,|+\rangle + e^{i(\phi_0-\gamma Bt)}\sin\frac{\theta_0}{2}\,|-\rangle\right).
\end{aligned} \tag{2.8}$$

By inspection we see that after a time t the spin vector has rotated to a new position defined by angles [2]

[2] The overall phase, of course, being irrelevant to the measurable state.

$$\begin{aligned}
\theta(t) &= \theta_0, \\
\phi(t) &= \phi_0 - \gamma Bt,
\end{aligned} \tag{2.9}$$

That is to say it has rotated an angle $-\gamma Bt$ about the z axis. By rotational symmetry, the magnetic field can be chosen to lie along any direction $\boldsymbol{n}$, with the corresponding Hamiltonian and time evolution operator

$$\hat{H}_S = -\gamma B\hat{S}_{\boldsymbol{n}} \text{ and} \tag{2.10}$$

$$\mathcal{U}(t,0) = \exp\left(-i\frac{(-\gamma Bt)\hat{S}_{\boldsymbol{n}}}{\hbar}\right) \tag{2.11}$$

rotating the spin vector about $\boldsymbol{n}$ at the angular *Larmor frequency* $\Omega_L = -\gamma B$ in a manner completely analogous to the classical precession of a magnetic moment under the influence of the torque exerted by the magnetic field.

2.2 Pseudo-spins and the Bloch sphere

In studying the specific example of a spin-1/2 particle in a magnetic field, we have given ourselves a toolbox with which to treat any general two-level Hamiltonian $\hat{H}$. Hermiticity dictates that such a Hamiltonian will take the most general form

$$\hat{H} = \begin{pmatrix} a & c \\ c* & b \end{pmatrix}, \quad a, b \in \mathbb{R}, c \in \mathbb{C}. \tag{2.12}$$

Without loss of generality, a, b and c can be redefined in terms of four real constants $r_{0,1,2,3}$ in the form

$$\begin{aligned} a &= \epsilon_0 + r_3, \\ b &= \epsilon_0 - r_3, \\ c &= r_1 - ir_2. \end{aligned} \tag{2.13}$$

Separating like this allows the Hamiltonian to be written in the suggestive manner

$$\hat{H} = \epsilon_0 \mathbb{I} + r_1 \hat{\sigma}_x + r_2 \hat{\sigma}_y + r_3 \hat{\sigma}_z. \tag{2.14}$$

We can then define a vector $\mathbf{\Omega}$ with components $\frac{2}{\hbar} r_{1,2,3}$, and hence length $\Omega = \frac{2}{\hbar}\sqrt{r_1^2 + r_2^2 + r_3^2}$, so that this can be rewritten as

$$\hat{H} = \epsilon_0 \mathbb{I} + \frac{\hbar}{2} \mathbf{\Omega} \cdot \hat{\boldsymbol{\sigma}}. \tag{2.15}$$

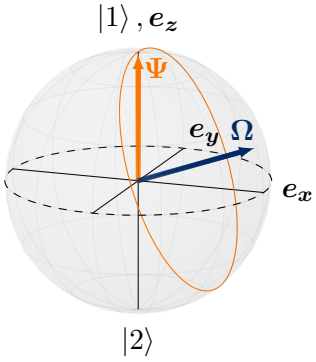

Figure 2.1: The Bloch sphere offers a natural way to visualise the state of a two-level system. A time-independent Hamiltonian, factorised as per Equation (2.15), causes the vector representation $\mathbf{\Psi}$ of the quantum state in this space to precess about the vector $\mathbf{\Omega}$ at an angular frequency determined by its amplitude $\Omega = |\mathbf{\Omega}|$.

Disregarding the global phase, any state of a two-level quantum system $|\Psi\rangle = c_1 |1\rangle + c_2 |2\rangle$ can be completely represented by three real numbers. We can use the expectation values of the Pauli spin matrices for this purpose,

$$\begin{aligned} u &\equiv \langle\Psi| \hat{\sigma}_x |\Psi\rangle = c_1^* c_2 + c_1 c_2^*, \\ v &\equiv \langle\Psi| \hat{\sigma}_y |\Psi\rangle = i(c_1 c_2^* - c_1^* c_2), \\ w &\equiv \langle\Psi| \hat{\sigma}_z |\Psi\rangle = c_1^* c_1 - c_2^* c_2, \end{aligned} \tag{2.16}$$

allowing the state to be described by a real-valued *pseudo-spin* vector $\mathbf{\Psi} \equiv (u, v, w)$. Assuming $|\Psi\rangle$ to be normalised, the normalisation condition $|c_1|^2 + |c_2|^2 = 1$ dictates that $\mathbf{\Psi}$ has unit length, mapping the accessible Hilbert space onto the surface of a sphere in the three-dimensional space it inhabits. This sphere is known as the *Bloch sphere*; it originates from NMR physics [74], and is reminiscent of the Poincaré sphere on which the Stokes polarisation parameters are often mapped [75]. The vector $\mathbf{\Psi}$ is occasionally referred to as the *Bloch vector* although, as $|\Psi\rangle$ and $\mathbf{\Psi}$ are different representations of the same quantum state, it is common to use the term *state vector* to describe them both interchangeably.

The z coordinate w determines the 'population difference' between the two basis states $|1\rangle$ and $|2\rangle$. When the Hamiltonian is of the form $\hat{H} = \hat{H}_0 + \hat{V}$ — with $\hat{H}_0$ being the Hamiltonian of a 'bare' system for which the eigenstates and energies are known and $\hat{V}$ representing additional interaction terms — then this basis can be chosen to be the eigenbasis of $\hat{H}_0$, giving w a clear physical significance. The physical significance of u and v will be revisited later, but for now we can just consider them coordinates in an abstract Cartesian space in which we can visualise the following, particularly simple, geometric picture for the time evolution of $|\Psi\rangle$ enacted by the Hamiltonian $\hat{H}$ [76].

We have seen that the Hamiltonian in Equation (2.10) causes a physical spin to precess about the magnetic field vector at a characteristic angular Larmor frequency Ω_L. By comparison, noting the definition of $\hat{S}_n$ from Equation (2.4), one can deduce that the Hamiltonian in Equation (2.15) will cause the pseudo-spin representation of the state vector $\mathbf{\Psi}$ to precess about $\mathbf{\Omega}$ at an angular frequency $\Omega = |\mathbf{\Omega}|$. Figure 2.1 visualises this and, indeed, in this frame the Schrödinger equation can be recast in the form of the equation of motion for a classical magnetic moment precessing due to the torque exerted by a magnetic field

$$\frac{\mathrm{d}\mathbf{\Psi}}{\mathrm{d}t} = \mathbf{\Omega} \times \mathbf{\Psi}. \tag{2.17}$$

For this reason, $\mathbf{\Omega}$ is referred to in the literature as the *field vector*.[3]

This reduces the problem of solving any time-independent two-level Hamiltonian to a matter of factorising it into the form given in Equation (2.14) and obtaining $\mathbf{\Omega}$. The Bloch sphere picture then makes the time evolution of a general state trivial to visualise, and the eigenstates of the Hamiltonian are just those of the operator $(\mathbf{\Omega}\cdot\hat{\boldsymbol{\sigma}})/\Omega$ — the projection of the pseudo-spin along the field vector — with eigenenergies $\epsilon_0 \pm \frac{\hbar}{2}\Omega$.

[3] We have defined the field vector with the opposite sign here to that found in a lot of other literature. Note that, while a real magnetic moment precesses clockwise around a magnetic field vector, our field vector is defined such that our pseudo-spin precesses counter-clockwise about it. This is the sign convention adopted by Metcalf and van der Straten among others, and is purely a matter of taste.

2.3 The interaction of light with a two-level atom

With all of this in mind, let us turn our attention to the specific problem of a two-level atom interacting with monochromatic light. An atom with just two internal energy states is a hypothetical convenience that can never be realised physically. However, the two-level model lends itself well to mathematical analysis and, in

the case of an atom interacting with near-resonant light, it provides a good approximation of physical phenomena.

So, consider a model atom that has just two eigenstates of its internal Hamiltonian $\hat{H}_0$, ground $|1\rangle$ and excited $|2\rangle$, with associated eigenenergies $-\epsilon_0/2$ and $\epsilon_0/2$ respectively.

$$\begin{aligned} \hat{H}_0 |1\rangle &= -\frac{\epsilon_0}{2} |1\rangle = -\hbar\frac{\omega_0}{2} |1\rangle , \\ \hat{H}_0 |2\rangle &= \frac{\epsilon_0}{2} |2\rangle = \hbar\frac{\omega_0}{2} |2\rangle . \end{aligned} \tag{2.18}$$

The most general wavefunction, including its time dependence, is then

$$|\Psi(t)\rangle = c_1(t) |1\rangle + c_2(t) |2\rangle , \tag{2.19}$$

where we have incorporated all of the time dependence of the state into the coefficients $c_{1,2}(t)$. In the absence of any perturbation, their evolution is then determined by the Schrödinger equation for the bare atom

$$i\hbar\frac{\partial}{\partial t} |\Psi(t)\rangle = \hat{H}_0 |\Psi(t)\rangle \tag{2.20}$$

yielding

$$c_{1,2}(t) = c_{1,2}(0)e^{\pm i\omega_0 t/2}. \tag{2.21}$$

A perturbation of the form

$$\hat{V}(t) = -\hat{\boldsymbol{d}} \cdot \boldsymbol{E}(t) \tag{2.22}$$

arises as a result of an optical electric field, large enough to be treated classically. This corresponds to the energy of the atom's quantum dipole moment $\hat{\boldsymbol{d}}$ in a classical electric field $\boldsymbol{E}(t)$. If we assume the dipole moment is due to the displacement of a single electron then we can rewrite $\hat{\boldsymbol{d}} = -e\hat{\boldsymbol{r}}$ where $\hat{\boldsymbol{r}}$ is the operator for the electron's position relative to the atomic centre of mass.

We have assumed that $\boldsymbol{E} = \boldsymbol{E}(t)$ is a function of time only; that is to say it is uniform over the region of space in which it interacts with the atom. This is the dipole approximation and is valid provided that the electromagnetic wavelength is much larger than the

atom's radius, as it is for a typical optical transition in a ground-state atom.[4]

[4]Note that this picture only describes the internal dynamics of the atom, but does not explain the effect of radiation pressure on the atom's centre of mass that we shall consider explicitly in Section 2.6.1.

For light linearly polarised in the x direction we have $\boldsymbol{E}(t) = E\cos(\omega t - \phi)\boldsymbol{e_x}$ so that

$$\hat{V}(t) = e\,\hat{x}E\cos(\omega t - \phi). \tag{2.23}$$

Here ϕ has been used to specify the phase of the light field and $\hat{x}$ is the operator for the component of the electron's displacement along $\boldsymbol{e_x}$. Note that $\hat{x}$ has negative parity so that, provided $|2\rangle$ and $|1\rangle$ have well-defined parity, the matrix elements $\langle 1|\,\hat{x}\,|1\rangle = \langle 2|\,\hat{x}\,|2\rangle = 0$. Hence the matrix form of $\hat{V}(t)$ in the $(|1\rangle, |2\rangle)$ basis can be written as

$$\begin{aligned}\hat{V}(t) &= \begin{pmatrix} 0 & \hbar\Omega_R\cos(\omega t - \phi) \\ \hbar\Omega_R\cos(\omega t - \phi) & 0 \end{pmatrix} \\ &= \hbar\Omega_R\cos(\omega t - \phi)\hat{\sigma}_x.\end{aligned} \tag{2.24}$$

This introduces the coupling constant

$$\Omega_R \equiv \langle 2|\,\hat{x}\,|1\rangle\,\frac{eE}{\hbar} = \langle 1|\,\hat{x}\,|2\rangle\,\frac{eE}{\hbar} \tag{2.25}$$

that we can take to be real for linearly polarised light that drives a $\Delta m = 0$ transition [77, §15.1].

The total Hamiltonian, including the perturbation, can then be written as

$$\hat{H} = \hat{H}_0 + \hat{V}(t) = \begin{pmatrix} -\hbar\omega_0/2 & \hbar\Omega_R\cos(\omega t - \phi) \\ \hbar\Omega_R\cos(\omega t - \phi) & \hbar\omega_0/2 \end{pmatrix} \tag{2.26}$$

that can be factorised as per Equation (2.15) to give

$$\hat{H} = -\frac{\hbar}{2}\omega_0\hat{\sigma}_z + \hbar\Omega_R\cos(\omega t - \phi)\hat{\sigma}_x. \tag{2.27}$$

This is in the form of Equation (2.15), but the time dependence is unedifying. In the Bloch sphere picture this corresponds to a field vector with a constant z component, whose x component is oscillating between $\pm 2\Omega_R$ at an angular frequency ω. Motivated

by the fact that this oscillation can be expressed as the sum of two counter-rotating vectors in the x–y plane, it is instructive to consider the form of the Hamiltonian in a frame rotating in phase with one of them.

Take the component that is rotating clockwise around the z axis. In part 2.1 it was demonstrated that the Hamiltonian in equation (2.6) causes a spin vector to precess about $\boldsymbol{e_z}$ clockwise at an angular frequency $-\gamma B$. Thus, borrowing from the time evolution operator in Equation (2.7), it is clear that the unitary operator

$$\mathcal{U}_R(t) \equiv \exp\left(-\frac{i\omega t}{2}\hat{\sigma}_z\right) \tag{2.28}$$

will act to rotate an arbitrary state vector $|\Psi(t)\rangle$ counter-clockwise about $\boldsymbol{e_z}$ at an angular frequency ω, just as it would in a frame rotating clockwise at the same frequency. We thus define our state vector in such a rotating frame to be

$$|\Psi_R(t)\rangle = \mathcal{U}_R(t)\,|\Psi(t)\rangle = \mathcal{U}_R(t)\mathcal{U}(t)\,|\Psi(0)\rangle\,, \tag{2.29}$$

where $\mathcal{U}(t)$ is the unitary time evolution operator for $\hat{H}$ in the rest frame. As this is not an inertial frame we can expect the physics to be different. However, plugging Equation (2.29) into the Schrödinger equation allows us to determine a new Hamiltonian $\hat{H}_R$ that is valid in this rotating frame,

$$\begin{aligned}
\hat{H}_R\,|\Psi_R(t)\rangle &= i\hbar\frac{\partial}{\partial t}\,|\Psi_R(t)\rangle \\
&= i\hbar\frac{\partial}{\partial t}\,[\mathcal{U}_R\mathcal{U}]\,|\Psi(0)\rangle \\
&= i\hbar\left(\frac{\partial\mathcal{U}_R}{\partial t}\mathcal{U} + \mathcal{U}_R\frac{\partial\mathcal{U}}{\partial t}\right)|\Psi(0)\rangle \\
&= i\hbar\frac{\partial\mathcal{U}_R}{\partial t}\mathcal{U}_R^\dagger\Big(\mathcal{U}_R\mathcal{U}\,|\Psi(0)\rangle\Big) + \mathcal{U}_R\Big(i\hbar\frac{\partial\mathcal{U}}{\partial t}\,|\Psi(0)\rangle\Big) \\
&= i\hbar\frac{\partial\mathcal{U}_R}{\partial t}\mathcal{U}_R^\dagger\,|\Psi_R(t)\rangle + \mathcal{U}_R\hat{H}\,|\Psi(t)\rangle \\
&= \left(i\hbar\frac{\partial\mathcal{U}_R}{\partial t}\mathcal{U}_R^\dagger + \mathcal{U}_R\hat{H}\mathcal{U}_R^\dagger\right)|\Psi_R(t)\rangle\,, \\
\Rightarrow \hat{H}_R &= i\hbar\frac{\partial\mathcal{U}_R}{\partial t}\mathcal{U}_R^\dagger + \mathcal{U}_R\hat{H}\mathcal{U}_R^\dagger.
\end{aligned} \tag{2.30}$$

Substituting (2.27) and (2.28) into (2.30) gives

$$\begin{aligned}\hat{H}_R &= i\hbar\left(-i\frac{\omega}{2}\hat{\sigma}_z\right) + \exp(-i\omega t\hat{\sigma}_z/2)\hat{H}\exp\left(i\omega t\hat{\sigma}_z/2\right) \\ &= \hbar\frac{\omega}{2}\hat{\sigma}_z - \hbar\frac{\omega_0}{2}\hat{\sigma}_z + \hbar\Omega_R\cos(\omega t-\phi)\mathcal{U}_R\hat{\sigma}_x\mathcal{U}_R^\dagger \\ &= \hbar\frac{(\omega-\omega_0)}{2}\hat{\sigma}_z + \hbar\Omega_R\cos(\omega t-\phi)\begin{pmatrix}0 & e^{-i\omega t}\\ e^{i\omega t} & 0\end{pmatrix}.\end{aligned} \tag{2.31}$$

A bit of algebra, expressing the cos function as a sum of complex exponentials, yields the following factorisation,

$$\begin{aligned}\hat{H}_R = \frac{\hbar}{2}(\omega-\omega_0)\hat{\sigma}_z + \frac{\hbar}{2}\Omega_R\left(\cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y\right) \\ + \frac{\hbar}{2}\Omega_R\left(\cos(2\omega t-\phi)\hat{\sigma}_x + \sin(2\omega t-\phi)\hat{\sigma}_y\right).\end{aligned} \tag{2.32}$$

This Hamiltonian still exhibits some time dependence, in a manner which could have been anticipated. Upon shifting into the frame of one rotating part of the linearly polarised oscillation, the counter-rotating part persists at twice the frequency.

Crucially, the $\hat{\sigma}_z$ component of the Hamiltonian — analogous to the magnetic field in Equation (2.6) — is reduced by the fictitious field arising from the shift to a non-inertial frame. On resonance, when $\omega = \omega_0$, the two cancel entirely and the term vanishes.

Defining the detuning $\delta \equiv \omega - \omega_0$ of the optical field from the angular transition frequency, the Hamiltonian can be separated into two parts corresponding to the two lines in Equation (2.32). The first part is time-independent and affects a precession of the state about a stationary field vector $\tilde{\boldsymbol{\Omega}} = \Omega_R\cos\phi\tilde{\boldsymbol{e}}_{\boldsymbol{x}} + \Omega_R\sin\phi\tilde{\boldsymbol{e}}_{\boldsymbol{y}} + \delta\tilde{\boldsymbol{e}}_{\boldsymbol{z}}$ at a pseudo-Larmor angular frequency of $\Omega = |\tilde{\boldsymbol{\Omega}}| = \sqrt{\Omega_R^2+\delta^2}$. We employ tildes to denote vectors in the rotating frame, but may drop them later on if the context is sufficiently obvious.

The remaining, time-dependent, part acts to rotate the state vector at angular frequency Ω_R about an axis that rotates in the x–y plane at 2ω. If $\omega \gg \Omega_R$ then this axis will rotate a full circle before the precession has advanced any considerable amount, and the net effect of the 'torque' will average out to 0. In this limit, it is a good approximation to drop the time-dependent terms from $\hat{H}_R$ in what is known as the *rotating wave approximation (RWA)*.

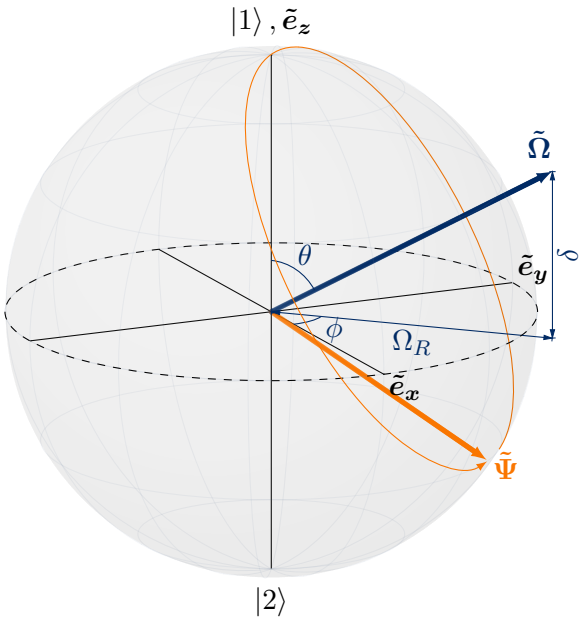

Figure 2.2: In the rotating frame, the state vector $\tilde{\boldsymbol{\Psi}}$ evolves according to $\frac{\partial \tilde{\boldsymbol{\Psi}}}{\partial t} = \tilde{\boldsymbol{\Omega}} \times \tilde{\boldsymbol{\Psi}}$, precessing about the field vector $\tilde{\boldsymbol{\Omega}}$ at an angular frequency $\Omega = \sqrt{\delta^2 + \Omega_R^2}$. On resonance $\delta = 0$, an atom starting in the ground state $|1\rangle$, will follow a great circle and reach the excited state $|2\rangle$ at the 'south' pole after a time $t = \pi/\Omega_R$. The further off resonance the optical field, the less far south the atom will get; the oscillations become faster while the enclosed area becomes smaller, with the speed of the state vector over the surface of the Bloch sphere only dependent on Ω_R.

Finally, then, the rotating frame Hamiltonian in the RWA reads

$$\begin{aligned} \hat{H}_R &= \frac{\hbar}{2}\left(\delta\hat{\sigma}_z + \Omega_R\cos(\phi)\hat{\sigma}_x + \Omega_R\sin(\phi)\hat{\sigma}_y\right) \\ &= \frac{\hbar}{2}\begin{pmatrix} \delta & \Omega_R e^{-i\phi} \\ \Omega_R e^{i\phi} & -\delta \end{pmatrix}. \end{aligned} \tag{2.33}$$

Sticking with the tildes to remind us that we are in the rotating frame, the corresponding evolution of the Bloch vector $\tilde{\boldsymbol{\Psi}} = \tilde{u}\tilde{\boldsymbol{e}}_{\boldsymbol{x}} + \tilde{v}\tilde{\boldsymbol{e}}_{\boldsymbol{y}} + \tilde{w}\tilde{\boldsymbol{e}}_{\boldsymbol{z}}$, as given by Equation (2.17), is

$$\frac{\mathrm{d}\tilde{u}}{\mathrm{d}t} = -\delta\tilde{v} + \Omega_R\sin(\phi)\tilde{w}, \tag{2.34a}$$

$$\frac{\mathrm{d}\tilde{v}}{\mathrm{d}t} = \delta\tilde{u} - \Omega_R\cos(\phi)\tilde{w}, \tag{2.34b}$$

$$\frac{\mathrm{d}\tilde{w}}{\mathrm{d}t} = -\Omega_R\sin(\phi)\tilde{u} + \Omega_R\cos(\phi)\tilde{v}. \tag{2.34c}$$

These are a form of the optical Bloch equations (OBEs), analogous to the Bloch equations employed in NMR spectroscopy, and they

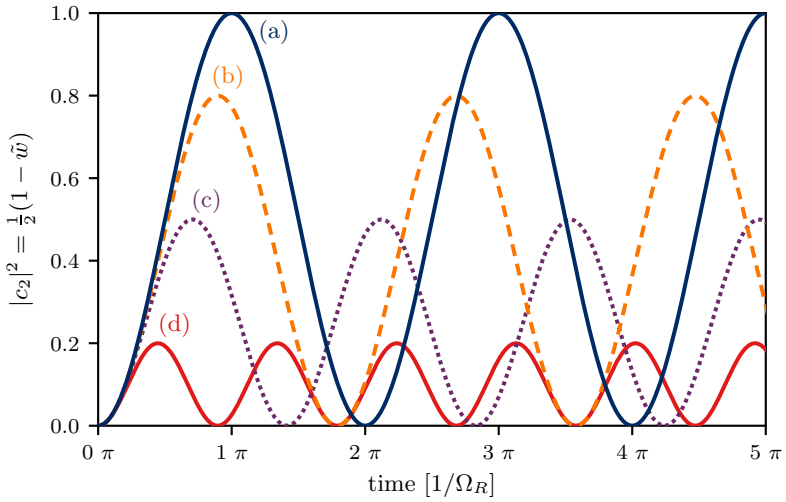

Figure 2.3: Rabi oscillations for different values of the detuning δ. (a) $\delta = 0$, (b) $\delta = \Omega_R/2$, (c) $\delta = \Omega_R$ and (d) $\delta = 2\Omega_R$.

encapsulate the time evolution of the Bloch vector represented pictorially in Figure 2.2.

2.4 Rabi oscillations, π-pulses and $\pi/2$-pulses

Consider the OBEs at resonance, that is $\delta = 0$. The field vector $\tilde{\mathbf{\Omega}}$ now lies in the x–y plane so that the Bloch vector for an atom initially in the ground state $|1\rangle$ at the north pole will follow a great circle, passing through the south pole before returning to where it started. The atom oscillates between the ground and excited states, and the angular frequency of this oscillation is Ω_R. In terms of the excited state probability $|c_2|^2$ we find

$$|c_2|^2 = \frac{1}{2}\left(1 - \cos(\Omega_R t)\right). \tag{2.35}$$

These oscillations are called *Rabi oscillations*, and Ω_R is correspondingly referred to as the *Rabi frequency*, after the early NMR work of Isidor Rabi [78, 79].

It is clear in Figure 2.2 that introducing a detuning means that the Bloch vector for an atom starting in the ground state will no longer make it all the way to the south pole. The amplitude of the oscillations is reduced but the field vector is longer, representing an increase in frequency. The length $\Omega = \sqrt{\Omega_R^2 + \delta^2}$ is referred to as the *generalised Rabi frequency*. When $\delta = \Omega_R$ the amplitude of the oscillations is reduced by a factor of two, peaking at $|c_2|^2 = |c_1|^2 = 1/2$. Clearly the ratio of δ to Ω_R is important:

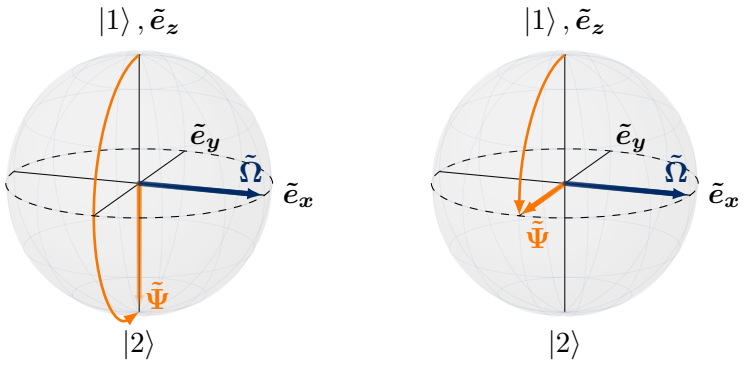

(a) A π- or mirror pulse. **(b)** A $\pi/2$- or beamsplitter pulse.

Figure 2.4: Finite duration, resonant laser pulses such as those shown here are the fundamental components of most optical atom interferometers, which often use two-photon resonances such as those described in Section 2.6 to minimise problematic spontaneous emission.

increasing Ω_R increases the range of detunings for which significant oscillations occur. As Ω_R is proportional to the amplitude of the optical field, this is an example of *power broadening*. Rabi oscillations for a variety of detunings are illustrated in Figure 2.3.

It is worth taking a moment now to introduce a couple of finite duration light–atom interactions that will prove useful later. On resonance, a pulse with a duration of half a Rabi period $T = \pi/\Omega_R$ will transfer an atom from the ground state to the excited state, or vice versa. Such a pulse is called a π-*pulse*, depicted on the Bloch sphere in Figure 2.4(a).

Similarly, a pulse lasting just a quarter of a Rabi period $T = \pi/2\Omega_R$ will prepare a state in an equal superposition of $|1\rangle$ and $|2\rangle$

$$\frac{1}{\sqrt{2}}\left(|1\rangle - ie^{i\phi}|2\rangle\right). \tag{2.36}$$

This $\pi/2$-*pulse* operation is depicted on the Bloch sphere in Figure 2.4(b) and, as the laser phase ϕ determines the orientation of the field vector in the x–y plane, we see that this phase is imprinted on the phase of the superposition. This is of particular interest for atom interferometry, along with the action of a π-pulse on such a state. A π-pulse acting on a superposition is often termed a *mirror pulse* as it serves to invert the state amplitudes.[5] A $\pi/2$-pulse in this context is often referred to as a *beamsplitter pulse*, and these operations are the fundamental components of most optical atom interferometers.

[5] With a phase factor, again dependent on the laser phase ϕ.

2.5 Light shifts and dressed states

Consider the expectation value of the atomic dipole moment in the laboratory (non-rotating) frame. Still considering linearly polarised light in the x direction, we shall assume that the atomic dipole will be solely in the direction of the external field $\hat{\boldsymbol{d}} = -e\hat{x}\boldsymbol{e_x}$, a reasonable assumption for one-electron atoms. Recalling that we have established $\langle 1|\,\hat{x}\,|1\rangle = \langle 2|\,\hat{x}\,|2\rangle = 0$ and $\langle 1|\,\hat{x}\,|2\rangle = \langle 2|\,\hat{x}\,|1\rangle$, the dipole moment of a general state $|\Psi\rangle = c_1\,|1\rangle + c_2\,|2\rangle$ is then[6]

$$\begin{aligned}
\langle\hat{\boldsymbol{d}}\rangle &= -\,e\left(c_1^* c_2\,\langle 1|\,\hat{x}\,|2\rangle + c_2^* c_1\,\langle 1|\,\hat{x}\,|2\rangle\right)\boldsymbol{e_x} \\
&= -\,e\,\langle 1|\,\hat{x}\,|2\rangle\left(e^{-i\omega t}\tilde{c}_1^*\tilde{c}_2 + e^{i\omega t}\tilde{c}_1\tilde{c}_2^*\right)\boldsymbol{e_x} \\
&= -\,e\,\langle 1|\,\hat{x}\,|2\rangle\Big([\tilde{c}_1^*\tilde{c}_2 + \tilde{c}_1\tilde{c}_2^*]\cos(\omega t) \\
&\qquad\qquad + i[\tilde{c}_1\tilde{c}_2^* - \tilde{c}_1^*\tilde{c}_2]\sin(\omega t)\Big)\boldsymbol{e_x} \\
&= -\,e\,\langle 1|\,\hat{x}\,|2\rangle\left(\tilde{u}\cos(\omega t) + \tilde{v}\sin(\omega t)\right)\boldsymbol{e_x},
\end{aligned} \tag{2.37}$$

[6] The letter e shoulders a lot of responsibility here, representing both the elemental unit of charge and Euler's number in addition to labelling the unit vectors. It is probably less confusing to let the context determine the meaning, rather than attempt to differentiate between them with additional notation.

where we have used $\tilde{c}_{1,2} = e^{\mp i\omega t/2}c_{1,2}$ to relate the lab-frame dipole moment to the Bloch vector coordinates in the rotating frame.

We see that the dipole oscillates at the angular frequency ω of the optical field that drives it. Without losing generality, let us set the laser phase to $\phi = 0$. We can identify in-phase and quadrature components whose amplitudes are determined respectively by the x coordinate $\tilde{u}$ and y coordinate $\tilde{v}$ of the Bloch vector in the rotating frame. As it is energetically favourable for a dipole to align with an electric field, the time-averaged potential energy of the dipole in the oscillating field is proportional to the in-phase term. The quadrature term results in absorption of energy from the driving field, and the corresponding dipole radiation in all directions results in a dissipative *scattering* force. This is a very classical picture, but remains qualitatively true in the quantum description. We will revisit this in the next chapter, where we consider the effect of spontaneous emission. This introduces dissipation into the OBEs (2.34) that allows steady-state solutions for $\tilde{u}$ and $\tilde{v}$ to be found.

For now, consider the resonant case $\delta = 0$ in the absence of such dissipation. We have set $\phi = 0$ meaning the field vector lies parallel to $\tilde{\boldsymbol{e}}_{\boldsymbol{x}}$ so that $\tilde{u} = 0$ and, for an atom starting in the ground state, $\tilde{v} = -\sin(\Omega_R t)$ oscillates between ± 1 at the Rabi frequency

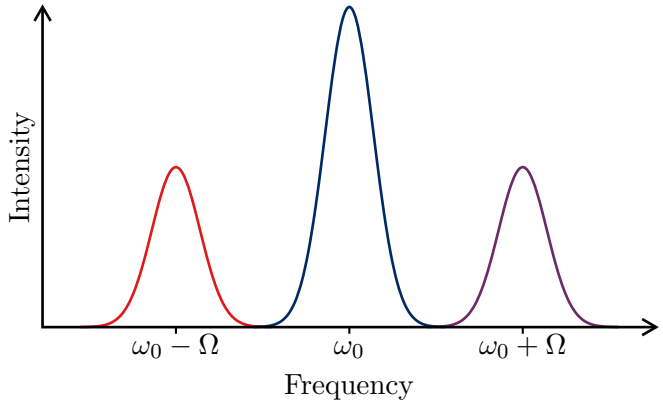

Figure 2.5: Illustration of the Mollow triplet observed in the fluorescence emitted from resonantly driven atoms. Radiation is observed at the transition angular frequency ω_0, with the modulation due to Rabi oscillations at the Rabi frequency Ω producing positive and negative sidebands.

Ω_R. The dipole, oscillating at ω, is thus modulated at Ω_R and we would expect the radiation from such a dipole to exhibit sidebands at the sum and difference frequencies $\omega \pm \Omega_R$. This is indeed apparent in the fluorescence light emitted from such a resonantly driven system; there we observe the classic Mollow triplet [80], an illustrative sketch of which is shown in Figure 2.5.

For a more quantum picture of this, and other effects, we turn to the treatment of Cohen-Tannoudji *et al.* [81] and reconsider the rotating frame Hamiltonian in Equation (2.33). As ever, its eigenstates are those that do not change with time, and these correspond to the Bloch vector being either parallel or anti-parallel to the field vector so that the cross product in $\frac{\partial \tilde{\boldsymbol{\Psi}}}{\partial t} = \tilde{\boldsymbol{\Psi}} \times \tilde{\boldsymbol{\Omega}}$ vanishes. With reference to Equations (2.5) we see that these eigenstates are

$$|+\rangle = \cos\frac{\theta}{2}|1\rangle + e^{i\phi}\sin\frac{\theta}{2}|2\rangle \quad \text{and} \tag{2.38a}$$

$$|-\rangle = e^{-i\phi}\sin\frac{\theta}{2}|1\rangle - \cos\frac{\theta}{2}|2\rangle\,, \tag{2.38b}$$

where θ and ϕ are the polar angles describing the orientation of the field vector. By convention, $0 \leq \theta \leq \pi$ is measured from the 'north' pole corresponding to $|1\rangle$ and is defined equivalently by $\cos\theta = \delta/\Omega$ or $\sin(\theta) = \Omega_R/\Omega$, while ϕ is measured counter-clockwise[7] from $\tilde{\boldsymbol{e}}_{\boldsymbol{x}}$ and conveniently equates to the laser phase with our choice of signs.

[7] To an observer at the 'north' pole.

These states form a natural basis in which to consider the light–atom system, and have eigenenergies

$$\epsilon_{\pm} = \pm\frac{\hbar}{2}\Omega = \pm\frac{\hbar}{2}\sqrt{\delta^2 + \Omega_R^2}. \tag{2.39}$$

These can be verified by substituting Equations (2.38) into the time-independent Schrödinger equation, but are also apparent from the fact that a superposition of these states will acquire phase at a rate Ω in what is nothing more than the Rabi oscillations detailed in Section 2.4.

In the absence of any light–atom coupling, $\Omega_R = 0$ and the energy separation between the states is just given by the detuning $\hbar\delta$. Why should this be? In shifting into a frame rotating at the optical frequency, we factored the photon energy $\hbar\omega$ out of the problem. The energy difference that remains is the difference between the the photon energy and the atomic transition. If we incorporate the photon energy in our lab frame picture then we observe an equivalence. Let us expand our lab frame Hamiltonian to include a term $\hat{H}_{\text{light}} = \hbar(\hat{a}^\dagger\hat{a} + 1/2)\omega$ where $\hat{a}^\dagger\hat{a}$ is the operator for the number of photons in the optical field,

$$\hat{H} = \hat{H}_0 + \hat{V} + \hat{H}_{\text{light}}. \tag{2.40}$$

When $\hat{V} = 0$, the eigenstates of the whole system can then be expressed as $|\alpha, n\rangle$ where $\alpha = 1, 2$ labels the internal state of the atom and n is the integer number of photons in the optical mode of the laser. Incrementing n increases the energy by $\hbar\omega$, resulting in the energy structure represented in Figure 2.6. For each value of n the atomic structure repeats, vertically shifted by $\hbar\omega$.

When the coupling is reintroduced, a transition from $|1\rangle$ to $|2\rangle$ is mediated by absorption of a photon from the laser mode with a corresponding reduction of n by 1, while a transition from $|2\rangle$ to $|1\rangle$ via stimulated emission returns a photon to the field and increases n by 1. Consequently, only levels $|1, n\rangle \leftrightarrow |2, n-1\rangle$ are coupled for any given value of n. Such couplings are labelled with blue arrows in Figure 2.6 and the coupled states, including the photon energy, are nearly degenerate except for the energy difference introduced by the laser detuning $\hbar\delta$. This is in agreement with the energy separation in the rotating frame; the result is the same however we account for the photon energy.

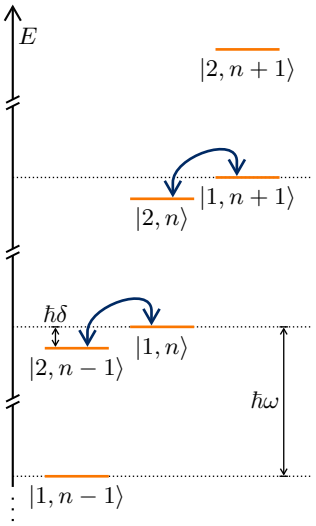

Figure 2.6: Including the energy in the optical field containing n photons, states $|1, n\rangle$ and $|2, n-1\rangle$ are near-degenerate, separated in energy by $\hbar\delta$ where δ is the detuning of the photon angular frequency ω from that of the atomic transition. The optical field loses a photon when it drives an atom from $|1\rangle$ to $|2\rangle$ and vice versa, so it is these near-degenerate levels that are coupled by the field.

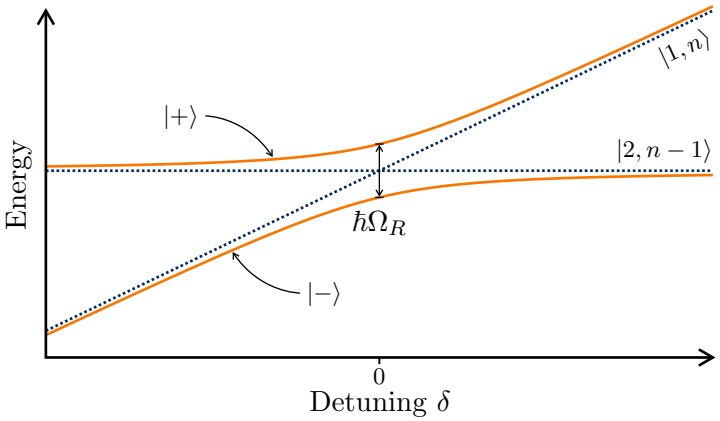

Figure 2.7: Dressed state energies as a function of detuning, relative to the energy of the bare state $|2, n-1\rangle$. When $\delta \ll 0$, $|+\rangle\,(|-\rangle) \approx |1, n\rangle\,(|2, n-1\rangle)$ with a slight shift up (down) in energy. On the other side of the resonance $|+\rangle\,(|-\rangle) \approx |2, n-1\rangle\,(|1, n\rangle)$ and, from the perspective of the bare states, the sign of the shift is reversed. On resonance, while the bare states (dashed lines) would have the same energy, the dressed states are separated by $\hbar\Omega_R$ in an archetypal avoided crossing.

It would be more accurate, though, to substitute $|1\rangle \to |1, n\rangle$ and $|2\rangle \to |2, n-1\rangle$ in Equations (2.38) to acknowledge the photons' critical role. It is common parlance to say that the optical field 'dresses' the atom, and to refer to $|\pm\rangle$ as *dressed states*.

Far from resonance and to the red, when $\delta < 0$ and $|\delta| \gg \Omega_R$, the field vector is pointing very close to the south pole so that $\theta \approx \pi \Rightarrow \sin(\theta/2) \approx 1$ and $\cos(\theta/2) \approx 0$. Consequently, ignoring phase factors, $|+\rangle \approx |2, n-1\rangle$ and $|-\rangle \approx |1, n\rangle$. In this limit the main effect of the light on the atom is to induce a shift in the state energies $\Delta\epsilon_\pm = \epsilon_\pm \mp \delta$ that we can calculate by taking a Taylor expansion of Equation (2.39) about $|\Omega_R/\delta| = 0$ to arrive at

$$\Delta\epsilon_\pm \approx \pm \left| \frac{\hbar\Omega_R^2}{4\delta} \right|, \tag{2.41}$$

taking the first nonzero term. We find a positive shift for $|+\rangle$, corresponding to the excited state $|2, n-1\rangle$, and a negative shift to the ground state $|1, n\rangle$.

Far to the blue side of the resonance, Equation (2.41) for the shift to the dressed states remains valid but now the field vector is pointing north. We have $\theta \approx 0$ so that now $|+\rangle \approx |1, n\rangle$ and $|-\rangle \approx |2, n-1\rangle$. The sign of the shift from the perspective of the bare states has thus reversed. This is immediately apparent

from inspection of Figure 2.7, where the dressed state energies are compared to those of the bare states as a function of detuning.

As the magnitude of $\Delta\epsilon_\pm$ scales in proportion to Ω_R^2, and this in turn is proportional to the optical intensity, this shift to the bare states is often referred to as the far-off-resonant *light shift*, or else the *ac Stark shift* as a counterpart to the similar shift induced by a dc electric field. We can use these shifts in order to gain considerable insights into many laser cooling and trapping processes, as we shall explore in the following chapter.

Figure 2.7 also shows the behaviour of the dressed states on resonance at $\delta = 0$. Here the atom is 'maximally dressed'; each dressed state is an equal superposition of the ground and excited state, with opposite phase relationships.[8] Consequently, while the bare states would here have equal energy, the dressed states are separated by $\hbar\Omega_R$ in an avoided crossing.

[8] In the lucid language of Cohen-Tannoudji, each dressed state is equally 'contaminated' by the ground and excited states.

It is apparent from the plot that if one were to sweep the detuning of the laser through the resonance from far off on one side to far off on the opposite side, one could adiabatically transfer atoms from one bare state to the other [82]. With a bit of thought, one would conclude that this sweep should be slow relative to the Rabi frequency Ω_R but also fast with respect to the decay rate of the upper state so that spontaneous emission is negligible and the form of the OBEs in Equations (2.34) is valid. This process therefore just about warrants the oxymoronic term *adiabatic rapid passage (ARP)* that it has acquired in the literature [83–85], and is particularly useful when full population inversion is desired in an ensemble of atoms that may exhibit a broad range of detunings and Rabi frequencies due to, for example, Doppler shifts or variations in laser intensity. We shall revisit this in another context in Chapter 9.

To close this section we return to the Mollow triplet with which we began it, armed with this more quantum picture. Note that spontaneously emitted fluorescence light is not emitted into the laser mode, but rather in a random direction. It therefore results in the loss of a photon from the system we are considering so that $|2, n-1\rangle$ decays to $|1, n-1\rangle$. With this in mind, Figure 2.8 shows two sets of near-degenerate states from Figure 2.6, this time in frequency space. On the left the bare states are shown, corresponding to the laser being off. Note that four levels are shown but, on resonance, these would be completely degenerate.

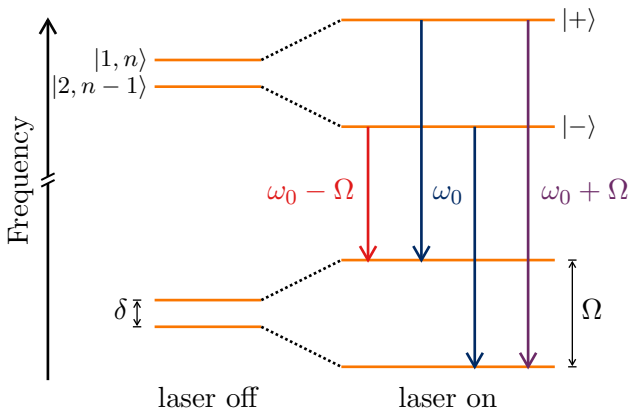

Figure 2.8: The Mollow triplet can be understood as transitions between dressed states. Spontaneous emission in the absence of the optical coupling only occurs from state $|2\rangle$ to $|1\rangle$ but, as the eigenstates of the optically dressed atom are 'contaminated' with both bare states in nearly equal proportion close to resonance, all transitions are possible giving rise to three possible frequencies of the emitted radiation.

Of course, fluorescence is not observed when the laser is off as the excited state is negligibly populated.

Upon turning the laser on, we enter the dressed state picture; the near-degenerate states in each manifold are then separated in angular frequency by Ω. Now, near resonance each dressed state is a roughly equal superposition of ground and excited states; spontaneous emission between either of the dressed states in the upper manifold to either of the dressed states in the lower manifold is thus dipole allowed and will be observed. From Figure 2.8 we see that the different decay routes will yield light at three separate frequencies that correspond to the components of the Mollow triplet in Figure 2.5 that we previously explained in classical terms as a modulation applied to an oscillating dipole at the Rabi frequency. Once again, quantum mechanics agrees with classical intuition.

2.6 Effective two-level model of a three-level Raman system

Let us take a step closer to real-world physics and consider now a three-level atom interacting with a pair of lasers. Specifically, we consider a ground level in an alkali atom that exhibits hyperfine splitting, resulting in two stable ground states $|1\rangle$ and $|2\rangle$ separated by an angular frequency ω_0 that is typically in the microwave (GHz) range. Each of these levels is then optically coupled to an

excited state $|3\rangle$ by lasers operating at optical angular frequencies ω_1 and ω_2 respectively.

This closely resembles the system experimentally realised in this thesis, and indeed in many atom interferometers. We will see that, under particular conditions, the excited state $|3\rangle$ effectively drops out of the problem and it can be considered a two-level system with an effective field vector coupling $|1\rangle$ and $|2\rangle$ just as we have considered in the previous sections. The advantage of this system is that typically there is no direct dipole coupling between the ground states $|1\rangle$ and $|2\rangle$. While in our previous discussions we conveniently ignored spontaneous decay from $|2\rangle \to |1\rangle$, here it is genuinely negligible and we can consider it a closed system that evolves coherently over timescales measured in seconds. Nonetheless, the fact that the two levels are coupled by optical photons allows a relatively large recoil momentum to be associated with the transition. This combination of long-term coherence and large momentum transfer is particularly useful for atom interferometry, as we will explore in Chapter 4.

2.6.1 Internal and external degrees of freedom and the transfer of momentum

Thus far we have not considered the momentum imparted to the atom by the optical field, but now would seem like an opportune moment to do so. In general we can separate an atomic system into two uncorrelated subsystems: an 'internal' electronic structure whose Hamiltonian takes the form $\hat{H}_{\text{int}} = \sum_\alpha \hbar\omega_\alpha |\alpha\rangle\langle\alpha|$ and has energy eigenstates denoted $|\alpha\rangle$ with corresponding eigenenergies $\epsilon_\alpha = \hbar\omega_\alpha$, and an 'external' centre-of-mass part whose Hamiltonian takes the form of that for a free particle $\hat{H}_{\text{ext}} = |\hat{\boldsymbol{p}}|^2/2m$, where $\hat{\boldsymbol{p}}$ is the operator for the momentum of the atom's centre-of-mass and m is the total mass of the atom. The eigenstates of this second part are thus momentum eigenstates $|\boldsymbol{p}\rangle$.

As the two parts are uncorrelated, the total Hamiltonian for the bare atom takes the form

$$\hat{H}_0 = \hat{H}_{\text{int}} \otimes \mathbb{I} + \mathbb{I} \otimes \hat{H}_{\text{ext}}, \tag{2.42}$$

with eigenstates $|\alpha, \boldsymbol{p}\rangle = |\alpha\rangle \otimes |\boldsymbol{p}\rangle$.[9] A general state can then be expressed in this basis as

[9] That we can divide the problem in this manner is often taken for granted, but is not a banal point by any means. The constituent particles of the atom are highly correlated; if we treat their positions and momenta individually in the lab frame then we find an entangled system. Employing this trick of considering the electronic structure in relative coordinates in the centre-of-mass frame, in addition to the combined centre-of-mass momentum in the lab frame, we find two uncorrelated systems. Entanglement is relative [86, 87].

$$\int \mathrm{d}^3\,\boldsymbol{p} \sum_\alpha c_{\alpha,\boldsymbol{p}}(t)\,|\alpha,\boldsymbol{p}\rangle\,. \tag{2.43}$$

Suppose an optical field couples levels $\alpha = a, b$ with an interaction term

$$\hat{V} = -\hat{\boldsymbol{d}}\cdot\hat{\boldsymbol{E}}(\boldsymbol{r},t), \tag{2.44}$$

where we explicitly include the dependence of the electric field on the spatial coordinate $\boldsymbol{r}$ and we take it to be a plane wave propagating in a direction determined by the wavevector $\boldsymbol{k}$

$$\begin{aligned}\hat{\boldsymbol{E}}(\boldsymbol{r},t) &= \hat{\boldsymbol{E}}_{\boldsymbol{0}}\cos\left(\boldsymbol{k}\cdot\hat{\boldsymbol{r}} - \omega t + \phi\right)\\ &= \frac{\hat{\boldsymbol{E}}_{\boldsymbol{0}}}{2}\left(e^{i\boldsymbol{k}\cdot\hat{\boldsymbol{r}}}e^{-i\omega t}e^{i\phi} + e^{-i\boldsymbol{k}\cdot\hat{\boldsymbol{r}}}e^{i\omega t}e^{-i\phi}\right).\end{aligned} \tag{2.45}$$

This interaction can be treated for the internal part as in previous sections — we still assume that the optical wavelength is large compared to the size of the atom — but in the momentum basis the $e^{\pm i\boldsymbol{k}\cdot\hat{\boldsymbol{r}}}$ terms act as translation operators [88]. Rewriting them using the closure relation

$$\begin{aligned}e^{\pm i\boldsymbol{k}\cdot\hat{\boldsymbol{r}}} &= \int \mathrm{d}^3\,\boldsymbol{p}\, e^{\pm i\boldsymbol{k}\cdot\hat{\boldsymbol{r}}}\,|\boldsymbol{p}\rangle\,\langle\boldsymbol{p}|\\ &= \int \mathrm{d}^3\,\boldsymbol{p}\,|\boldsymbol{p}\pm\hbar\boldsymbol{k}\rangle\,\langle\boldsymbol{p}|\,,\end{aligned} \tag{2.46}$$

we retrieve the unsurprising and well-known result: absorbing or emitting a photon with wavevector $\boldsymbol{k}$ changes an atom's centre-of-mass momentum by that of the photon $\hbar\boldsymbol{k}$. In practice this means that the state $|a,\boldsymbol{p}\rangle$ is only coupled to the single state $|b,\boldsymbol{p}+\hbar\boldsymbol{k}\rangle$ by the optical interaction. We can denote the energies of these states

$$\begin{aligned}\hbar\omega'_a(\boldsymbol{p}) &= \hbar\omega_a + \frac{|\boldsymbol{p}|^2}{2m}, \text{ and}\\ \hbar\omega'_b(\boldsymbol{p}) &= \hbar\omega_b + \frac{|\boldsymbol{p}+\hbar\boldsymbol{k}|^2}{2m}.\end{aligned} \tag{2.47}$$

This introduces a modified resonance condition. While in the previous sections we considered the optical field to be on resonance when the detuning $\delta = \omega - (\omega_b - \omega_a) = 0$ — corresponding to the optical frequency being equal to the frequency difference between

the two internal states it coupled — now we see that we should include the external kinetic energy as well. We thus introduce the modified detuning

$$\begin{aligned}\delta'(\boldsymbol{p}) &= \omega - \left(\omega_b'(\boldsymbol{p}) - \omega_a'(\boldsymbol{p})\right) \\ &= \delta - \frac{\boldsymbol{k}\cdot\boldsymbol{p}}{m} - \frac{\hbar|\boldsymbol{k}|^2}{2m},\end{aligned} \tag{2.48}$$

with the associated resonance condition $\delta' = 0$.

We see the resonance is modified by two shifts that we can immediately identify as the *Doppler shift* $\boldsymbol{k}\cdot\boldsymbol{p}/m = \boldsymbol{k}\cdot\boldsymbol{v}$ associated with an atom's initial velocity $\boldsymbol{v}$ in the frame of the laser, and the *recoil shift* $\hbar|\boldsymbol{k}|^2/2m$ originating from the kinetic energy difference between the two states coupled by absorption or emission of a photon.

Conservation of energy and momentum makes this inevitable, and it could all have been derived by drawing a simple picture, such as that in Figure 2.9, where the total energy of each internal state is plotted as a function of the component of atomic momentum p along the propagation axis of the laser. Photons are represented by arrows drawn from a point $\omega_a'(p)$ with a slope c until they reach $p+\hbar k$. The detuning δ' is then just the vertical separation between the top of the arrow and the parabola representing the energy of the coupled state, up to a factor a $\hbar$.

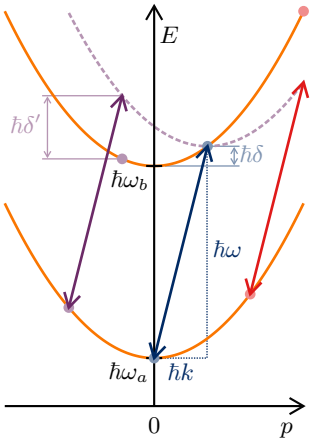

Figure 2.9: Total energy of the atomic states of a two-level atom, plotted as solid orange parabolas as a function of the component p of atomic momentum parallel to the laser wavevector $\boldsymbol{k}$. Photons, represented by arrows with slope c, couple a closed pair of states, represented by like-coloured circles on the dispersion curves, for a given value of p. The central, on-resonance, photon is detuned from the transition angular frequency $\omega_b - \omega_a$ by an amount δ to account for the recoil shift, while the momentum dependence of the detuning δ' due to the Doppler shift is represented by the dashed parabola and measured from the curve representing the upper state energy.

2.6.2 And back to the problem at hand

So let us return to the three level atom. Motivated by the discussion in the previous interlude, we introduce the wavevectors for our two lasers $\boldsymbol{k_{1,2}}$ in addition to their angular frequencies $\omega_{1,2}$. Henceforth we will let these lie along the z axis so that $\boldsymbol{k_{1,2}} \equiv k_{1,2}\boldsymbol{e_z}$ and only consider the z component of the atomic momentum that we denote p. If we assume that the first laser only couples the states $|1\rangle \leftrightarrow |3\rangle$ and the second only couples $|2\rangle \leftrightarrow |3\rangle$ then, again, for a ground state momentum p we find that the lasers couple a closed ladder of momentum-inclusive states

$$\begin{aligned}|1'\rangle &\equiv |1, p\rangle\,, \\ |2'\rangle &\equiv |2, p + \hbar(k_1 - k_2)\rangle \quad \text{and} \\ |3'\rangle &\equiv |3, p + \hbar k_1\rangle\end{aligned} \tag{2.49}$$

with associated energies

$$\begin{aligned}\hbar\omega_1'(p) &= \frac{p^2}{2m},\\ \hbar\omega_2'(p) &= \hbar\omega_0 + \frac{\big(p+\hbar[k_1-k_2]\big)^2}{2m}, \text{ and}\\ \hbar\omega_3'(p) &= \hbar\omega_3 + \frac{(p+\hbar k_1)^2}{2m}\end{aligned} \tag{2.50}$$

respectively, relative to the energy of the $p = 0$ ground state.[10] Note that if we forgo this assumption and consider the coupling between all levels by all lasers then this ladder is no longer closed and the problem spans a much larger basis. However, we will consider the case where the two-photon detuning

$$\delta'(p) \equiv (\omega_1 - \omega_2) - \big(\omega_2'(p) - \omega_1'(p)\big) \tag{2.51}$$

is much smaller than the single photon detuning

$$\Delta'(p) \equiv \omega_1 - \big(\omega_3'(p) - \omega_1'(p)\big)\,. \tag{2.52}$$

In this limit the dynamics are dominated by two-photon processes and the contributions from other couplings are negligible, with the exception of an additional light shift that is derived in [89], making it reasonable to restrict our analysis to this basis of just three momentum-inclusive states.

In terms of the 'bare' one- and two-photon detunings, respectively defined as $\Delta \equiv \omega_1 - \omega_3$ and $\delta \equiv (\omega_1 - \omega_2) - \omega_0$, Equations (2.51) and (2.52) become

$$\delta' = \delta - \frac{p(k_1-k_2)}{m} - \frac{\hbar(k_1-k_2)^2}{2m}, \text{ and} \tag{2.53a}$$

$$\Delta' = \Delta - \frac{pk_1}{m} - \frac{\hbar k_1^2}{2m}. \tag{2.53b}$$

We see that the single-photon detuning exhibits the Doppler and recoil shifts identified in Equation (2.48), while the two-photon detuning has equivalent shifts associated with the *effective wavevector*

$$\boldsymbol{k} \equiv \boldsymbol{k_1} - \boldsymbol{k_2} = (k_1 - k_2)\boldsymbol{e_z} \equiv k\boldsymbol{e_z}. \tag{2.54}$$

[10] Here we have shunted the kinetic energy due to the transverse atomic motion into our choice of energy reference. As the dynamics are only governed by frequency *differences* — and the transverse momentum is unaffected by interaction with the lasers — this means that, while this reference may vary from atom to atom, the end result will not.

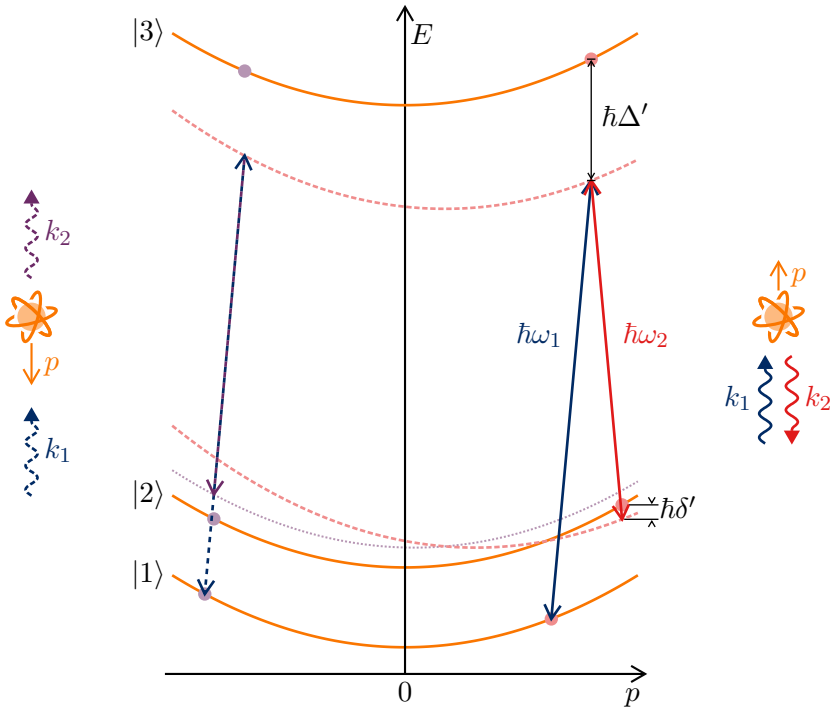

Figure 2.10: Total energies, not to scale, of the three atomic states described in the main text, plotted in solid orange as a function of the z component of atomic momentum p. We consider photons, shown as arrows, of angular frequencies $\omega_{1,2}$ and wavenumbers $k_{1,2}$ propagating in the z direction that couple a closed set of states, represented by like-colour circles on the dispersion curves. *Left:* co-propagating, dashed, photons have the same slope and thus couple states with very similar momenta. *Right:* counter-propagating case, solid, photons have equal but opposite slopes. The one- and two-photon detunings Δ' and δ' are labelled for this arrangement, and their momentum dependence is shown by red, dashed parabolas. Δ' is large enough that the fractional shift due to this momentum dependence is negligible, but the same is not true of δ'. On the other hand, the co-propagating photons have the same single-photon detuning but the momentum dependence of the two-photon detuning, represented by the dotted purple line, is negligible owing to the very small effective wavenumber $k = k_1 - k_2$. We thus term the counter- and co-propagating arrangements Doppler sensitive and insensitive respectively.

We will consider the limit where Δ is much greater than any of these shifts so that $\Delta' \approx \Delta$ is effectively constant. However, we will find in later chapters that the shifts to δ' are central to atom interferometry experiments. We note that these shifts are considerable when the lasers are counter-propagating. In this arrangement $k_2 \approx -k_1 \Rightarrow k \approx 2k_1$. By contrast, when the lasers are co-propagating $k_2 \approx k_1$ and $k \approx 0$.

These detunings and shifts are represented graphically in Figure 2.10 for both the co- and counter-propagating cases. The co-propagating photons, represented by dashed arrows, have the same slope and couple states $|1, p\rangle$ and $|2, p + \hbar k\rangle$ with very similar momenta; the counter-propagating photons, represented by solid lines, have slopes with opposite signs and couple states with a correspondingly large momentum separation $\hbar k$.

Let us proceed to treat the system in a similar manner to the way we approached the two-level problem in Section 2.3. In our restricted three-state basis, the most general state is

$$|\Psi\rangle = c_1(t)\,|1,p\rangle + c_2(t)\,|2,p+\hbar(k_1-k_2)\rangle + c_3(t)\,|3,p+\hbar k_1\rangle\,, \quad (2.55)$$

which we express in column vector form

$$|\Psi(t)\rangle = \begin{pmatrix} c_1(t) \\ c_2(t) \\ c_3(t) \end{pmatrix}, \quad (2.56)$$

allowing us to represent the bare atomic Hamiltonian with the diagonal matrix

$$\hat{H}_0 = \begin{pmatrix} \hbar\omega_1' & 0 & 0 \\ 0 & \hbar\omega_2' & 0 \\ 0 & 0 & \hbar\omega_3' \end{pmatrix}. \quad (2.57)$$

Here we have omitted the explicit dependence of the eigenenergies on p for brevity.

As before, the operator for the interaction energy is given by $\hat{V} = -\hat{\boldsymbol{d}} \cdot \hat{\boldsymbol{E}}$ that we represent with the matrix

$$\hat{V} = \begin{pmatrix} V_{11} & V_{12} & V_{13} \\ V_{12}^* & V_{22} & V_{23} \\ V_{13}^* & V_{23}^* & V_{33} \end{pmatrix}. \quad (2.58)$$

The electric field operator from the combined laser fields is

$$\hat{\boldsymbol{E}} = \sum_{l=1,2} \mathrm{Re}\left(\hat{\boldsymbol{E}}_{\boldsymbol{l}} e^{i(k_l \hat{z} - \omega_l t + \phi_l)}\right), \quad (2.59)$$

where l enumerates the laser fields, $\hat{\boldsymbol{E}}_{\boldsymbol{l}}$ encodes their respective polarisations and ϕ_l their phases. Recall that the effect of the spatial variation of the electric field on the momentum part of the states has already been accounted for in our choice of basis and the corresponding eigenenergies. The optical wavelength is still large

with respect to the spatial extent of the atom, so we can proceed to treat the interaction in the dipole approximation which we do by evaluating at $z = 0$, making the matrix elements of $\hat{V}$

$$V_{\alpha\beta} = \langle \alpha' | - \hat{\boldsymbol{d}} \cdot \sum_{l=1,2} \mathrm{Re}\left(\hat{\boldsymbol{E}}_{\boldsymbol{l}} e^{-i(\omega_l t - \phi_l)} \right) | \beta' \rangle \,, \tag{2.60}$$

with $\alpha, \beta = 1, 2, 3$ and the primed kets defined in Equation (2.49).

Now, as before, the odd parity of the dipole operator and the well defined parity of the basis states ensure that the diagonal elements $V_{\alpha\alpha}$ vanish; states are not coupled to themselves. We intimated at the start of this section that for many systems of interest the $|1'\rangle \leftrightarrow |2'\rangle$ transition is not dipole allowed. This is the case for the two hyperfine ground levels of Rubidium, as used in this work, and is the case we shall consider here. Thus the V_{12} matrix element also vanishes.

The only nonzero matrix elements left to consider are then

$$\begin{aligned} V_{\alpha 3} &= \langle \alpha' | - \hat{\boldsymbol{d}} \cdot \sum_{l=1,2} \mathrm{Re}\left(\hat{\boldsymbol{E}}_{\boldsymbol{l}} e^{-i(\omega_l t - \phi_l)} \right) | 3' \rangle \\ &= \sum_{l=1,2} \hbar \, \mathrm{Re}\left(\Omega_{l\alpha} e^{-i(\omega_l - \phi_l)} \right) , \end{aligned} \tag{2.61}$$

where $\alpha = 1, 2$ and $\Omega_{l\alpha} = \langle \alpha' | - \hat{\boldsymbol{d}} \cdot \hat{\boldsymbol{E}}_{\boldsymbol{l}} | 3' \rangle / \hbar$ is the Rabi frequency for the coupling of state $|\alpha'\rangle$ to the excited state $|3'\rangle$ by the lth laser. We have neglected at this stage to make any assumptions about the polarisation of either laser, and so we note that the $\Omega_{l\alpha}$ may, in general, be complex.

We now explicitly make the simplification, justified above and implicit in our choice of a closed basis, that laser $l = 1, 2$ only couples states $|l'\rangle \leftrightarrow |3'\rangle$. This corresponds to setting $\Omega_{12} = \Omega_{21} = 0$, making the total Hamiltonian $\hat{H} = \hat{H}_0 + \hat{V}$ including the interaction

$$\hat{H} = \hbar \begin{pmatrix} \omega_1' & 0 & \mathrm{Re}\left(\Omega_{11} e^{-i(\omega_1 t - \phi_1)} \right) \\ 0 & \omega_2' & \mathrm{Re}\left(\Omega_{22} e^{-i(\omega_2 t - \phi_2)} \right) \\ \mathrm{Re}\left(\Omega_{11} e^{-i(\omega_1 t - \phi_1)} \right) & \mathrm{Re}\left(\Omega_{22} e^{-i(\omega_2 t - \phi_2)} \right) & \omega_3' \end{pmatrix} . \tag{2.62}$$

To proceed further, we transform into a new frame given by a carefully chosen unitary transformation

$$\hat{\mathcal{U}}_R(t) = \begin{pmatrix} e^{i\omega_1' t} & 0 & 0 \\ 0 & e^{i(\omega_2'+\delta')t} & 0 \\ 0 & 0 & e^{i(\omega_3'+\Delta')} \end{pmatrix}. \tag{2.63}$$

State vectors and operators in this frame then take the respective forms

$$|\tilde{\Psi}(t)\rangle = \begin{pmatrix} \tilde{c}_1(t) \\ \tilde{c}_2(t) \\ \tilde{c}_3(t) \end{pmatrix} = \hat{\mathcal{U}}_R(t)\,|\Psi(t)\rangle = \begin{pmatrix} e^{i\omega_1' t} c_g(t) \\ e^{i(\omega_2'+\delta')t} c_e(t) \\ e^{i(\omega_3'+\Delta')t} c_3(t) \end{pmatrix}, \text{ and} \tag{2.64a}$$

$$\tilde{\hat{\mathcal{O}}}(t) = \hat{\mathcal{U}}_R(t)\hat{\mathcal{O}}(t)\hat{\mathcal{U}}_R^\dagger(t). \tag{2.64b}$$

This transformation is designed to eliminate the time dependence from the Hamiltonian and is akin to our translation to a rotating frame in Section 2.3. Plugging Equations (2.62) and (2.63) into Equation (2.30) to find the effective Hamiltonian $\hat{H}_R$ in this new frame we get

$$\hat{H}_R = \begin{pmatrix} 0 & 0 & \tilde{V}_{13} \\ 0 & -\hbar\delta' & \tilde{V}_{23} \\ \tilde{V}_{13}^* & \tilde{V}_{23}^* & -\hbar\Delta' \end{pmatrix}, \tag{2.65}$$

where the interaction terms are

$$\begin{aligned} \tilde{V}_{13} &= \hbar\, e^{-i(\Delta'+\omega_3'-\omega_1')t}\,\mathrm{Re}\left(\Omega_{11} e^{-i(\omega_1 t-\phi_1)}\right), \text{ and} \\ \tilde{V}_{23} &= \hbar\, e^{-i(\Delta'-\delta'+\omega_3'-\omega_2')t}\,\mathrm{Re}\left(\Omega_{22} e^{-i(\omega_2 t-\phi_2)}\right). \end{aligned} \tag{2.66}$$

By expanding the real values in terms of complex values plus their conjugates we can identify terms $e^{i\left(\omega_1-[\omega_3'-\omega_{1,2}']-\Delta'\right)t}$ and $e^{i\left(\omega_2-[\omega_3'-\omega_2']-[\Delta'-\delta']\right)t}$ which equate to unity, with the time dependence cancelling out as hoped. The remaining terms oscillate at approximately double the optical angular frequencies $\omega_{1,2}$. Provided $w_{1,2} \gg |\Omega_{11}|, |\Omega_{22}|, \Delta'$, we can neglect these fast, *counter-rotating*, terms in a generalisation of the RWA that we made back in Section 2.3. In doing so, the rotating frame Hamiltonian finally becomes

$$\hat{H}_R = \frac{\hbar}{2}\begin{pmatrix} 0 & 0 & \Omega_{11}e^{-i\phi_1} \\ 0 & -2\delta' & \Omega_{22}e^{-i\phi_2} \\ \Omega_{11}^* e^{i\phi_1} & \Omega_{22}^* e^{i\phi_2} & -2\Delta' \end{pmatrix}, \tag{2.67}$$

leading to three, coupled, differential equations for the state amplitudes $\tilde{c}_\alpha$ in the rotating frame

$$\frac{\partial \tilde{c}_1}{\partial t}(t) = -\frac{i}{2}\Omega_{11}e^{-i\phi_1}\tilde{c}_3(t) \tag{2.68a}$$
$$\frac{\partial \tilde{c}_2}{\partial t}(t) = i\delta'\tilde{c}_2(t) - \frac{i}{2}\Omega_{22}e^{-i\phi_2}\tilde{c}_3(t) \tag{2.68b}$$
$$\frac{\partial \tilde{c}_3}{\partial t}(t) = -\frac{i}{2}\Omega_{11}^* e^{i\phi_1}\tilde{c}_1(t) - \frac{i}{2}\Omega_{22}^* e^{i\phi_2}\tilde{c}_2(t) + i\Delta'\tilde{c}_3(t). \tag{2.68c}$$

These look formidable initially, but can be reduced by employing the dark art of *adiabatic elimination*. This has been rationalised in different ways [90–92] but usually relies on the fact that we are considering the case of large single-photon detuning $\Delta' \gg \delta'$ in conjunction with the additional assumption that the coupling is sufficiently weak so that $\Delta' \gg |\Omega_{11}|, |\Omega_{22}|$ as well.

This allows us to identify the final term in Equation (2.68c) as a driver of oscillations in $\tilde{c}_3(t)$ at a much faster rate than anything else in the system. For an initial condition $\tilde{c}_3(0) = 0$ these will act to keep the value of $\tilde{c}_3(t)$ small and it is then common practice to set its time derivative to zero. This may or may not be immediately intuitive;[11] for a bit of intuition, we therefore identify the cycle period of the fast oscillations as $T = {}^{2\pi}/_{\Delta'}$ and proceed to calculate the cycle average $\langle \tilde{c}_3(t) \rangle = {}^{1}/_{T}\int_t^{t+T} \tilde{c}_3(t')\,\mathrm{d}t'$ by integrating Equation (2.68c):

[11] One could get hooked up on the fact that, for very rapid oscillations, the time derivative could become instantaneously rather large.

$$\frac{1}{T}\int_t^{t+T} \mathrm{d}t' \frac{\partial \tilde{c}_3}{\partial t'}(t') = -\frac{i}{2}\frac{1}{T}\int_t^{t+T} \mathrm{d}t' \Big(\Omega_{11}^* e^{i\phi_1}\tilde{c}_1(t') + \Omega_{22}^* e^{i\phi_2}\tilde{c}_2(t')\Big) + i\Delta'\frac{1}{T}\int_t^{t+T} \mathrm{d}t'\tilde{c}_3(t'). \tag{2.69}$$

Here, the integral on the left-hand side evaluates to $\tilde{c}_3(t+T) - \tilde{c}_3(t)$. As the cycle period T is short then — in a first order approximation — we can assume that after a full cycle $\tilde{c}_3(t)$ will return to where it started and we can, indeed, set this term to zero. By the

same logic, we assume that the other amplitudes do not vary considerably over the cycle so that we can make the approximation $\tilde{c}_{1,2}(t') \approx \tilde{c}_{1,2}(t)$ and take them out of the integral.

We thus arrive at a first order approximation for the cycle average

$$\langle \tilde{c}_3(t) \rangle \approx \frac{1}{2\Delta'} \left(\Omega_{11}^* e^{i\phi_1} \tilde{c}_1(t) + \Omega_{22}^* e^{i\phi_2} \tilde{c}_2(t) \right) \tag{2.70}$$

that we can substitute for $\tilde{c}_3(t)$ in Equations (2.68a) and (2.68b), turning them into coarse-grained equations for the slow dynamics that smooth over the fine-grained ripples induced by the rapid oscillations. This first-order approximation is far from completely watertight, for example it varies depending on the particular choice of rotating frame, but in practice it is sufficient for everything we wish to illustrate here. More discussion and higher order approximations can be found in [92–94], but for now we arrive at the effective two-level Hamiltonian for $\tilde{c}_{1,2}$ that arises from this coarse-grained treatment,[12]

$$\hat{H}_2 = \frac{\hbar}{2} \begin{pmatrix} \frac{|\Omega_{11}|^2}{2\Delta'} & \frac{\Omega_{11}\Omega_{22}^*}{2\Delta'} e^{i(\phi_2 - \phi_1)} \\ \frac{\Omega_{11}^*\Omega_{22}}{2\Delta'} e^{-i(\phi_2 - \phi_1)} & \frac{|\Omega_{22}|^2}{2\Delta'} - 2\delta' \end{pmatrix}. \tag{2.71}$$

[12] Note that, for those not entirely satisfied, the method of integrating factors can be employed to solve Equation (2.68c) and get an exact integral expression for $\tilde{c}_3$ in terms of $\tilde{c}_{1,2}$. Substituting this into Equations (2.68a) and (2.68b), it should be possible to calculate the exact dynamics of the field vector driving the evolution of the effective two-level system on the Bloch sphere in order to gain a more geometric insight.

Defining new quantities

$$\delta_L \equiv \delta' - \frac{|\Omega_{11}|^2}{4\Delta'} + \frac{|\Omega_{22}|^2}{4\Delta'} \text{ and} \tag{2.72}$$

$$\phi_L \equiv \phi_1 - \phi_2 - \arg(\Omega_{11}\Omega_{22}^*), \tag{2.73}$$

this can be factorised as per Equation (2.15)

$$\hat{H}_2 = \frac{\hbar}{2} \Bigg(\left(\frac{|\Omega_{11}|^2}{4\Delta'} - \frac{|\Omega_{22}|^2}{4\Delta'} \right) \mathbb{I} + \frac{|\Omega_{11}\Omega_{22}^*|}{2\Delta'} \cos(\phi_L) \hat{\sigma}_x + \frac{|\Omega_{11}\Omega_{22}^*|}{2\Delta'} \sin(\phi_L) \hat{\sigma}_y + \delta_L \hat{\sigma}_z \Bigg), \tag{2.74}$$

with the associated field vector

$$\tilde{\mathbf{\Omega}} = \left(\frac{|\Omega_{11}\Omega_{22}^*|}{2\Delta'} \cos(\phi_L), \frac{|\Omega_{11}\Omega_{22}^*|}{2\Delta'} \sin(\phi_L), \delta_L \right), \tag{2.75}$$

where we resume employing the tilde as a reminder that we are in a rotating frame. We see that the two stable states are coupled by the two-photon *stimulated Raman* process in a manner directly analogous to the simple two-level, single-photon, picture we considered in Section 2.3. We can identify the on-resonance Rabi frequency

$$\Omega_R = \frac{|\Omega_{11}\Omega_{22}|}{2\Delta'} \approx \frac{|\Omega_{11}\Omega_{22}|}{2\Delta} \tag{2.76}$$

and a shifted resonance condition $\delta_L = 0$. The shift to the resonance, defined in Equation (2.72), has a pleasing interpretation: by comparison to Equation (2.39) we see that it is just the result of the far-off-resonant light shift experienced by the states $|1'\rangle$ and $|2'\rangle$ due to their couplings to the excited state by lasers $l = 1$ and $l = 2$ respectively. This is where the coupling terms that we neglected from Equation (2.62) — and indeed additional couplings to even-more-off-resonant states not previously considered — reenter the problem: they manifest as a shift to the resonance. In general we can clump these shifts together in a single parameter δ_{ac} and, redefining

$$\delta_L \equiv \delta' + \delta_{\text{ac}}, \tag{2.77}$$

our solution still holds.

The role of the laser phase from the single-photon case is played here by ϕ_L, the relative phase between the two lasers that corresponds to the phase of the beatnote that would exist between them. This can be varied by phase modulation of just one of the beams, as will become important in later chapters.

Less coherent manipulation of atoms with laser light

3

Guided by confused light
I am guided by confused light.

— The Front Bottoms, *Tie-dye dragon*

Absorption, stimulated emission and spontaneous emission. A physicist will be aware of these three processes from pretty much the start of their career, and will likely spend a great deal of time thinking about them regardless of their specialism.

They will likely be aware of Einstein's rate equations and his A and B coefficients long before any formal studies in quantum mechanics. A course in statistical mechanics will tell them that the rate of absorption B_{12}, corresponding to the excitation of an atom from energy level 1 to 2, is equal to the rate of stimulated emission B_{21} multiplied by the ratio g_2/g_1 of the level degeneracies. In the same course they will perhaps derive the relationship

$$A_{21} = \frac{\hbar\omega^3}{\pi^2 c^3} B_{21}, \tag{3.1}$$

that links the rate of spontaneous emission A_{21} to that of stimulated emission B_{21} (at a given angular frequency of radiation ω), from thermodynamic principles [79]. This may well be several years before they are proficient enough in quantum mechanics to derive an expression for B_{12} using time-dependent perturbation theory after which, for many, the picture will be complete. Some may go further, studying quantum field theory and coming to understand A_{21} quantitatively as arising from coupling of the atom to modes of the quantised electromagnetic vacuum.

If they go into experimental atomic physics, on the other hand, they will likely take home the message that the A and B coefficients are properties of the atom that result in the probabilistic decay of any excited state — and the corresponding emission of a photon — at a characteristic rate γ that has an associated lifetime $\tau = 1/\gamma$ and can be measured experimentally.[1] It is from this basis that we will launch into this chapter, phenomenologically

[1] When there are several possible states to which it can decay then there will be a characteristic rate for each one and γ will equate to their sum. The ratios of the individual decay rates to γ are then known as *branching ratios*. Again, these are measurable quantities, but for now we will concern ourselves with *closed* atomic transitions that have only one decay path.

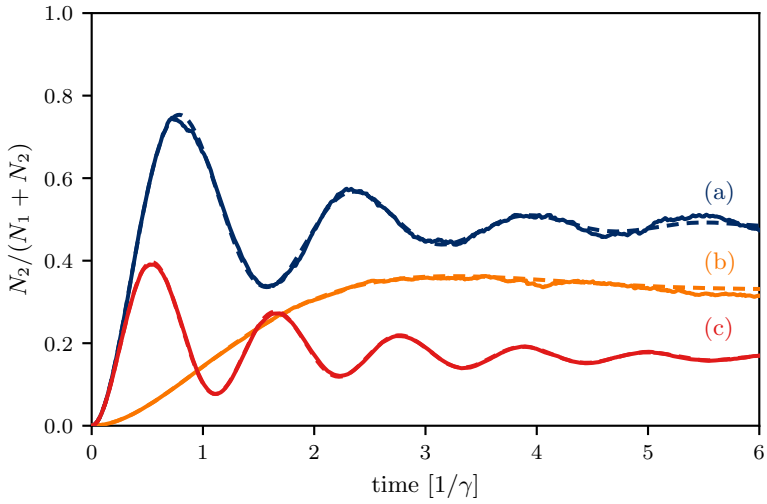

Figure 3.1: Monte Carlo simulations of the fraction of N=1000 atoms in the excited state during a laser excitation (solid lines) compared to a numerical integration of Equations (3.2) (dashed lines) for (a) $\Omega_R = 4\gamma$, $\delta = 0$, (b) $\Omega_R = 1\gamma$, $\delta = 0$, (c) $\Omega_R = \delta = 4\gamma$.

incorporating our understanding of spontaneous emission into the coherent dynamics that we discussed last chapter in order to determine how they are disrupted and what new physics emerges. This will lead to an understanding of the laser cooling and trapping processes used in all of the experiments that follow, and that are central to the work in Part III.

3.1 Spontaneous emission in the OBEs

Consider an ensemble of atoms, all prepared in the ground state and driven by a resonant laser. Each atom will undergo Rabi oscillations at the same rate, and these could be observed by measuring the fraction of atoms in the excited state after interaction with the laser for varying times t. That is, until spontaneous emission is considered. This is a stochastic process that an individual atom will instantaneously return to the ground state at a random time so that it no longer oscillates in sync with the rest of the ensemble.

After many spontaneous emission events, the relative phases between the Rabi oscillations of the atoms in the ensemble will be completely randomised and, on measuring the fraction of atoms $N_2/(N_1 + N_2)$ in the excited state, one would expect to measure a value that is independent of t.

The rate of this randomisation process is determined by γ and one might infer that, for $\Omega_R \gg \gamma$, many Rabi oscillations would be observed before $N_2/(N_1 + N_2)$ eventually settles to a steady-state

value of 1/2. On the other hand, for $\Omega_R \ll \gamma$, the probability of an atom completing a full Rabi cycle before a decay occurs is small. Atoms in the ensemble will tend to remain near the 'north' pole of the Bloch sphere — corresponding to the ground state $|1\rangle$ — no oscillations would be discernible and, for decreasing Ω_R, one would find $N_2/(N_1 + N_2) \to 0$.

We can show this by conducting a Monte Carlo simulation of an ensemble of atoms, incorporating these stochastic decays, to produce the solid lines in Figure 3.1 that represent the fraction of the atoms measured in the excited state as a function of time for different values of Ω_R and the detuning δ. The dashed lines are a numerical integration of the full form of the OBEs that incorporate decay terms to account for these new dynamics:

$$\frac{\mathrm{d}\tilde{u}}{\mathrm{d}t} = -\delta\tilde{v} + \Omega_R \sin(\phi)\tilde{w} - \frac{\gamma}{2}\tilde{u} \tag{3.2a}$$

$$\frac{\mathrm{d}\tilde{v}}{\mathrm{d}t} = \delta\tilde{u} - \Omega_R \cos(\phi)\tilde{w} - \frac{\gamma}{2}\tilde{v} \tag{3.2b}$$

$$\frac{\mathrm{d}\tilde{w}}{\mathrm{d}t} = -\Omega_R \sin(\phi)\tilde{u} + \Omega_R \cos(\phi)\tilde{v} + \gamma\,(1 - \tilde{w})\,. \tag{3.2c}$$

A derivation, or at least justification, of these equations is given in Appendix A. It is important to highlight a shift here: $\tilde{\mathbf{\Psi}} = (\tilde{u}, \tilde{v}, \tilde{w})$ now represents the average state vector for all of the atoms in the ensemble. Formally, the ensemble would be represented by a *density matrix* that arises from a reformulation of quantum mechanics to describe such statistical mixtures [81]. In practice, at least for this thesis, it is sufficient to continue thinking of the individual atomic state vectors as lying on the surface of the Bloch sphere, while noting that — as they may point in different directions — their average $\tilde{\mathbf{\Psi}}$ may reside inside it.

3.2 Dephasing versus decoherence

As a brief aside, it is worth noting that spontaneous emission damping is not the only reason that we observe decay in Rabi oscillations. In many atom interferometry experiments, including those in this thesis, two-photon Raman transitions such as those considered in Section 2.6 are employed that result in an effective two-level system with negligible spontaneous emission. Nonetheless, Rabi oscillations will decay over time periods much shorter than the effective lifetime of the dipole-forbidden transition.

These decays are caused by variations in laser detuning — typically Doppler induced — and Rabi frequency among the atoms in the ensemble. Each atom undergoes uninterrupted, sinusoidal Rabi oscillations for the duration of the interaction, but the spread of different frequencies and amplitudes within the ensemble results in them dephasing over time. However, this is still a coherent process; the system retains a memory of its initial state and in principle the interaction could be reversed so that all of the oscillations rephase again. This is beautifully demonstrated in the spin echo Rabi oscillations presented in [89].[2]

[2]This remains true on timescales over which the Rabi frequency and detuning for a single atom remain constant. Transverse motion of an atom in a Gaussian beam, for instance, would result in an irreversible change of the Rabi frequency. While not strictly decoherence, from an experimental perspective it might as well be. If it looks like a duck and quacks like a duck...

By contrast, the random nature of spontaneous emission acts to erase any dependence of the system on the initial conditions and is inherently an irreversible, incoherent, process. It is thus an example of *decoherence*, as opposed to the *dephasing* just described. While dephasing can be mitigated with coherent control techniques like those we investigate in Chapter 9, decoherence can only be combated by minimising the coupling of the system to its environment.

3.3 Forces on a two-level atom

The damping terms introduced to the OBEs to account for spontaneous emission allow them to reach a steady state. Physically, each atom will plot a cyclic trajectory on the surface of the Bloch sphere that, after its first spontaneous emission, is guaranteed to pass through the ground state, but with a repeatedly randomised phase. For times $t \gg \tau$ the state of the ensemble is completely independent of its initial state, and is characterised by an average field vector determined by setting the time derivatives in Equations (3.2) to zero. If we let the laser phase $\phi = 0$ then we find the components of this averaged field vector to be

$$\tilde{u} = 2\frac{\delta}{\Omega_R}\frac{s_0}{1+\left(2\delta/\gamma\right)^2+s_0} \tag{3.3a}$$

$$\tilde{v} = -\frac{\gamma}{\Omega_R}\frac{s_0}{1+\left(2\delta/\gamma\right)^2+s_0} \tag{3.3b}$$

$$\tilde{w} = \frac{1+\left(2\delta/\gamma\right)^2}{1+\left(2\delta/\gamma\right)^2+s_0}. \tag{3.3c}$$

Here

$$s_0 = 2\left(\frac{\Omega_R}{\gamma}\right)^2 \tag{3.4}$$

is called the *saturation* parameter. When it is large $\tilde{w} \to 0$, corresponding to half of the ensemble being in the excited state, and when it is small $\tilde{w} \to 1$ and most of the ensemble is found in the ground state. This quantifies what we inferred from intuition in Section 3.1: it is the ratio of the Rabi frequency to the decay rate that determines the fraction of excited atoms in the steady state.

In Section 2.5 we identified that, when $\phi = 0$, $\tilde{u}$ and $\tilde{v}$ are respectively proportional to the amplitudes of the in-phase and quadrature components of the atomic dipole moment relative to the laser. Now, in the steady state limit, Equations (3.3) represent the components of both the *instantaneous* ensemble-averaged state vector, and the *time-averaged* state vector of any individual atom. Using the steady state values of $\tilde{u}$ and $\tilde{v}$ in our expression for the dipole moment, then, we can calculate the time-averaged force the laser exerts on an atom from the instantaneous force

$$\boldsymbol{F} = -\nabla\left(-\langle\hat{\boldsymbol{d}}\cdot\hat{\boldsymbol{E}}\rangle\right), \tag{3.5}$$

the spatial derivative of the expectation value for the energy of the atomic dipole moment in an electric field $\boldsymbol{E}$.

In common with our treatment in Chapter 2, we will continue to consider the electric field to be linearly polarised in the x direction and assume that the atomic dipole moment aligns to it. Furthermore, we will consider a travelling wave in the positive z direction, again with amplitude E, giving $\boldsymbol{E} = E\cos(\omega t - kz)\,\boldsymbol{e_x}$. Combined with the expression for the dipole moment in Equation (2.37), we arrive at an expression for the instantaneous force at $z = 0$ [95]:[3]

$$\begin{aligned}\boldsymbol{F} = -\hbar\left(\tilde{u}\cos(\omega t) + \tilde{v}\sin(\omega t)\right)\\ \times\left(\nabla\Omega_R\cos(\omega t) + \Omega_R k\sin(\omega t)\boldsymbol{e_z}\right),\end{aligned} \tag{3.6}$$

where we account for the fact that the Rabi frequency may vary spatially due to amplitude variations of the electric field within, for instance, a focused Gaussian beam.

[3] Recalling that, for our linearly polarised light, $e\langle 1|\hat{x}|2\rangle = \hbar\Omega_R/E$ and is real. We also persist with the dipole approximation, assuming that the dimension of the atom is small with respect to the optical wavelength.

Explicitly averaging over the fast oscillations at the optical angular frequency ω then only the products of in-phase terms remain, picking up a factor of $1/2$ from the cycle average,

$$\langle \boldsymbol{F} \rangle = -\frac{1}{2} \left(\tilde{u} \hbar \nabla \Omega_R + \tilde{v} \Omega_R \hbar k \boldsymbol{e_z} \right). \tag{3.7}$$

By equating $\tilde{u}$ and $\tilde{v}$ to their steady-state values given by Equations (3.3), we implicitly consider the force to be averaged over timescales corresponding to the decay rate γ, that is typically much slower than the optical oscillations. Doing so, we can immediately make a couple of qualitative observations.

Firstly, when $\tilde{u}$ is large and negative then the atomic dipole moment is in phase with the electric field, aligning to it. This is energetically favourable, and the atom will reduce its energy still further if the electric field is increased. The first term then results in a force proportional to $\nabla \Omega_R$ and the atom climbs the hill to the point where the electric field amplitude is greatest. On the other hand, if $\tilde{u}$ is large but positive, the dipole moment is π out of phase with the driving field and the atom feels a force towards lower amplitudes. From Equation (3.3a) we see that the sign of $\tilde{u}$ corresponds to that of the detuning δ. We conclude that, on average, atoms will be attracted to regions of high-intensity red-detuned light and repelled from regions of high-intensity blue-detuned light.

Secondly, when $\tilde{v}$ is large and negative then the phase of the oscillating dipole moment is retarded by $\pi/2$ relative to that of the electric field.[4] Classically, we would identify this as the condition in which the electric field does maximal work, transferring the most energy to the dipole with each cycle. Quantum mechanically, we might infer that this is when the atom exchanges the most photons with the field. As spontaneous emission in all directions breaks the symmetry between the rate of absorption and stimulated emission in the direction of the travelling wave, over many absorption cycles this exchange results in a force along the direction of the beam. The second term in Equation (3.7) encapsulates this and, noting the presence of the photon momentum $\hbar k \boldsymbol{e_z}$, we might infer — and will confirm shortly — that $\tilde{v}\Omega_R/2$ corresponds to the photon scattering rate.

[4] By inspection of Equation (3.3b), we see that the steady state value of $\tilde{v}$ is never positive.

The two terms in Equation (3.7) are therefore commonly referred to as the *dipole force* and *scattering force* respectively, $\langle F \rangle = \langle F_{\text{dip}} \rangle + \langle F_{\text{scatt}} \rangle$. The ratio $\tilde{u}/\tilde{v} = -2\delta/\gamma$ confirms that the

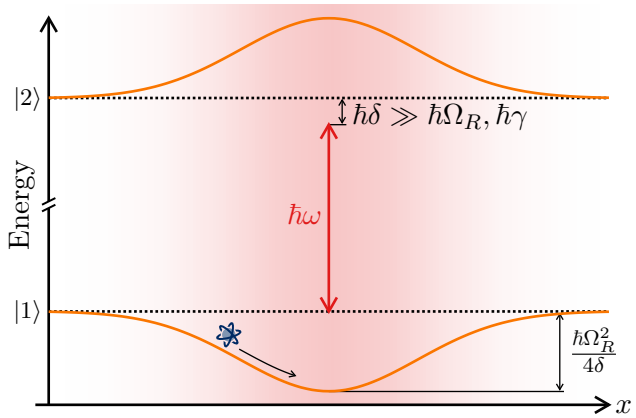

Figure 3.2: A red-detuned Gaussian beam provides a conservative trapping potential for two-level atoms in the ground state, and this can be explained by considering the spatial dependence of the light shift. It is worth noting that the excited state exhibits anti-trapping, and so off-resonant excitation is a potential loss, or at least heating, mechanism.

dipole term dominates at large detuning, while the scattering term dominates near resonance.

3.3.1 The dipole force

A significant result arises if we rewrite the dipole term as the spatial derivative of an effective potential:

$$\begin{aligned}\langle F_{\text{dip}}\rangle &= -\nabla\left[\frac{\hbar}{2}\int \tilde{u}\,\mathrm{d}\Omega_R\right] \\ &= -\nabla\left[\frac{\hbar}{2}\delta\ln\left(\gamma^2+4\delta^2+2\Omega_R^2\right)\right] \\ &= -\nabla\left[\frac{\hbar}{2}\delta\ln\left(1+\frac{2\Omega_R^2}{\gamma^2+4\delta^2}\right)\right] \approx \quad -\nabla\left[\frac{\hbar\Omega_R^2}{4\delta}\right],\end{aligned} \tag{3.8}$$

in the limit of large detuning $\delta \gg \gamma, \Omega_R$. We can identify the potential in its final form as the far off-resonant light shift from Equation (2.41); in this limit the probability of excitation is negligible and the atoms seek a minimum of the spatially varying energy landscape arising from light shift to the ground state. This has become second nature to atomic physicists since the concept was introduced by Dalibard and Cohen-Tannoudji [96], and diagrams such as that in Figure 3.2 have become common. This shows the spatial dependence of the light shift to the states of a two-level atom induced by a red-detuned Gaussian laser beam, demonstrating that the shift to the ground state then provides a

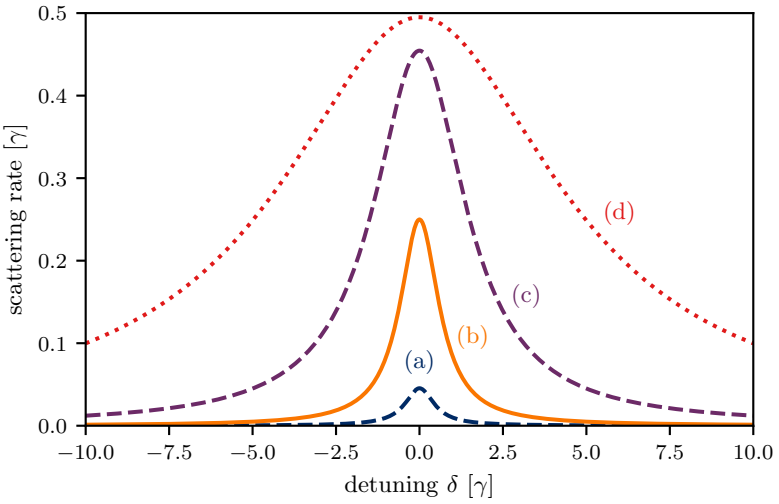

Figure 3.3: Scattering rate as a function of detuning for (a) $s_0 = 0.1$ (b) $s_0 = 1$ (c) $s_0 = 10$ (d) $s_0 = 100$. For $s_0 > 1$ the width of the resonance is significantly broadened. Figure based on one from Metcalf and van der Straten [97].

conservative potential that can trap atoms with kinetic energies smaller than the maximum light shift induced at the centre. This in the governing principle of the far off-resonant optical dipole trap (FORT), the optimal loading of which is the principle focus of the work in Part III.

3.3.2 The scattering force

The decay rate of the excited state is just $\gamma|c_2|^2 = \gamma/2(1-\tilde{w})$. In the steady state, the excitation rate and decay rate are balanced, and so this is equivalent to the rate at which photons are absorbed from the laser and scattered into the vacuum. As the photons from the laser all propagate in the same direction while the scattered photons are emitted in all directions, this results in a force in the direction of the laser beam which, noting that $1/2\Omega_R\tilde{v} = \gamma/2(1-\tilde{w})$, is in quantitative agreement with the $\langle F_{\text{scatt}}\rangle$ term in Equation (3.7).

The scattering rate as a function of detuning is plotted in Figure 3.3 for different values of the saturation parameter $s_0 = 2\left(\Omega_R/\gamma\right)^2$. It reaches its peak value of $\gamma/2$ for very large values as anticipated, but we also observe that for $s_0 > 1$ the linewidth is significantly broadened in another example of *power broadening*.

As the saturation parameter scales with Ω_R^2 that is proportional to the optical intensity I, it is often expressed in the form $s_0 = I/I_{\text{sat}}$ where the *saturation intensity* I_{sat} is given by

$$I_{\text{sat}} = \frac{4\pi^2}{3}\frac{\hbar\omega}{\lambda^2}\gamma. \tag{3.9}$$

We shall not derive this here but, written in this form, it is clear that at this intensity an average of one photon impinges on an area $\sim \lambda^2$ during the lifetime of the excited state. Above such an intensity one might expect each atom to start 'seeing' more than one photon per decay cycle.

3.3.3 Concluding remarks

One could argue that this section describes everything there is to know about the force light exerts on atoms under typical laboratory conditions. Again we have eschewed rigour in favour of physical insight, but Equation (3.7) is just a specific form of a more general equation for the force that is summed up concisely in Equation (3.5). Combined with the steady-state solutions to the OBEs in Equations (3.3), the picture is complete.

But that is like being given the standard model and being expected to understand everything about how the observable universe works. In practice, predicting the effects of different alignments and polarisations of light on real, multi-level atoms with complex selection rules is not simple. A host of surprising and useful effects emerge from these fundamental principles that are by no means obvious.[5]

We will shift gear now and, for the remainder of this chapter, take a whistle-stop tour of just a few of these that are relevant to later chapters, focusing on the cooling and trapping of neutral atoms. Books can be, and have been, filled with beautiful and subtle analyses of these effects [97]; here we shall just summarise the results with a focus on qualitative explanations.

[5] As in any field, it is often unexpected results yielded by experiments that inform theoretical advances, in turn leading to new insights that inspire future experiments. When it was theoretically suggested, nobody expected laser cooling to be as efficient as it proved to be, and this opened up a rich seam of theoretical insights into a host of sub-Doppler cooling mechanisms, a couple of which we shall outline in Section 3.6.

3.4 Optical molasses

One might initially expect the scattering force on an atom from two near-resonant, counter-propagating laser beams with equal intensity to exactly cancel. This is to neglect the velocity dependence of the force that enters through the contribution of the Doppler shift to the detuning δ. An atom travelling in the opposite direction to a negatively detuned beam will see it Doppler shifted closer to

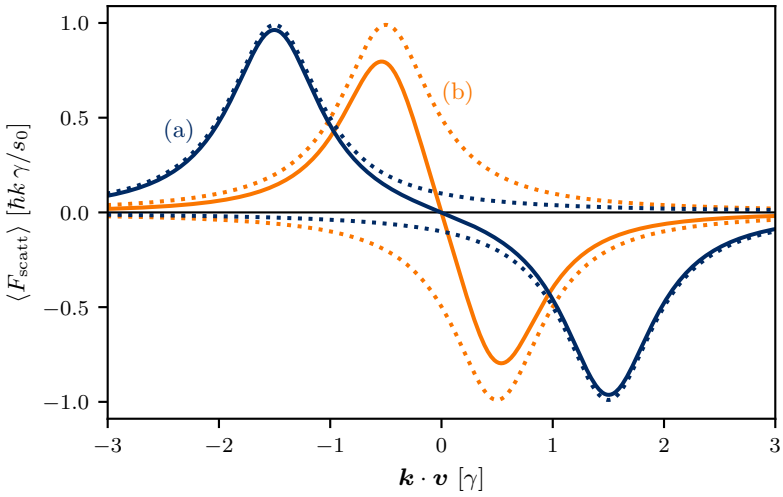

Figure 3.4: Force exerted by two counter-propagating laser beams with $s_0 = 0.01$ as a function of atomic velocity for (a) $\delta = -3\gamma/2$ (b) $\delta = -\gamma/2$. Dotted lines show the force from each individual beam, solid lines show their sum.

resonance while $|\boldsymbol{k} \cdot \boldsymbol{v}|$ is less than the detuning. As the scattering force is maximal on resonance, one can envisage that an atom subjected to two counter-propagating red-detuned beams will scatter more photons from the one that it is travelling against and feel a resultant force that opposes its velocity, slowing it down.

More precisely, we can engineer a gradient $\frac{\partial \langle F_{\text{scatt}} \rangle}{\partial v} < 0$ about $v = 0$ that acts to damp the atomic motion. This gradient is typically not trivial to calculate but, for $I \ll I_{\text{sat}}$ we can treat the total force as just the sum of the individual forces that would be exerted by each beam in isolation.[6] This is shown in Figure 3.4, where it is clear that a non-negligible gradient about $v = 0$ is manifest for detunings on the order of the natural linewidth γ. The force diminishes again for $|\delta - \boldsymbol{k} \cdot \boldsymbol{v}| \gtrsim \gamma$, but atoms moving slower than this *capture velocity* feel the desired damping towards $v = 0$. It is the atomic analogue of trying to move through treacle, and so this technique is known as an *optical molasses*. With beams aligned along all three Cartesian axes, motion in all dimensions in thus damped.

This would appear to imply that the affected atoms will be brought to a complete halt, a clearly nonphysical result. In fact there will be a minimum spread in velocities that tends to be characterised by the temperature T_D that this spread would correspond to if the atoms were in an equilibrium state

$$k_B T_D = \frac{\hbar \gamma}{2}. \tag{3.10}$$

[6] The treatment can be generalised to the higher intensities usually employed in experiments by including a saturation parameter that takes into account the total intensity of both beams, or all 6 beams in a 3-D case. However, if the counter-propagating beams have non-orthogonal polarisations then there will be a standing wave present, the intensity variations in which affect the cooling non-trivially [98]. This complicates the analysis beyond the scope of this work, and the low intensity case offers a good qualitative description.

This is called the *Doppler cooling limit* and can be derived in a variety of ways. It arises from the fluctuations in $\langle F_{\text{scatt}} \rangle$ about the average value due to the momentum — or indeed energy — fluctuations of the spontaneously emitted photons. A thorough derivation can be found in, for example, [79] but, as the lower limit on the frequency spread of these photons is set by the natural linewidth γ, the form of Equation (3.10) seems eminently reasonable at face value.

3.5 The magneto-optical trap

So far we have limited our discussion to two-level atoms under the influence of linearly polarised light. A wealth of possibilities opens up once we extend these basic principles to the multi-level atoms and range of polarisations found in real experiments. One such possibility is the magneto-optical trap (MOT) [13], that employs a magnetic field gradient to add a restorative nature to the optical molasses force for atoms that exhibit the requisite Zeeman structure. This simple and robust technique allows the preparation of the cold, dense and localised clouds of atoms that form the initial stage of most cold atom experiments.

We will qualitatively outline this here for the simplest case of cooling and trapping on a $J = 0 \rightarrow J' = 1$ atomic transition. A magnetic field $\boldsymbol{B}$ will then break the degeneracy of the $m'_J = 0, \pm 1$ sublevels of the excited state. Consider the case, illustrated in Figure 3.5, where a linear magnetic field gradient is engineered in the z direction with $\boldsymbol{B} = 0$ at $z = 0$. As an atom moves in the positive z direction, the $m'_J = 1$ level experiences a positive energy shift proportional to the distance it moves from the origin and the $m'_J = -1$ level experiences an equal and opposite shift. These shifts are reversed in the negative z direction.

Now let us realise an optical molasses along the z axis with red-detuned, right-circularly polarised light. Along the quantisation axis z the light propagating in the positive direction is then σ^+ polarised, driving the $\Delta m_J = +1$ transition, while the σ^- light propagating in the negative direction drives the $\Delta m_J = -1$ transition. As Figure 3.5 illustrates, for $z < 0$ the magnetic field shifts the $\Delta m_J = 1$ transition towards resonance, and the $\Delta m_J = -1$ transition further from resonance making it more probable that an atom will absorb photons from the positive-propagating beam and feel a net force towards $z = 0$. The further it gets from the origin,

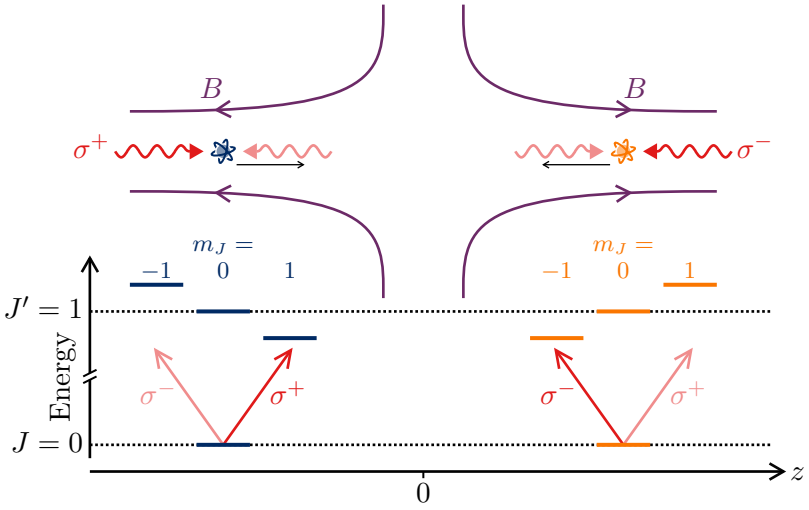

Figure 3.5: Operating principle of the MOT. *Above:* schematic representation of the optical and magnetic fields acting on the atoms; *Below:* structure of the Zeeman shifted atomic states on either side of the trap with the optical couplings illustrated, showing that on either side of the trap and atom will scatter more photons from the beam pushing it towards the origin $z = 0$.

the greater this effect,[7] and again the opposite is true on the other side of what we can now call the trap.

[7] For as long as the magnetic shift is smaller than the detuning.

A pair of anti-Helmholtz coils with current circulating in opposite directions, as illustrated in Figure 3.6(a), can be used to generate a magnetic field such as that sketched in Figure 3.5 with rotational symmetry about the z axis. At the centre of such an arrangement, a linear magnetic field gradient is then found in all directions. Three pairs of counter-propagating laser beams with appropriate polarisations can thus be employed to cool and trap atoms at the centre, in a *3-D MOT*.

Alternatively, four coils arranged as per Figure 3.6(b), or else four magnets in a quadrupole configuration, can be used to realise the magnetic field in Figure 3.5 with translational symmetry along the axis directed out of the page. In this case, counter-propagating beams in the horizontal and vertical directions can be used to realise 2-D confinement; the atomic density is, of course, low in this arrangement as atoms are lost in 'beams' emerging from either end of the trap.

These beams, however, provide a particularly directional flux of transversely cooled atoms. Such a *2-D MOT* [99–101] can thus be used — often in conjunction with a near-resonant 'push' laser beam aligned along the trap axis, the scattering force from which enhances atomic flux in a particular direction — as a source of

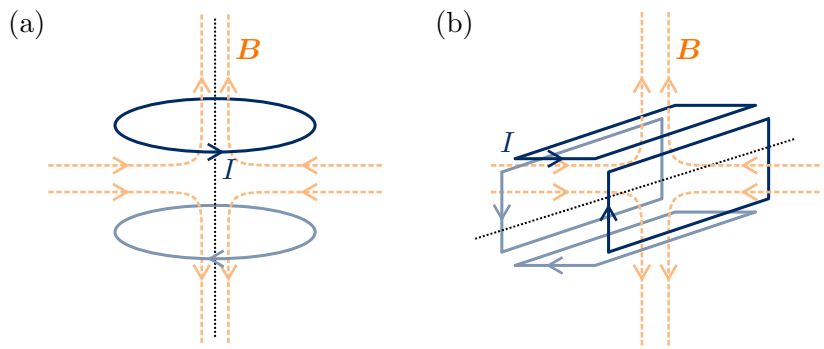

Figure 3.6: (a) 3-D and (b) 2-D MOT coil geometries. A current I flowing in the configurations shown by the solid (blue) lines produces a magnetic quadrupole field $\boldsymbol{B}$, dashed (orange) lines, with (a) rotational or (b) translational symmetry about the respective dotted (black) axes.

atoms for the loading of a 3-D MOT; this is particularly useful in experiments where long coherence times are desired for atoms in the vicinity of the 3-D MOT, negating the possibility of loading it from a background vapour, as is the case in the experiments described in Part III.

This magneto-optical cooling and trapping of atoms works far better than these arguments imply that it should. Even the very first MOTs produced a trapped cloud of atoms with a momentum spread much lower than that implied by the Doppler cooling limit in Equation (3.10). This fortuitous result stems from a subtle interplay between the dipole and scattering forces, similar to that outlined below.

3.6 Sub-Doppler cooling mechanisms

3.6.1 Sisyphus cooling

Consider the $J = 1/2 \rightarrow J' = 3/2$ transition shown in Figure 3.7. Driven by σ^+ polarised light, the ground state population can be *optically pumped* into the $m_J = 1/2$ sublevel by virtue of the fact that the pump light only drives $\Delta m_J = +1$ transitions while spontaneous decay can occur on any dipole-allowed transition. Similarly, driving with σ^- polarised light pumps the ground state population into the $m_J = -1/2$ state.

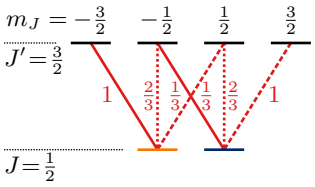

Figure 3.7: Zeeman structure of a $J = 1/2 \rightarrow J = 3/2$ transition conducive to Sisyphus cooling. Solid, dashed and dotted red lines show σ^-, σ^+ and π transitions, with their relative coupling strengths labelled.

Now consider two laser beams of the same frequency, slightly red-detuned from this $J = 1/2 \rightarrow J = 3/2$ transition, counter-propagating along the z axis with orthogonal linear polarisations. The result is a form of standing wave — illustrated in the lower

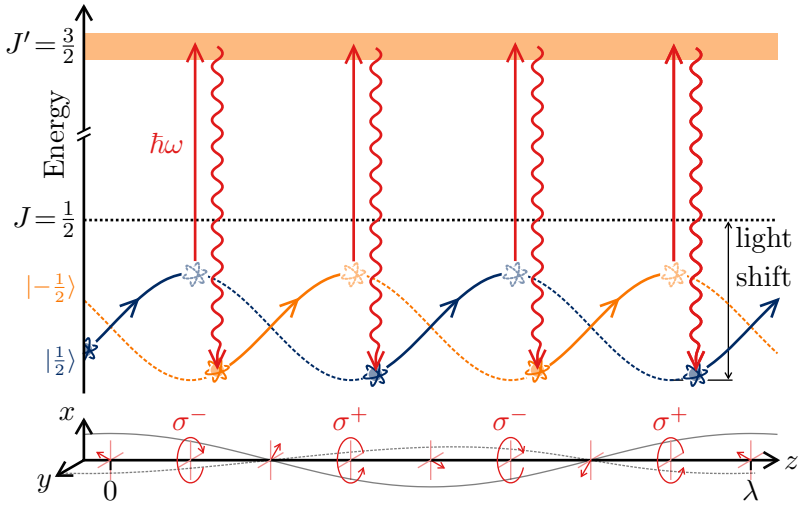

Figure 3.8: A simple example of Sisyphus cooling. *Below:* the stationary polarisation gradient that arises from two counter-propagating beams with orthogonal linear polarisations; *Above:* The resulting variation in the differential light shift between the $m_J = \pm 1/2$ ground states of the level scheme shown in Figure 3.7, in combination with optical pumping, means that atoms are always climbing 'uphill' in the energy landscape as the propagate along the beam. The dissipation of kinetic energy that follows can narrow the atomic momentum distribution below the Doppler cooling limit.

part of Figure 3.8 — in which, over a distance of half a wavelength, the polarisation will vary from σ^+ to σ^- and back again, interspersed with regions where the polarisation is linear. A slow-moving atom propagating along this standing wave will thus be alternately pumped between the $m_J = \pm 1/2$ ground levels.

The light shift from the red-detuned light will act to reduce the energy of these levels in proportion to the coupling strength. The ratios between coupling strengths of the various transitions are given by the Clebsch-Gordan coefficients [102], labelled in Figure 3.7, and we see that for σ^+ polarised light (dashed lines) the light shift of the $m_J = 1/2$ ground state is greater than that of the $m_J = -1/2$ ground state by a factor of 3. For σ^- polarised light the opposite is true and, as shown in the upper part of Figure 3.8, the optical pumping thus always drives atoms into a state of lower energy.

As they propagate along the beam, the lower energy state becomes the higher energy state as the polarisation reverses. We see that each atom is continuously climbing up a potential energy gradient and so losing kinetic energy, that gets dissipated via the energy difference between the absorbed and spontaneously emitted photons represented respectively by straight and wobbly arrows in Figure 3.8.

This cooling mechanism is an example of a family of methods known as *polarisation gradient cooling*, or else *Sisyphus cooling*,[8] and the natural timescale is the time it takes for atoms to be pumped between ground states. At low laser power, this is much longer than the excited state lifetime that is responsible for the Doppler cooling limit, allowing much colder effective temperatures to be reached [102]. As, in principle, the laser power can always be reduced, the fundamental limit to the momentum spread is thus set by the single-photon recoil momentum, defined in Section 2.6.1, that is an unavoidable consequence of any exchange of energy between an atom and the optical field. In a typical experiment, sub-Doppler molasses cooling of rubidium atoms such as those used in this thesis can yield sample temperatures as much as two orders of magnitude below the Doppler cooling limit.

[8] Named after the king from Greek mythology who was condemned to eternally push a great boulder up a hill, only for it to roll to the bottom again upon nearing the top.

3.6.2 VSCPT and grey molasses

Consider again the Raman system described in Section 2.6, and in particular the rotating frame Hamiltonian in Equation (2.67). If, for simplicity, we assume that $\phi_1 = \phi_2 = 0$ and that Ω_{11} and Ω_{22} are real then, via a further rotation, we can recast the system in a basis made up of states $|\psi_D\rangle$, $|\Psi_B\rangle$ and $|3, p + \hbar/2(k_1 + k_2)\rangle$ where

$$|\psi_D\rangle = \frac{1}{\Omega}\left(\Omega_{22}\,|1, p - \hbar k/2\rangle - \Omega_{11}\,|2, p + \hbar k/2\rangle\right), \text{ and} \tag{3.11a}$$

$$|\psi_B\rangle = \frac{1}{\Omega}\left(\Omega_{11}\,|1, p - \hbar k/2\rangle + \Omega_{22}\,|2, p + \hbar k/2\rangle\right). \tag{3.11b}$$

Here $\Omega = \sqrt{\Omega_{11}^2 + \Omega_{22}^2}$ [103]. For $\delta' = -kp/m = -kv = 0$ we find that $\langle 3, p + \hbar/2(k_1 + k_2)|\,\hat{H}_R\,|\psi_D\rangle = \langle\psi_B|\,\hat{H}_R\,|\psi_D\rangle = 0$, and so $|\psi_D\rangle$ becomes a *dark state* that does not couple to either of the other states. An atom in this state therefore does not absorb photons or experience a light shift, in contrast to the *bright state* $|\psi_B\rangle$ that does.[9]

In general, the coupling of the dark state to the bright state is proportional to the average atomic momentum p,

$$\langle\psi_B|\,\hat{H}_R\,|\psi_D\rangle = -\frac{\Omega_{11}\Omega_{22}}{\Omega^2}\hbar k\frac{p}{m}, \tag{3.12}$$

[9] Note the shift in basis here to one where p represents the average momentum of the closed momentum family. With counter-propagating beams $k_1 \approx -k_2$ and this is essentially equivalent to the momentum of the excited state $|3\rangle$. The detuning δ' is then correspondingly redefined to be relative to the recoil shifted resonance. This feels like a more natural basis [104], and is commonly employed in the context of the sort of cooling mechanisms we are describing here, but the basis used in Section 2.6 seems to be more common in atom interferometry literature.

so that for $p \neq 0$ the dark state will evolve into the bright state and scatter photons. This scattering, mediated by spontaneous emission from the excited state $|3, p + \hbar/2(k_1 + k_2)\rangle$, allows atoms to jump between momentum classes until they arrive in the $p = 0$ dark state from which they no longer scatter. This process of atoms collecting in a specific momentum state that no longer 'sees' the light is a cooling mechanism known as velocity-selective coherent population trapping (VSCPT) [105, 106] that has a couple of interesting properties.

In this description, an atom does not feel a velocity dependent force. Indeed, when the lower levels $|1\rangle$ and $|2\rangle$ are degenerate and the laser beams are counter-propagating then, on average, there is no force at all. The process is driven by momentum fluctuations induced by the spontaneous emission, the magnitude of which are governed by the natural linewidth of the excited state. In the $p = 0$ dark state these fluctuations vanish as the atom no longer scatters light. We thus find ourselves in the curious position in which the very fluctuations that were responsible for the Doppler cooling limit in Section 3.4 are here responsible for distinctly sub-Doppler cooling!

As the $p = 0$ dark state is a superposition of momentum states $|1, -\hbar k/2\rangle$ and $|2, \hbar k/2\rangle$, a measurement of the momentum distribution thus yields two distinct peaks. Again considering degenerate states with counter-propagating beams, a nice picture emerges to explain the lack of scattering from the dark state [97]: spatially, the two momentum states that comprise it are counter-propagating plane waves so that the state is a completely delocalised standing wave whose fixed phase relationship to the standing wave formed by the laser beams ensures that the spatial integral of the dipole transition matrix element vanishes.

In the presence of a polarisation gradient, the assertion that there is no velocity-dependent force exerted is no longer true and a Sisyphus-type mechanism can again be employed to enhance the efficiency of the VSCPT process [107, 108]. The bright state is subject to a spatially dependent light shift from its coupling to the excited state that is periodic at half the optical wavelength, while the energy of the dark state remains fixed.

For $p \neq 0$ the dark state is again motionally coupled to the bright state and undergoes a non-adiabatic transfer, but now the spatially varying light shift needs to be taken into account using first-order

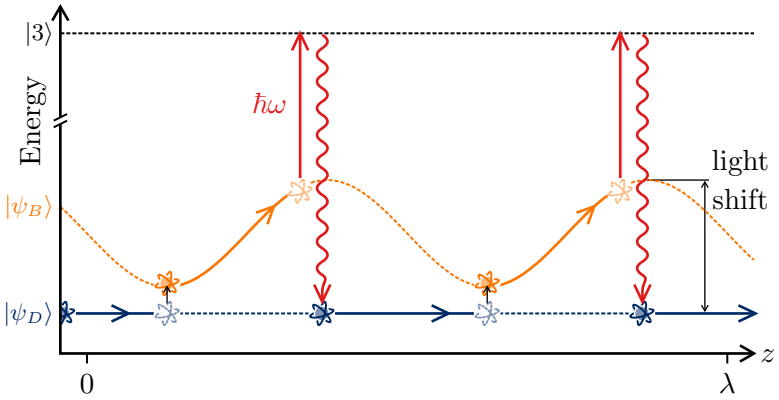

Figure 3.9: In the presence of a polarisation gradient such as that shown in Figure 3.8, another Sisyphus-type mechanism acts to introduce a velocity-dependent force to VSCPT in what is known as a grey molasses. The bright state couples to the excited state and experiences a spatially varying light shift, that is periodic at half the optical wavelength λ. At $p \neq 0$, the dark state undergoes non-adiabatic transfer to the bright state when the light shift is minimal, before being optically pumped back to the dark state when the light shift is close to maximal. For blue-detuned light this results in a damping force.

perturbation theory. The result is a transfer rate that is proportional to the atomic velocity and inversely proportional to the light shift, so that moving atoms tend to enter the bright state at points in the standing wave where the light-shift is minimal.

If the laser detuning is positive then the light shift to the bright state is positive and so, as shown in Figure 3.9, the atoms gain potential energy as they propagate along the polarisation gradient and the light shift increases. Near the maximum light shift they are optically pumped back into the dark state, dissipating kinetic energy as in conventional Sisyphus cooling. One differentiating feature of this form of VSCPT, enhanced by a polarisation gradient, relative to a standard Sisyphus molasses is that in this case it is a blue rather than a red detuning that results in a cooling force.

This simple treatment for a Λ-type Raman system extends well to real atomic transitions [106, 109], although in real systems the dark states are often comprised of basis states with different kinetic energies. The associated dephasing results in an eventual coupling of such states to the optical field so that they are not truly dark over finite timescales. This has lead to polarisation gradient cooling with such states being referred to as *grey molasses* [110].

Such techniques have recently been employed to sub-Doppler cool a variety of atomic species and achieve phase space densities sufficient for effective evaporative cooling where conventional Sisyphus

cooling proves slow of ineffective [111–117]. A key component of the work described in Part III involved realising such a scheme on the ^{87}Rb D_2 line.

Atom interferometry

4

Any form of transport which involved tearing you apart atom by atom, flinging those atoms through the sub-etha, and then jamming them back together again just when they were getting their first taste of freedom for years had to be bad news.

— Douglas Adams, *The Restaurant at the End of the Universe*

Detailed descriptions of the atom interferometry conducted during this thesis are given in the relevant chapters, with Section 7.1 being particularly pertinent. However, it would be remiss not to give it some treatment during these introductory chapters. In the following sections, then, we give a brief overview of the standard picture of atom interferometry presented in many texts, before considering a simpler and slightly less common description that is pleasingly classical and will lead us nicely into the next chapter, where we can have some fun with boats and clocks.

4.1 A standard picture

We will discuss in particular atom interferometers that are employed to measure acceleration and, by extension, rates of rotation. These are particularly common and are employed, among other things, to make precision gravity measurements [22–25, 27, 118] and highly sensitive inertial sensors [37, 38, 41, 42, 119]. They rely on a spatial separation of the quantum atomic wavefunction and it has become standard to achieve this with light–atom interactions, taking advantage of the recoil momentum associated with absorption of a photon. As discussed in Section 2.6.1, coupling two internal atomic states with a laser also couples two distinct momentum states of the atom. Employing a $\pi/2$-pulse, as described in Section 2.4, to prepare a superposition of internal states therefore also amounts to preparing a superposition of external momentum states. A suitable time after the interaction, the atomic wavefunction can thus be considered as being split into two discrete wavepackets propagating along distinct spatial paths [21, 120].

As suggested in Section 2.6, it is common to use alkali atoms for atom interferometry and employ transitions between hyperfine ground levels. Using two counter-propagating lasers to drive two-photon Raman transitions between two such states results in an effective two-level system where the coupled states are separated considerably in momentum thanks to the short wavelength of the optical photons, while the long-term stability of the basis states means that superpositions between them also remain stable and evolve coherently for meaningful experimental timescales.

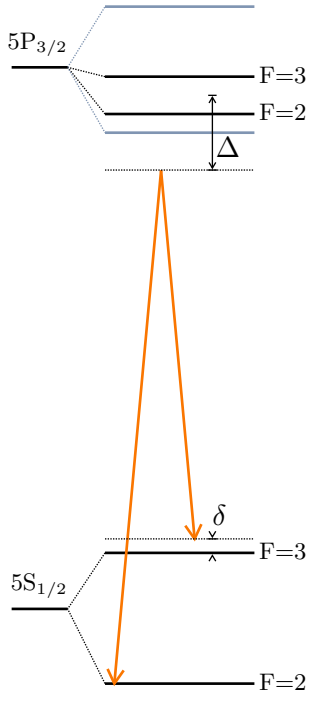

Figure 4.1: Energy levels of ^{85}Rb, coupled by two-photon Raman transitions.

In this particular thesis we employ ultra-cold ^{85}Rb atoms for the interferometry experiments and drive two-photon transitions between the $|5S_{1/2}, F = 2\rangle$ and $|5S_{1/2}, F = 3\rangle$ ground states, mediated by the $|5P_{3/2}, F' = 2, 3\rangle$ excited states as illustrated in the level structure shown in Figure 4.1.

A typical atom interferometry experiment starts with an ensemble of ultra-cold atoms prepared in a single hyperfine ground state. A $\pi/2$-pulse then puts each atom into an equal superposition of hyperfine states, and thus also momentum states, effectively acting as an atomic beamsplitter. Similarly a π-pulse, in reversing the phase of the superposition, also reverses the associated momenta and is analogous to a mirror.

In this manner, an atomic equivalent of a Mach-Zehnder optical interferometer can be realised by combining an initial $\pi/2$ beamsplitter pulse with a π mirror pulse and a final beamsplitter pulse, as illustrated in Figure 4.2. With an equal time T between each pulse the atomic wavefunction is divided spatially and then recombined to enclose a finite area. For a given atom, a phase accumulates between the wavepackets propagating along each 'arm' of the interferometer. The phase accumulated by each wavepacket can be determined using the Feynman path integral formalism, described in Section 8.3.1, of evaluating the quantum mechanical action integral along their respective paths [121].

The relative phase between the two wavepackets is the quantity of interest, and the final beamsplitter pulse acts to project this phase onto the probability of measuring the atom in one or other of the basis states just as the final beamsplitter in an optical Mach-Zehnder maps the phase acquired between each path onto the probability of a photon getting transmitted or reflected. Thus, by measuring the number of atoms from the ensemble in each of the internal states at the end of the interferometer sequence, the

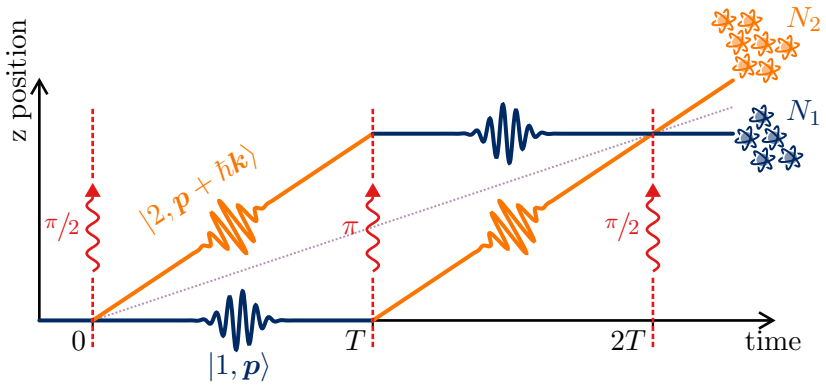

Figure 4.2: Schematic representation of the atomic equivalent of an optical Mach-Zehnder interferometer.

difference in phase acquired by the wavepackets in each arm can be determined.

A series of rules compiled by Friedberg and Hartmann [122] for the phase acquired between the two arms during free-space propagation and interaction with pulsed laser fields in the presence of various external potentials have been refined and specialised through subsequent literature [33, 123, 124], culminating in a concise but thorough treatment by Bongs *et al.* [125].

To first order, a rotation of the apparatus at angular velocity $\boldsymbol{\Omega}$ and an acceleration $\boldsymbol{a}$ would induce a phase difference at the output port given by [119]

$$\Phi = \left[\boldsymbol{k} \cdot \boldsymbol{a} + 2\boldsymbol{v} \cdot (\boldsymbol{k} \times \boldsymbol{\Omega})\right] T^2, \tag{4.1}$$

where $\boldsymbol{k}$ is the effective wave vector given by Equation (2.54) corresponding to the difference between the wavevectors of the two Raman beams, $\boldsymbol{v}$ is the mean velocity of an atom throughout the sequence and T is the duration between pulses.

Clearly, if $\boldsymbol{k}$ is aligned vertically and the rate of rotation $\boldsymbol{\Omega}$ can be zeroed then the first term can be used to measure the vertical acceleration due to gravity. Applications for measuring rotation are also immediately apparent [37, 126]. In state-of-the-art interferometers, however, T is generally large and more complicated pulse sequences are employed to impart more momentum and increase the wavepacket separation. In these cases the higher order phase terms can become non-negligible, and indeed of metrological interest, and so need to be calculated.

4.2 A reductionist picture

Consider a two-level atom in its centre-of-mass frame, prepared in an equal superposition of its two internal states. Let us introduce a resonant laser beam that we take to be a plane wave propagating with a wavevector $\boldsymbol{k}$. With our minds still in the rest frame of the atom and thinking about the electromagnetic field that the atom 'sees', let us nonetheless — as a matter of convenience — introduce the position $\boldsymbol{r}$ of the atom in an inertial frame in which, for now, we consider the lasers to be static. If we define the coordinate system such that the atomic superposition and laser oscillations are in phase at the origin then the phase between the superposition and the laser field that the atoms sees will be

$$\phi = \boldsymbol{k} \cdot \boldsymbol{r}. \tag{4.2}$$

If the atom is moving in this inertial frame with a velocity $\boldsymbol{v}$ then, assuming time $t = 0$ corresponds to the atom being at the origin, its position at time t is just $\boldsymbol{r}(t) = \boldsymbol{v}t$ and the phase between the atom and laser becomes time dependent,

$$\phi = \boldsymbol{k} \cdot \boldsymbol{v}t, \tag{4.3}$$

which is just the Doppler shift that, in Section 2.6, we incorporated into the laser detuning.

If we suppose, in addition, that the laser, and thus the electric field, is subjected to a constant acceleration $\boldsymbol{a}$ and angular velocity $\boldsymbol{\Omega}$ about the origin then it is harder to determine what laser phase the atom sees at any particular instant. The lasers are now in a non-inertial frame, but a coordinate $\boldsymbol{x}$ in the inertial frame can be recast in the frame of the laser as

$$\begin{aligned} \boldsymbol{x}'(t) &= \boldsymbol{x} - \int_0^t \boldsymbol{\Omega} \times \boldsymbol{x} \, \mathrm{d}t' - \frac{1}{2}\boldsymbol{a}t^2, \\ &\approx \boldsymbol{x} - \boldsymbol{\Omega} \times \boldsymbol{x}t - \frac{1}{2}\boldsymbol{a}t^2 \end{aligned} \tag{4.4}$$

for small rotation angles Ωt. A graphical representation of this transformation is given in Figure 4.3. Plugging in the atomic position $\boldsymbol{r} = \boldsymbol{v}t$ in the inertial frame we get[1]

[1] Note that the t here is distinct from the t' integrated over in Equation (4.4) and can be taken out of the integral. The atom will reach the same location in the inertial frame at time t regardless of the motion of the laser.

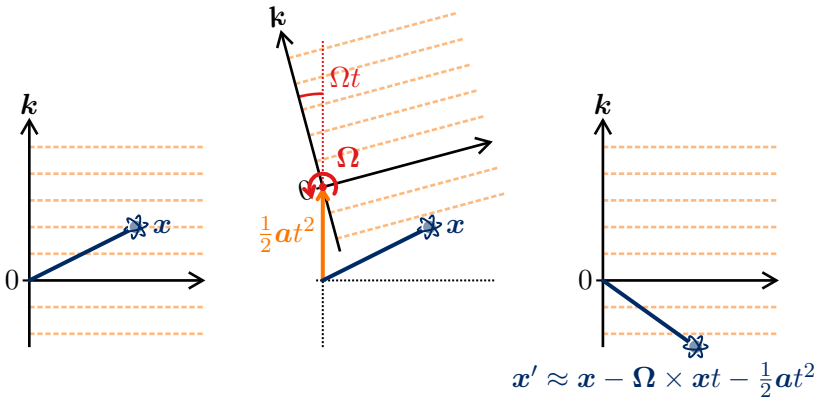

Figure 4.3: Acceleration of the lasers relative to an atom in an inertial frame can be recast as an acceleration of the atom in the non-inertial frame in which the lasers, and hence wavefronts, are static.

$$\boldsymbol{r}'(t) \approx \boldsymbol{v}t - \boldsymbol{\Omega} \times \boldsymbol{v}t^2 - \frac{1}{2}\boldsymbol{a}t^2, \tag{4.5}$$

making the time-dependent laser phase observed by the atom

$$\begin{aligned} \phi(t) &= \boldsymbol{k} \cdot r'(t) \\ &\approx \boldsymbol{k} \cdot \left(\boldsymbol{v}t - \boldsymbol{\Omega} \times \boldsymbol{v}t^2 - \frac{1}{2}\boldsymbol{a}t^2 \right). \end{aligned} \tag{4.6}$$

Now, consider how this phase enters into the output of the $\pi/2 - \pi - \pi/2$ interferometer shown in Figure 4.2 as it plays out on the Bloch sphere, assuming that the atom behaves as a point-like particle and follows a trajectory given by its average velocity $\boldsymbol{v}$ that corresponds to the central dotted line in Figure 4.2. Recall that the laser phase observed by the atom determines the angle of the field vector $\tilde{\boldsymbol{\Omega}}$ — about which the atomic state vector $\tilde{\boldsymbol{\Psi}}$ rotates — in the x–y plane of the Bloch sphere, and that the $\pi/2-$ and $\pi-$pulses discussed in Section 2.4 correspond to 90° and 180° rotations respectively.

At $t = 0$, then, the $\pi/2$-pulse rotates the state vector down onto the equatorial plane, preparing a superposition with a fixed phase relationship relative to the laser which we can be take to be $\phi_0 = 0$, corresponding to the field vector lying along the x axis, without loss of generality. This rotation in shown in Figure 4.4(a), showing that after the interaction the state vector is aligned orthogonal to the field vector in the equatorial plane.

When the atom next sees the laser, at time $t = T$, its phase has advanced by

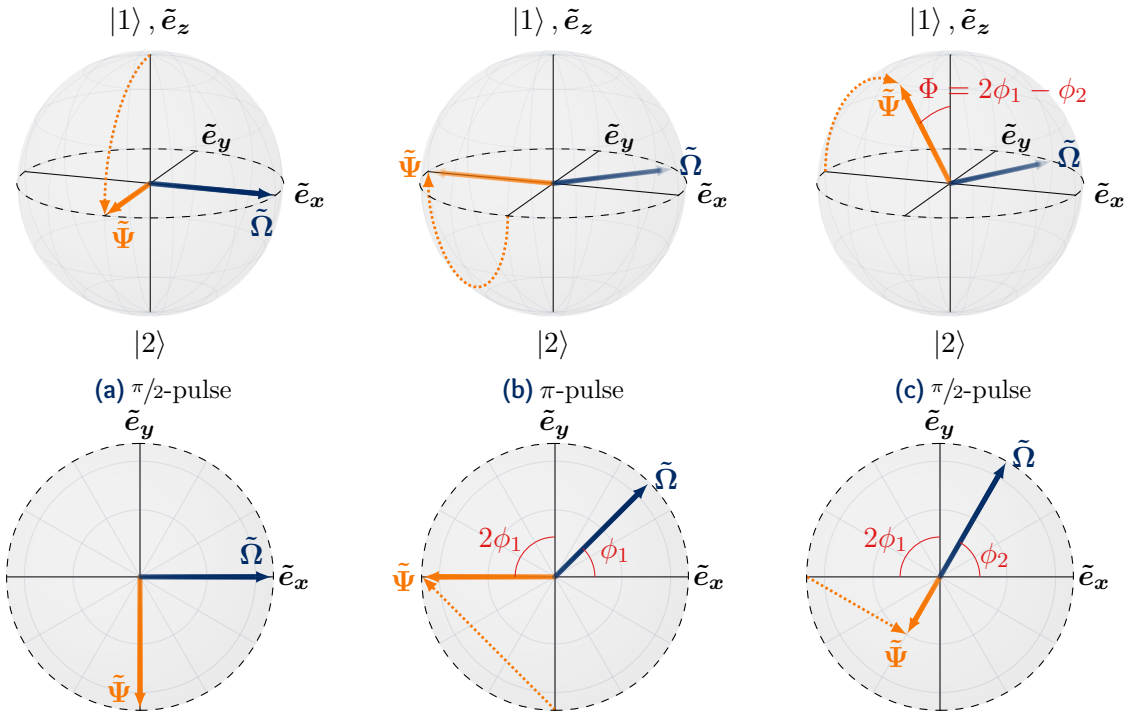

Figure 4.4: $\pi/2 - \pi - \pi/2$ interferometer phase evolution on the Bloch sphere. The laser phase seen by the atom corresponds to the angle of the field vector $\boldsymbol{\Omega}$ in the equatorial plane. Dotted orange lines represent the rotation of the atomic state about the axis of the field vector effected by interaction with the laser during pulses. The top frames show a 3-D projection, and the bottom frames show a top-down view of the equatorial plane.

$$\phi_1 = \boldsymbol{k}_{\text{eff}} \cdot \left(\boldsymbol{v}T - \boldsymbol{\Omega} \times \boldsymbol{v}T^2 - \frac{1}{2}\boldsymbol{a}T^2 \right), \tag{4.7}$$

relative to the atomic superposition and the field vector is rotated accordingly. As shown in Figure 4.4(b), the rotation by π about the field vector effected by the π-pulse maps this phase onto the angle $2\phi_1$ that the state vector ends up at relative to the y axis.

At the time $t = 2T$ of the final $\pi/2$-pulse, the atom sees the laser with the new phase

$$\begin{aligned} \phi_2 &= \boldsymbol{k}_{\text{eff}} \cdot \left(\boldsymbol{v}(2T) - \boldsymbol{\Omega} \times \boldsymbol{v}(2T)^2 - \frac{1}{2}\boldsymbol{a}(2T)^2 \right) \\ &= \boldsymbol{k}_{\text{eff}} \cdot \left(2\boldsymbol{v}T - 4\boldsymbol{\Omega} \times \boldsymbol{v}T^2 - 2\boldsymbol{a}T^2 \right) \end{aligned} \tag{4.8}$$

and, in general, the $\pi/2$ rotation of the state vector about the field vector at this new angle lifts it off the equator. As shown in Figure 4.4(c), the angle Φ that it makes relative to the 'north' pole that corresponds to the initial ground state $|1\rangle$ is then

$$\begin{aligned}\Phi &= 2\phi_1 - \phi_2, \\ &= 2\boldsymbol{k}_{\text{eff}} \cdot \left(\boldsymbol{v}T - \boldsymbol{\Omega} \times \boldsymbol{v}T^2 - \frac{1}{2}\boldsymbol{a}T^2\right) \\ &\qquad - \boldsymbol{k}_{\text{eff}} \cdot \left(2\boldsymbol{v}T - 4\boldsymbol{\Omega} \times \boldsymbol{v}T^2 - 2\boldsymbol{a}T^2\right) \\ &= \boldsymbol{k}_{\text{eff}} \cdot (2\boldsymbol{\Omega} \times \boldsymbol{v} + \boldsymbol{a})\, T^2.\end{aligned} \tag{4.9}$$

This is the interferometer phase, equivalent to Equation (4.1) arrived at by Bordè, Riedl, Barrett *et al.* [41, 119, 127], mapped onto the z axis that reflects the difference in probability of finding an atom in the ground $|1\rangle$ or excited $|2\rangle$ state. Again, it can thus be determined by measuring the number of atoms $N_{1,2} = \left(1 \pm \cos(\Phi)\right)/2$ in each state after an ensemble is subjected to the interferometer sequence.

This is the central path approximation [128]. In this picture each atom can be thought of as a stable clock, ticking at the same frequency as the laser oscillations and moving relative to them. At each pulse the phase of the atomic clock is compared to the local phase of the laser. The phase accrued between the two increases by 2π each time the atom traverses one optical wavelength.[2]

[2] Or, equivalently, if the optical field moves one wavelength relative to the atom.

Note that this particular interferometer sequence corresponds to a differential measurement. The first $\pi/2$-pulse synchronises the atomic clock to the laser phase. At the π-pulse, a record of the phase difference that has accumulated over the first *dwell period* of duration T is made. A further phase accrues during the second dwell period, and the final pulse effectively measures the difference between these two phases.

This is made particularly explicit in Equation (4.9); if the phase accrued during the first dwell period is equal to that accrued in the second due to, for example, the atom propagating at a constant speed across the wavefronts, then $\phi_2 = 2\phi_1$ and $\Phi = 0$. In effect this is a differential measurement of the atomic velocity along the laser beam and so is only non-zero in the presence of an acceleration.

4.2.1 A few remarks

As derived in Section 2.6, when two lasers are used to perform the interferometry with Raman transitions between stable hyperfine levels then it is the relative phase between the two lasers that

plays the role of the single-photon laser phase and determines the angle of the field vector on the Bloch sphere.

This makes sense given that this corresponds to the phase of the beatnote between the two lasers and this is the quantity that oscillates at the frequency of the atomic transition. How this quantity varies spatially is encoded in the effective wavevector $\boldsymbol{k}$ and depends on how the lasers are aligned relative to each other. The more the phase of the beatnote changes for a small spatial translation then the more sensitive the interferometer is.

Consider the scenarios where two Raman beams propagate parallel or anti-parallel to each other. In the co-propagating arrangement, the sign of the phase shift for the individual lasers upon translation along the propagation axis is the same and so the relative phase between the lasers varies very little and is effectively static over regions of space on the scale of millimetres. [3]

[3] For two lasers separated by the microwave frequency $\sim 2\pi \times 3\,\mathrm{GHz}$ separating the hyperfine ground levels in ^{85}Rb, a $2\ \pi$ phase shift occurs over the microwave wavelength, ~ 10cm.

By contrast, in the counter-propagating arrangement translation along the beam axis results in phase shifts for each laser that have opposing signs such that the relative phase between them varies rapidly. For 780 nm lasers a 2π phase shift occurs over half the optical wavelength $\sim 390\,\mathrm{nm}$. In the limiting case where both lasers have exactly the same frequency, if they are co-propagating then they have a fixed phase relationship all along the propagation axis, while counter-propagating beams form a standing wave with nodes (anti-nodes) where the two lasers are exactly out of phase (in phase) every half wavelength.

It is clear then that, with co-propagating Raman beams, an atom's motion relative to the wavefronts is not going to induce a noticeable shift in phase.[4] On the other hand, for counter-propagating Raman beams where the spatial period is just a few hundred nanometres, atoms rapidly accrue a phase relative to the laser if they move along the beam path or, indeed, if the lasers move in relation to them. This reinforces our assertion that counter-propagating beams result in a sensitive interferometer, previously justified by relating the large momentum imparted by the corresponding two-photon transition to a large separation between the two 'arms' of the interferometer in the more conventional picture presented in the previous section.

[4] The r.m.s. speed of an atom in a cloud of temperature $\mathcal{T}$ is $\sqrt{k_B\mathcal{T}/M} \approx 10\sqrt{\mathcal{T}}$ for ^{85}Rb. For a typical cloud temperature of $\mathcal{T} \sim 50\,\mu\mathrm{K}$, it would take about a second for an atom to traverse the microwave wavelength and pick up a 2π phase shift, a long time compared to a typical interferometer period on the order of ms.

It is important to be aware of the approximations and omissions made in this picture. Apart from assuming small rotation angles ΩT — an assumption also made in other models — by taking the

field vector to lie in the equatorial plane and have a fixed angle throughout each pulse, we effectively assume that the interactions are short enough that an atom does not propagate a considerable fraction of a wavelength along the beam axis during them. Equivalently, we assume that the Rabi frequency Ω is much greater than the Doppler shift $\boldsymbol{k} \cdot \boldsymbol{v}$.

As intimated in the previous section, there are cases where the separation of the atomic wavepacket is metrologically significant and cannot be neglected. For example, if the atomic wavepackets are macroscopically separated vertically for a long enough duration then the interferometer becomes sensitive to the differential force experienced by each wavepacket due to gravity gradients [118]. Another beautiful experiment [129] uses two $\pi/2$-pulses to put atoms into a superposition in which there are spatially separated wavepackets corresponding to the same internal state. These wavepackets are then held at different positions in a vertically orientated optical lattice for upwards of 20 s. The motion of the atoms along the beam axis is effectively frozen and the dominant phase that accrues between the wavepackets is due to the difference in the gravitational potential energy between the vertically separated lattice sites.

However, these effects only become observable when the spatial separation between the wavepackets is large, the atoms remain in a superposition for a long time, or both. This is not true of the interferometers considered in this thesis and so these effects are negligible and it is sufficient to think of the atoms as point-like clocks, leading us nicely into our following discussion of the Longitude problem.

Some musings on boats and clocks

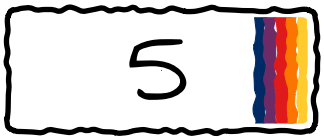

So we beat on, boats against the current, borne back ceaselessly into the past.

— F. Scott Fitzgerald, *The Great Gatsby*

And now for something completely different: some heavy analogy, hand-drawn figures and tales of nautical adventure. This chapter — in conjunction with Chapter 8 — is loosely based on a pedagogical article, currently in preparation, designed to introduce the classical navigators who will one day be the users of atom interferometric inertial sensors to some of the ideas behind quantum technologies. While it is sincerely hoped that something can be found here to please and inform all but the most salty of quantum seadogs, it should also be said that there is nothing here that cannot be skipped over without affecting one's enjoyment or understanding of the rest of this thesis.

5.1 The longitude problem

Before the global positioning system (GPS) now incorporated in digital chart plotters made nautical navigation a triviality, exploring the world by sea was a dangerous pursuit that, up until the nineteenth century, had many of the greatest minds in the world turned to the problem of enabling sailors to establish where, in fact, they were. Since antiquity we have employed latitudinal and longitudinal coordinates, measured north-south and east-west respectively, for navigation. In modern times latitude is understood to be a natural concept, with the Earth's rotation axis defining the north and south poles as fixed points that any navigator can agree on.[1] Lines of constant latitude, generally marked on nautical charts every degree north and south of the equator, remain happily fixed in place as the Earth rotates below them.

Longitude is harder to define, and indeed measure. Throughout history nations have defined arbitrary meridian lines that run north-south from which navigators can, in principle, measure their

[1] I should imagine that navigators adhering to the view that the Earth is flat have more serious problems to contend with, but the centre and circumference of a circle are similarly well defined and they appear to have devised a system that allows them to avoid falling off the edge.

distance in order to establish some sort of longitudinal reference. Again since antiquity, it has been understood that the further one is from a defined meridian, the greater will be the difference in local, solar, time. When it is noon at a given meridian, it will be midnight 180° east or west from it.

This is moot unless you have a reference of the time at the meridian to which to compare your local solar time. This necessitates a timepiece that can be synchronised at the meridian and hold good time over the timescales associated with global navigation while subjected to the particular mechanical stresses also associated with it. These are particularly acute at sea, as anybody who has ever felt at all seasick with attest. With such a timepiece, however, a navigator can determine their longitude by comparing its value to their locally measured time, determined by the Sun's elevation from the horizon, and noting the difference. This is most naturally done at local noon, where the Sun reaches its highest point in the sky making it easy to determine with accuracy.

A timepiece that kept sufficiently good time at sea was eventually developed by John Harrison, earning him the Longitude Prize in 1773 and revolutionising nautical navigation [130]. Where sailors used to cross the Atlantic ocean by rigidly maintaining a constant latitude until they sighted land, they could now embark on unconstrained travel and accurately track their position. In 1884 the meridian passing through the Royal Observatory in Greenwich, England, was internationally agreed upon as the 'Prime Meridian' for global navigation, establishing a globally recognised 'universal time' and longitude standard.[2]

[2]The French, of course, abstained and clung to their own Paris meridian for a further 30 years as their navigational standard. Even after adopting the time at Greenwich as the universal time reference for global navigation in 1914, they would continue to refer to it as 'Paris mean time, retarded by 9 minutes and 21 seconds' until the late 20th century.

5.2 As an analogue to atom interferometry

A marine chronometer inspired by one of Harrison's early clocks is a stable oscillator designed to rotate in time with the Earth's rotation about its axis. Synchronised to universal time, it will be 'in phase' with the Earth's rotation, apparent from the local solar time, when it is at the same longitude as the Royal Observatory in Greenwich, England. As the ship in possession of the chronometer sails east or west around the globe then a phase shift will accrue between the timepiece and the local solar time that can be used to unambiguously determine how far the ship has travelled longitudinally, modulo one complete circumnavigation.

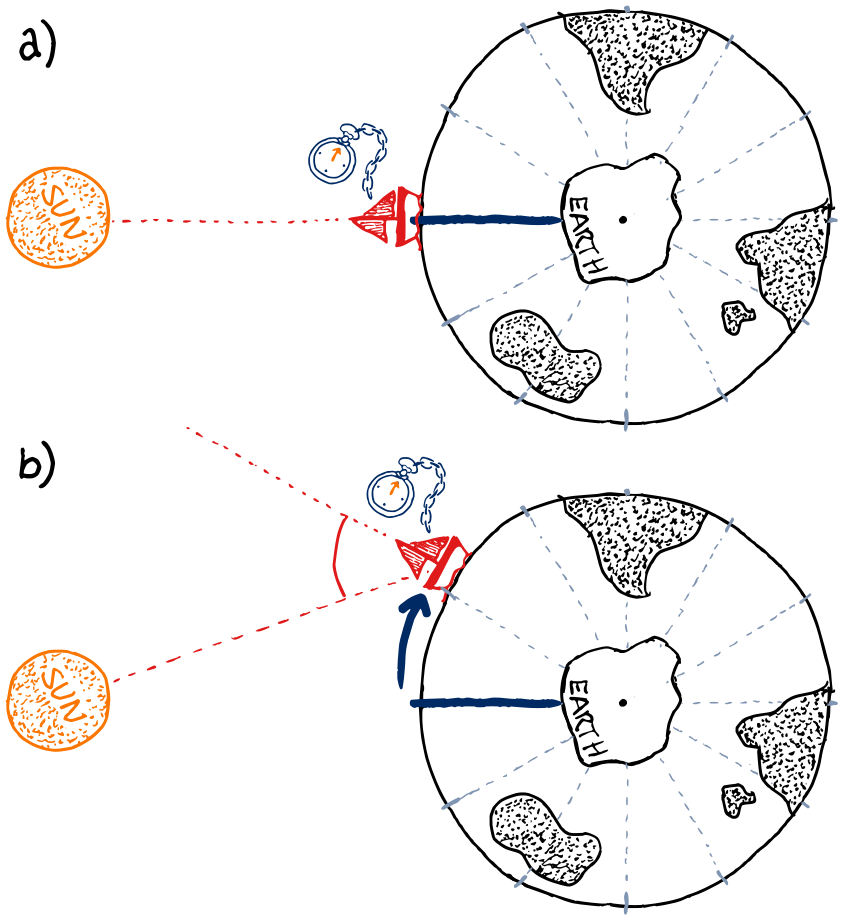

Figure 5.1: The principle of determining longitude with a marine chronometer. In (a), a navigator synchronises his ship's clock with the solar time at the prime meridian, in this case midday where the sun's elevation from the horizon is maximal. The further east or west the ship then sails from the meridian, the greater the difference between the local solar time and the meridian time, as measured on the clock. In (b) it is again midday at the meridian, but the boat has travelled far enough to note approximately a two-hour shift in solar time.

This is much the same principle as atom interferometry, where an oscillating atomic superposition plays the role of the chronometer. Instead of the Earth's rotation, its frequency ticks in time with oscillations of the electromagnetic field in a resonant laser beam. A pulse of the laser synchronises the phase of the atomic superposition to that of the laser at the atom's initial location, much as a ship's navigator might synchronise his clock at the Greenwich meridian. The phase accrued between the atom and the laser then provides a measure of how far the atom has travelled along the laser beam, modulo one laser wavelength.

Such interferometers are also sensitive to fields that might alter the frequency of the moving oscillator relative to the reference one. Suppose, in circumnavigating the globe, our ship sailed over a particularly large mass that altered the acceleration due

to Earth's gravity g enough to noticeably increase the periodicity of the chronometer's pendulum for an appreciable duration.[3]

[3]Clearly pendulum clocks do not hold good time at sea, hence the need for Harrison's invention of the marine chronometer that employed a temperature-compensated spiral spring driving a fast-beating balance wheel as its stable oscillator, but humour me.

Upon arriving back at Greenwich — or perhaps the Marina Greenwich near Mascarat, Spain, if they favoured a pleasant climate and safe port over hours of tacking up the Thames river — the sailors would find their ship's clock to be slightly ahead of local solar time. The change in gravitational acceleration would register as a slight shift in phase of the ship's chronometer from local time after a complete circumnavigation. Of course, anybody who has tried to build a stable mechanical oscillator or tell the time from the position of celestial bodies will appreciate the impossibility of such a measurement.

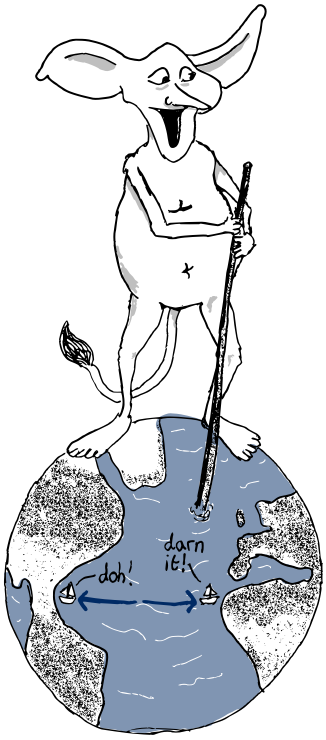

Figure 5.2: Two boats fitted with marine chronometers could, in principle, detect the rate at which a nefarious demon was stirring the Earth's seas.

Imagine, now, that it was seriously suspected that some nefarious demon was rotating the mass of water that forms the Earth's seas at some undetermined rate relative to the land masses. One could dispatch two ships from, say, Marina Greenwich with orders to maintain constant speed through the water, one travelling due east and one due west. If the water was rotating at the same rate as the land then the boats should meet at a longitude of exactly 180° many weeks later.

If, instead, the malevolent demon depicted in Figure 5.2 (or indeed a celestial body) was causing the seas to rotate at a different rate than the Earth's land masses then the boats would meet at some different longitude. In knowing both how long they had been travelling and their longitude from their onboard chronometers, they could determine the rate at which the seas were being stirred.[4]

Anybody who has tried to maintain a constant speed through the water on a boat for any length of time, to sail along a constant latitude without hitting land at some point, or indeed to consider the effect of tides on such a measurement would appreciate that this is perhaps the worst possible way to establish the demon's disruptive intentions although, in principle, the large enclosed area of such an 'interferometer' would make it attractively sensitive to even a small perturbation by said demon if these issues could be addressed.

It is worth noting, however, that the chronometer method does not offer a sensitive measurement of longitude on more local length scales, as the absolute distance one needs to travel in order to see a measurable shift in phase between the ship's chronometer and local solar time is large. Similarly, the frequency of the Earth's

rotation is slow, so that it would take a long time to detect any difference in frequency between the Earth's rotation rate and that of the clock hand.

By contrast, the distance an atom must travel relative to the laser in an atom interferometer in order for a full 2π phase shift to accrue will typically be measured in just hundreds of nanometres. Similarly, the frequency of the oscillations will be on the order of gigahertz[5] so that any shift of the atomic frequency relative to the laser quickly becomes measurable. Such measurements of rotation rates with atomic clocks, then, have the potential to be very sensitive [133]. Introducing an atom interferometer in these terms — rather than invoking the momentum superpositions of Section 4.1 — strongly suggests that they do not lend their supreme sensitivity to anything intrinsically 'quantum', but rather that it is all a matter of scale.

[4] There is a comparison to be drawn here to an experiment [131, 132] in which atomic clocks were flown on eastwards and westwards around the world on jet planes, and their times were compared to a local reference at the end to test the principle of special relativity that "a moving clock runs slow". In an inertial frame, the westbound clock — travelling against the Earth's rotation — moved slower, and hence gained time relative to the eastbound one. This can be considered a coarse measurement of the Earth's rotation rate. As the concept of "local time" is valid in a non-inertial reference frame in which the Earth's rotation is cancelled, our demon provides a rotation our ships could actually measure.

[5] Or indeed terahertz in the case of a single photon resonance.

Part II

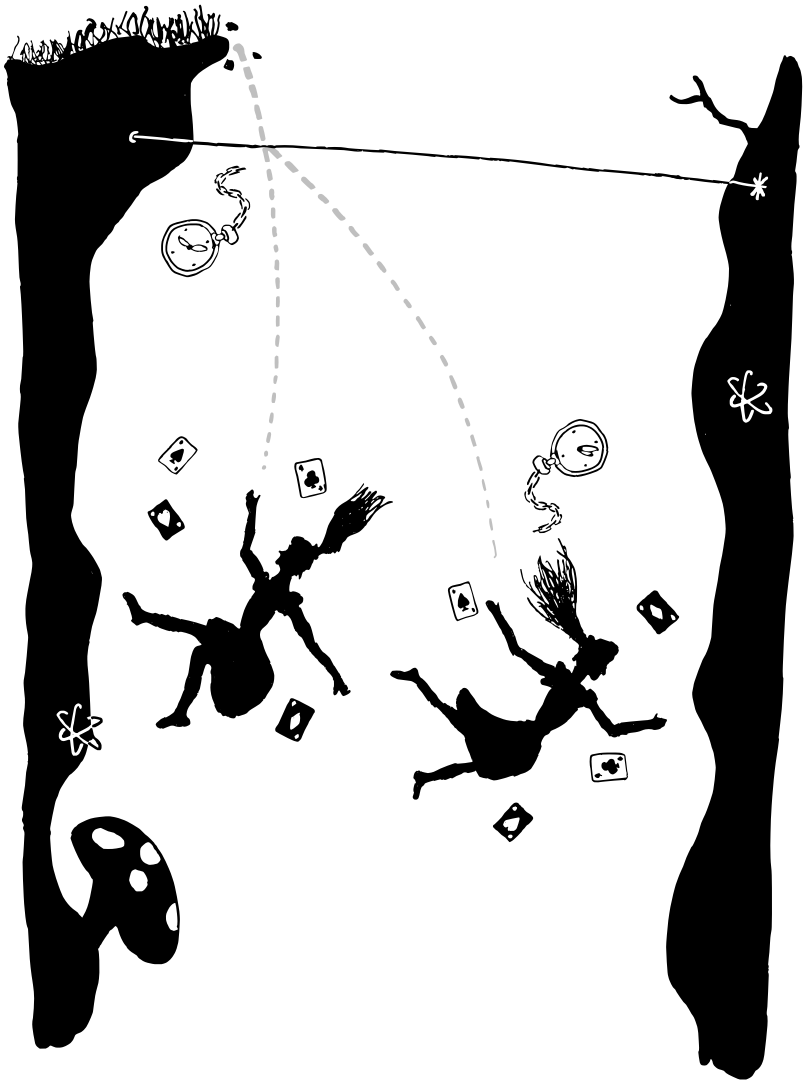

A basement in Southampton

Experimental setup

6

The major difference between a thing that might go wrong and a thing that cannot possibly go wrong is that when a thing that cannot possibly go wrong goes wrong it usually turns out to be impossible to get at or repair.

— Douglas Adams, *Mostly Harmless*

The setup on which most of the experiments in the thesis were conducted consists of a MOT of ^{85}Rb atoms positioned at the centre of counter-propagating horizontal Raman beams in the 'basement'[1] of the Southampton University physics department. The optical and magnetic MOT fields are extinguished prior to pulsing the Raman beams in a variety of experiments detailed in the following chapters, and the output of each experimental run is a measure of the fraction of atoms in the upper ground state $|5S_{1/2}, F = 3\rangle$ made by pulsing the MOT beams and capturing the fluorescence light on a photo-multiplier tube (PMT).

[1]The main entrance to the Southampton Physics building is on the first floor at the top of a small hill. This probably explains why the ground floor at the bottom of the hill is colloquially referred to as 'the basement', despite possessing exterior walls and a ground-level goods entrance.

The horizontal Raman beam configuration eliminates any shift to the Raman resonance from gravitational acceleration, although it offers a limited time window in which to perform experiments before the atoms have fallen out of the beams. The unique aspect of the experiment is that the Stokes and anti-Stokes frequency components for the Raman interactions are in spatially discrete, phase-coherent beams. One component is generated by a 310 MHz acousto-optic modulator (AOM) and taken to one side of the chamber through a polarisation maintaining (PM) fibre, while the other is generated by a 2.7 GHz electro-optic modulator (EOM) and taken to the other side of the chamber through another fibre. This is in contrast to many other experiments in which both frequency components are present in each beam and discriminated between by, for example, a gravitational Doppler shift [23].

In this chapter we describe this setup in more detail, although the definitive version remains in the thesis of Alex Dunning [89].[2] In Section 6.1 we briefly describe the specifics of the MOT, in Section 6.2 we describe the generation and control of the Raman interaction beams, and in Section 6.3 we explain the state preparation and read-out process common to all of our experiments. In Sections 6.4

[2]I inherited the coherent control experiment in a disassembled state, and it is a largely due to the clarity of Alex's thesis that I was able to reconstruct it. I emphasise here only the elements of the experiment that have been changed or characterised further.

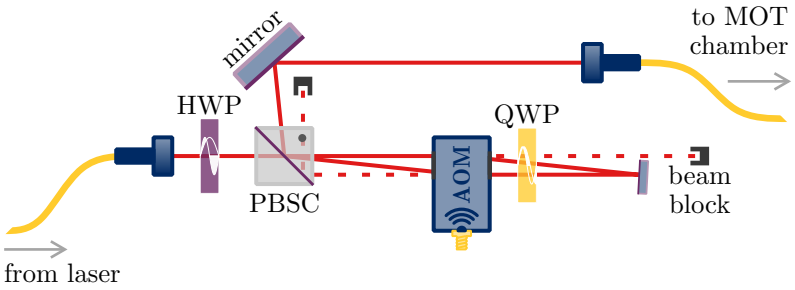

Figure 6.1: Double-pass AOM shutter setup used to control the amplitude of the cooler and repump beams. After the first pass, the undiffracted beam is dumped and the diffracted order is reflected back through the AOM where it is diffracted with the opposite frequency shift, having passed twice though a quarter waveplate (QWP) to ensure that it has has a linear polarisation orthogonal to the input beam such that it can be separated by a polarising beamsplitter cube (PBSC). After the two diffractions, approximately 50 % of the input light is coupled into an output fibre for delivery to the MOT chamber.

and 6.5 we describe the use of the Raman beams to characterise the magnetic field experienced by the atoms and how the MOT temperature varies as a function of different cooling parameters.

6.1 Magneto-optical trap

A laser detuned slightly to the red of the $|5S_{1/2}, F = 3\rangle \rightarrow |5P_{3/2}, F' = 4\rangle$ transition provides the cooling light for the MOT. In principle this is a closed, cycling transition suitable for laser cooling; in practice, off-resonant excitation to the $|5P_{3/2}, F' = 3\rangle$ state is unavoidable, an open transition whereupon spontaneous decay to the lower ground state $|5S_{1/2}, F = 2\rangle$ makes atoms 'dark' to the cooling light and free to leave the trap.

This 'optical pumping' to the dark state occurs with a $1/e$ time constant on the order of 100 µs, necessitating an additional repump beam resonant with the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition to repopulate the 'light' upper ground state $|5S_{1/2}, F = 3\rangle$, again via spontaneous decay, and return atoms to the cooling cycle.[3]

These cooling and repumping beams are produced by two *Vescent D2-100-DBR* distributed Bragg reflector (DBR) lasers. A peak lock servo keeps the repump laser locked to the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition with an error signal derived from frequency modulation (FM) spectroscopy of a saturated absorption setup. The cooling laser is maintained at a fixed

[3] Of course, in the absence of the cooling light, this laser will causes atoms to collect in the upper ground state $|5S_{1/2}, F = 3\rangle$ and, as it resonantly drives an open transition, this process is much more efficient and happens orders of magnitude faster than the off-resonant pumping driven by the cooling light. This ability to use the MOT lasers to pump atoms between the two ground states, and the fluorescent light from the spontaneous decays that facilitate it, form the basis of the state read-out mechanism that is described in detail later.

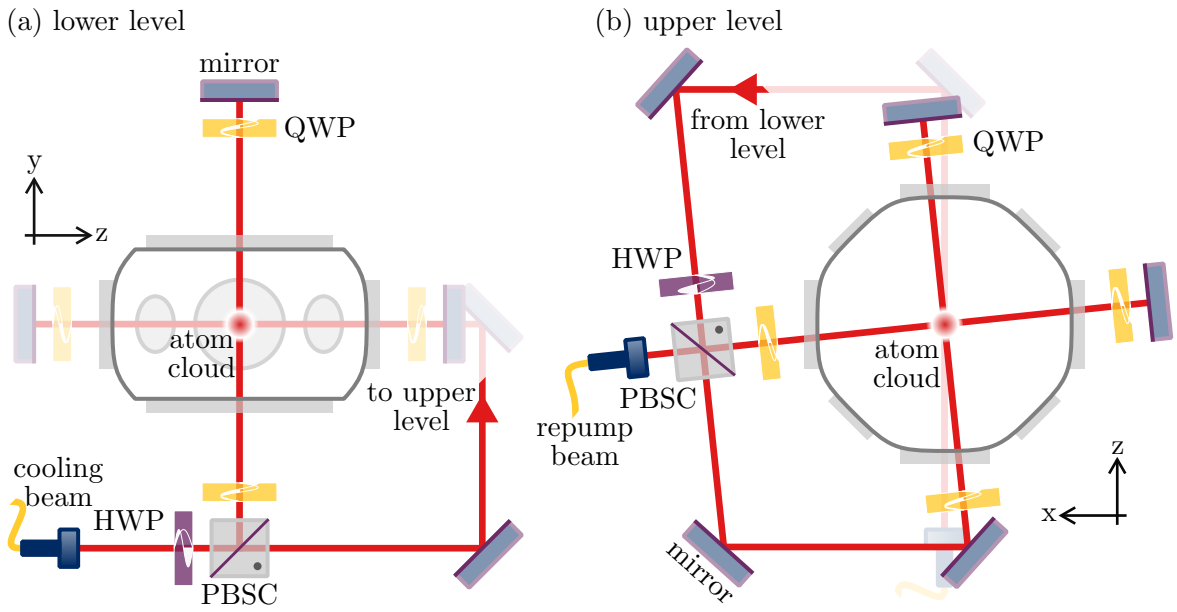

Figure 6.2: Schematic of the free-space optics that divide and retroreflect the cooling and repump beams along the MOT axes. On the lower level (a), one third of the cooling beam is reflected vertically up through the chamber by a polarising beamsplitter cube (PBSC), with a quarter waveplate (QWP) circularising the polarisation. Above the chamber the beam is retroreflected by a mirror, and a double pass through a second QWP ensures that the light returns with polarisation of the same handedness. The other two thirds of the beam is directed to an elevated platform (b), where it is evenly split by a second PBSC between orthogonal horizontal arms that are similarly retroreflected through the chamber to form a MOT where the three arms intersect. The repump beam is injected in the adjacent port of the PBSC on this level and is evenly divided between the two horizontal arms. Non polarisation maintaining, single mode, fibres are used to deliver the cooling and repump light, but the polarisation is stable enough that the half waveplates (HWPs) that precede the PBSCs to control the balance of the cooling beam between the three arms rarely require adjustment throughout the day.

frequency from the repump by an offset phase lock servo that compares a beat note obtained by combining the two lasers on a fast photodiode to an adjustable internal reference.[4]

The beams are fibre coupled, and each one is taken into free space to pass twice through its own AOM in order to produce a diffracted order with no net frequency shift that is coupled back into fibre for delivery to the MOT chamber as shown in Figure 6.1. These AOMs are used to control the amplitude of the cooler and repump quickly and independently, with the double-passed configuration allowing the light to be extinguished with up to 75 dB suppression, while exhibiting just 3 dB fibre-to-fibre losses from stray diffraction and imperfect coupling when the shutter is fully 'open'. Each beam is delivered to the MOT in a single fibre, and out-coupled as collimated beams with $1/e^2$ diameters of 7.5 mm that are divided along three Cartesian axes to form the optical field for the MOT by an arrangement of free space optics, shown in Figure 6.2.

[4] The optical and electronic components for this setup are all provided by Vescent Photonics in a package that is sufficiently 'plug-and-play' that I feel justified in omitting from this work the detailed description of stabilising a laser to an atomic frequency reference that may once have been a lengthy section in the thesis of any student of the field, primarily because of how much of each day they would spend cursing it.

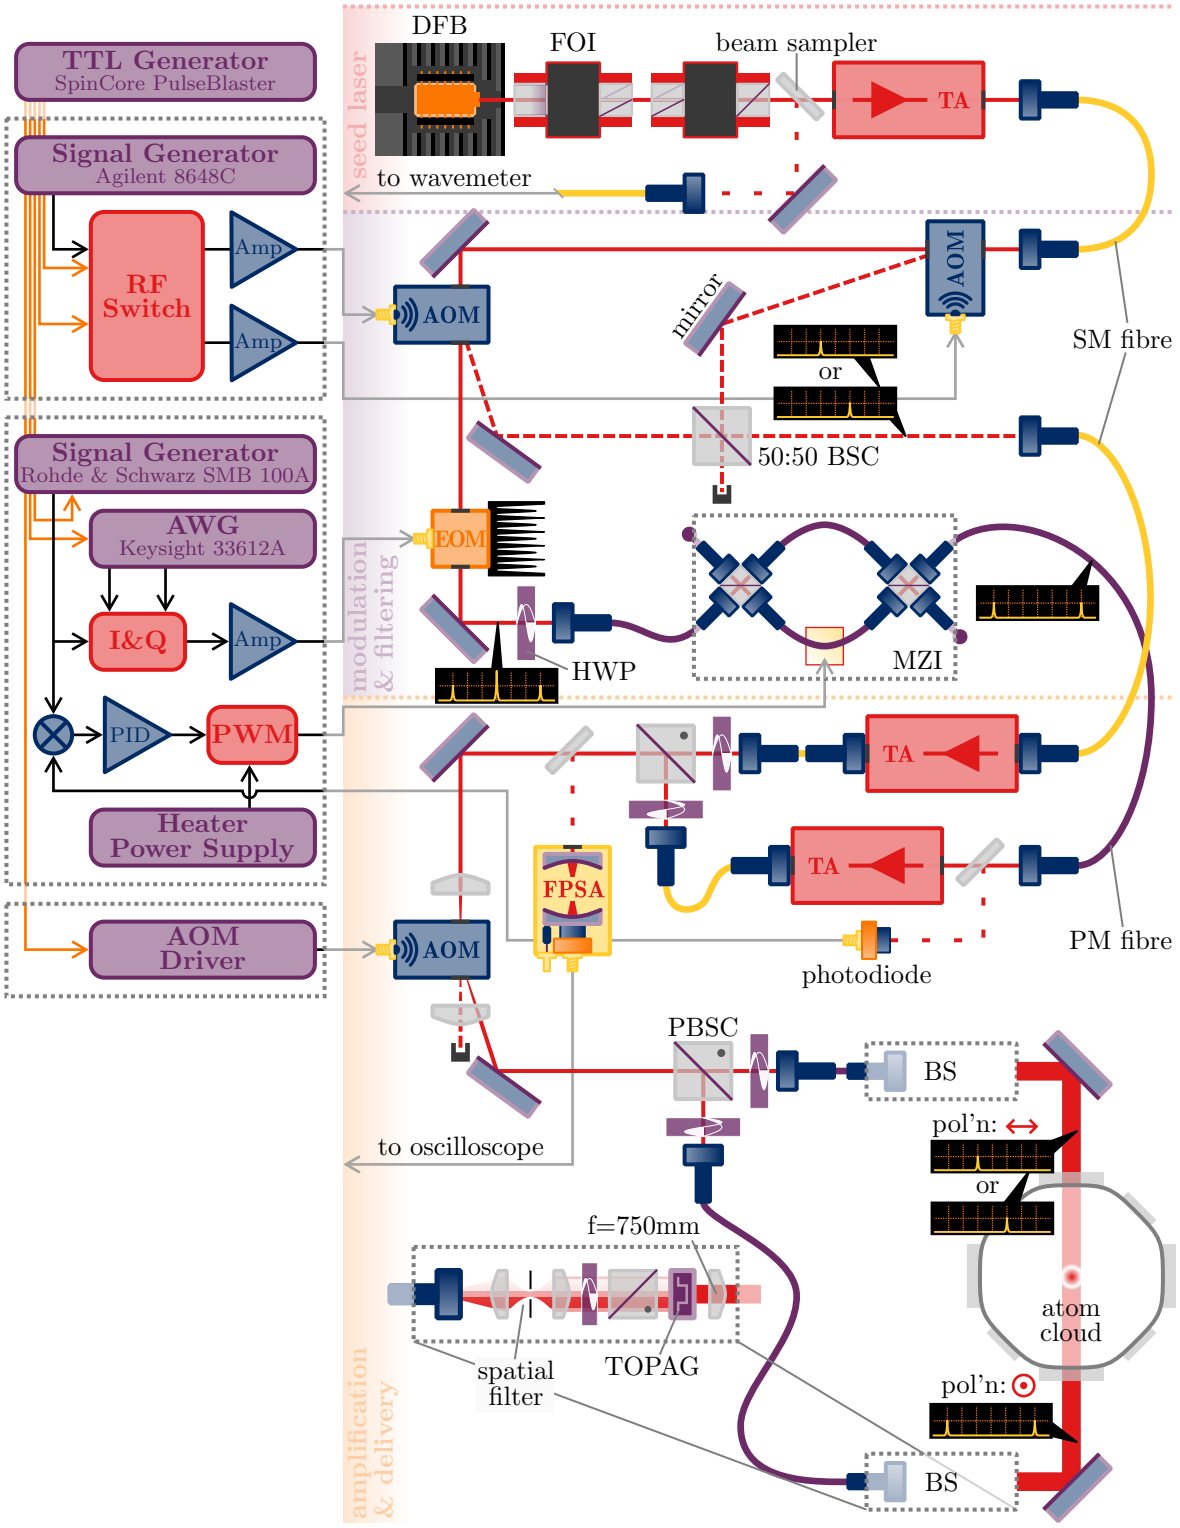

Figure 6.3: Schematic of optics and electronics used to generate Raman interaction pulses. Symbols are labelled as follows, with only one instance labelled for some repeated symbols: Distributed feedback laser diode (DFB); Faraday optical isolator (FOI); tapered amplifier (TA); acousto-optic modulator (AOM); electro-optic modulator (EOM); half waveplate (HWP); polarising (PBSC) and non-polarising (BSC) beamsplitter cubes; fibre Mach-Zehnder filter (MZI), one arm of which is temperature controlled by a heater whose current is pulse width modulated (PWM) within a feedback loop (PID), details of which are provided in the text, to suppress the carrier from the EOM output; single-mode (SM) and polarisation-maintaining (PM) fibres; Fabry-Perot spectrum analyser (FPSA); beam-shaping optics (BS), including Topag GTH-4-2.2 refractive beam shapers (TOPAG); arbitrary waveform generator (AWG); RF amplifiers (Amp) and in-phase and quadrature phase modulator (I&Q).

6.2 Raman interaction beams

A schematic of the optical apparatus and control electronics for generating the two spatially discrete, phase-coherent beams to drive the two-photon Raman transitions used for our coherent control and interferometry experiments is shown in Figure 6.3.

The beams are generated from a single 780 nm *Eagleyard* distributed feedback (DFB) laser diode. After passing through two Faraday optical isolators (FOIs), providing 60 dB of isolation from back-propagating light,[5] a microscope slide samples a small amount of the light to be fibre coupled and monitored on a wavelength meter, and the rest of the beam is amplified by a Toptica BoosTA tapered amplifier and fibre coupled to clean the beam and take it to the modulation and filtering stage.

[5] This laser, along with many DFB and DBR lasers used by the author, exhibits a particularly high sensitivity to optical feedback. When combined with a tapered amplifier (TA) and fibre-coupled, two optimally aligned FOIs were required in order to prevent frequency jumps and linewidth broadening.

6.2.1 Spatially discrete Raman frequency components

To drive two-photon transitions between the two hyperfine ground states of ^{85}Rb we require two phase-coherent laser beams separated in angular frequency by the hyperfine splitting $\omega_0 = 2\pi \times 3.035\,732\,44\,\text{GHz}$.

One of these beams is produced from the diffracted order of AOMs placed at the beginning of the modulation and filtering stage. The two AOMs are driven at $\omega_\text{A} \approx 2\pi \times 320\,\text{MHz}$, one aligned to diffract into the first positive order and one aligned to diffract into the first negative order. These diffracted orders are combined on a non-polarising beamsplitter cube (BSC) and re-coupled into fibre to go to the amplification and delivery stage. Only one AOM is driven at a time, selected with a radio frequency (RF) switch, determining whether this beam goes on to become the Stokes or anti-Stokes light.

The second beam is modulated by an EOM, driven by a frequency synthesiser locked to a common reference with the AOM driver, at $\omega_\text{E}+\delta$, with $\omega_0 = \omega_\text{A}+\omega_\text{E}$. This gives it two sidebands at $\pm(\omega_\text{E}+\delta)$, and the carrier is then suppressed by a fibre Mach-Zehnder filter that we describe in more detail below.

At the amplification and delivery stage, the fibre-coupled AOM and EOM beams are separately amplified by TA diodes and combined on a polarising beamsplitter cube (PBSC) to pass through a common AOM for fast $\sim 50\,\text{ns}$ shuttering and future amplitude

modulation. The beams are then separated by polarisation and coupled into PM fibres for delivery to the MOT chamber. Before entering the chamber, each beam passes through a beam shaper, consisting of a spatial filter to ensure a high mode quality, a PBSC to maintain polarisation purity[6] and a Topag GTH-4-2.2 refractive element that results in a $\sim 1.2\,\mathrm{mm}$ square beam of uniform intensity — to within $\pm 5\,\%$ — at the focus of a 750 mm focal length lens where the MOT cloud is located. The PBSCs in each beam shaper are aligned orthogonally, ensuring that the beams are delivered with orthogonal linear polarisations, and they enter the chamber in a counter-propagating alignment.

[6] Half waveplates (HWPs) prior to entering the PM fibres and before the PBSCs at the output ensure the light polarisation remains aligned to one of the transverse fibre axes and that the PBSC transmission is maximised.

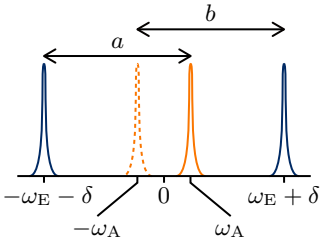

Figure 6.4: Illustration of frequency components present in the AOM modulated (orange) and EOM modulated (blue) beams, relative to the carrier frequency. Two sidebands at $\pm(\omega_E + \delta)$ are present in the EOM beam and, by switching between the positive or negative order AOM, either a $+\omega_A$ or $-\omega_A$ component is present in the AOM beam. This results in two possible Raman pairs, a or b, depending on which AOM is active. Both are separated in angular frequency by $\omega_0 + \delta$ but have opposite effective k-vectors as seen by the atoms.

This results in two beams arriving at the atoms from opposite directions with the frequency components shown in Figure 6.4. The sign of the frequency shift of the AOM beam (shown in orange) is selected by the RF switch. When positive, it forms a Raman pair a with the lower frequency sideband in the EOM beam (shown in blue), and when negative it forms Raman pair b with the higher frequency EOM sideband. Although two sidebands are present in the EOM beam, one is always far enough off-resonance that its only considerable contribution to the interaction is to induce a slight light-shift to the resonance.

The RF switch can alternate between Raman pairs a and b in under $5\,\mu\mathrm{s}$, providing a way to quickly reverse the sign of the effective k-vector, and hence momentum kick, the atoms experience when the beams are pulsed on.

6.2.2 EOM carrier suppression

The carrier is suppressed from the EOM beam by a bespoke fibre Mach-Zehnder interferometer produced by OZ Optics, described in detail in [134] and represented in Figure 6.5. Two PM fibre inputs are combined on a 50:50 BSC, divided along two PM fibre arms with a path length difference $\Delta l = 38.5(5)\,\mathrm{mm}$ and recombined on a second 50:50 BSC into two PM fibre outputs. The path length difference means that an optical field with an angular frequency ω injected at one of the input ports will recombine on the second BSC with one path having accrued an additional phase factor of

$$\begin{aligned} \phi &= \frac{\Delta l \eta}{c}\omega \\ &= \frac{2\pi\omega}{\mathcal{F}}, \end{aligned} \tag{6.1}$$

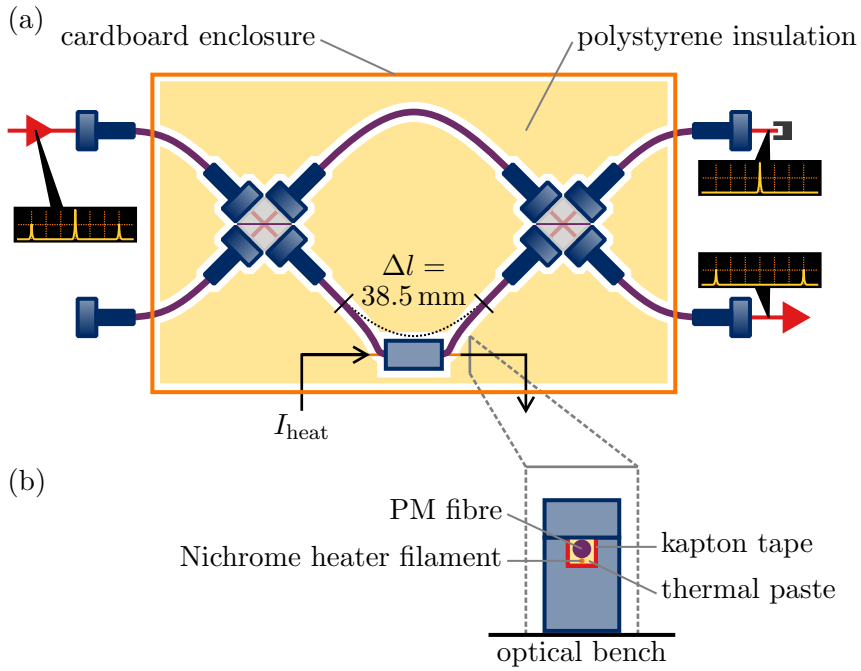

Figure 6.5: Schematic of stabilised fibre Mach-Zehnder interferometer used to suppress the carrier from the EOM output. The input light is divided evenly between two fibres by a free-space 50:50 BSC, that are then recombined on a second free-space BSC whose output ports are again fibre coupled. A path difference of $\Delta l = 38.5(5)\,\mathrm{mm}$ gives the interferometer a FSR of approximately double the angular frequency of the modulation ω_E so that, when the carrier transmission is minimised at one output port, the sideband transmission is close to maximal. The device is encased in a cardboard enclosure with polystyrene insulation to minimise thermal drift, and a $\sim 50\,\mathrm{mm}$ length of one of the fibre arms is encased in a heat-sinking aluminium block with a nichrome heating filament running parallel to it that is used to actively stabilise the optical path length and maintain the carrier transmission at a minimum.

where η is the refractive index of the fibre that we assume does not vary with ω, c is the speed of light in vacuum and we define $\mathcal{F} = 2\pi c/(\Delta l \eta)$ as the free spectral range (FSR) of the interferometer.

This means that the transmitted field at one of the output ports will be

$$\begin{aligned} E_\mathrm{out} &= \frac{1}{2}E_0 e^{i\omega t} + \frac{1}{2}E_0 e^{i(\omega t + 2\pi\omega/\mathcal{F})} \\ &= T(\omega)E_0 e^{i\omega t}, \end{aligned} \tag{6.2}$$

where E_0 is the amplitude of the input field and we define the transmission function $T(\omega) = \frac{1}{2}(1 + e^{i2\pi\omega/\mathcal{F}})$.[7] For an input intensity $I_\mathrm{in} = |E_0|^2$, the transmitted intensity $I_\mathrm{out} = |E_\mathrm{out}|^2$ then varies sinusoidally with $2\pi\omega/\mathcal{F}$ as

[7] Conservation of energy dictates that an extra factor of π accrues in phase between the two terms of the transmission function at the other output port so that, when there is complete destructive interference at one port, there is complete constructive interference at the other.

$$\begin{aligned} I_{\text{out}} &= |T(\omega)|^2 I_{\text{in}} \\ &= \frac{1}{2}\left[1 + \cos\left(\frac{2\pi\omega}{\mathcal{F}}\right)\right] I_{\text{in}}. \end{aligned} \tag{6.3}$$

For our interferometer, the path length difference between the two fibre arms is $\Delta l = 38.5(5)\,\text{mm}$, corresponding to a free spectral range (FSR) $\mathcal{F} \approx 2\pi \times 5.45\,\text{GHz}$ of approximately twice the EOM modulation angular frequency ω_{E}. This means that, when the carrier transmission $|T(\omega)|^2$ is minimised, the first order sidebands produced by the modulation are close to maximally transmitted.

In the limit of small modulation depth β, the phase modulated output of the EOM that we take to have unit amplitude

$$E_{\text{mod}} = e^{i(\omega t + \beta \sin(\omega_{\text{E}} t))} \tag{6.4}$$

can be approximated, to first order in β, as

$$E_{\text{mod}} \approx e^{i\omega t} + \frac{\beta}{2} e^{i(\omega+\omega_{\text{E}})t} - \frac{\beta}{2} e^{i(\omega-\omega_{\text{E}})t}, \tag{6.5}$$

comprising the carrier component with angular frequency ω and two sidebands at $\omega \pm \omega_{\text{E}}$. After passing through the Mach-Zehnder interferometer this becomes

$$\begin{aligned} E_{\text{T}} =& T(\omega) e^{i\omega t} + \frac{\beta}{2} T(\omega+\omega_{\text{E}}) e^{i(\omega+\omega_{\text{E}})t} \\ & - \frac{\beta}{2} T(\omega-\omega_{\text{E}}) e^{i(\omega-\omega_{\text{E}})t}, \end{aligned} \tag{6.6}$$

resulting in intensity modulation at the EOM angular frequency ω_{E},

$$\begin{aligned} I_{\text{T}} \propto |E_{\text{T}}|^2 =& |T(\omega)|^2 + \frac{\beta^2}{4}\left[|T(\omega+\omega_{\text{E}})|^2 + |T(\omega-\omega_{\text{E}})|^2\right] \\ & - \beta \sin\left(\frac{\pi\omega_{\text{E}}}{\mathcal{F}}\right) \sin\left(\frac{2\pi\omega}{\mathcal{F}}\right) \cos\left(\omega_{\text{E}} t + \frac{\pi\omega_{\text{E}}}{\mathcal{F}}\right) \\ & + (\text{terms oscillating at } 2\omega_{\text{E}} t). \end{aligned} \tag{6.7}$$

As depicted in Figure 6.3, this intensity modulation is detected by a suitably fast photodiode and the amplitude of the modulation

$$A = \beta \sin\left(\frac{\pi\omega_{\mathrm{E}}}{\mathcal{F}}\right) \sin\left(\frac{2\pi\omega}{\mathcal{F}}\right) \tag{6.8}$$

is extracted by mixing that signal with a fraction of the RF signal driving the EOM with an appropriate phase delay[8] and running the output through a low-pass filter.

[8] The ∼1 m wavelength of the RF signal corresponds rather nicely with readily available, pre-prepared lengths of coaxial cable.

As $\mathcal{F} \approx 2\omega_{\mathrm{E}}$, the term $\sin\left(\pi\omega_{\mathrm{E}}/\mathcal{F}\right) \approx 1$ and, as per the transmitted intensity of the carrier $|T(\omega)|^2$, the amplitude of the intensity modulation varies sinusoidally with $2\pi\omega/\mathcal{F}$. However, the signals are $\pi/2$ out of phase so that A passes through 0 when the carrier transmission is minimal or maximal. This makes it an appropriate error signal for stabilising the interferometer to a point of minimal carrier transmission.

To achieve this, the signal is passed through a proportional-integral (PI) control circuit and an Arduino micro-controller produces a pulsed logic signal whose pulse width is proportional to the PI output. This signal drives a field-effect transistor (FET) circuit that modulates the current I_{heat} through a ∼50 mm nichrome heater filament running parallel to one of the fibre arms of the interferometer within an aluminium heat-sinking block. By changing the temperature of the length of fibre, the optical path length difference $\Delta l\eta \propto 1/\mathcal{F}$ is varied to maintain $\sin(2\pi\omega/\mathcal{F})$ at the zero crossing and minimise the carrier transmission.

To minimise drifts in the optical path lengths due to variations in laboratory temperature, the entire interferometer sits within a cardboard enclosure to mitigate air currents and is thermally isolated with polystyrene insulation.[9] The thermal nature of the feedback gives it a low bandwidth so that any modulation it imposes on the light can be considered constant over the duration of a typical experimental sequence and the effect of the phase modulation of the EOM carrier frequency on the error signal occurs too quickly for it to affect the stabilisation.

[9] An A4 box file the department receptionist had no further use for proved the perfect size to form the enclosure, and the disproportionately liberal use of packing foam by the majority of equipment suppliers ensures that thermal insulation is never in short supply.

The result is ∼30 dB suppression of the carrier component from the output port in question, effectively just leaving the modulation sidebands to be amplified and delivered to the MOT. After fibre-couping losses, ∼ 15 % of the input light is emitted in the sidebands and amplified while, as shown in Figure 6.5 where incoming and outgoing beams are labelled with illustrative spectral profiles, the carrier is emitted from the other output port and dumped on a beam block.

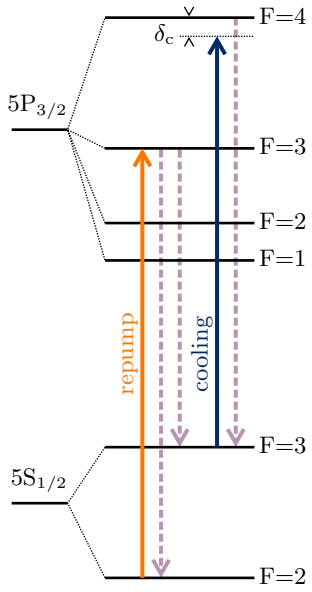

Figure 6.6: Level scheme showing the optical transitions driven by the cooling and repump beams during the preparation and read-out phases. Dashed arrows show possible spontaneous decay routes. The cooling beam drives the closed $|5S_{1/2}, F = 3\rangle \rightarrow |5P_{3/2}, F' = 4\rangle$ transition, but off-resonant coupling to the $|5P_{3/2}, F' = 3\rangle$ state results in slow $\sim$ 100 µs optical pumping to the $|5S_{1/2}, F = 2\rangle$ ground state. The repump beam drives the open $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition, quickly pumping atoms back into the $|5S_{1/2}, F = 3\rangle$ excited state.

6.3 Cooling, state preparation and read-out

As briefly touched upon in Section 6.1, atoms excited into the upper state manifold by the cooling and repump lasers will occasionally spontaneously decay back to the lower manifold and emit a photon in a random direction. The possible decay routes are shown as dashed arrows in Figure 6.6, while the excitations caused by absorbing photons from the cooling and repump lasers are represented with solid arrows.

The repump laser excites atoms from the $|5S_{1/2}, F = 2\rangle$ state to the $|5P_{3/2}, F' = 3\rangle$ state. As atoms in this state can decay to either the $|5S_{1/2}, F = 2\rangle$ state or the dark $|5S_{1/2}, F = 3\rangle$ state — where they are no longer resonant with the repump laser — then, when just the repump light is activated, atoms will quickly accumulate in the dark state via optical pumping.

The cooling light excites atoms from the $|5S_{1/2}, F = 3\rangle$ state to the $|5P_{3/2}, F' = 4\rangle$ state from which they can only decay back down to the $|5S_{1/2}, F = 3\rangle$ state and remain coupled by the cooling laser. The closed nature of the cooling transition means that atoms interacting with this laser will experience repeated spontaneous decays, with the photons they emit in random directions making the MOT visible and the associated recoil force ultimately providing the cooling mechanism by unbalancing the otherwise symmetric forces associated with stimulated emission and absorption.

However, in practice this cycle does not continue in perpetuity. The cooling laser is detuned by a factor δ_c on the order of the natural linewidth of the transition, so is already driving off-resonant excitations. There is also a far weaker off-resonant coupling to the $|5P_{3/2}, F' = 3\rangle$ state and, in the relatively rare event that an atom ends up here, it has a chance to decay down to the $|5S_{1/2}, F = 2\rangle$ state that is dark to the cooling laser. All atoms will go dark within a few hundred µs, and so the repump laser needs to be active in order to sustain the MOT long enough to load an appreciable number of atoms.

A telescopic lens setup collects a fraction of the fluorescence photons associated with all these decays onto a *Hamamatsu H7422-50* PMT. A pinhole at the shared focus of the telescope lenses and a 780 nm band-pass filter serve to block out the majority of any other light. Figure 6.7 shows how the signal from this PMT varies during a typical experimental cycle.

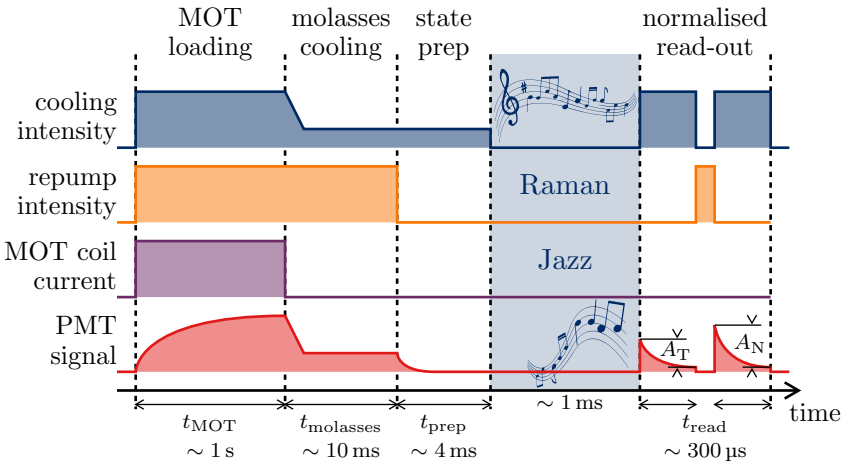

Figure 6.7: Representation, not to scale, of how the MOT fields are varied during the preparation and read-out phases, and the corresponding fluorescence signal. The cooling, repump and MOT coils are all on for the MOT loading phase that typically lasts for $t_{\mathrm{MOT}} \approx 1\,\mathrm{s}$. The coil current is then quenched, and the cooling beam is reduced in power over $\sim 1\,\mathrm{ms}$ so that the atoms are cooled in an optical molasses for $t_{\mathrm{molasses}} \approx 10\,\mathrm{ms}$. The repump beam is extinguished $t_{\mathrm{prep}} \approx 4\,\mathrm{ms}$ prior to the cooling beam in order to optically pump the cooled atoms into the $|5S_{1/2}, F = 2\rangle$ ground state. After an experimental pulse sequence — or 'Raman jazz' — lasting $\sim 1\,\mathrm{ms}$, the fraction of atoms excited into the $|5S_{1/2}, F = 3\rangle$ state in then measured by fitting exponential curves to the decaying fluorescence signal associated with two pulses of the cooling beam, each lasting $t_{\mathrm{read}} = 300\,\mu\mathrm{s}$, separated by $t_{\mathrm{repump}} = 100\,\mu\mathrm{s}$ during which the repump beam is pulsed in order to quickly transfer all of the atoms to the excited $|5S_{1/2}, F = 3\rangle$ state. The amplitude measured for the second 'normalisation' pulse A_{N} then corresponds to all the atoms being excited, and the ratio between the two amplitudes $A_{\mathrm{T}}/A_{\mathrm{N}}$ represents the fraction of atoms excited during the experimental sequence.

Each experiment starts with a MOT loading phase, during which the MOT quadrupole coils, cooling and repump lasers are all on. This phase lasts for $t_{\mathrm{MOT}} \sim 1\,\mathrm{s}$ during which the number of trapped atoms increases and the PMT signal asymptotically rises to a steady-state value.

This is followed by a molasses cooling phase that lasts for $t_{\mathrm{molasses}} \sim 10\,\mathrm{ms}$, during which the current in the quadrupole coils is quenched and the power of the cooling laser is ramped down to a lower level within $\sim 1\,\mathrm{ms}$. Varying this level offers an effective mechanism for controlling the temperature of the cloud. As the PMT signal is dominated by fluorescence on the cyclic cooling transition, it reduces in line with the power of the cooling laser during this phase.

The repump laser is then extinguished $t_{\mathrm{prep}} \sim 4\,\mathrm{ms}$ prior to the cooling laser. During this time the cooling laser pumps atoms into the $|5S_{1/2}, F = 2\rangle$ state, continuing to provide a cooling force

while it does so. This is visible as an exponential decay of the PMT signal as the atoms go dark. As the cooling laser power is maintained at a low level following the molasses cooling phase, this pumping occurs slower than it would with the cooling laser at full power, necessitating the relatively long t_{prep} duration as increasing the cooling power for faster pumping results in unwanted heating.

After the state preparation is completed, the cooling laser is extinguished and the actual coherent control experiments are performed by various sequences of pulses with the Raman lasers that are typically completed in $< 1\,\mathrm{ms}$. The measured output of these experiments is the fraction of atoms prepared in the $|5S_{1/2}, F = 2\rangle$ state that have been transferred to the $|5S_{1/2}, F = 3\rangle$ state. This is determined by pulsing the cooling laser at full power for 300 µs. This results in a decaying fluorescence signal to which an exponential function is fitted to extract an amplitude A_{T}. All of the atoms are then pumped into the excited state by a short 100 µs pulse of the repump laser.[10] The amplitude A_{N} of the fluorescence signal associated with a second pulse of the cooling laser then serves as a normalisation factor, with the ratio $A_{\mathrm{T}}/A_{\mathrm{N}}$ providing the measure of the fraction of atoms transferred.

[10] While the cooling laser drives a cycling transition, meaning that each atom can spontaneously emit many thousands of photons before being pumped to the dark state after off-resonant excitation, the repump drives an open transition. This means that the pumping is very fast and the fluorescence signal very weak. With a lot of atoms it is occasionally visible as a small spike following the first cooling pulse, but it is not uncommon for it not to be visible at all.

6.4 Magnetic field switching and optimisation

Transitioning atoms from the magnetic gradient required to keep them trapped in the MOT to a homogeneous magnetic environment conducive to the experiments that follow is not a trivial part of any atomic physics experiment.

In the Southampton setup, large anti-Helmholtz coils that are external to the vacuum chamber provide the quadrupole field for the atoms inside. When on, each coil draws $\sim 2\,\mathrm{A}$ of current and they double up as 60 W heaters that serve to desorb enough rubidium from the walls of the chamber to negate the need for the fitted alkali metal dispensers when loading the MOT. Around 0.2 J of energy is stored in the quadrupole field when it is active and so, in order to quench it on ms timescales, on the order of 100 W of power needs to be dissipated somewhere.

This strain is taken by power FETs and Zener diodes, and the current through the coils is extinguished within 1 ms. However, eddy currents sustain the magnetic field for somewhat longer. A Hall probe external to the chamber suggests that a steady state

is reached in $\sim 3\,\text{ms}$, although it is the field the atoms experience that is of interest.

6.4.1 A brief look at two-photon Zeeman spectra (measuring the field experienced by the atoms)

In the presence of a magnetic field, the magnetically sensitive sublevels of the $|5S_{1/2}, F = 2\rangle$ and $|5S_{1/2}, F = 3\rangle$ states, between which we drive two-photon Raman transitions, will receive an energy shift proportional to their m_f value. The constant of proportionality, or g-factor, for this splitting is equal and opposite for each state to within 1%. This results in a frequency shift to the Raman resonance that is dependent on the initial m_f state and, as our state preparation does not preferentially pump into any particular sublevel, sub-Doppler Raman spectroscopy of the cloud will reveal multiple peaks at different values of the two-photon detuning δ.

We can perform sub-Doppler Raman spectroscopy by modifying the delivery system described in Section 6.2 to combine both Raman components in one of the two PM fibres with orthogonal polarisations and removing the PBSC from the output beam shaper. This results in co-propagating beams, such that a moving atom will see both frequency components Doppler shifted with the same sign and the two-photon resonance condition is unaffected.

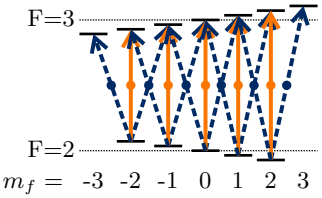

Figure 6.8: Allowed Raman transitions between Zeeman shifted m_f states.

When the Raman beams are propagating parallel to the magnetic field vector, which provides a quantisation axis for the Zeeman splitting, the two-photon transitions impart no net angular momentum along the magnetic field axis and so they only drive $\Delta m_f = 0$ transitions between states with equal m_f values, shown as solid orange lines in Figure 6.8. This results in a five-peak Zeeman spectrum. In contrast, when the Raman beams propagate perpendicular to the magnetic field vector they impart a unit of angular momentum along the quantisation axis, driving the $\Delta m_f = \pm 1$ transitions shown as dashed blue lines in Figure 6.8. There are 10 such transitions but adjacent ones, illustrated by solid blue circles at the crossover points, receive an equal frequency shift[11] resulting in a six peak Zeeman spectrum. As the frequency shifts to the $\Delta m_f = 0$ and $\Delta m_f = \pm 1$ are interleaved, when the Raman propagation axis has components both parallel and perpendicular to the magnetic field vector then both sets of peaks are visible resulting in eleven evenly spaced peaks at the frequencies marked by solid circles in Figure 6.8 [135].

[11] Equal in the sense that the g-factors are equal, to within 1%.

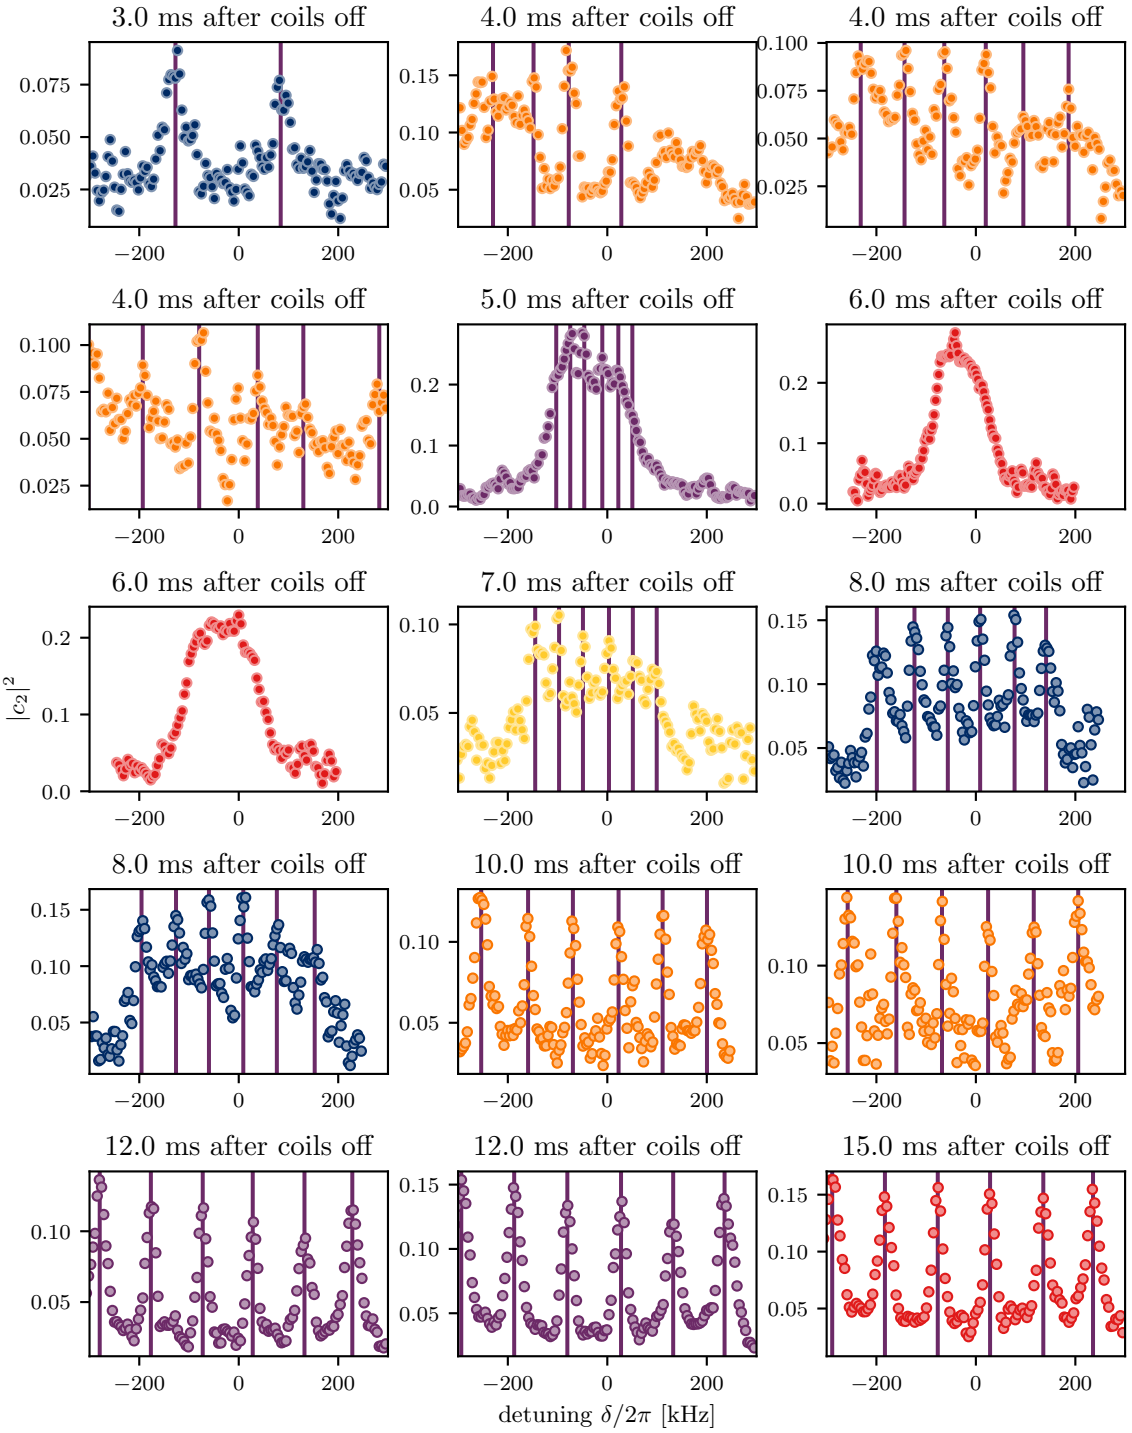

Figure 6.9: Two-photon Zeeman spectra — showing the fraction of atoms transferred to the upper state $|c_2|^2$ as a function of the detuning of an applied, co-propagating, Raman pulse — measured for a range of t_{molasses} values, labelled in terms of the duration between extinguishing the MOT coil current and pulsing the Raman lasers. The locations of peaks, represented by solid vertical lines, are determined by a peak finding algorithm that identifies peaks based on carefully chosen threshold values and interpolates with Gaussian fitting to data points about the identified maximum to estimate the peak centre. The thresholds and the number of data points used in the Gaussian fit were varied between data sets in order to achieve results that appeared sensible to a human eye, and some identified peaks that were deemed to be falsely identified have been omitted.

6.4.2 Magnetic field switching transients

Using such two-photon Zeeman spectra as a probe, the variation in the magnetic field vector after the MOT quadrupole field in quelled can be monitored. Employing low power Raman pulses ($\sim 1\,\mathrm{mW}$ per beam) with a long ($\sim 300\,\mu\mathrm{s}$) duration to reduce power broadening, the fraction of atoms excited to the $|5S_{1/2}, F = 3\rangle$ state was measured as a function of the two-photon detuning δ to take Zeeman spectra for a range of t_{molasses} values. For these sub-Doppler measurements, a cold cloud was not essential and the cooling beam was maintained at a relatively high power during the molasses cooling so that $t_{\mathrm{prep}} = 2\,\mathrm{ms}$ was a sufficient duration for the state preparation. The duration of the cooling beam power ramp was kept fixed at 1 ms, acting as a lower limit to t_{molasses} and meaning that the minimum delay between quenching the coil current and pulsing the Raman beams was 3 ms.

After this duration the coil current has completely diminished and the magnetic field on the coil axis exterior to the chamber has reduced to $< 10\%$ of its maximum value as measured on a Hall probe, but significant variation is still experienced by the atoms. Spectra for a range of t_{molasses} values, labelled in terms of the delay between turning the coils off and pulsing the Raman beams, are presented in Figure 6.9.

Prior to taking these spectra, three Helmholtz coil pairs had been used to zero the magnetic field at $\sim 6\,\mathrm{ms}$ after switching the coils off in a manner akin to that described in the next section. 6 ms

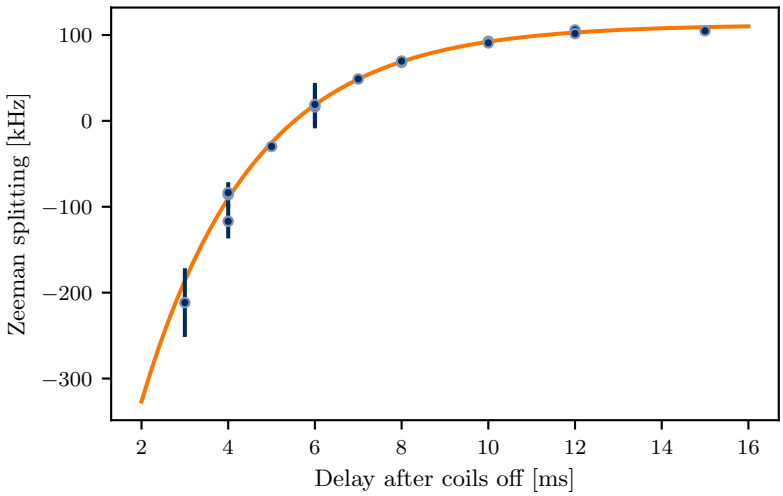

Figure 6.10: Mean splitting between adjacent $\Delta m_f = \pm 1$ transitions as a function of the duration after extinguishing the coil current, taking into account a change in sign about the point where the states are degenerate.

after switching the coils off, the Zeeman peaks are thus sufficiently overlapped so as to be unresolvable. The spectra taken at shorter delays are noisy and hard to interpret, likely because the magnetic field is varying fast relative to the 300 µs Raman pulses. Nonetheless, some $\Delta m_f = \pm 1$ peaks are still resolvable with possible signs of $\Delta m_f = 0$ peaks between them.

For delays longer than 6 ms, $\Delta m_f = \pm 1$ peaks become clearly visible and diverge, while $\Delta m_f = 0$ peaks remain suppressed. The MOT coils are aligned orthogonally to the Raman beam axis, and this is perhaps an indication that, while the quadrupole field rescinds relatively quickly, eddy currents sustain a field along the coil axis for much longer.[12] The temporal evolution of the mean separation between adjacent $\Delta m_f = \pm 1$ states taken from these spectra is shown in Figure 6.10, taking into account a change in sign as the states pass through degeneracy, and is well fitted by an exponential function with a $1/e$ time constant of ~ 2.6 ms that approaches a steady-state asymptote only after ~ 15 ms.

[12] The MOT quadrupole coils are located along the y axis of Figure 6.2, while the Raman beams are injected along the z axis.

6.4.3 Achieving Zeeman degeneracy

The Zeeman splitting evident in sub-Doppler spectra will cause unwanted broadening in Doppler-sensitive experiments and so three pairs of Helmholtz coils located along the x, y and z axes of

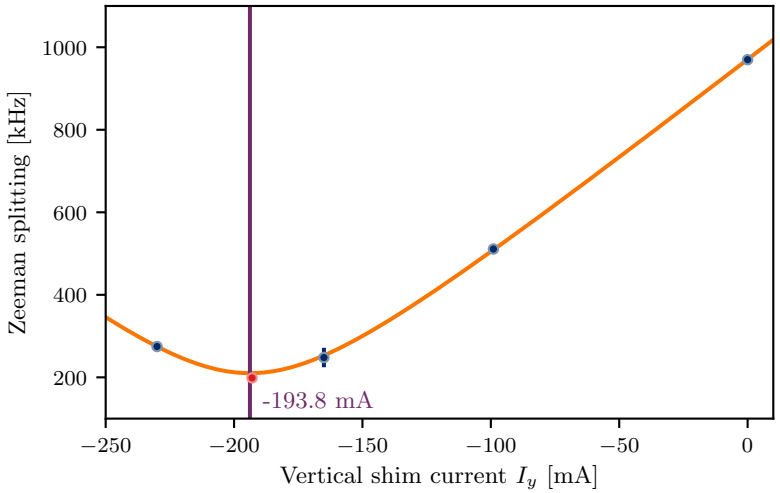

Figure 6.11: Mean frequency separation between adjacent $\Delta m_f = 0$ and $\Delta m_f = \pm 1$ peaks as a function of the Helmholtz coil current along the vertical y axis. Four measurements are made, and a function of the form $\Delta = A\sqrt{(I_y - I_0)^2 + c}$ is fitted to find the current value I_0, marked with a vertical line, that minimises the splitting and zeroes the magnetic field along that axis. A verification measurement, taken at $I_y = I_0$ and shown in a different colour, agrees with the expected minimum value.

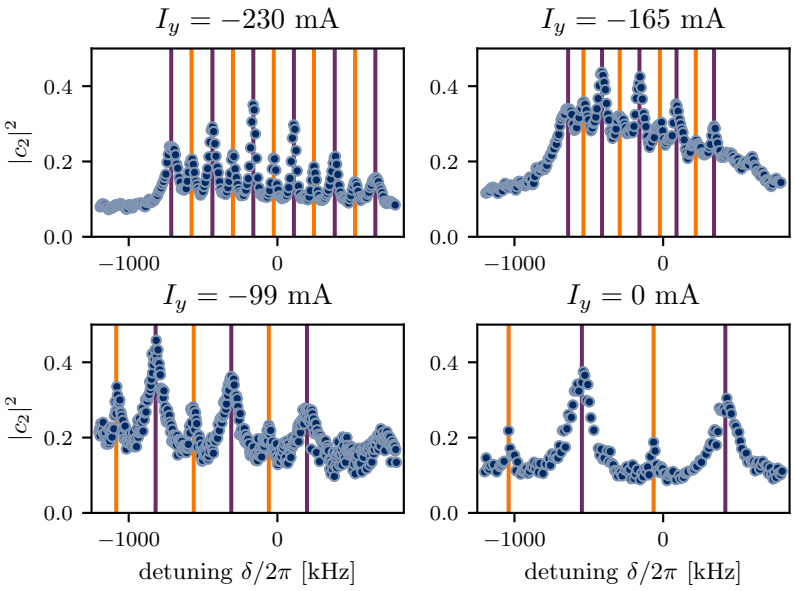

Figure 6.12: Two-photon Zeeman spectra taken at various values of the current in the vertical Helmholtz coils I_y corresponding to the data points in Figure 6.11. Purple (dark) and orange (light) vertical lines indicate resolvable $\Delta m_f = \pm 1$ and $\Delta m_f = 0$ peaks respectively, identified with the same peak fitting algorithm — with individually tailored threshold values — employed in Figure 6.9.[13]

Figure 6.2 are used to zero the magnetic field so that the Zeeman sublevels overlap once the magnetic transient has died out.

This is achieved by measuring the mean splitting between adjacent pairs of $\Delta m_f = 0$ and $\Delta m_f = \pm 1$ peaks for a range of values for a particular coil current. The magnitude of this splitting scales with the scalar magnetic field magnitude $B = \sqrt{B_x^2 + B_y^2 + B_z^2}$, where $B_{x,y,z}$ are the components of the magnetic field along each axis. As these scale linearly with the Helmholtz coil currents $I_{x,y,z}$ along the same axes, the splitting as a function of coil current varies with the form

$$\Delta = A\sqrt{(I - I_0)^2 + c}, \tag{6.9}$$

where Δ is the splitting, I is the current through the coil pair in question and A, c and I_0 are fitting parameters with I_0 representing the current that zeroes the magnetic field along that particular axis. An example of this is shown in Figure 6.11, where the mean splitting between adjacent $\Delta m_f = 0$ and $\Delta m_f = \pm 1$ peaks is plotted as a function of the Helmholtz coil current I_y along the vertical y axis. Four measurements from the spectra shown in Figure 6.12 are plotted, to which Equation (6.9) is fitted to find the current value I_0 that minimises the splitting and zeroes the

[13] These spectra were actually taken $\sim$ 6 ms after the MOT coil current was quenched, prior to determining how long-lived the switch-off transient is, and the fact that the magnetic field was still varying may contribute to the asymmetry evident in the peak heights.

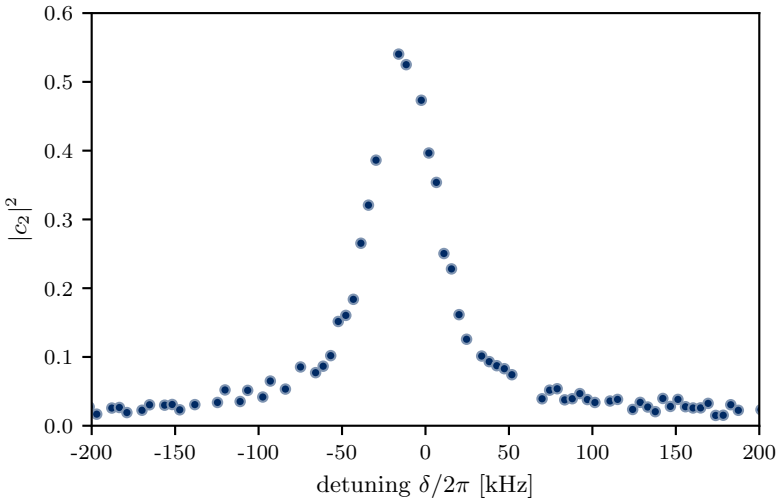

Figure 6.13: Sub-Doppler spectrum taken $> 15\,\text{ms}$ after the MOT coils were turned off and once the magnetic field has been zeroed along all three Cartesian axes. It shows a single peak, with all of the Zeeman sublevels degenerate to well within $2\pi \times 50\,\text{kHz}$. The detuning is measured relative to the bare hyperfine splitting ω_0, from which the centre of the resonance is shifted due to the light shifts induced by the Raman beams.

magnetic field along that axis. A verification measurement, taken at $I_y = I_0$ agrees with the expected minimum value.

Repeating this process for each Cartesian axis allows the magnetic field that the atoms experience to be zeroed sufficiently that all of the Zeeman sublevels are degenerate to well within 50 kHz, as shown in the sub-Doppler spectrum presented in Figure 6.13.

6.5 MOT temperature characterisation

The power of the cooling beam is controlled by an AOM shutter as shown in Figure 6.1. The carrier frequency of the AOM is varied with a control voltage produced by a *Keysight 33612A* arbitrary waveform generator (AWG), changing the diffraction efficiency and hence power to vary it in the manner illustrated in Figure 6.7.

The value to which this control voltage is ramped during the molasses sequence affects the efficiency of the sub-Doppler cooling process and hence provides a suitable knob for varying the cloud temperature during the experimental sequence.

The cloud temperature can be measured with Raman Doppler spectroscopy [136, 137]. By reducing the power of the Raman beams and increasing the two-photon detuning, the Rabi frequency can be reduced sufficiently so that the spectral width of a π-pulse is much smaller than the Doppler width of the cloud. In this case,

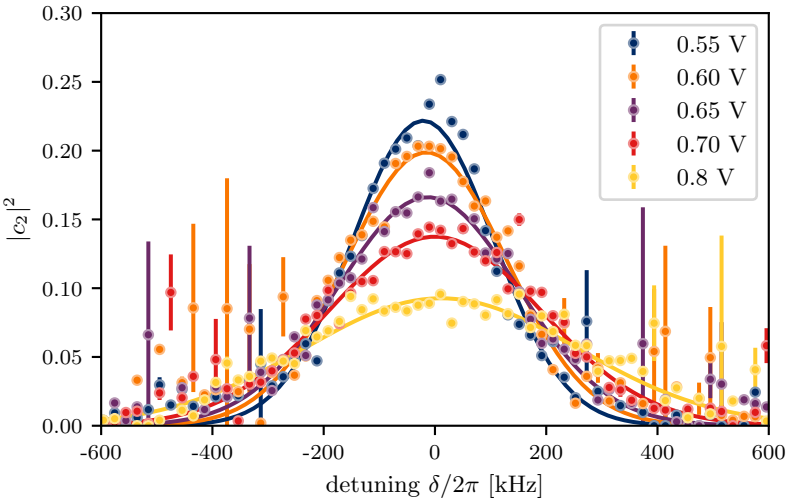

Figure 6.14: Maxwell-Boltzmann distributions measured with 13 µs π-pulses for a range of voltages that control the cooling beam power.

measuring the number of atoms transferred by a π-pulse into the upper ground state as a function of the two-photon detuning δ will produce a Gaussian profile whose width is dominated by the 1-D Maxwell-Boltzmann velocity distribution of the atoms along the Raman propagation axis.

The temperature can be estimated by fitting a Gaussian function to these data, such as in Figure 6.14. Depending on the MOT alignment, and other experimental parameters that drift from day to day, the temperature can be tuned in a range between ∼20 µK and ∼140 µK. A typical plot of temperature as a function of the voltage controlling the cooling beam power, taken from the Gaussian fits in Figure 6.14, is shown in Figure 6.15.

Colder and hotter samples can be achieved by varying the detuning of the cooling beam. However, the ability to vary the detuning during an experimental run has not yet been implemented and so the same detuning is used during both the MOT loading and molasses cooling. Detuning further to the red improves the efficiency of the sub-Doppler cooling and results in a cooler sample but also reduces the number of atoms loaded and adversely affects the signal-to-noise ratio (SNR). Adjusting the detuning to maximise the number of atoms and varying just the beam power to control the cloud temperature has proved to be flexible enough so far but, for future experiments where a narrower velocity distribution may be preferable, we plan to incorporate a detuning ramp into the loading sequence presently.

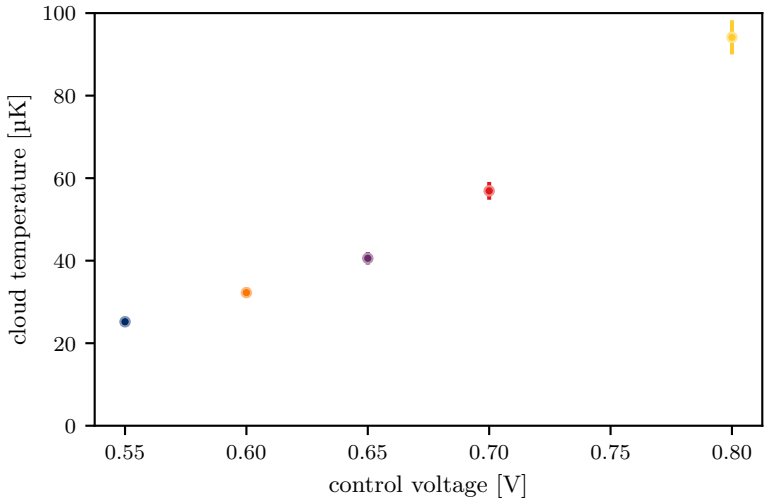

Figure 6.15: Estimated temperature as a function of the voltage controlling the cooling beam power derived from the Gaussian fits in Figure 6.14.

6.6 Summary of experimental parameters

Now would seem an opportune moment to collate, in one place, some key experimental parameters that could prove useful to those who work on this experiment in the future. Table 6.1 lists the typical conditions of the MOT fields used throughout this thesis, along with some representative powers and detunings of the Raman interaction beams used in the experiments that are detailed later.

Approximate Rabi frequencies and light shifts associated with different Raman beam parameters are also listed. Light shifts were determined by identifying the detuning that resulted in peak population transfer after applying a π-pulse, while Rabi frequencies were determined by detuning the Raman transition to the light-shifter resonance, applying pulses of varying durations — in order to observe Rabi flopping — and associating the duration of a π-pulse to the first peak in population transfer. This process is somewhat iterative — the Rabi flopping signal is required to determine the duration of a π-pulse, while the π-pulse duration is needed in order to determine the light shift and tune the lasers to resonance for the Rabi flopping — but in practice the pulse duration does not need to be too exact for the light shifted resonance to be identified sufficiently.

Parameter	Typical value
Magneto-optical trap fields	
Cooling beam total power (during MOT loading)	12 – 16 mW
Cooling beam total power (during molasses cooling)	0.5 – 12 mW (adjusted to vary cloud temperature)
Repump beam total power	2 – 4 mW
MOT beams $1/e^2$ diameter	7.5 mm
Cooling beam detuning	$\sim 2\pi \times 6 - 18$ MHz ($1 - 3\gamma$)
Magnetic field gradient (along coil axis)	$\sim 5\,\mathrm{G\,cm^{-1}}$
Loading duration	$\sim 0.5 - 2$ s
Molasses cooling duration	∼15 ms (inc. 1 ms to ramp cooling power)
Raman transition parameters (for velocimetry experiments)	
EOM beam power (inc. off-resonant sideband)	∼200 mW
AOM beam power	∼100 mW
Carrier wavelength	∼780.256 nm ($\Delta \approx -2\pi \times 7$ GHz)
Rabi frequency	$\sim 2\pi \times 650$ kHz
Light shift	$\sim -2\pi \times 1$ MHz
Raman transition parameters (for interferometry / optimal control experiments)	
EOM beam power (inc. off-resonant sideband)	∼100 mW
AOM beam power	∼50 mW
Carrier wavelength	∼780.270 nm ($\Delta \approx -2\pi \times 14$ GHz)
Rabi frequency	$\sim 2\pi \times 250$ kHz
Light shift	$\sim -2\pi \times 200$ kHz
Raman transition parameters (for Doppler spectroscopy / Zeeman spectra)	
EOM beam power (inc. off-resonant sideband)	∼10 mW
AOM beam power	∼5 mW
Carrier wavelength	∼780.270 nm ($\Delta \approx -2\pi \times 14$ GHz)
Rabi frequency	$\sim 2\pi \times 40$ kHz
Light shift	$\sim -2\pi \times 16$ kHz (comparable to the two photon recoil shift)

Table 6.1: Summary of typical experimental parameters.

Velocimetry of cold atoms with matter-wave interferometry

7

It's quantum, baby.

— Keanu Reeves *(Hello to Jason Isaacs)*

At its heart, the Mach-Zehnder interferometer sequence that is the basis of atom interferometric sensors of rotation, gravitational acceleration and gradients thereof, can be considered as comprising two measurements of atomic velocity and being sensitive to the difference between them. While much has been written about these interferometers — from reporting state-of-the-art achievements such as macroscopic wavefunction separation in interferometers on 10 m length scales [27] and proposals to perform interferometry with beams of atoms travelling between satellites in the vacuum of space [138, 139], to technical analyses of higher-order phase terms acquired in the presence of external fields [125] and at short interrogation times [140] — there have been few investigations into the underlying velocity measurements.

This is perhaps because only a relative velocity measurement between the atoms and the apparatus can be made and, as they begin an experiment in the same inertial frame, this negates any obvious applications as a velocity sensor. However, Weitz and Hänsch proposed utilising the velocity-dependent interferometer output to cool atoms and molecules with carefully timed laser pulse sequences, relaxing the constraints on laser frequency that most laser cooling mechanisms are subject to [43], and this was later demonstrated by Dunning *et al.* [44] while Weiss *et al.* [33] employed the technique to measure the photon recoil.

In this chapter we show that atom interferometry can be used to measure the distribution of velocities within an atom cloud itself. We study, in detail, the velocimetry process that underlies all atom interferometric sensors, and it proves a unique and informative lens through which to consider the field in general. This culminates in the development of a precision, interferometric technique for cold atom velocimetry.

In Section 7.1 we consider the velocity sensitivity of a two-pulse Ramsey interferometer sequence in both the laboratory and atomic reference frames. Based on ideas that were developed jointly with co-authors in [1], this serves both to introduce the velocimetry concept and provides a fresh context in which to consider atom interferometry in general.

In Section 7.2 we present results of a velocimetry experiment, based upon the principles we have established, that show fundamental shortcomings of the technique in a non-ideal physical experiment where finite length laser pulses exhibit intra-pulse Doppler sensitivity that is non-negligible. Section 7.3 presents a three-pulse velocimetry experiment, published in [3], that mitigates these shortcomings and proves to be an effective technique.

7.1 Conceptual overview: velocimetry with a Ramsey interferometer

The phase measured at the output of a matter-wave interferometer is independent of any chosen reference frame, but it is nonetheless instructive to consider the contribution of an atom's velocity to the interferometric phase in both the laboratory and atomic rest frames, where it defines both the evolution of the atom's quantum mechanical state and its classical trajectory.

7.1.1 Laboratory frame

Depicted in the laboratory frame, in which atoms move relative to fixed lasers, interaction of the atoms with the optical field prepares and resolves a quantum mechanical superposition of two atomic states with different momenta. The residual energy difference between these superposed states upon subtraction of the laser frequency then determines the rate at which interferometric phase is accrued. The lasers thus play the role of a frequency reference for the experiment and a source of the $\pi/2$- and π-pulses that act as the matter-wave beamsplitters and mirrors.

The velocity dependence of the interferometric phase accrued following a $\pi/2$ beamsplitter pulse becomes evident when the kinetic energy, and its modification by the photon recoil, is included in the atomic Lagrangian [37]. For illustrative purposes we consider the simple example of an atom with two electronic states $|1\rangle$ and $|2\rangle$,

radiatively coupled upon absorption or emission of a single photon with angular frequency ω and wavevector $\boldsymbol{k}$.[1] Denoting the energies of the internal electronic states $\mathcal{E}_{1,2}$ and writing the combined internal and external motional states as $|1, \boldsymbol{p}_1\rangle$ and $|2, \boldsymbol{p}_2\rangle$, conservation of energy and momentum dictate that

$$\hbar\omega = \left(\mathcal{E}_2 + \frac{|\boldsymbol{p}_2|^2}{2m}\right) - \left(\mathcal{E}_1 + \frac{|\boldsymbol{p}_1|^2}{2m}\right) \text{ and} \tag{7.1}$$

$$\hbar\boldsymbol{k} = \boldsymbol{p}_2 - \boldsymbol{p}_1, \tag{7.2}$$

respectively, where m is the atomic mass. Writing $\boldsymbol{p}_{1,2} \equiv \boldsymbol{p} \mp {}^1\!/_2\hbar\boldsymbol{k}$ [104], so that the photon couples states $|1, \boldsymbol{p} - {}^1\!/_2\hbar\boldsymbol{k}\rangle$ and $|2, \boldsymbol{p} + {}^1\!/_2\hbar\boldsymbol{k}\rangle$ to satisfy Equation (7.2), Equation (7.1) becomes

$$\hbar\Delta = \hbar\omega - (\mathcal{E}_2 - \mathcal{E}_1) = \frac{\boldsymbol{p}\cdot\hbar\boldsymbol{k}}{m}, \tag{7.3}$$

where Δ is the detuning from resonance. This is simply the classical Doppler shift, illustrated graphically in Figure 7.1 where a photon is represented by a line of slope c that must connect two parabolas indicating the combined electronic and kinetic energy of each atomic state so that the length of the line must increase with the component p of the atomic momentum in the direction of the photon propagation.

In the absence of any external field, the Lagrangian for each of the photon-coupled states is simply

$$\mathcal{L}_{1,2} = \frac{{\boldsymbol{p}_{1,2}}^2}{2m}, \tag{7.4}$$

so that

$$\mathcal{L}_2 - \mathcal{L}_1 = \hbar\Delta. \tag{7.5}$$

The phase ϕ accrued by a superposition prepared by a ${}^\pi\!/_2$-pulse from the laser at $t = t_1$ after a period $\tau \equiv t_2 - t_1$ is then just [141]

[1] This is readily extended to the two-photon Raman transitions we employ in our experiment, where the difference in frequency between the two photons plays the role of the frequency reference and an effective wavevector is defined by the difference between the individual wavevectors, as detailed in Section 2.6.

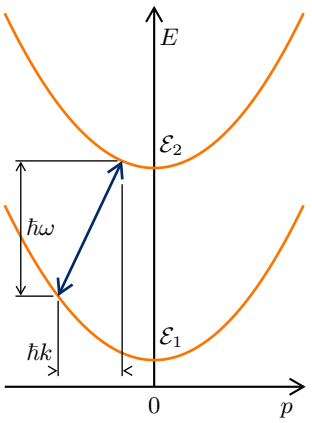

Figure 7.1: Graphical representation of the conservation of energy and momentum during absorption and emission of a photon with angular frequency ω. The parabolic dispersion curves represent the combined kinetic energy ${}^{p^2}\!/_m$ and electronic energy $\mathcal{E}$ of an atom in each of its two internal states, and the photon is represented by the line of slope c whose length varies, according to the classical Doppler shift, with the atomic momentum p.

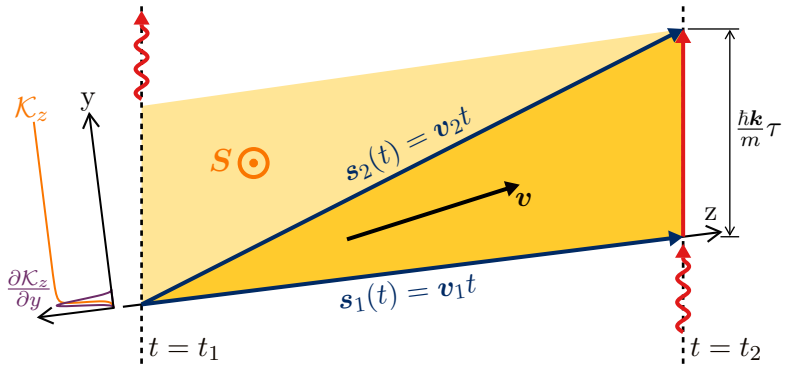

Figure 7.2: Interferometer paths $s_{1,2}(t)$ and area S used for calculating the path integral of the quantum mechanical action.

$$\begin{aligned}\phi &= \frac{1}{\hbar}\int_{t_1}^{t_2}(\mathcal{L}_2 - \mathcal{L}_1)\,\mathrm{d}t \\ &= \int_{t_1}^{t_2}\Delta\,\mathrm{d}t \\ &= \frac{\boldsymbol{p}\cdot\boldsymbol{k}}{m}\tau = \boldsymbol{k}\cdot\boldsymbol{v}\tau\end{aligned} \tag{7.6}$$

in terms of the atomic velocity $\boldsymbol{v}$. A subsequent $\pi/2$-pulse at $t = t_2$ can then allow this phase to be measured by mapping it onto the probability for the atom to be detected in one or other of the superposed states.

Defining

$$\mathcal{L} = \mathcal{K}\cdot\boldsymbol{p} \tag{7.7}$$

allows the path integral [121] of Equation (7.6) to be rewritten as

$$\begin{aligned}\phi &= \frac{1}{\hbar}\left(\int_{t_1}^{t_2}\mathcal{K}\cdot\boldsymbol{p}_2\,\mathrm{d}t - \int_{t_1}^{t_2}\mathcal{K}\cdot\boldsymbol{p}_1\,\mathrm{d}t\right) \\ &= \frac{m}{\hbar}\left(\int_{\boldsymbol{s}_2(t_1)}^{\boldsymbol{s}_2(t_2)}\mathcal{K}_2\cdot\mathrm{d}\boldsymbol{s}_2 - \int_{\boldsymbol{s}_1(t_1)}^{\boldsymbol{s}_1(t_2)}\mathcal{K}_1\cdot\mathrm{d}\boldsymbol{s}_1\right) \\ &= \frac{m}{\hbar}\left(\oint\mathcal{K}\cdot\mathrm{d}\boldsymbol{s} + \int_{\boldsymbol{s}_1(t_2)}^{\boldsymbol{s}_2(t_2)}\mathcal{K}\cdot\mathrm{d}\boldsymbol{s}\right) \equiv \phi_{\mathrm{loop}} + \phi_{12},\end{aligned} \tag{7.8}$$

where $\boldsymbol{s}_{1,2}$ is the path followed by state (1,2) from $t = t_1$ to $t = t_2$, as shown in Figure 7.2, and the final term ϕ_{12} is the shift due to the separation of the wavepackets at $t = t_2$ [125].

The Kelvin-Stokes theorem then allows the first term in Equation (7.8) to be cast as a surface integral over $\boldsymbol{S}$, the area enclosed by the interferometer, giving

$$\phi_{\text{loop}} = \frac{m}{\hbar} \iint (\boldsymbol{\nabla} \times \boldsymbol{\mathcal{K}}) \cdot \mathrm{d}\boldsymbol{S}. \tag{7.9}$$

External magnetic, electronic or gravitational fields, or equivalent non-inertial motion, will be manifest in the quantity $\boldsymbol{\nabla} \times \boldsymbol{\mathcal{K}}$ so that the interferometer is sensitive to them, and this formulation of the integral shows explicitly that the sensitivity of this simple two-pulse interferometer scales with the enclosed area. In the common Mach-Zehnder interferometer configuration [22], two such measurements are combined so that the wavepackets overlap at the output and the ϕ_{12} term is eliminated.

For the Lagrangian of Equation (7.4), $\boldsymbol{\mathcal{K}} = \boldsymbol{p}/2m = \boldsymbol{v}/2$ and is uniform except during the beamsplitter interaction. To evaluate the surface integral, we consider the impulse to occur over a small but finite time and distance. Defining a z axis $\boldsymbol{z} \equiv \boldsymbol{v}_1 t$ to lie colinear with $\boldsymbol{s}_1$, and a y axis in the $\boldsymbol{k}$–$\boldsymbol{v}$ plane, the first term in the interferometer phase can be written as

$$\begin{aligned} \phi_{\text{loop}} &= \frac{m}{\hbar} \iint (\boldsymbol{\nabla} \times \boldsymbol{\mathcal{K}}) \cdot \boldsymbol{e_x} \, \mathrm{d}y \, \mathrm{d}z \\ &= \frac{m}{\hbar} \iint \left(\frac{\partial \mathcal{K}_z}{\partial y} - \frac{\partial \mathcal{K}_y}{\partial z} \right) \mathrm{d}y \, \mathrm{d}z, \end{aligned} \tag{7.10}$$

where the integral is over the triangular area enclosed by the two interferometer paths. As the initial state $|1, \boldsymbol{p} - {}^1\!/_2\hbar\boldsymbol{k}\rangle$ receives no impulse, ${}^{\partial\boldsymbol{\mathcal{K}}}/_{\partial z} = 0$ and, as the beamsplitter impulse on the other state occurs around $z = y = 0$, ${}^{\partial \mathcal{K}_z}/_{\partial y}$ will be zero in most of the triangle so that the area can be extended to a trapezium without affecting the result and the integrals over z and y can be separated.

The integral over y then follows from conservation of momentum,

$$\int \frac{\partial \mathcal{K}_z}{\partial y} \, \mathrm{d}y = \frac{\hbar k_z}{2m} = \frac{\hbar \boldsymbol{k} \cdot \boldsymbol{v_1}}{2m v_1}, \tag{7.11}$$

and the first term in Equation (7.8) is thus

$$\phi_{\text{loop}} = \frac{m}{h} \int_0^{v_1 \tau} \frac{\hbar \boldsymbol{k} \cdot \boldsymbol{v_1}}{2m v_1} \, \mathrm{d}z = \frac{1}{2} \boldsymbol{k} \cdot \boldsymbol{v}_1 \tau. \tag{7.12}$$

As the path separation at $t = t_2$ will be $\hbar\boldsymbol{k}\tau/m$, the second term may be written, to first order in $\hbar\boldsymbol{k}/m$, as

$$\phi_{12} = \frac{m}{\hbar}\int_{\boldsymbol{s}_1(t_2)}^{\boldsymbol{s}_1(t_2)+\hbar\boldsymbol{k}\tau/m} \frac{\boldsymbol{v}_2}{2}\cdot \mathrm{d}\boldsymbol{s} = \frac{m}{\hbar}\frac{\boldsymbol{v}_2}{2}\cdot\frac{\hbar\boldsymbol{k}\tau}{m} = \frac{1}{2}\boldsymbol{k}\cdot\boldsymbol{v}_2\tau, \quad (7.13)$$

so that the total interferometer phase, the sum of Equations (7.12) and (7.13), is

$$\phi = \boldsymbol{k}\cdot\boldsymbol{v}\tau, \quad (7.14)$$

reproducing Equation (7.6).

7.1.2 Atomic frame

In Chapter 5 we used an analogy to the longitude problem [130] — in which nautical navigators would compare the phase of a ship's clock to the periodic variations in the sun's elevation from the horizon in order to determine their longitudinal progress — to describe the functioning of an atom interferometer in an atom's inertial frame.

In this picture, the first pulse of a Ramsey interferometer [14] synchronises the phase of the atomic superposition that acts as a local clock to that of the optical field at a particular point in space. The change in position and hence relative phase during a given measurement interval τ thus reveals the clock's velocity relative to the optical field.

This phase is measured by a second pulse that maps it onto the atomic state probabilities so that, with ideal pulses that perform this mapping exactly, the probability that a given atom is in the second atomic state $|2\rangle$ after the second pulse is

$$|c_2|^2 = \frac{1}{2}[1+\cos(kv_z\tau)], \quad (7.15)$$

where $v_z \equiv \boldsymbol{k}\cdot\boldsymbol{v}/k$ is the component of the atom's velocity vector along the optical propagation axis.

7.1.3 Atom interferometric velocimetry

For a statistical ensemble of atoms with a distribution $P(v_z)$ of velocities v_z in the beam direction, the fraction of atoms in state $|2\rangle$ after the second pulse will be

$$\mathcal{S}(\tau) = \int_{-\infty}^{\infty} P(v_z)\frac{1}{2}[1+\cos(kv_z\tau)]dv_z. \tag{7.16}$$

Each velocity class v_z hence contributes to the interferometer signal a component that varies sinusoidally with τ, with angular frequency kv_z and amplitude proportional to $P(v_z)$, akin to the contributions from different wavelengths of light to the signal produced when varying the arm length of a Michelson interferometer in Fourier transform spectroscopy [142, 143]. The velocity distribution is thus mapped onto the frequency domain of the signal but, owing to the symmetry of the cosine function, positive and negative velocities cannot be distinguished.

If the laser phase is advanced by ϕ between the two interferometer pulses, this phase is mapped onto the output signal such that

$$\mathcal{S}(\phi,\tau) = \int_{-\infty}^{\infty} P(v_z)\frac{1}{2}[1+\cos(kv_z\tau-\phi)]dv_z. \tag{7.17}$$

The absolute value of the Fourier transform of the quantity

$$\mathcal{S}_I + i\mathcal{S}_Q \equiv \mathcal{S}(0,\tau) + i\mathcal{S}(\pi/2,\tau) \tag{7.18}$$

is then proportional to the velocity distribution $P(kv_z)$, expressed as a function of the angular frequency kv_z.[2]

[2] In practice the time series data are unavoidably truncated and the time origin determines how the velocity distribution is divided between the real and imaginary parts of the transform, but the absolute value always offers a good measure.

7.2 Ramsey velocimetry experiment

Figure 7.3 shows the in-phase and quadrature interferometer traces, $\mathcal{S}_I(\tau)$ and $\mathcal{S}_Q(\tau)$, for our atom cloud. By initially adjusting the Raman detuning δ to maximise the population transferred by a π-pulse, we cancel the light shift during the interferometer pulses but incur a detuning of $2\pi \times 400\,\text{kHz}$ in between, giving the traces the form of damped oscillations. The Raman Rabi frequency Ω is around $2\pi \times 450\,\text{kHz}$.

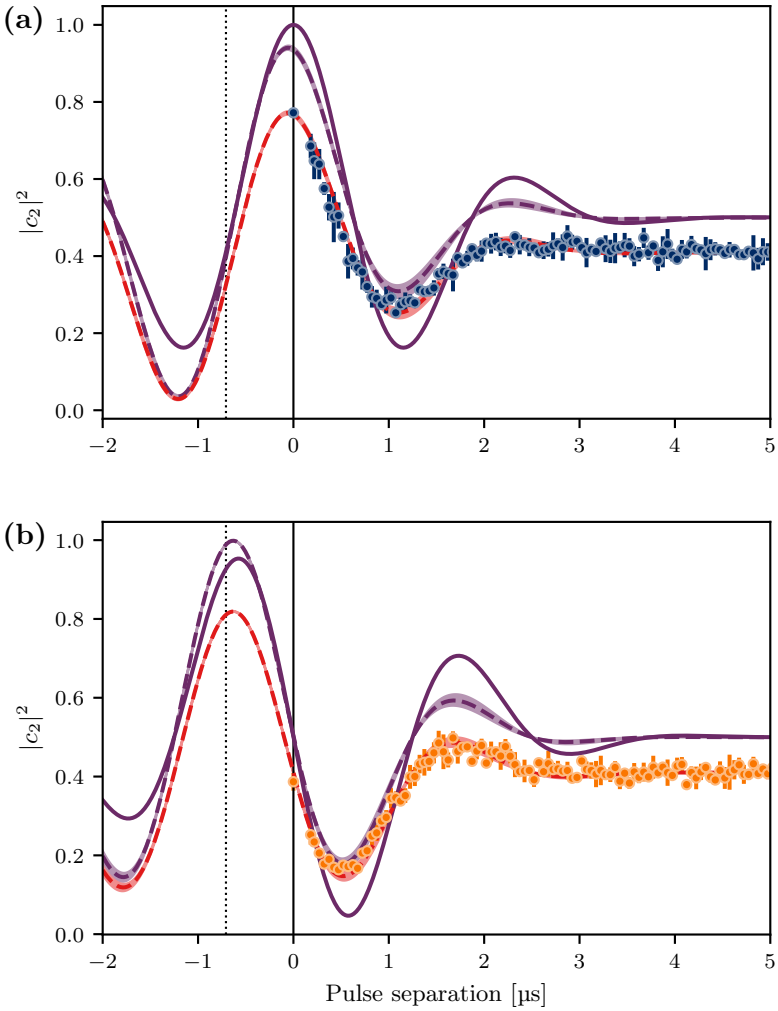

Figure 7.3: (a) In-phase $\mathcal{S}_I(\tau)$ and (b) quadrature $\mathcal{S}_Q(\tau)$ components of the interferometer signal for a Rabi frequency of $2\pi \times 450\,\text{kHz}$ and detuning of $2\pi \times 400\,\text{kHz}$ to cancel the light shift during the interferometer pulses. Circles show experimental data; solid purple curves show predictions assuming perfect beamsplitter pulses; dashed purple curves are predictions including the phase and amplitude corrections described in Section 7.3.1; dashed red curves are further scaled by a factor of 0.82, which we attribute to atom loss from the Raman beams [89]. Shaded regions about the dashed theory curves represent $\pm 10\,\%$ variations in cloud temperature. The dotted lines indicate the effective time origin $\tau = -2/\Omega$.

The distribution of velocities within our atom cloud, measured by Raman Doppler velocimetry (2.5 mW per Raman beam for 100 μs), is shown in Figure 7.4, and fits well a Gaussian distribution with a temperature of 21(1) μK superimposed upon a broad background that we have previously attributed to inhomogeneous sub-Doppler cooling [44, 144]. This velocimetry method is itself subject to inaccuracies, and at longer exposure times yields higher temperatures, perhaps for the reasons discussed in [145].

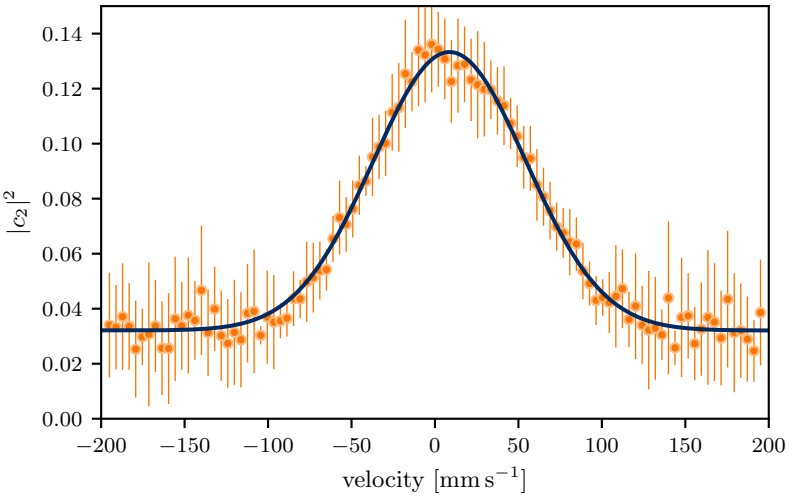

Figure 7.4: Velocity distribution measured by low intensity Raman Doppler velocimetry, orange (filled) circles, with error bars showing the standard deviation of 10 measurements. A thermal distribution with a temperature of 21(1) µK, solid (blue) fit, sits atop a broad background attributed to inhomogeneous sub-Doppler cooling [44, 144] or off-resonant excitation [145].

The solid lines in Figure 7.3 show the expected form of the interferometry fringes for this velocity distribution for infinitely short pulses. In practice the Raman pulses have a finite duration and so exhibit a residual Doppler sensitivity. This means that the amplitude and phase of the periodic fringes that each velocity class contributes to the signal in Equations (7.16) and (7.17) are modified by a velocity dependent factor that we will derive in Section 7.3.1. The dashed purple curves in Figure 7.3 include these amplitude and phase modifications to leading order. The red dashed curves incorporate an additional empirical scaling factor of 0.82 that accounts for the effect of atoms being lost from the region illuminated by the Raman beams but remaining within the readout beams [89] and brings the theoretical curves into excellent agreement with the experimental data.

The theoretical curves are extended into the hypothetical $\tau < 0$ region that has no physical meaning. The physical restriction of $\tau \geq 0$ can be expressed mathematically by multiplying Equation (7.18) by the Heaviside step function $\Theta(\tau)$ prior to taking the Fourier transform to extract the velocity distribution.

However, the velocity dependent amplitude and phase modifications due to Doppler-sensitive pulses complicates matters. While the fringe amplitude varies little across the velocity distribution, the leading order term in the phase shift

$$\beta \approx 2\left(\frac{kv_z}{\Omega}\right) \tag{7.19}$$

is linear with respect to the Doppler-shifted angular frequency kv_z. This has the effect of shifting the effective time origin for the interferometry to $\tau_0 = -2/\Omega$, indicated by the dotted vertical lines in Figure 7.3.[3] Mathematically, the result is multiplication of Equation (7.18) by an offset step function $\Theta(\tau + \tau_0)$ prior to transformation so that the output of the transform is then the convolution of the velocity distribution $P(kv_z)$ with

[3]Physically this corresponds to the phase accrued when an atom propagates a non-negligible fraction of the effective wavelength of the optical field during the pulses themselves. Some difference in phase is then inevitably observed between the first and second pulse, even when $\tau = 0$ and one immediately follows the other.

$$\tilde{\Theta}(kv_z) = \frac{i}{\sqrt{2\pi}(kv_z)} e^{ikv_z\tau_0} + \sqrt{\frac{\pi}{2}}\delta(kv_z), \tag{7.20}$$

the transform of the offset step function, where $\delta(kv_z)$ is the Dirac delta function. When the spread of the velocity distribution is small with respect to $1/\tau_0$, the first term remains approximately orthogonal to the second during the convolution so that the absolute value of the transform is broadened but the velocity distribution remains unaltered in the real part.

However, when the Doppler spread of the velocity distribution is a considerable fraction of the Rabi frequency, the velocity dependent phase factor $e^{ikv_z\tau_0}$ in the first term becomes significant so that it broadens both the real and imaginary parts of the transform and the velocity distribution is irretrievable from a single measurement.

The absolute value of the FFT of the data from Figure 7.3 after baseline subtraction is shown in Figure 7.5 in units of atomic velocity after accounting for the $2\pi \times 400\,\text{kHz}$ detuning of the Raman frequency from the known hyperfine frequency that cancels the light shift during the interferometry pulses. If the underlying velocity distribution is assumed to have the Gaussian form associated with a 1-D Maxwell-Boltzmann distribution, that is

$$P(kv_z) \propto \exp\left(-\frac{1}{2}\frac{(kv_z - kv_0)^2}{k^2{\sigma_v}^2}\right) \tag{7.21}$$

[4]The atomic wavepacket is separated into two states, one receiving a recoil kick and one not. The mean impulse imparted by the interferometer thus offsets the centre of the distribution by half of the recoil velocity.

where v_0 is the non-zero centre of the distribution that will turn out to be half of the two-photon recoil velocity[4] and $\sigma_v = \sqrt{k_B T/m}$ represents the thermal distribution of velocities for atoms of mass m and temperature T, then an analytical form for the FFT data can be determined.

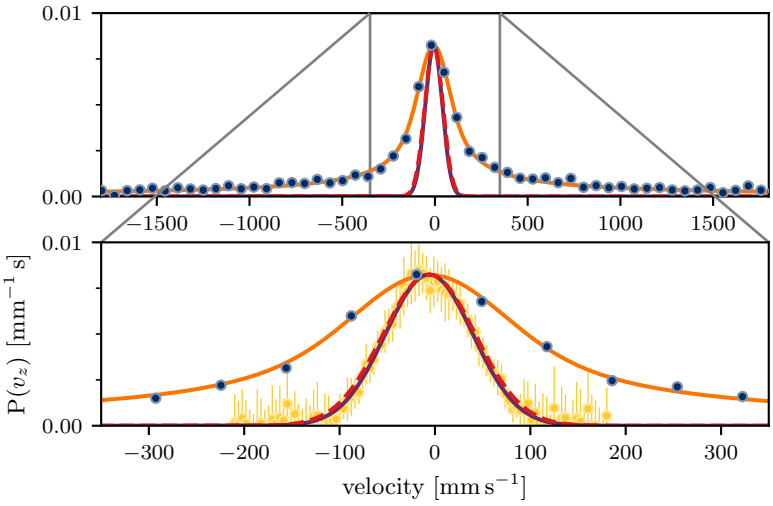

Figure 7.5: The absolute value of the fast Fourier transform (FFT) of interferometric measurements (blue circles), corrected for the $2\pi\times400\,\mathrm{kHz}$ detuning from the known hyperfine frequency, is a version of the underlying atomic velocity distribution that is broadened due to the Doppler sensitivity inherent in the finite duration interferometer pulses. If the velocity distribution is assumed to be Gaussian, an analytical function (described in the main text) can be fitted (solid orange line) from which the assumed Gaussian distribution can be determined (dashed red line). The established temperature of this distribution is 24(2) µK, in agreement with the 21(1) µK Gaussian distribution measured by Raman Doppler spectroscopy that, for comparison, is represented here (light yellow circles and solid purple Gaussian fit) with the baseline subtracted and centred on the interferometrically derived velocity distribution that is offset from $v_z = 0$ due to the recoil associated with interacting with the interferometry beams.

Convolution of the Gaussian distribution with Equation (7.20) is equivalent to multiplying it by a complementary error function, yielding a signal of the form

$$\tilde{\mathcal{S}}(kv_z) \propto P(kv_z)\,\mathrm{erfc}\left(-\frac{1}{\sqrt{2}}\left[k\sigma_v\tau_0 + i\frac{kv_z - kv_0}{k\sigma_v}\right]\right). \qquad (7.22)$$

The solid orange line in Figure 7.5 shows a fit of the absolute value of this function to the transformed data, with σ_v, v_0 and the overall amplitude as free parameters and $\tau_0 = -2/(2\pi\times450)\,\mathrm{ms}$ set by the empirically determined Rabi frequency. The dashed red line is the Gaussian velocity distribution determined from the fitting parameters, with an offset of $v_0 = -6(2)\,\mathrm{mm\,s^{-1}}$ consistent with half of the two-photon recoil velocity $v_\mathrm{R} = 12\,\mathrm{mm\,s^{-1}}$ and a velocity spread σ_v corresponding to a cloud temperature of 24(2) µK. This is in agreement with the 21(1) µK measurement made by Raman Doppler spectroscopy that is represented in Figure 7.5 by a solid

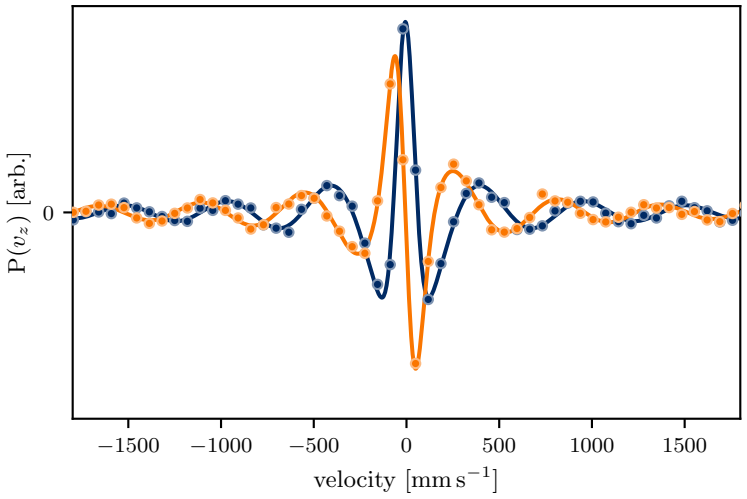

Figure 7.6: Real (blue circles) and imaginary (orange circles) parts of the FFT of data from Figure 7.3 after baseline subtraction and padding with zeroes for negative times up to the effective time origin $\tau = \tau_0$. The data show good agreement with the real and imaginary parts of the fitted Equation (7.22), represented as solid lines.

purple Gaussian centred upon the interferometric measurement for comparison.

It is fun, though not particularly informative, to pad out the data from Figure 7.3 with zero-valued data points for negative times up to the effective time origin of $\tau = \tau_0$. This then introduces oscillations between the real and imaginary parts of the transform, periodic in v_z with a period of $2\pi/k\tau_0$. The real and imaginary parts of the transform then fit well to the real and imaginary parts of the fitted Equation (7.22), as shown in Figure 7.6. As the periodicity of the oscillations in both the fitted function and the data are determined by the fixed parameter τ_0, this tells us little but looks pretty. It does suggest the possibility of making τ_0 a free parameter and using it to determine the Rabi frequency.

While this demonstrates a practical method of determining the atomic cloud temperature, several factors contribute to the large uncertainty in the measurement. In our experiment, the delay between when the AOM shutter is (de)activated and when the optical field is turned on (off) at the atoms differs for each of the counter-propagating Raman beams. This means that, between two sequential pulses, there are short periods where only one beam is activated, contributing a light-shift, whilst a 'pulse' is considered to be active only while both beams illuminate the atoms. This means that, when the AOM is turned off for a given duration, the actual duration between pulses is larger by a factor of the delay $\tau_{\text{delay}} = 75(10)$ ns between the response of one field relative to the

other.[5] This is illustrated in Figure 7.7 and accounted for in the data presented in Figure 7.3 to the tens of ns level, but the ±10 ns uncertainty in the spacing between the data point at $\tau = 0$ ns and the rest of the data contributes to the uncertainty in the final temperature distribution while the extra light-shift when only one beam is on is unaccounted for and contributes to the uncertainty in v_0.

Additionally, the function in Equation (7.22) that is fitted to the data tends to zero in the wings upon the multiplication of a very large number by a very small number. This becomes difficult to handle computationally so that only those data within a range about the centre of the distribution, where the function can be numerically evaluated, can be fitted to. This range is approximately equivalent to that shown in the top axes of Figure 7.5, while the data extend several times further.

These issues, along with many of the complications attributed to the intra-pulse Doppler sensitivity can be overcome by employing a three-pulse interferometer to make the same measurement.

7.3 Practical velocimetry with a three-pulse interferometer[6]

When $\tau_0 = 0$, depicted in Figure 7.8(a) and equivalent to having instantaneous interferometer pulses of infinite power, the first term in Equation (7.20) remains orthogonal to the second for all atomic velocities and a Ramsey interferometer can be used to measure the true atomic velocity distribution. While we have demonstrated that, with a good understanding of the intra-pulse Doppler sensitivity and its effect on the Fourier transformed signal, this technique can still be powerful in a realistic situation with pulses of finite duration like those in Figure 7.8(b), in practice this requires some initial assumptions about the underlying velocity distribution.

These issues can be overcome by interleaving a 'mirror' π-pulse between the 'beamsplitters' of our interferometer, as shown in Figure 7.8(c). In its time-symmetrical form ($T_2 = T_1$), atoms divide their time equally between the two interferometer states, and the interferometer forms a basic *composite pulse* [56] in which systematic contributions to the phase accrued during the first evolution period T_1 are reversed during the second period T_2.

[5] This was likely due to a difference in alignment between the beams as they passed through the shuttering AOM that was subsequently reduced with a corresponding reduction of τ_{delay} in later experiments.

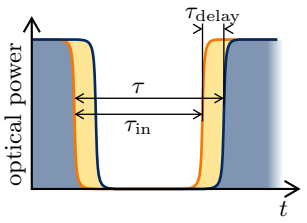

Figure 7.7: Illustration (not to scale) of how a delay $\tau_{\text{delay}} = 75(10)$ ns in the arrival time between each of the Raman beams at the atoms introduces an uncertainty in the Ramsey velocimetry measurement. Blue (dark) shaded regions indicate periods during which both beams are illuminated and a pulse in considered 'on'. Yellow (light) shaded regions indicate periods during which only one beam is active, inducing a light-shift but not driving the two-photon transition. This introduces an uncertainty in the superposition phase acquired in between pulses, affecting the measurement of v_0. The actual time between pulses τ is longer than the programmed time τ_{in} by an extra τ_{delay}, introducing an uncertainty in the separation between the data point at $\tau = 0$ and the subsequent points that affects precisely the temperature of the atoms can be determined.

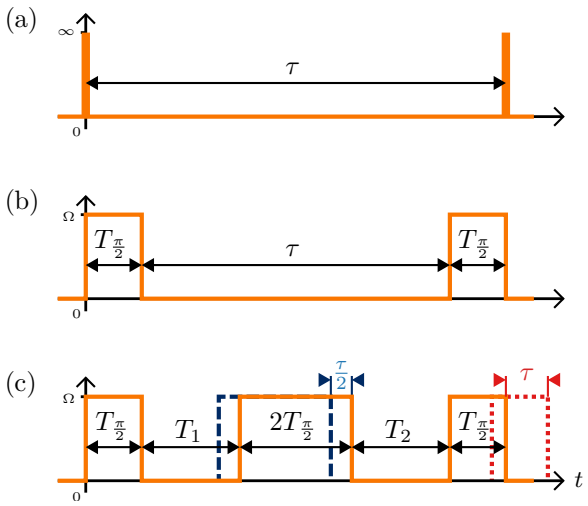

Figure 7.8: Temporal pulse profiles for the interferometer sequences considered here. a) Ramsey sequence with ideal, zero-length, pulses. b) Ramsey sequence with realistic, finite-length, pulses. c) Asymmetric Mach-Zehnder sequence with realistic pulses.

The absolute velocity-dependent phase shifts cancel, but those due to changes in velocity remain; this *Mach-Zehnder* arrangement thus forms the basis of atom interferometric inertial sensors [22, 125, 146].

When the interferometer is asymmetric, however, we retain the velocity sensitivity according to the temporal asymmetry $\tau = T_2 - T_1$, which can be varied continuously over both negative and positive values, while taking advantage of partial cancellation of phase shifts accrued during the pulses themselves [146].

7.3.1 Theoretical interferometer output

For a quantitative understanding of this we derive analytical expressions for the output of the physically realisable interferometers in Figure 7.8, calculating the excited state probability $|c_2|^2$ at the end of each sequence for an atom starting in the ground state $c_1 = 1$.

With a suitable shift of energy reference, the effective two-level Hamiltonian for the Raman system given in Equation (2.74) can be written as

$$\hat{H} = \frac{\hbar}{2}\begin{pmatrix} \delta_L & \Omega_R e^{-i\phi_L} \\ \Omega_R e^{i\phi_L} & -\delta_L \end{pmatrix}, \tag{7.23}$$

where Ω_R is the (real) Rabi frequency and ϕ_L is the effective laser phase defined in Equations (2.76) and (2.73) respectively. We can separate the laser detuning δ_L into three terms

$$\delta_L = \delta_{\text{laser}} - \delta_{\text{ac}} - \delta_{\text{Doppler}}. \tag{7.24}$$

Here $\delta_{\text{Doppler}} = kv_z$ is ultimately the quantity of interest, corresponding to the Doppler shift of an atom whose average velocity along the beam axis is v_z. The effect of the light shift on the resonance condition is encapsulated in δ_{ac} — that is thus only non-zero when the lasers are on during the pulses — while δ_{laser} incorporates all of the other detuning terms that are assumed to be constant.

Further defining $\delta_0 \equiv \delta_{\text{laser}} - \delta_{\text{Doppler}}$, the propagator that acts on a state $\binom{c_1}{c_2}$ at a time t in order to give the state at a later time $t + \Delta t$ in the absence of any interaction with the lasers is then

$$\hat{\mathcal{U}}(\Delta t) = \begin{pmatrix} e^{i\delta_0 \cdot \Delta t/2} & 0 \\ 0 & e^{-i\delta_0 \cdot \Delta t/2} \end{pmatrix}. \tag{7.25}$$

During the pulses the evolution is governed by the OBEs and, assuming completely coherent evolution, the appropriate propagator can be written as

$$\hat{\Omega}(\Delta t, \phi_L) = \begin{pmatrix} C^*(\Delta t) & -iS^*(\Delta t, \phi_L) \\ -iS(\Delta t, \phi_L) & C(\Delta t) \end{pmatrix}, \tag{7.26}$$

where we have employed the formalism of Stoner *et al.* [140] in which

$$C(\Delta t) = \cos\left(\frac{\Delta t}{2}\sqrt{\delta_L^2 + \Omega_R^2}\right) + i\frac{\delta_L}{\sqrt{\delta_L^2 + \Omega_R^2}} \sin\left(\frac{\Delta t}{2}\sqrt{\delta_L^2 + \Omega_R^2}\right), \tag{7.27a}$$

and

$$S(\Delta t, \phi_L) = \frac{\Omega_R e^{i\phi_L}}{\sqrt{\delta_L^2 + \Omega_R^2}} \sin\left(\frac{\Delta t}{2}\sqrt{\delta_L^2 + \Omega_R^2}\right). \tag{7.27b}$$

We denote the nth pulse with subscripts $\hat{\Omega}_n, C_n, S_n$ etc. so that the output of a 3-pulse interferometer with pulse separations $T_{1,2}$ is

$$\begin{aligned}
|c_2|^2 &= \left|\left(\begin{smallmatrix}0&0\\0&1\end{smallmatrix}\right)\cdot\hat{\Omega}_3\cdot\hat{\mathcal{U}}\left(T_2\right)\cdot\hat{\Omega}_2\cdot\hat{\mathcal{U}}\left(T_1\right)\cdot\hat{\Omega}_1\cdot\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)\right|^2, \\
&= |S_1|^2|S_2|^2|S_3|^2 + |C_1|^2|S_2|^2|C_3|^2 \\
&\quad - 2\,\mathrm{Re}\left[e^{i\delta_0(T_2-T_1)}C_1S_1\left(S_2^*\right)^2C_3^*S_3\right] \\
&\quad + |S_1|^2|C_2|^2|C_3|^2 + |C_1|^2|C_2|^2|S_3|^2 \\
&\quad + 2\,\mathrm{Re}\left[e^{i\delta_0(T_1+T_2)}C_1^*S_1^*\left(C_2^*\right)^2C_3^*S_3\right] \\
&\quad + 2\left(|C_3|^2-|S_3|^2\right)\mathrm{Re}\left[e^{-i\delta_0\cdot T_1}C_1S_1C_2S_2^*\right] \\
&\quad + 2\left(|C_1|^2-|S_1|^2\right)\mathrm{Re}\left[e^{i\delta_0\cdot T_2}C_2^*S_2^*C_3^*S_3\right].
\end{aligned} \tag{7.28}$$

To proceed we note that, for the Mach-Zehnder interferometers under consideration,

$$C_3 = C_1, \tag{7.29a}$$
$$S_3 = e^{i\phi}S_1, \tag{7.29b}$$
$$\begin{aligned}\arg(S_1) &= \arg(S_2) \\ \Rightarrow S_1^*S_2 &= S_1S_2^* = |S_1||S_2|,\end{aligned} \tag{7.29c}$$

where ϕ is an advance in the laser phase introduced prior to the final pulse. This allows us to write Equation (7.28) as

$$\begin{aligned}
|c_2|^2 &= |S_1|^4|S_2|^2 + |C_1|^4|S_2|^2 + 2|S_1|^2|C_2|^2|C_1|^2 \\
&\quad - 2|C_1|^2|S_1|^2|S_2|^2\cos[\delta_0(T_2-T_1)+\phi] \\
&\quad + 2|S_1|^2\,\mathrm{Re}[C_1^2C_2^2]\cos[\delta_0(T_1+T_2)+\phi] \\
&\quad + 2|S_1|^2\,\mathrm{Im}[C_1^2C_2^2]\sin[\delta_0(T_1+T_2)+\phi] \\
&\quad + 2(|C_1|^2-|S_1|^2)|S_1||S_2|\,\mathrm{Re}[C_1C_2] \\
&\qquad\qquad \times[\cos(\delta_0\cdot T_1)+\cos(\delta_0\cdot T_2+\phi)] \\
&\quad + 2(|C_1|^2-|S_1|^2)|S_1||S_2|\,\mathrm{Im}[C_1C_2] \\
&\qquad\qquad \times[\sin(\delta_0\cdot T_1)+\sin(\delta_0\cdot T_2+\phi)].
\end{aligned} \tag{7.30}$$

For velocimetry purposes, we want to introduce an asymmetry $\tau = T_2 - T_1$ into the interferometer sequence. We will first consider the interferometer output when this is introduced by keeping

$T_1 = T$ constant and varying $T_2 = T + \tau$. This is the case illustrated by the red (dotted) line in Figure 7.8(c), and can also be used to model the Ramsey interferometer shown in 7.8(b) by setting $T = 0$ and $C_2 = 1, S_2 = 0$.

By expanding the trigonometric functions of Equation (7.30), we can isolate the τ dependence into terms of $\sin(\delta_0\tau + \phi)$ and $\cos(\delta_0\tau + \phi)$ such that the interferometer output can be written in the form

$$|c_2|^2 = \frac{1}{2}\left\{\mathcal{C}' + \mathcal{A}'\cos\left(\delta_0\tau + \phi\right) + \mathcal{B}'\sin\left(\delta_0\tau + \phi\right)\right\} \tag{7.31}$$

or, equivalently,

$$|c_2|^2 = \frac{1}{2}\left\{\mathcal{C}' + \alpha\cos\left(\delta_0\tau + \phi + \beta\right)\right\}, \tag{7.32}$$

where $\alpha = \sqrt{(\mathcal{A}')^2 + (\mathcal{B}')^2}$ and $\beta = \arctan\left(-\mathcal{B}'/\mathcal{A}'\right)$. When $\mathcal{C}' = \alpha = 1$ and $\beta = 0$ then this closely resembles the integrand in Equation (7.16) and the atomic velocity distribution is well mapped onto the frequency domain with an offset given by δ_{laser}. As discussed in Section 7.2, provided $|\beta| \ll 1$ then the velocity distribution can be retrieved from the real part of the FFT, even in the Ramsey configuration where time domain data can only be collected for $\tau \geq 0$.

In practice, explicitly evaluating the form of the coefficients $\mathcal{A}'$, $\mathcal{B}'$ and $\mathcal{C}'$, we find

$$\begin{aligned}
\mathcal{A}' =& -a + b\cos(2\delta_0 T) + c\sin(2\delta_0 T) \\
&+ d\cos(\delta_0 T) + f\cos(\delta_0 T), && \text{(7.33a)} \\
\mathcal{B}' =& -b\cos(2\delta_0 T) + c\sin(2\delta_0 T) \\
&- d\cos(\delta_0 T) + f\cos(\delta_0 T), && \text{(7.33b)} \\
\mathcal{C}' =& 2|S_1|^4|S_2|^2 + 2|C_1|^4|S_2|^2 + 4|S_1|^2|C_2|^2|C_1|^2 \\
&+ d\sin(\delta_0 T) + f\sin(\delta_0 T), && \text{(7.33c)}
\end{aligned}$$

using

$$
\begin{aligned}
a =& 4|C_1|^2|S_2|^2|S_1|^2, \\
b =& 4|S_1|^2\,\mathrm{Re}(C_1^2C_2^2), \\
c =& 4|S_1|^2\,\mathrm{Im}(C_1^2C_2^2), \\
d =& 4(|C_1|^2-|S_1|^2)|S_1||S_2|\,\mathrm{Re}(C_1C_2), \\
f =& 4(|C_1|^2-|S_1|^2)|S_1||S_2|\,\mathrm{Im}(C_1C_2).
\end{aligned}
\tag{7.34}
$$

For $T \gg 0$, beating between the terms oscillating at T and $2T$ results in rapid oscillations of $\mathcal{A}'$ and $\mathcal{B}'$ in frequency space that restrict the usefulness of the interferometer to a very narrow velocity range.

The contrast α and phase shift β of the Ramsey interferometer are plotted in Figure 7.9(a) as a function of $\delta_{\mathrm{Doppler}} = kv_z$ in units of the on-resonance Rabi frequency Ω_R for the case where

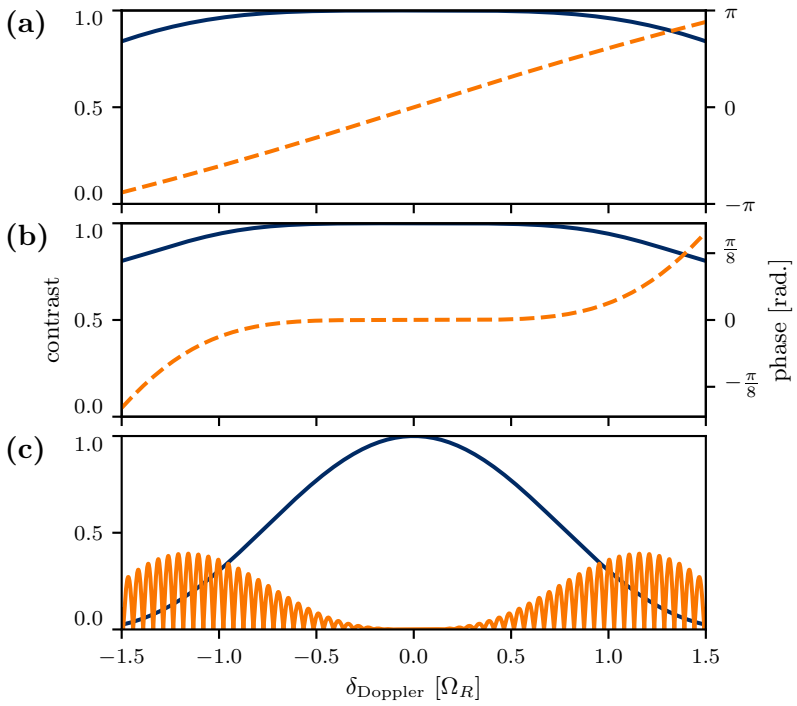

Figure 7.9: Contrast and phase shifts to frequency components in interferometer outputs as a function of Doppler detuning $\delta_{\mathrm{Doppler}}$, in units of the on-resonance Rabi frequency Ω_R. In these plots $\delta_{\mathrm{ac}} = \delta_{\mathrm{laser}}$; when this is not the case, the profile is shifted along the x axis by the difference. Contrast α and phase shift β are plotted as solid and dashed lines respectively for (a) a Ramsey interferometer and (b) a $T_1 = T, T_2 = T+\tau$ Mach-Zehnder interferometer with $T = \pi/2\Omega_R$. In plot (c) , the contrast of the harmonic $|\mathcal{A}|$ and subharmonic $|\mathcal{B}|$ terms are shown as solid blue (dark) and orange (light) lines respectively for a $T_{1,2} = T \mp \tau$ Mach-Zehnder interferometer with $T = 20\pi/\Omega_R$. For larger T, the oscillations in the contrast of the subharmonics increases but the envelope remains constant.

$\delta_{\text{laser}} = \delta_{\text{ac}}$ is tuned to be on-resonant during the pulses. The contrast is approximately flat about $\delta_{\text{Doppler}} = 0$, but the phase shift is dominated by a linear term $\beta \approx {}^{2\delta_{\text{Doppler}}}/_{\Omega_R}$ that, as discussed in Section 7.2, acts to shift the effective time origin of the interferometer. Physically, when $\tau = 0$, a velocity-dependent phase shift is still apparent due to the finite duration of the pulses. This means that the true velocity distribution cannot be extracted from the FFT given the physical restriction $\tau \geq 0$.

In Figure 7.9(b), the same parameters are plotted for a $T_1 = T, T_2 = T + \tau$ Mach-Zehnder configuration with $T = {}^{\pi}/_{2\Omega_R}$. The additional pulse and short periods of free evolution act to undo the first-order phase shift accrued during the pulses, flattening out the phase profile about $\delta_{\text{Doppler}} = 0$. The negative range of $\tau \geq -T$ is still severely limited in such an interferometer but the negligible phase shift means that the velocity distribution can be retrieved from the real part of the FFT when data are collected for $\tau \geq 0$. However, like the Ramsey interferometer, this configuration is similarly sensitive to the absolute timing of the pulses discussed in Section 7.2.

The phase shift is eliminated entirely if the total free evolution time $T_1 + T_2$ of the interferometer is kept constant and the asymmetry is instead introduced by setting $T_{1,2} = T \mp \tau/2$, as shown by the blue (dashed) line in Figure 7.8(c). Noting that

$$\sin\left[\delta_0\left(\frac{\tau}{2}+T\right)+\phi\right]+\sin\left[\delta_0\left(T-\frac{\tau}{2}\right)\right]= 2\sin\left(\delta_0\cdot T+\frac{\phi}{2}\right)\cos\left(\delta_0\frac{\tau}{2}+\frac{\phi}{2}\right), \text{ and} \tag{7.35a}$$

$$\cos\left[\delta_0\left(\frac{\tau}{2}+T\right)+\phi\right]+\cos\left[\delta_0\left(T-\frac{\tau}{2}\right)\right]= 2\cos\left(\delta_0\cdot T+\frac{\phi}{2}\right)\cos\left(\delta_0\frac{\tau}{2}+\frac{\phi}{2}\right), \tag{7.35b}$$

the output in this case can be shown to be

$$|c_2|^2 = \frac{1}{2}\Big\{\mathcal{C} - \mathcal{A}\cos\left(\delta_0\tau+\phi\right) + \mathcal{B}\cos\left[\frac{1}{2}\left(\delta_0\tau+\phi\right)\right]\Big\}, \tag{7.36}$$

with

$$\mathcal{A} = 4|C_1|^2|S_1|^2|S_2|^2, \tag{7.37a}$$

$$\begin{aligned}\mathcal{B} = 8(|C_1|^2 - |S_1|^2)|S_1||S_2| \\ \times \big[\mathrm{Re}(C_1C_2)\cos(\delta_0 \cdot T + \phi/2) \\ + \mathrm{Im}(C_1C_2)\sin(\delta_0 \cdot T + \phi/2)\big],\end{aligned} \tag{7.37b}$$

$$\begin{aligned}\mathcal{C} = 4|C_1|^2|C_2|^2|S_1|^2 + 2|C_1|^4|S_2|^2 + 2|S_1|^4|S_2|^2 \\ + 4|S_1|^2\big[\mathrm{Re}(C_1^2C_2^2)\cos(\delta_0 \cdot T + \phi) \\ + \mathrm{Im}(C_1^2C_2^2)\sin(\delta_0 \cdot T + \phi)\big].\end{aligned} \tag{7.37c}$$

The output has sinusoidal components in τ with velocity-dependent angular frequency δ_0, and amplitude $\mathcal{A}(\Omega_R, \delta_L)$, purely constructed from elements of the pulse matrices and thus only dependent on the offset from the light-shifted resonance (with the overall scale determined by the on-resonance Rabi frequency). There is no detuning-dependent phase shift to these harmonic components, though parasitic subharmonics of amplitude $\mathcal{B}(\Omega_R, \delta_0, \delta_L, \phi, T)$ become significant at large detunings.

With $\mathcal{A} = \mathcal{C} = 1$, $\mathcal{B} = 0$, and $\delta_{\mathrm{laser}} = 0$, Equation (7.36) again resembles the analogous Ramsey output in the integrand of Equation (7.17), albeit with an inversion arising from the additional rotation by π. As long as $\mathcal{A} \gg \mathcal{B}$ then, as in Equation (7.17), the atomic velocity distribution is well mapped onto the frequency domain and the scaling by $\mathcal{A}$ can be corrected for by multiplying through by its reciprocal.

The magnitudes of $\mathcal{A}$ and $\mathcal{B}$ are plotted as functions of $\delta_{\mathrm{Doppler}}/\Omega_R$ in Figure 7.9(c) which shows that this criterion is satisfied for $-0.4 < \delta_{\mathrm{Doppler}}/\Omega_R < 0.4$, where $|\mathcal{B}/\mathcal{A}| < 0.1$. Doppler profiles falling within this window will thus incur little distortion from subharmonics, which would act to artificially narrow broader distributions. As the interferometer time is made longer by increasing T, the oscillations in $\mathcal{B}$ become more rapid but remain within the same envelope.

It should be noted that the velocity distribution will be centred about δ_{laser} in the frequency domain, so the ambiguity of the DC component $\mathcal{C}(\Omega_R, \delta_0, \delta_L, \phi, T)$ can be negated by setting δ_{laser} much larger than the width of the Doppler profile and subtracting it off in analysis. When δ_{laser} is measured relative to the 'bare' hyperfine splitting ω_0, subtracting it off leaves the velocity distribution

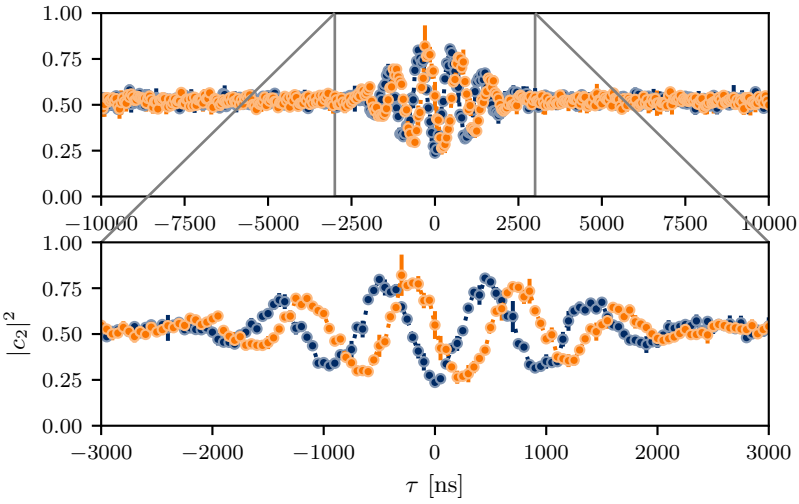

Figure 7.10: Measurements of fractional population transfer $|c_2|^2$ as a function of temporal asymmetry τ for interferometers with (light orange circles) and without (dark blue circles) a $\pi/2$ phase shift before the final recombination pulse. Data are taken in a pseudo-random order at a rate of 2 Hz to facilitate MOT loading. Each point is an average of two measurements, with error bars representing the standard deviation. A detuning of $\delta_{\text{laser}} = -2\pi \times 1050\,\text{kHz}$ from the hyperfine splitting results in oscillations which appear within an envelope whose shape is governed by the velocity distribution.

centred about the average velocity v — depicted in Figure 7.2 and to which we previously alluded — corresponding to half the two-photon recoil velocity.

7.3.2 Experimental results

On balance, the final interferometer geometry considered in the previous section, in which the total interferometer duration $T = T_1 + T_2$ is kept constant and the asymmetry τ is varied between $-T$ and T, shows the most promise for our setup. We can realise Rabi frequencies that are suitably large compared to the Doppler width of our atomic sample so that the amplitude of subharmonics that would act to enhance the apparent probability of lower velocities is negligible and the residual velocity-dependent modification of the fringe amplitude $\mathcal{A}$ can be corrected for in subsequent analysis.

Figure 7.10 shows typical output from the interferometer, with a detuning $\delta_{\text{laser}} = -2\pi \times 1050\,\text{kHz}$ introduced from the two-photon Raman resonance to offset the AC Stark shift during the pulses. The in-phase and quadrature fringes correspond to $-\mathcal{S}_I$ and $\mathcal{S}_Q$ from Equation (7.18), with the inversion of the in-phase component

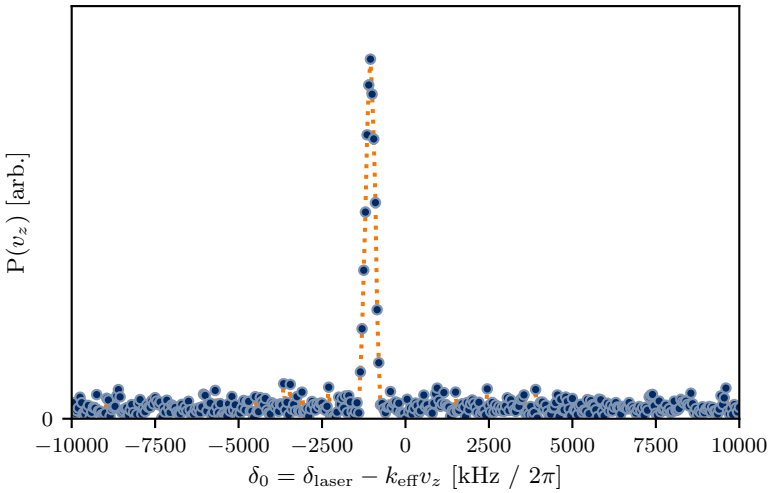

Figure 7.11: Absolute value of raw FFT of data in Figure 7.10 after baseline subtraction, prior to subtracting the $\delta_{\text{laser}} = -2\pi \times 1050\,\text{kHz}$ laser detuning from the hyperfine splitting, which manifests as a $-410\,\text{mm}\,\text{s}^{-1}$ shift to the centre of the velocity distribution. The quadrature measurement is able to resolve the sign of the displacement, so there is no component with an opposite shift.

arising from the additional rotation by π. The FFT of these data is shown, as a function of frequency, in Figure 7.11.

The velocity distribution can be directly extracted from these data, as shown in Figure 7.12(a) where they are reproduced as a function of velocity with the two-photon detuning δ_{laser} subtracted. Here they are overlaid upon a measurement made by conventional Raman Doppler spectroscopy. Fitted Gaussian profiles, with temperatures of 31.0(15) µK and 33.6(15) µK respectively, are shown.

The FFT profile, shown in Figure 7.12(a), requires a small correction to account for a slight dependence of the fringe amplitude $\mathcal{A}(k_{\text{eff}} v_z)$ upon atomic velocity. Figure 7.12(b) shows this velocity dependence (dotted line), together with the corrected velocity distribution which yields the same temperature as the Doppler measurements, with notably enhanced signal-to-noise ratio. The profile in this instance is slightly displaced to account for a difference of $2\pi \times 40\,\text{kHz}$ between the assumed laser detuning δ_{laser} and the AC Stark shift recorded for this experiment.

The Gaussian fit to the corrected data is centred at $v_z = -4.5(13)\,\text{mm}\,\text{s}^{-1}$. This is in agreement with a second measurement, shown in Figure 7.14, centred at $v_z = -5.1(7)\,\text{mm}\,\text{s}^{-1}$. The two-photon recoil velocity for ^{85}Rb is $12\,\text{mm}\,\text{s}^{-1}$ and there could be an impulse imparted to the cloud as the magnetic field is terminated.

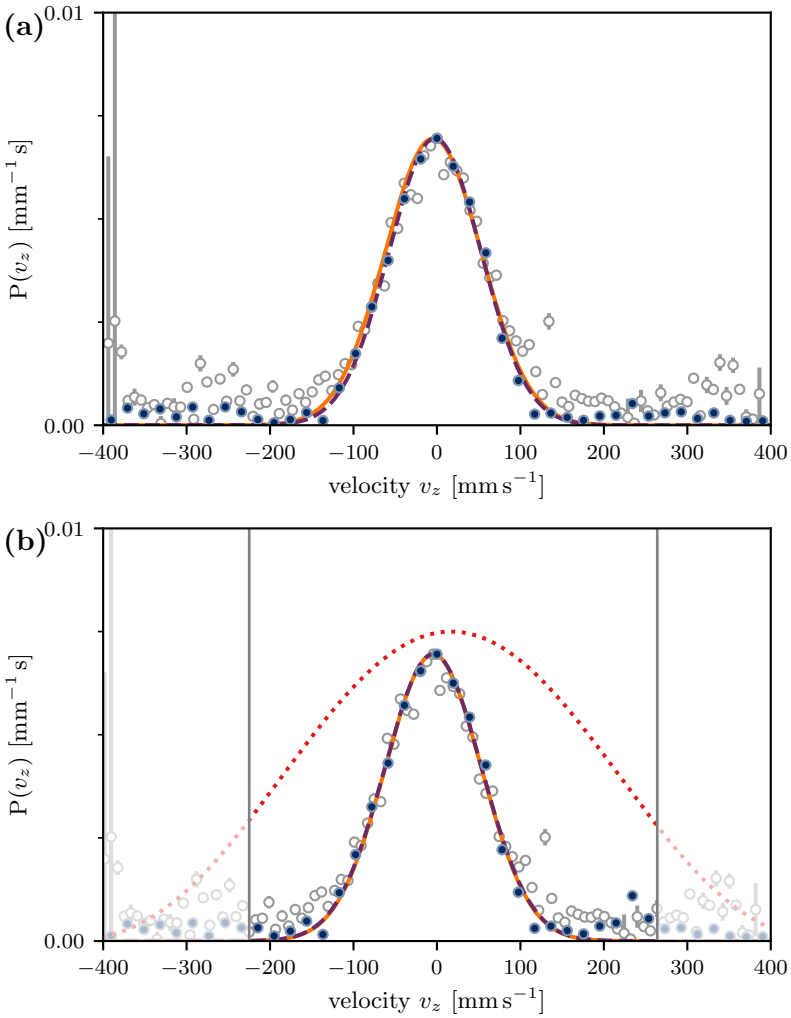

Figure 7.12: Absolute value of the FFT of data in Figure 7.10, filled (blue) circles, overlaid upon a velocity profile determined by Raman Doppler spectroscopy, empty (grey) circles, in units of velocity. Solid (orange) lines are a Gaussian fit to the spectroscopic data, with a temperature of 33.6(15) µK, and dashed (purple) lines show Gaussian fits to the interferometric data. Plot (a) uses raw FFT data and the Gaussian fit is narrower than the spectroscopic measurement with a temperature of 31.0(15) µK. Plot (b) has a correction factor applied, multiplying each point by the reciprocal of the theoretical amplitude $\mathcal{A}(k_{\text{eff}} v_z)$ based on our experimental parameters, overlaid as a dashed (red) line. Only points in the $\mathcal{A} > 0.4$ range (highlighted) were corrected to avoid amplifying noise at the extremities, bringing the fitted temperature to 34.1(16) µK in agreement with the fit to the spectroscopic data. The AC Stark shift-induced offset from the spectroscopic data has been subtracted, centring it on the interferometric data at $v_z = -4.5\,\text{mm}\,\text{s}^{-1}$ to better compare their shapes.

The velocity resolution of the FFT is determined by the range of τ, limited in principle to the time it takes atoms to leave the interaction region. Our data, in the range $|\tau| < 10\,\mu s$ (spanning $\Delta\tau = 20\,\mu s$), give a velocity resolution of $\delta v = 1/(k_{\text{eff}}\Delta\tau) \approx 20\,\text{mm}\,\text{s}^{-1}$ that is approximately twice the recoil velocity. The value $\delta\tau$ by which τ is incremented between adjacent data points, in this case $\delta\tau = 50\,\text{ns}$, determines the range of velocities that can be measured $\Delta v = 1/(k_{\text{eff}}\delta t)$, although oversampling reduces the sensitivity of the measurement to the noise on any individual data point.

7.3.3 Discussion

The velocity distribution determined from our interferometric measurements agrees well for a warm sample with that obtained by conventional Doppler spectroscopy using low power Raman pulses (single-photon detuning $\sim 2\pi \times 15\,\text{GHz}$, Rabi frequency $\Omega_R \approx 2\pi \times 25\,\text{kHz}$). Whereas the interferometric measurements are well represented by a single Gaussian, however, the Doppler measurements show an additional, broader, component. This is particularly noticeable when a cooler sample, such as that shown in Figure 7.13, is probed with a strong Raman field (single-photon detuning $\sim 2\pi\times7\,\text{GHz}$, Rabi frequency $\Omega_R \approx 2\pi\times50\,\text{kHz}$), yielding

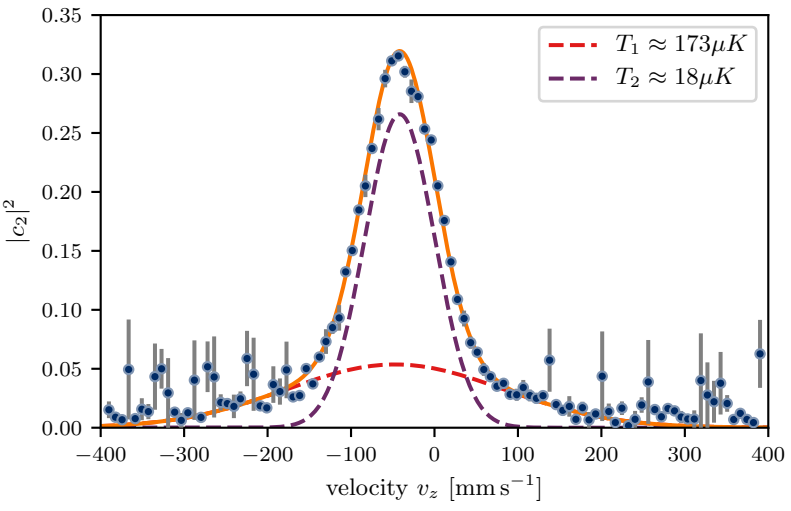

Figure 7.13: Doppler spectroscopy velocity profile taken with Raman π-pulses with a single-photon detuning $\sim 2\pi\times7\,\text{GHz}$ and a Rabi frequency $\Omega_R \approx 2\pi \times 50\,\text{kHz}$. A single Gaussian does not make a good fit, but a two Gaussian distribution fits well. The solid (orange) line is the sum of the two dashed Gaussians, a cold central peak and a broader background which we have previously attributed to inhomogeneous sub-Doppler cooling. The resonance is AC Stark shifted, resulting in the large offset of the distribution from $v_z = 0$.

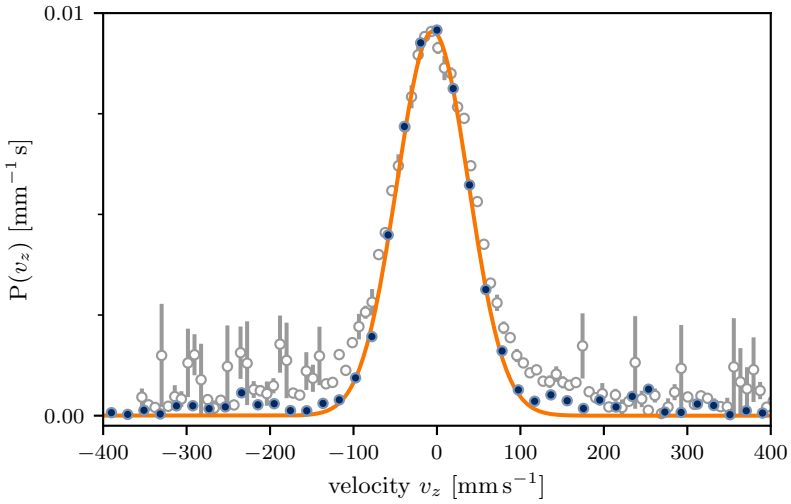

Figure 7.14: Interferometric velocimetry measurement, filled circles (blue), overlaid upon the Doppler spectroscopy profile from Figure 7.13, empty circles (grey). The measurements were taken under the same conditions, with the offset subtracted from the spectroscopic data to centre them on the interferometric profile at $v_z = -5.1\,\mathrm{mm\,s^{-1}}$. The solid (orange) line shows a Gaussian fit to the interferometric data, with a temperature of 18.7(6) µK corresponding closely to the 17.8(9) µK colder Gaussian fitted in Figure 7.13. The signal-to-noise ratio of the interferometric measurement is good, and does not show signs of the broad background evident in the spectroscopic measurement.

a central Gaussian distribution with a temperature of 17.8(9) µK superimposed upon a much broader background.

Interferometric measurement under the same conditions, shown in Figure 7.14, does not display this broad component, but modelling suggests that this is not a limitation of the interferometric technique. We have previously attributed the broad background to inhomogeneous sub-Doppler cooling [144]; such a distribution might also result if the Doppler technique detected warmer, untrapped atoms outside the region interrogated by the interferometer. The dependence upon the strength of the probe laser in the Doppler measurements, however, suggests that the broadening is an artefact of conventional Doppler methods, perhaps due to off-resonant excitation [145]. This is consistent with several determinations by other researchers of atom cloud temperatures from measurements of the coherence length of the atomic wavepacket by measuring the fringe contrast as a function of wavepacket separation, in each case yielding a temperature below that estimated by Doppler [147] or time of flight (TOF) [148–150] methods.

TOF measurements are often used for colder samples and condensates [151]. While these can be taken in a single shot, they are

limited by the physical extent of the cloud and the imaging resolution [152]. In practice this means that the time of expansion required to measure the coldest distributions is typically $\sim 10\,\mathrm{ms}$, limiting its usefulness when studying dynamic behaviour such as in [153]. Both interferometric and Doppler measurements require multiple experimental runs, but can consequently be performed faster; such measurements with a Fourier transform limited resolution equivalent to Figure 7.14 for a 10 nK cloud could be made in as little as 500 µs. However, the Doppler measurement requires a continuous interaction for this time, increasing the probability of the artefacts we have observed while limiting the resonance to a small number of atoms and reducing the signal-to-noise ratio.

Interferometric measurement, in contrast, ideally involves interactions that last for a small fraction of the total measurement time and interact uniformly with the entire velocity and spatial distribution of the atom cloud so that, on average, half of the atoms contribute to the signal, limited by the finite range that can be addressed in practice. Interferometric velocimetry is hence a particularly effective complement to existing methods and is particularly suitable for colder atom samples in which artefacts such as off-resonant excitation, saturation and scattering force heating would otherwise distort the measured velocity distributions. It uses techniques, apparatus and, in some cases [149, 154], datasets that are often already to hand.

7.4 Conclusion

We have described the use of Ramsey matter-wave interferometry for the measurement of the velocity distribution, and hence translational temperature, of ultracold rubidium atoms. By using an asymmetrical three-pulse arrangement with switchable pulse phases, we record quadrature signals over both positive and negative effective interferometer durations. The Fourier transform, with correction for the residual Doppler effect within the interferometer pulses themselves, then reveals the atomic velocity distributions with good fidelity because the whole atomic sample contributes to each data point. The technique can distinguish between positive and negative velocities with a quadrature measurement, and is more effective at lower cloud temperatures. It is in many ways complementary to conventional techniques of Doppler-sensitive spectroscopy and TOF measurement, as it is not subject to distorting artefacts arising from off-resonant excitation

[145] or the physical extent of the atom cloud and measurement beam [152].

Principles of optimal control: some more musings on boats and clocks

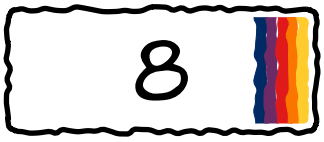

"Can you row?" the Sheep asked, handing her a pair of knitting-needles as she spoke.
"Yes, a little – but not on land – and not with needles – " Alice was beginning to say, when suddenly the needles turned into oars in her hands, and she found they were in a little boat, gliding along between banks: so there was nothing for it but to do her best.

— Lewis Carroll, *Through the Looking Glass*

In setting the scene for the next chapter, in which we present some results of optimal control theory, we have an excuse to return to the world of hand-drawn boats and clocks that we last visited in Chapter 5. Based, once more, on a pedagogical article in preparation, we introduce the *calculus of variations* — and the theory of *optimal control* that derives from it — by studying a classic problem in (aero)nautical navigation. Along the way, we will find that this problem is closely tied to the fundamental nature of quantum mechanics. Again, while this chapter offers some additional insight into the work that follows, it can be quite comfortably skipped by readers prone to seasickness.

8.1 Zermelo's navigation problem

I am no stranger to sailing myself. My pastime has taken me on many a passage across the English channel, heading roughly northwards or southwards to cross a waterway in which the tide flows predominantly east or west. Planning such a voyage entails looking at tidal predictions to ascertain how far east or west the tide is likely to carry your vessel over the duration of your voyage so that you can adjust your heading to aim an equal but opposite amount east or west of your destination in compensation. If I calculate that the tide will carry me 15 nautical miles east of where I want to be drinking my beer that night over the course of my

crossing then I maintain a bearing that, in the absence of tide, would take me 15 nautical miles west of my destination and hope I got my sums right.

This simple approach has served me, and mariners the world over, perfectly well — even earning me the odd trophy — but I have often wondered whether it is the optimal solution to the problem. In fact it is a version of Zermelo's navigation problem, a classic conundrum in the field of the calculus of variations, first considered by Ernst Zermelo in 1930 in the context of calculating the optimal course to steer an airship to a target destination when the wind experienced is a function of the ship's position [155–157]. The calculus of variations is the study of finding the extrema of 'functionals': functions that map a function, such as the trajectory one might steer across the channel as a function of time, to a scalar, such as the total time that said trajectory might take. We assume the optimal course across the the channel to be the one that maximises valuable drinking time (VDT) and thus minimises the duration of the crossing.

8.2 A solution using the calculus of variations

In Appendix B I demonstrate a method for calculating the optimal route across a waterway with a time-dependent tide loosely based

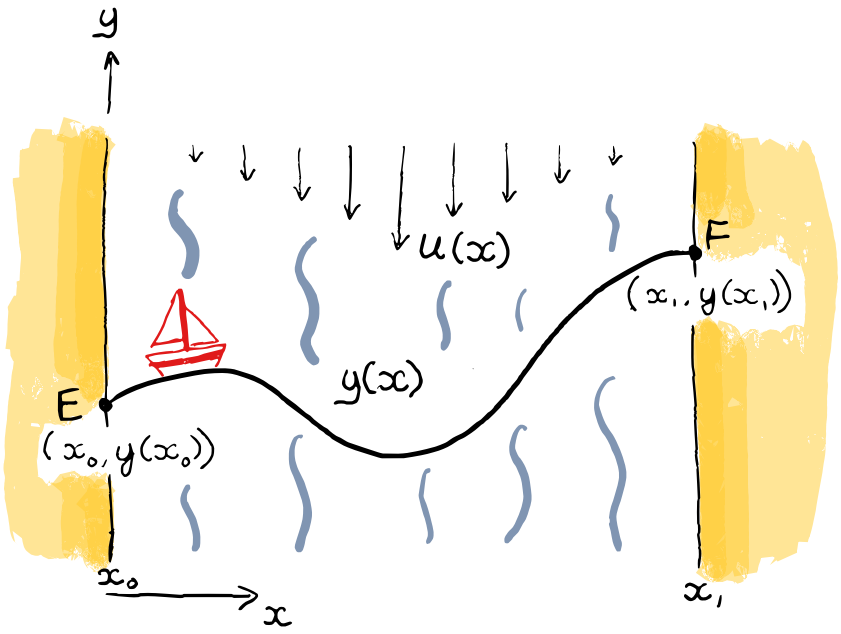

Figure 8.1: A simplified version of the optimal channel crossing problem, with a tide that is independent of time.

on the English channel,[1] but the time dependence of the tide necessitates the use of somewhat advanced techniques from the field of optimal control that occlude a lot of the elegance of the calculus of variations that underpins them. To illustrate these principles here, let us instead consider crossing a stretch of water such as a river in which the current does not vary with time.

[1] Spoiler: sailors have been getting it right for all these years.

As illustrated in Figure 8.1, let us define a coordinate system with the y axis lying parallel to the edges of the waterway given by the lines $x = x_0$ and $x = x_1$. The current $u(x)$ flows in the negative y direction and is a function of x. Assuming our boat can travel through the water with unit boat speed, we seek the fastest path to sail between a point $\mathbf{E} = (x_0, y_0)$ on the $x = x_0$ side and point $\mathbf{F} = (x_1, y_1)$ on the $x = x_1$ side (in keeping with our England-France theme). Assuming it will always be worth our while to make some progress towards the opposite bank, so that our x coordinate can be assumed to increase monotonically, we can describe our path across the water as a function $y(x)$.

Our boat's total velocity is then described by the pair of equations

$$\begin{aligned} \frac{\mathrm{d}x}{\mathrm{d}t} &= \cos(\theta) \text{ and} \\ \frac{\mathrm{d}y}{\mathrm{d}t} &= \sin(\theta) - u(x), \end{aligned} \tag{8.1}$$

and it is our task to find the function $y(x)$ that minimises the total duration of the crossing

$$T[y(x)] = \int_{x_0}^{x_1} \left(\frac{\mathrm{d}x}{\mathrm{d}t} \right)^{-1} \mathrm{d}x \tag{8.2}$$

when subject to the boundary conditions

$$\begin{aligned} y(x_0) &= y_0 \text{ and} \\ y(x_1) &= y_1 \end{aligned} \tag{8.3}$$

that ensure that the path starts at $\mathbf{E}$ and ends at $\mathbf{F}$.

$T[y(x)]$ is our functional, with the square brackets indicating that it operates on the function $y(x)$ as a whole to return a scalar denoting the duration of the crossing. The right-hand side of Equation (8.2) needs to be expanded to make the $y(x)$ dependence explicit.

From Equations (8.1) we can see that

$$\left(\frac{\mathrm{d}x}{\mathrm{d}t}\right)^2 + \left(\frac{\mathrm{d}y}{\mathrm{d}t} + u(x)\right)^2 = 1. \tag{8.4}$$

Writing $\frac{\mathrm{d}y}{\mathrm{d}t} = \frac{\mathrm{d}y}{\mathrm{d}x}\frac{\mathrm{d}x}{\mathrm{d}t}$ gives us a quadratic equation for $\frac{\mathrm{d}x}{\mathrm{d}t}$ in terms of $\frac{\mathrm{d}y}{\mathrm{d}x}$ and u, both functions of x, with the solution

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \frac{-u\frac{\mathrm{d}y}{\mathrm{d}x} + \sqrt{1 + \left(\frac{\mathrm{d}y}{\mathrm{d}x}\right)^2 - u^2}}{1 + \left(\frac{\mathrm{d}y}{\mathrm{d}x}\right)^2}, \tag{8.5}$$

where we take the positive root because we have already assumed that x will increase monotonically and this should be true in particular when $u(x) = 0$. Inserting this solution into Equation (8.2) we get

$$\begin{aligned} T[y(x)] &= \int_{x_0}^{x_1} \frac{1 + \left(\frac{\mathrm{d}y}{\mathrm{d}x}\right)^2}{-u\frac{\mathrm{d}y}{\mathrm{d}x} + \sqrt{1 + \left(\frac{\mathrm{d}y}{\mathrm{d}x}\right)^2 - u^2}}\,\mathrm{d}x \\ &\equiv \int_{x_0}^{x_1} \mathcal{L}\left(x, y(x), \frac{\mathrm{d}y}{\mathrm{d}x}(x)\right)\mathrm{d}x. \end{aligned} \tag{8.6}$$

While, in this particular case, the integrand $\mathcal{L}(x, y(x), \frac{\mathrm{d}y}{\mathrm{d}x}(x))$ depends on the path just through the derivative $\frac{\mathrm{d}y}{\mathrm{d}x}(x)$ and not $y(x)$ itself, we choose to consider the possibility of such a dependence anyway as the following steps are quite general. We will obtain a famous result that we can then specialise to the case at hand.

We are seeking a function $\bar{y}(x)$ that minimises the value of $T[y(x)]$. Just as in regular calculus, where the value of a function $f(x)$ does not vary to first order upon small variations about a value $\bar{x}$ that minimises (or maximises) it, $T[y(x)]$ should not vary to first order for small variations about $\bar{y}(x)$. Mathematically, making the substitution $y(x) \to \bar{y}(x) + \varepsilon\eta(x)$ and suppressing the function arguments for brevity,

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}\varepsilon} T[\bar{y} + \varepsilon\eta]\bigg|_{\varepsilon=0} = 0, \\ \Rightarrow\ &\frac{\mathrm{d}}{\mathrm{d}\varepsilon} \int_{x_0}^{x_1} \mathcal{L}\left(x, \bar{y} + \varepsilon\eta, \frac{\mathrm{d}\bar{y}}{\mathrm{d}x} + \varepsilon\frac{\mathrm{d}\eta}{\mathrm{d}x}\right)\mathrm{d}x\bigg|_{\varepsilon=0} = 0. \end{aligned} \tag{8.7}$$

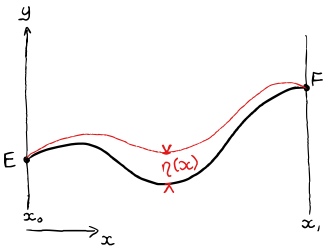

Figure 8.2: An arbitrary modification to the optimal path $\bar{y}(x)$ is represented by $\eta(x)$. $\bar{y}(x)$ is an extremum of the functional $T[y(x)]$ if $T[\bar{y}(x) + \varepsilon\eta(x)]$ does not vary to first order about $\varepsilon = 0$ for any $\eta(x)$.

Here $\eta(x)$ is an arbitrary, differentiable, function that, when added to $\bar{y}(x)$, generates another path that still goes from point **E** to **F**. Consequently it must vanish at the boundaries $\eta(x_0) = \eta(x_1) = 0$, as shown in Figure 8.2, but we place no further conditions on it.

Differentiating under the integral sign and setting $\varepsilon = 0$, Equation (8.7) becomes

$$\int_{x_0}^{x_1} \frac{\partial \mathcal{L}}{\partial y}\,\eta + \frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\,\frac{\mathrm{d}\eta}{\mathrm{d}x}\,\mathrm{d}x = 0. \tag{8.8}$$

The second term can be integrated by parts to give

$$\int_{x_0}^{x_1} \frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\,\frac{\mathrm{d}\eta}{\mathrm{d}x}\,\mathrm{d}x = \frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\,\eta\Bigg|_{x_0}^{x_1} - \int_{x_0}^{x_1} \frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\right)\eta\,\mathrm{d}x. \tag{8.9}$$

The boundary term vanishes by virtue of $\eta(x_0) = \eta(x_1) = 0$ and so, substituting back into Equation (8.8), we find

$$\int_{x_0}^{x_1} \left(\frac{\partial \mathcal{L}}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\right)\right)\eta\,\mathrm{d}x = 0. \tag{8.10}$$

At a path that minimises $T[y(x)]$ this must be true for any $\eta(x)$, and so we conclude that the multiplying factor in brackets must vanish.

$$\frac{\partial \mathcal{L}}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\partial \mathcal{L}}{\partial \frac{\mathrm{d}y}{\mathrm{d}x}}\right) = 0. \tag{8.11}$$

Equation (8.11) is called the Euler-Lagrange equation [158]. Solutions to this differential equation are known as critical functions and are guaranteed to be extrema of functionals defined as in Equation (8.6). In our case, $\mathcal{L}$ has no explicit dependence on y and so the first term is zero and we are left to solve

$$
\begin{aligned}
-\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\partial\mathcal{L}}{\partial\frac{\mathrm{d}y}{\mathrm{d}x}}\right) = \\
-\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{u\left(1-\frac{\mathrm{d}y}{\mathrm{d}x}^2\right)-\frac{\frac{\mathrm{d}y}{\mathrm{d}x}\left(1-2u^2+\frac{\mathrm{d}y}{\mathrm{d}x}^2\right)}{\sqrt{1-u^2+\frac{\mathrm{d}y}{\mathrm{d}x}^2}}}{\left(u\,\frac{\mathrm{d}y}{\mathrm{d}x}+\sqrt{1-u^2+\frac{\mathrm{d}y}{\mathrm{d}x}^2}\right)^2}\right) = 0.
\end{aligned}
\tag{8.12}
$$

This is not pretty, but it yields a second-order differential equation for $y(x)$ that can, in principle, be solved for the boundary conditions defined in Equation (8.3).

By way of example, let us choose $x_0 = -1$, $x_1 = 1$ and

$$u(x) = c(1 - x^2), \tag{8.13}$$

dictating that the current in the waterway has a quadratic profile, vanishing at the edges and flowing most strongly in the deep water at the centre with a maximum speed parameterised by c. This can now be inserted into Equation (8.12) and solved for $y(x)$. As an experimental physicist, my normal approach to solving complicated differential equations when I don't have an experimental apparatus to do it for me is to numerically integrate them. That is the approach we take here. To make things interesting, let us choose the boundary conditions $y_0 = 0$ and $y_1 = 3$ so that our goal is to get to the other side of the waterway having made some progress with or against the tide. A simple shooting method is employed to find optimal paths for a range of values of the parameter c, shown in Figure 8.3.

The results can be interpreted intuitively. When there is no tide flowing ($c = 0$) then the fastest path is a straight line to our destination. However, if our destination is down tide ($c > 0$) then it is beneficial to take a sigmoid path that allows us to benefit from the fast-flowing tide in the centre of the waterway for longer. Conversely, if our destination is up tide ($c < 0$) then we should take a sigmoid path of the opposite sign in order to cross the foul tide as quickly as possible.

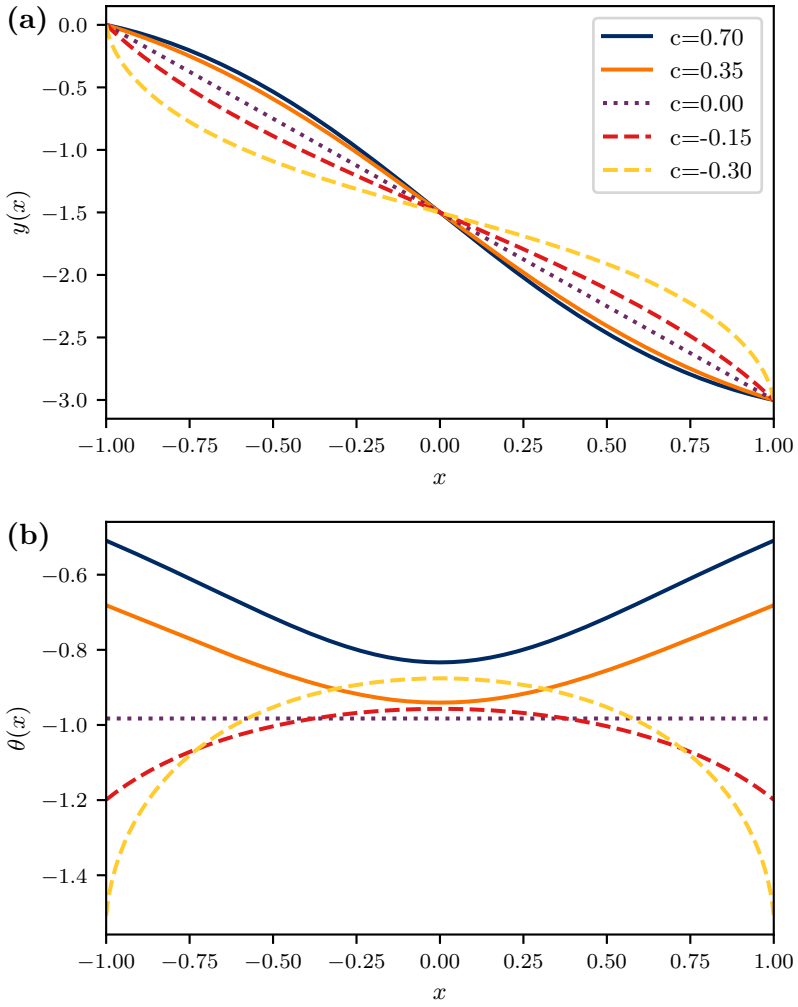

Figure 8.3: Results of numerical calculations for optimal (a) paths $\bar{y}(x)$ and (b) steering angles $\theta(x)$ for a range of values of the parameter c. The steering angle θ determines the direction we should point our boat and does not incorporate the component of the velocity vector due to the tide $u(x)$, as defined in Equation (8.1). The results follow intuition in showing that, when the tide is favourable and carries our boat towards our target ($c > 0$, solid lines), it is beneficial to sail more directly to the centre of the waterway in order to benefit from the fastest flow. Conversely, when the tide is acting against us ($c < 0$, dashed lines), we should make most of our progress against the flow in the slow current near the edges and cross the fast-flowing tide in the middle as quickly as possible. In the case of no current ($c = 0$, dotted line), a straight line is the fastest path to the target.

8.3 Relations to this work

This application of the calculus of variations to channel crossings is very interesting — if, perhaps, overly complicated — but appears to bear very little relation to the work of this thesis. In fact, while the ability to plan a channel crossing is not high on the list of priorities for most atomic physicists,[2] the calculus of variations is central to much of what we do.

[2] Although it should be said that it did prove genuinely useful during one return trip to England after working in the Bordeaux laboratory. See Figure 8.4.

Figure 8.4: Your author using the sun as a navigational aid (to measure latitude, not longitude, for which my sextant skills are not nearly precise enough) in order to calculate an updated course to steer while sailing back to England, following a stint working in Bordeaux.

8.3.1 Variational principles and the path integral formulation of quantum mechanics

Feynman's path integral formulation is a mathematically formidable approach to quantum mechanics that nonetheless stems from attractively simple and intuitive principles that can be effectively conveyed to a general audience, as Feynman himself does in his beautifully clear text 'QED: the strange theory of light and matter' [159]. Feynman postulates that the probability for a quantum mechanical particle to travel from a point **A** to a point **B** is calculated by summing the 'probability amplitudes' for every possible path the particle could take and then squaring the total. In a sense the particle, and indeed the universe, does everything it possibly can and, upon measurement, the probability of a given outcome is determined by summing over the contributions from all of its many histories.

The probability amplitudes that we sum over are complex numbers, so that they can combine constructively or destructively depending on their phase, and this concept of combining complex amplitudes and squaring to get a total probability is a very familiar one to quantum physicists. For the general reader, Feynman describes them as spinning arrows that should be combined head to tail in order to add them, with the square of the total length

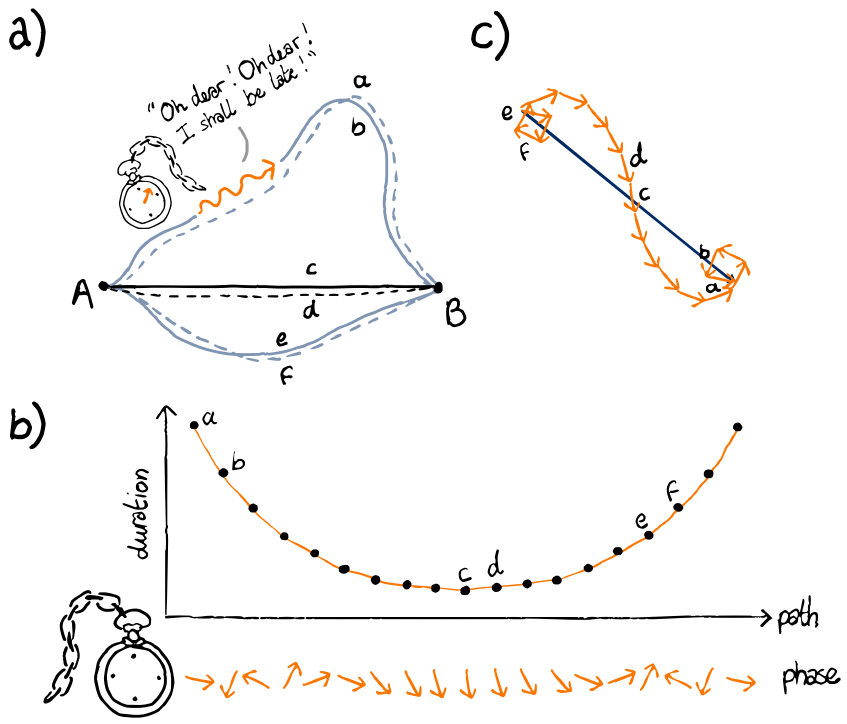

Figure 8.5: A selection of different paths that a photon can take between points **A** and **B** (a), and a schematic plot of their durations and how this affects the phase of the associated probability amplitudes (b). When the amplitudes are summed to give a final probability (c) it is the amplitudes for paths close to the extremum that contribute most to the final length, as the rapidly varying phase of the amplitudes for paths far from the extremum ensure they sum in ever decreasing circles and effectively cancel out.

representing the probability. In this chapter, perhaps it is appropriate to consider them as hands of a clock.

It remains to be said exactly how the amplitudes for each possible path are to be calculated. For photons, particles of light, the picture of a clock hand is very close to the truth. To a good approximation the amplitude for a particular path has a fixed length and the phase, or angle of the clock hand, can be thought of as that reached by a stopwatch that is started when the photon leaves point **A** and stopped when it arrives at point **B**. The rate at which the hand rotates is just the frequency of the light. Clearly, then, the phase of the amplitude depends on the duration a photon takes to propagate along the path at the speed of light.

Figure 8.5(a) shows a selection of different paths that a photon could take from a point **A** to **B**. The path that takes the shortest time, labelled c, is a straight line. Now, from our channel crossing example, we know that a small variation to this path will not affect the duration to first order. Consequently, the stopwatch timing the

duration along path d will stop at a very similar angle to that for path c. Conversely, the stopwatches for paths a and b, or e and f, far from the extremal path, could stop at very different angles.

Figure 8.5(b) schematically shows the duration of various paths around the minimum, and the corresponding stopwatch angles (or phases of the probability amplitudes).[3] The further a path is from the minimum, the more rapidly the stopwatch angle varies. When the amplitudes are summed by aligning them end to end, as in Figure 8.5(c) the rapidly varying phases from paths of increasing durations spiral in on themselves and do not contribute much to the length of the total amplitude, which is dominated by contributions from stopwatches around the minimum whose hands roughly align.

[3]This is an illustrative sketch. In reality the path of least time is a minimum of an infinite dimensional landscape that we reduce to one dimension for illustrative purposes by assuming that light travels in a straight line between **A** and an intermediate point **C**, and then again in a straight line from **C** to **B**. The horizontal axis of Figure 8.5(b) corresponds to translation of point **C** to vary the total path length.

As it is the square of this total length that dictates the ultimate probability that a photon leaving point **A** will arrive at point **B** we can see that, if we insert barriers to obstruct the paths at the top and bottom of Figure 8.5(a) then the likelihood of photons reaching point **B** will not change much. In effect, we would chop off some of the spiralling at the ends of Figure 8.5(c) but the overall length would stay largely the same. However, if we were to obstruct the straight-line, or line-of-sight, path then, as shown in Figure 8.6, *only* the spiralling ends would remain; the total amplitude would become much shorter. Point **B** would go dark as the probability of photons arriving there drastically diminishes. Feynman's formulation of quantum mechanics agrees, light travels in straight lines.

It agrees on much more than that, however. In chapter 2 of 'QED' Feynman proceeds to explain many of the familiar, and not so familiar, behaviours of quantum light using just the idea that anything that can happen does happen, and contributes the phase of its spinning stopwatch to the outcome of any subsequent measurement. In the case of photons, it is the paths of extremal duration that contribute most dominantly. This we can understand having looked under the hood of the calculus of variations, and it generally results in classical behaviour. The contributions from the other, non-classical, paths become important at a quantum level, however, and the sum-over-histories approach has proved supremely effective in determining, for example, small quantum corrections to the electron's magnetic moment. These are continually being calculated and then measured to higher and higher precision, making QED one of the most thoroughly and precisely tested theories in physics [34–36, 160, 161].

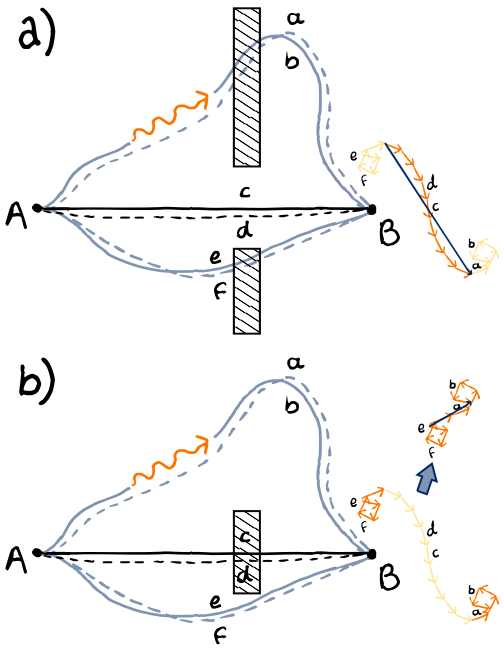

Figure 8.6: When the number of available paths from **A** to **B** is limited by blocking certain routes with impenetrable barriers, the probability of light reaching **B** from **A** is reduced. When line-of-sight path is unoccluded (a), the probability is not changed much as the occluded paths did not contribute much to the size of the total amplitude. Conversely, when paths around the line-of-sight route are blocked (b), the size of the total amplitude is greatly reduced and photons do not reach point **B**, in agreement with classical intuition.

The photon is a special, massless, case of a more general formula for calculating the probability amplitudes. For a general particle, the amplitude for a path $x(t)$ starting at a location $x(0) = x_0$ at time $t = 0$ and ending at a location $x(t') = x'$ at $t = t'$ is

$$e^{\frac{i}{\hbar}\int_0^{t'} \mathcal{L}\left(t, x(t), \frac{\mathrm{d}x}{\mathrm{d}t}(t)\right) \mathrm{d}t}. \tag{8.14}$$

The total probability amplitude (that should be squared to get the total probability) for a particle located at x_0 at $t = 0$ to be at x' by $t = t'$ is then

$$\langle x', t' | x_0, 0 \rangle = A(t) \int_{\text{all } x(t)} e^{\frac{i}{\hbar}\int_0^{t'} \mathcal{L}\left(t, x(t), \frac{\mathrm{d}x}{\mathrm{d}t}(t)\right) \mathrm{d}t}, \tag{8.15}$$

where the integral is over all possible paths and $A(t)$ is a normalisation factor independent of any particular path and therefore just dependent on t. The term $\mathcal{L}(t, x(t), \frac{\mathrm{d}x}{\mathrm{d}t}(t))$ integrated over in the exponential is reminiscent of the integrand in Equation (8.6) and is called the Lagrangian. It quantifies the difference between the kinetic and potential energy of the particle at each point along the

path $x(t)$. For a photon this is just $\hbar\omega$, where ω is the optical angular frequency, but for a massive particle it depends on the particle's momentum and the potential energy landscape.

In a sense, the particle's energy determines the rate that its stopwatch is ticking at any point on a given trajectory. It is, in the general case, paths about the extrema of the integral

$$S[x(t)] = \int_0^{t'} \mathcal{L}\left(t, x(t), \tfrac{\mathrm{d}x}{\mathrm{d}t}(t)\right) \mathrm{d}t \tag{8.16}$$

in the exponent that result in probability amplitudes of stationary phase and that hence contribute most to the total probability amplitude. As in Equation (8.6), this integral is a functional of the path $x(t)$ and we represent it with the capital S that the same quantity is denoted with in classical mechanics, where it is given the special name of the 'action'.

That critical functions of the action result in the most probable outcomes of quantum experiments should not come as a surprise to classical physicists. Indeed, the 'principle of least action' has been known to determine the equations of motion for classical systems for centuries. In retrospect, the existence of such a variational principle in classical mechanics is considered by some to have been a clear indicator of the wave-like nature of matter available to physicists long before the atom was discovered and quantum theory was developed to explain its properties [158]. Indeed, in developing his path integral formulation, Feynman derived the Schrödinger equation — an earlier, equivalent formulation of quantum mechanics, taught at undergraduate level — from not much more than the principle of least action and analogy to classical optics.

8.3.2 Optimal control

If we could model our channel crossing system as some sort of medium with a cleverly varying refractive index that we could shine light through, or else a potential energy landscape that we could fire a beam of particle through, then we could get the fundamental optimisation mechanisms built into nature to determine the optimal path for us, without having to resort to computation.[4]

This would be one approach to solving what is, at its heart, a problem of optimal control. Optimal control theory, loosely speaking, is the branch of mathematics that deals with finding the

[4] Indeed, when Zermelo published his papers presenting the navigation problem, physicists writing in the German literature were quick to draw analogies to the propagation of light in such a medium [157, 162, 163].

'best' way to drive a system from an initial state $\mathbf{x}_0$ to a target state $\mathbf{x}_1$.[5] More concretely, the definition of an optimal control problem generally starts with the definition of a dynamical system in the form

[5] Our collaborator Ilya Kuprov describes it as the theory of how to achieve the most, in the least amount of time, while expending the least amount of effort.

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{f}\left(\mathbf{x}(t), \mathbf{w}(t)\right), \tag{8.17}$$

where $\mathbf{x}(t)$ is a vector representing the state of a controllable system at time t and $\mathbf{w}(t)$ is a vector of variable control parameters that affect the evolution of the system. For our version of Zermelo's navigation problem, Equation (8.1) is in such a form, with $\mathbf{x}(t) = (x(t), y(t))$ representing the x and y coordinates of our boat, and the only control we have at our disposal is the steering angle $\mathbf{w}(t) = (\theta(t))$.

An optimal form of the control signal $\mathbf{w}(t)$ can then be sought that takes the solution of Equation (8.17) from $\mathbf{x}_0$ to $\mathbf{x}_1$ and minimises a cost functional

$$J(\mathbf{w}(t), \mathbf{x}_0, \mathbf{x_1}) = \int_0^T F(\mathbf{x}(t), \mathbf{w}(t), t)\,\mathrm{d}t, \tag{8.18}$$

where T is the time taken to evolve the system from $\mathbf{x}_0$ to $\mathbf{x}_1$ that is not constrained but is determined by the form of $\mathbf{w}(t)$. In Zermelo's case, a solution was sought that minimises T, and so one would set $F = 1$. This is indeed our approach in Equation (8.2) onward, but we write $\mathrm{d}t = (\frac{\mathrm{d}x}{\mathrm{d}t})^{-1}\,\mathrm{d}x$ and constrain the allowed states of the system by defining y as a function of x in order to find a solution by solving the Euler-Lagrange equation.

Clearly optimal control has its roots in the calculus of variations, but it has evolved a great deal and is now a very active field of study in its own right. In the spirit of Zermelo, it has found many applications in the field of navigation, particularly in trajectory determination of autonomous vehicles and the planning of space missions and manoeuvres. You would not thank me for going into the details of the various abstractions and computational methods that comprise modern optimal control theory here, although I give a flavour in demonstrating the application of Pontryagin's maximum principle to a more complicated channel crossing in Appendix B. Instead, using what we already know, I will describe a quantum mechanical problem in atom interferometry that we apply these techniques to in Chapter 9.

One of the strange consequences of quantum mechanics is that it is an analogue system with a digital output. In atom interferometry, we use lasers to manipulate atoms between two states with different momenta. Interaction with a laser for an appropriate amount of time will transfer an atom from a state that we label $|0\rangle$ to one that we label $|1\rangle$ and vice versa. When we look at an atom we will find it either in state $|0\rangle$ or state $|1\rangle$ and nothing in between. Such 'two-level' systems are often referred to as qubits, the quantum analogue of the digital bit that can be either off (0) or on (1) and forms the basis of digital computers. Quantum computers will be based upon qubits, but so far it all sounds very much the same. Putting funny angled brackets around the values that a bit can assume does not seem to make a computer any more useful.

But the wonderful thing about qubits is that they are, in fact, analogue. A single qubit can be in an infinite number of states that are 'superpositions' of $|0\rangle$ and $|1\rangle$. The general state of a qubit can be written as

$$|\psi\rangle = c_0 |0\rangle + c_1 |1\rangle \tag{8.19}$$

where $c_{0,1}$ are complex numbers with amplitudes in the interval $[0, 1]$. That these coefficients are complex means that, not only can qubits assume values between $|0\rangle$ and $|1\rangle$, they can have different relative phases that facilitate a rich set of quantum logic operations that are much more powerful than their digital equivalents.

The less wonderful thing about qubits is the aforementioned digital read-out. Reading a qubit in a state $|\psi\rangle = c_0 |0\rangle + c_1 |1\rangle$, we will sometimes measure it in state $|0\rangle$ and sometimes in state $|1\rangle$. The coefficients just determine with what probability we will measure one state or the other. The probability that we will measure state $|0\rangle$ ($|1\rangle$) is $|c_0|^2$ ($|c_1|^2$). We can not determine anything about the amplitude of c_1 relative to c_0 from a single measurement, so we must repeat the measurement many times to get a picture of the probability distribution. Unfortunately, when we measure my qubit as being in state $|0\rangle$ or $|1\rangle$ once, then subsequent measurements will yield the same result. The act of measurement is said to have 'collapsed' the superposition into a single basis state.

The answer is to perform the measurement on many qubits that have all been prepared identically. This is what we do in atom

interferometry. A cloud of atoms are prepared in identical superpositions that are acted upon by the environment. Further manipulation of the atoms is performed so that the effect of the environment, which acts on the relative phase between c_0 and c_1, can be determined by measuring the state of all of the atoms at the end and counting how many we find in state $|0\rangle$ and how many we find in state $|1\rangle$. The sensitivity of the final measurement depends on how identical we can make the response of each atom to our interactions in practice.

Now, as probabilities must sum to unity, clearly $|c_0|^2 + |c_1|^2 = 1$. This means that the possible states a qubit can occupy can be mapped onto the surface of a sphere, called the Bloch sphere after the physicist Felix Bloch, as illustrated in Figure 8.7. The z coordinate of state $|\psi\rangle$ represented on this sphere then determines the probability to be measured in state $|0\rangle$ or $|1\rangle$, with the north and south poles representing $|c_0|^2 = 1$ and $|c_1|^2 = 1$ respectively and the equator representing equal probability $|c_0|^2 = |c_1|^2 = 0.5$. The relative phase between c_0 and c_1 is then encoded in the state's x and y coordinates.

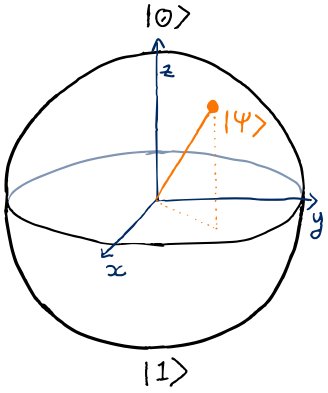

Figure 8.7: The state of a qubit can be represented on the surface of a sphere called the Bloch sphere. The z coordinate of a state represents the probability for it to be measured in each of the two basis states, and the x and y coordinates encode the superposition phase.

Manipulation of the qubit state is achieved by interaction with a control field, in the case of atom interferometry a resonant laser beam, and can be represented on the Bloch sphere as a rotation of the state about an effective field vector $\mathbf{\Omega}$ as depicted in Figure 8.8. Problems arise in atom interferometry, however, because the control field will never be completely homogeneous across all of the atoms that we wish to manipulate. Similarly, the atoms are in motion relative to the laser and they are all moving along the beam with different velocities.

This means that the effective field vector experienced by each atom is different; it will be orientated at differing angles and the rotations will occur at differing rates. In a simple laser pulse in which the field vector is kept static, this results in each atom being rotated into a slightly different state. How much the states differ from each other is determined by how well you can prepare your cloud of atoms to reduce the spread of velocities and how quickly you can perform your manipulations. Unfortunately, reducing the spread of velocities in your atom cloud generally means throwing away atoms and sacrificing the strength of your signal, while the speed with which you can perform the manipulations is limited by the available laser power.

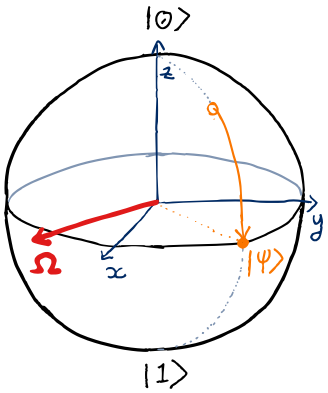

Figure 8.8: Interaction of a qubit with a control field corresponds to a rotation of the state on the surface of the Bloch sphere about an effective field vector $\mathbf{\Omega}$.

This is a problem to which, in Chapter 9, we have applied optimal control techniques with significant success. Rather than keeping the field vector fixed during the manipulations, we move it around by varying the phase of the laser light during a pulse. The optimal control algorithm treats the laser phase as a control parameter and considers the effect that varying it in time will have on a collection of atoms with a spread of velocities experiencing a range of laser intensities. The algorithm optimises the time-dependent form of the laser phase $\phi(t)$ to brings all the atoms considered as close as possible to a target state $|\psi_T\rangle$ from a single initial state.

The effect that the phase variation will have on the field vector experienced by each atom is not the same. The problem is akin to having a fleet of boats of differing abilities, some old and slow and some new and fast. You then seek a particular course to steer labelled 'A to B' that you can photocopy and distribute to every captain in your fleet, safe in the knowledge that whenever they are leaving port A and want to arrive at port B then they can do so by just steering the course you have given them, regardless of the the time of day, the particular currents flowing at the time or the capabilities of their boat.

This sounds impossible, and indeed it is if you want that course to be time optimal. However, if you allow for the passage to take a longer time then the problem starts to seem more tractable. Consider a channel crossing; you could imagine designing a course that would take exactly 12 hours to complete — a complete tidal cycle — so that the effect of tides was negligible, and that would result in the faster boats sailing a longer distance in that time but eventually ending up in the same place as slower boats who travel a shorter distance.

This is how we apply optimal control to perform the same operation for multiple atoms that experience different interactions. We increase the duration of our interactions to give ourselves more time to play with, and then use optimal control to design a laser pulse that sends atoms that are moving faster around the Bloch sphere off on a more circuitous route than those that are moving slower, but ensures that eventually they all end up in the same place.

It should be said that, as it stands, this optimisation is what we call 'open-loop'. The atomic system is modelled on a computer by making some approximations to reduce it to a tractable set of differential equations that are then solved numerically to arrive at

an optimal solution, but the actual atoms in the experiment do not form part of the optimisation process. When applying a similar approach to calculating an optimal course for a channel crossing in Appendix B, the end result essentially boiled down to saying that the approach sailors learn the world over is the correct one. This should hardly be surprising; sailors have been navigating tidal waterways for centuries, testing new approaches and pooling their collective knowledge in a giant feedback loop that has culminated in the principles and methods we are taught today. Community feedback has arrived at the solution that works the best, in the least possible time, and that requires the least amount of effort.

We are interested in pursuing a future in which this same approach will be employed in atom interferometers. The atom-interferometric inertial sensors that ultimately make their way into the navigation systems of boats and aeroplanes could be 'closed-loop' hybrid devices in which the shape of the laser pulses is constantly being tuned based on feedback from the output of both fast but coarse classical sensors and the slow but precise quantum signal in order to maximise the sensitivity of the measurement [57, 164, 165].

Optimised Raman pulses for atom interferometry

9

Capital letters were always the best way of dealing with things you didn't have a good answer to.

— Douglas Adams, *Dirk Gently's Holistic Detective Agency*

Optimal control, introduced in the last chapter in the context of Zermelo's navigation problem, is a powerful application of the calculus of variations to determining the best way to vary a set of controls at one's disposal in order to achieve a desired outcome. In Zermelo's case, we determined the optimal way to vary a ship's course over time in order to reach a target destination at the other side of a waterway with known currents in the shortest possible time.

It is appealing to consider whether, by defining the problem appropriately, similar techniques can be applied to determining the optimal way in which to vary the inputs of an atom interferometer in order to extract the most sensitive measurement.

Optimal control has been successfully employed in many quantum physics experiments, including nitrogen-vacancy (NV) centre magnetometry [166], the stabilisation of ultra-cold molecules [167], and the control of BECs [57–60]. One area of physics to which optimal control has been applied with particularly great success is NMR spectroscopy [50], in which it is employed to design control pulses that can perform desired manipulations of two-level spin systems even when subjected to variations in the strength of the control field and its detuning from resonance.

Many aspects of these interactions are directly analogous to the atom–laser interactions of atom interferometry. Indeed, NMR optimal control pulses are a natural extension of composite pulses [52] — that replace state vector rotations traditionally performed by fractional Rabi oscillations with pulses that comprise of a series of steps with different phases and durations, tailored to compensate for inhomogeneities in the interaction — and the applicability of such pulses to atom interferometry has already been established [54–56].

In this chapter we apply optimal control techniques [47–49, 158] to design tailored Raman pulses for atom interferometry that achieve high-fidelity state transfer in a hot atomic sample. The superior velocity acceptance and resistance to coupling strength variations when compared with existing composite and shaped pulses also make these pulses potential candidates for use as augmentation pulses in large momentum transfer (LMT) atom interferometers that extend the Doppler separation between the superposed atomic states to increase the interferometer area and sensitivity [55, 168]. We also consider the application of optimal control techniques to all three pulses of a Mach-Zehnder type interferometer.

The theoretical and numerical work was mainly conducted by Jack Saywell using the *Spinach* software suite for spin dynamics [48], and a more in-depth description of the methodology can be found in [2]. I was able to offer some advice on how the NMR toolset could best be put to use in the context of atom interferometry, and conducted the experimental work presented here to demonstrate and validate the computational designs. This chapter presents the results of our collaboration, soon to be published in a paper that we have co-authored [51].

9.1 Computational optimal control methods for atom interferometry

For the purposes of optimising the Raman interactions in our experiment we assume that the single-photon detuning of the interferometry beams is sufficiently large that the excited state can be adiabatically eliminated and the system can be treated effectively as a two-level one. We furthermore assume coherent (unitary) evolution of the system throughout the interferometry sequences.

As derived in Section 2.6, in a rotating frame with an appropriate choice of energy reference, the interaction Hamiltonian then becomes

$$\hat{H} = \frac{\hbar}{2}\begin{pmatrix} \delta_L & \Omega_R e^{-i\phi_L} \\ \Omega_R e^{i\phi_L} & -\delta_L \end{pmatrix} \tag{9.1}$$

in the $\binom{c_1}{c_2}$ basis, where δ_L is the laser detuning from the light shifted Raman resonance, Ω_R is the (real) Rabi frequency and ϕ_L is

the relative phase between the Raman beams that is controllable in our experiment by I&Q phase modulation of the EOM RF signal.

The time-dependent Schrödinger equation is then exactly soluble and, as in Section 7.3.1, the change in state effected by a pulse of constant intensity and phase lasting for a duration Δt can be represented by the propagator

$$\hat{\Omega}(\Delta t, \delta_L, \Omega_R, \phi_L) = \begin{pmatrix} C^* & -iS^* \\ -iS & C \end{pmatrix}, \tag{9.2}$$

with $C(\Delta t, \delta_L, \Omega_R)$ and $S(\Delta t, \delta_L, \Omega_R, \phi_L)$ defined in Equations (7.27).

Atom interferometers typically employ rectangular $\pi/2$- and π-pulses — introduced in Section 2.4 — with respective durations of $\tau_{\pi/2} = \pi/2\Omega_R$ and $\tau_\pi = \pi/\Omega_R$ so that, when $\delta_L = 0$, they perform $\pi/2$ and π rotations about an axis in the x–y plane of the Bloch sphere. Variations in δ_L and Ω_R, known in the NMR literature as *off-resonance* and *pulse length* errors respectively, translate into shifts to both the axis and angle of these rotations that ultimately reduce the contrast of the interferometer output.

In this chapter we consider pulses in which the relative laser phase ϕ_L is varied over the pulse duration. Since, for computation and experiments, the timebase is discretised into timesteps of equal duration $\mathrm{d}\tau$, the action of such a pulse on a state then takes the form of a time ordered product of propagators $\hat{\Omega}_n(\mathrm{d}\tau, \delta_L, \Omega_R, \phi_n)$

$$\hat{\Omega}_{\text{pulse}} = \hat{\Omega}_N \hat{\Omega}_{N-1} ... \hat{\Omega}_2 \hat{\Omega}_1, \tag{9.3}$$

where N is the total number of timesteps. The $\phi_n, (n = 1, 2, ..., N)$ then become our control parameters that we use to steer our atom interferometry pulse, with the aim of steering a course that is as close as possible to a desired rotation for a specified range of Rabi frequencies Ω_R and detunings δ_L.

The control parameters are optimised using the *gradient ascent pulse engineering (GRAPE)* algorithm. Given an initial guess for the control parameters, GRAPE efficiently calculates the derivatives of a chosen fidelity function with respect to them and ascends the steepest path to a local maximum.

The fidelity is a function of the optimised pulse propagator $\hat{\Omega}_{\text{pulse}}$, and establishing a form for this function is a crucial step in optimising pulses for a particular application. We consider two possible types of fidelity function in the following sections. One is the *point-to-point (PP)* state transfer fidelity where, given an initial state $|\psi_0\rangle$, the fidelity $\mathcal{F} = |\langle\psi_T|\hat{\Omega}_{\text{pulse}}|\psi_0\rangle|^2$ is defined as the overlap with a target state $|\psi_T\rangle$ after application of the pulse propagator. Alternatively, a *universal rotation (UR)* pulse is yielded by choosing a fidelity $\mathcal{F} = {}^1/_2 \operatorname{Tr}\left(\hat{\Omega}_T{}^\dagger \hat{\Omega}_{\text{pulse}}\right)$ that quantifies the overlap between the optimised pulse and a target propagator $\hat{\Omega}_T$.

9.2 Broadband pulse for large momentum transfer atom interferometry

Atom interferometers are the matter-wave analogue of optical interferometers in which coherent light is divided between separate spatial paths and recombined to produce interference. Atom interferometers replace the photons of optical interferometry with slow, massive, atoms for an initial sensitivity gain but, as with an optical interferometer, sensitivity is improved by increasing the separation between the interfering wavepackets. While ultrastable lasers and optical fibre technology make extremely sensitive optical interferometry possible over kilometre length scales, obtaining a macroscopic path separation in an atom interferometer by either increasing the interferometer duration (and size) or the initial momentum difference remains technologically challenging.

In general, attempts to increase the path separation in an atom interferometer for more sensitivity will be to the detriment of fringe visibility, limiting their usefulness [55, 65, 66, 169]. This is particularly true of large momentum transfer (LMT) interferometers that employ extended pulse sequences to separate atomic trajectories by multiple recoil momenta, resulting in a larger enclosed interferometer area — and hence sensitivity — for the same interferometer duration [64]. LMT atom optics typically rely on an atomic sample with a narrow initial momentum distribution [170], with Bloch oscillations [169] and Bragg diffraction [25, 65] demonstrating the greatest separation.

In a large, laboratory-based, experiment with a Bose-Einstein condensed (BEC) atomic sample, such techniques can produce momentum superpositions where the wavepacket is separated on the half-meter length scale [27, 118]. However, BECs require long

preparation times and, in order to make atom interferometers a viable alternative to optical devices for applications such as inertial sensing where a high repetition rate is important, there is a renewed interest in techniques for high-sensitivity atom interferometers with less stringent requirements on the atomic source. Such techniques include LMT interferometers using adiabatic state transfer [67, 85] and composite Raman pulses [55].

Adiabatic transfer, while robust, is necessarily a slow process and is most effective at high-fidelity transfer from one quantum state to another, not preparing superpositions [171]. Composite Raman pulses are a technically simple alternative that have shown promise [55], and here we seek to significantly increase their performance and robustness with optimal control.

9.2.1 Pulse design and implementation

One way to increase the momentum separation between the arms of LMT interferometers is to 'augment' the beamsplitter and mirror operations with multiple *augmentation pulses* that have alternating effective wavevectors designed to swap the population

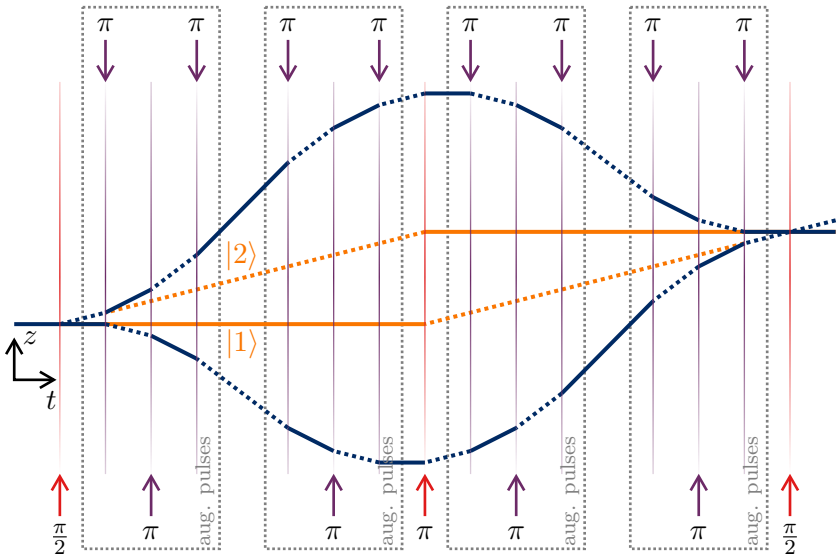

Figure 9.1: Illustration of a LMT interferometer sequence. Vertical arrows represent the effective k-vector of interferometry pulses. The red pulses correspond to the standard $\pi/2 - \pi - \pi/2$ Mach-Zehnder interferometer sequence, with the corresponding atomic trajectories shown with thick orange lines with solid (dotted) lines representing time spent in state $|1\rangle$ ($|2\rangle$). The purple pulses in the designated boxes show additional augmentation π-pulses where the direction of the effective k-vector is alternated in order to increase the momentum difference between the arms of the interferometer and increase the enclosed area, as demonstrated by the atomic trajectories illustrated with thick blue lines.

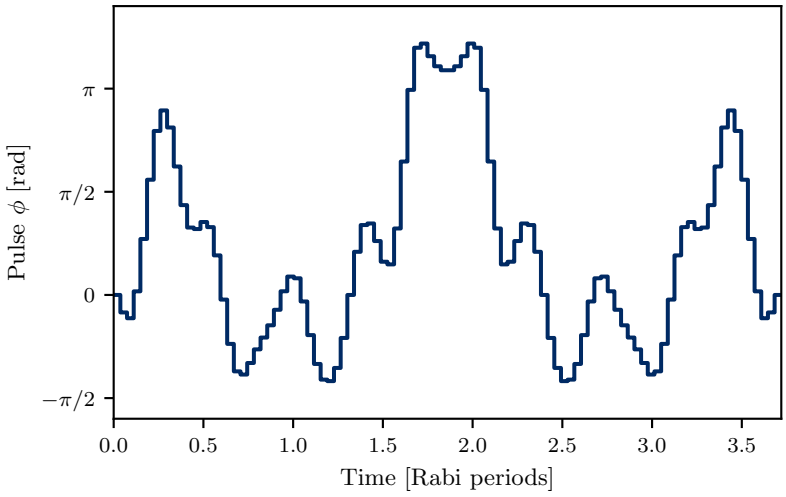

Figure 9.2: Phase profile of BAT pulse, with the time axis represented in units of the Rabi period $^{2\pi}/_{\Omega_0}$.

of the internal states whilst imparting additional momentum as depicted in Figure 9.1 [55, 64]. These pulses require a large velocity acceptance, equivalent to working over a large range of detunings, as they must remain resonant with both interferometer arms as they are Doppler shifted in opposite directions. However, in order to optimise these augmentation pulses it is sufficient to consider the point-to-point fidelity

$$\mathcal{F}_{\mathrm{A}} = |\langle 2|\,\hat{\Omega}_{\mathrm{A}}\,|1\rangle\,|^2 \tag{9.4}$$

without concern for the relative phase introduced between the two states. This is because the augmentation pulses appear in pairs within the extended pulse sequence [55, 56, 64] so that the interferometer phase introduced by each pulse is, to first order, cancelled out by that introduced by a subsequent one. Relaxing this constraint on the optimisation effectively gives the GRAPE algorithm a larger target to shoot at, allowing a greater fidelity to be achieved within a given pulse duration.[1]

To that end, the GRAPE algorithm was employed to optimise a pulse with 100 time steps and a duration of 3.72 Rabi periods, maximising $\mathcal{F}_{\mathrm{A}}$ when averaged over a Gaussian distribution of detunings δ_L with a standard deviation of $\sim 0.9\Omega_R$ and a linear spread of Rabi frequencies of $\pm 10\%$ about an assumed central frequency Ω_0. These values were chosen based on a typical experimental Rabi frequency of $\Omega_R = 2\pi \times 310\,\mathrm{kHz}$, giving the pulse a duration of 12 µs and meaning that the detuning spread would match that expected within a cloud with a Maxwell-Boltzmann temperature

[1] In practice this cancellation is such that the phase introduced by the first augmentation pulse in the interferometer sequence is cancelled out by that introduced by the final augmentation pulse. This leaves the interferometer sensitive to time-dependent variations in Rabi frequency and detuning. While point-to-point pulses can give a higher fidelity than universal rotations for the same pulse duration, it is feasible that augmentation pulses that perform universal rotations may yield greater contrast at the interferometer output when such time-dependent variations are present. It is worth noting that adiabatic rapid passage (ARP) pulses are essentially a form of point-to-point transfer.

of 120 µK. In practice, the length of each timestep is scaled if the empirically determined experimental Rabi frequency varies.

To enforce waveform smoothness, a penalty term proportional to the difference between adjacent phase steps was added to the fidelity function. No symmetry constraints were imposed, but the algorithm nonetheless converged to the approximately symmetric phase profile shown in Figure 9.2 that we affectionately christen the broadband atomic transfer (BAT) pulse.[2]

[2]It is commonplace in the NMR literature to assign acronymic monikers to pulses with varying levels of contrivance (CORPSE, WALTZ and SCROFULOUS being just a few examples). It would be a shame to go against such a precedent, and broadband atomic transfer (BAT) proves to be a concise and apt description of both our pulse's function and the visual appearance of its phase profile.

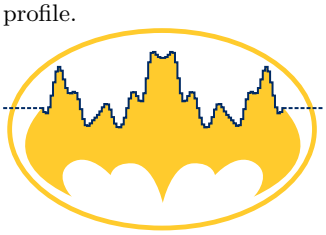

Such a phase profile $\phi_n, (n = 1, 2, ..., N)$ is realised experimentally by modulating the phase of the RF signal driving the EOM that produces one of the Raman frequency components with a *Miteq SDM0104LC1CDQ* I&Q modulator. The I and Q inputs are controlled by the dual outputs of a *Keysight 33612A* AWG that are programmed with the waveforms $I_n = V_0 \sin\phi_n$ and $Q_n = V_0 \cos\phi_n$ and configured to maintain the final phase value ϕ_N until a hardware trigger is received. The response of the I&Q exhibits a degree of nonlinearity in this phase mapping [89] that increases with the modulation amplitude V_0. Rather than attempt to correct for this, we instead operate at a low modulation amplitude where the response is approximately linear.

The sample rate at which the waveforms are played is varied in order to set the total duration of the modulation τ_{mod}. In practice, the time the AWG takes to respond to a hardware trigger is determined by the sample rate and so we oversample, with each phase step ϕ_n comprising multiple subsamples of equal phase, so that the sample period is much less than the ~ 100 ns rise time of the AOM that shutters the Raman beams.

Having scaled the modulation duration τ_{mod} to roughly correspond to the appropriate number of empirically determined Rabi periods, the pulse performance is optimised by setting the Raman detuning δ_L to resonance and measuring the temporal evolution of the excited fraction $|c_2|^2$ of a thermal atom cloud throughout the pulse in order to identify the point of peak transfer. This is achieved by concurrently triggering the phase modulation and the AOM that shutters the Raman beams, then measuring $|c_2|^2$ once the AOM is turned off again after a variable time τ. The hardware trigger delay and sample rate are then adjusted to respectively adjust the start time τ_{offset} and duration τ_{mod} of the phase modulation, as shown in Figure 9.3 in order to maximise the height of the peak transfer in these scans.

9.2.2 Performance characterisation

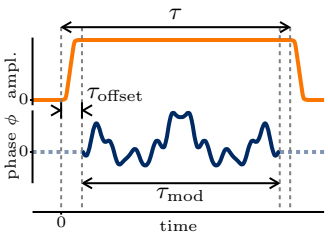

Figure 9.3: Illustration, not to scale, of the timing parameters optimised to realise optimal BAT pulses. The Raman light (top) is turned on at $t = 0$, and off again at $t = \tau$ by an AOM with a rise-time of $\sim 100\,\text{ns}$. The sample rate of the AWG controlling the phase (bottom) is adjusted to set the total duration of the phase waveform τ_{mod}, and this is adjusted together with the trigger delay τ_{offset} to achieve the optimal peak transfer when τ is scanned. Once the phase waveform has finished, the phase is maintained at ϕ_N, the value of the final step.

Such a temporal scan for this broadband atomic transfer (BAT) pulse, on resonance and having optimised the timing parameters τ_{mod} and τ_{offset}, is shown in Figure 9.4. It shows a peak in $|c_2|^2$ at the end of the phase sequence, represented by the vertical dashed line, after which damped Rabi oscillations are observed as the phase is fixed. We find that optimised pulses demonstrate a considerable resilience to variations in the trigger delay τ_{offset}, with efficient transfer still observed even when the phase modulation is commenced up to a quarter of a π-pulse duration after the Raman light is turned on.

The peak transfer is achieved when the light is kept on for slightly longer than the duration of the phase modulation, with $\tau > \tau_{\text{mod}}$ by $\sim 200\,\text{ns}$. The optimisation algorithm assumes that the light is turned on instantaneously as the phase modulation begins, without consideration of the AOM rise time, which probably accounts for this. Nonetheless, the form of the temporal profile shows excellent agreement with a numerical simulation that also assumes instantaneous illumination, shown as a solid line in the bottom panel of Figure 9.4.

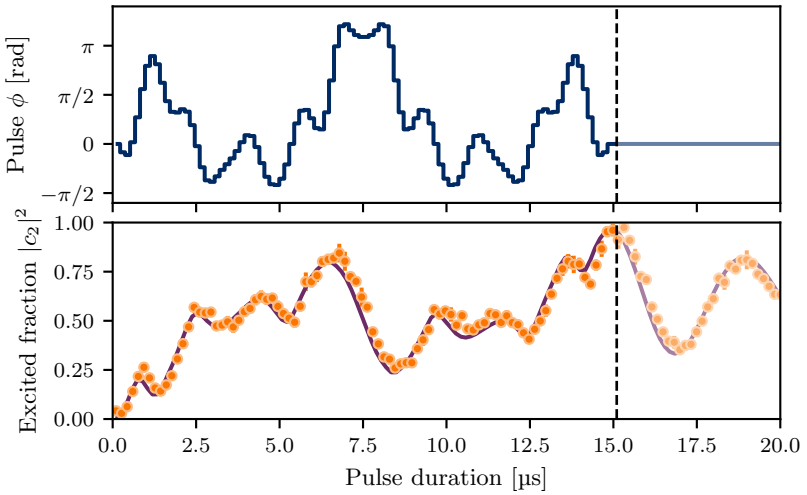

Figure 9.4: *Above*: BAT pulse phase profile designed using GRAPE to transfer atoms between levels $|1\rangle$ and $|2\rangle$ after scaling the duration of the modulation to give maximal transfer. *Below*: Measured fraction (circles) of atoms in the excited state $|c_2|^2$ after the Raman light is extinguished at various times during a pulse. The solid line is a theoretical curve produced by the model from [56], in which the two-level Hamiltonian is numerically integrated over the range of detunings and coupling strengths present in a thermal cloud of ^{85}Rb atoms, and assumes that the light reaches full intensity instantaneously and concurrently with the start of the phase modulation. Excellent agreement is observed for a simulated temperature of $35\,\mu\text{K}$.

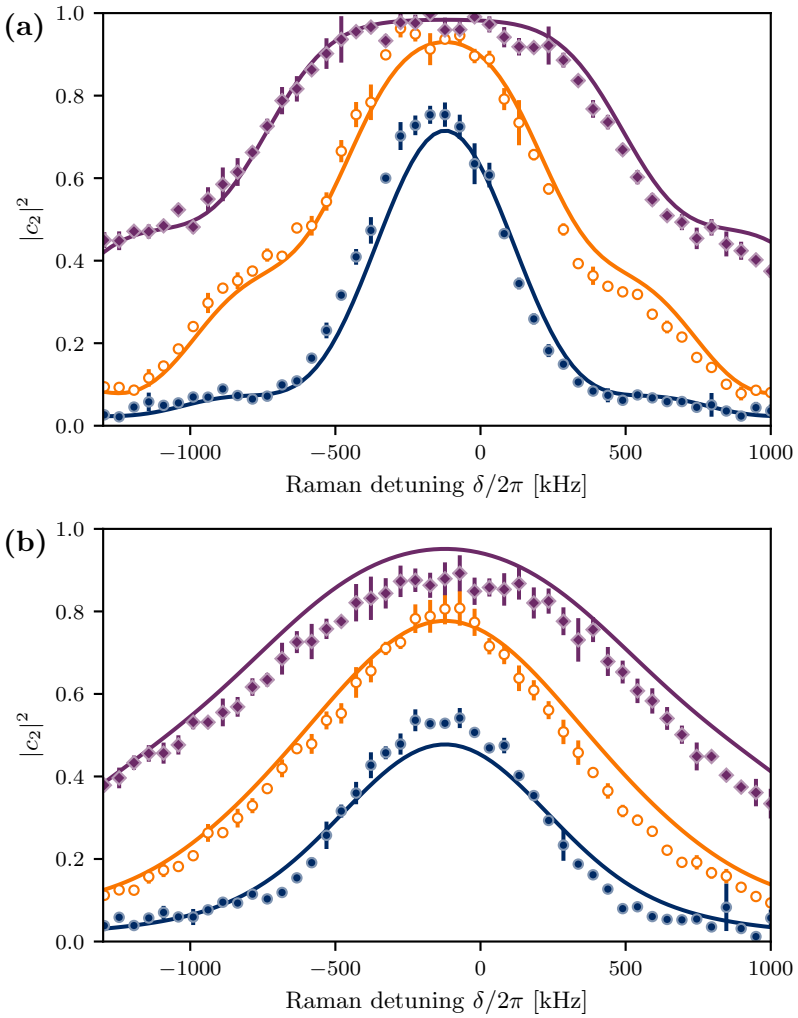

Figure 9.5: Fraction of atoms transferred to the excited state $|c_2|^2$ as a function of laser detuning δ_L in (a) a $\sim$ 35 µK atom cloud and (b) a $\sim$ 150 µK cloud. Data are shown for BAT pulse (diamonds), WALTZ pulse (empty circles) and a rectangular π-pulse (filled circles). The effective Rabi frequency was $2\pi \times 270$ kHz. solid lines show theory curves produced from the model used by Dunning *et al.* [56], which assumes a Maxwell-Boltzmann atomic velocity distribution. While, for cold atoms, the difference in transfer efficiency between BAT and WALTZ pulses is small on resonance, the greater range of detunings for which the BAT pulse remains efficient illustrates its potential utility for LMT interferometry.

In order to demonstrate the potential of the BAT pulse as an augmentation pulse for LMT, we compare the fraction of thermal atoms it transfers from state $|1\rangle$ to state $|2\rangle$ to the WALTZ pulse — the best composite Raman pulse previously employed in LMT interferometry [55] — and a rectangular π-pulse over a range of Raman detunings δ_L.

In Figure 9.5(a) this comparison is shown for a 35 µK sub-Doppler cloud where, nonetheless, the Doppler width of the velocity distribution is broader than the spectral width of a rectangular π-pulse

so that it transfers just 75(3) % of the atoms. In contrast, the BAT and WALTZ pulses achieve 99.8(3) % and 96(2) % transfer respectively. The BAT pulse earns its moniker, however, demonstrating its broadband nature by maintaining $> 90\,\%$ transfer over a $\pm 2\pi \times 380\,\mathrm{kHz}$ range compared to $\pm 2\pi \times 100\,\mathrm{kHz}$ for WALTZ.

This broad spectral profile is a signature of the large velocity acceptance required for LMT interferometry. For the ^{85}Rb atoms used here, each augmentation pulse would introduce $\sim 2\pi \times 60\,\mathrm{kHz}$ difference in the Raman resonance condition for each interferometer arm so it can be inferred that, starting on resonance, BAT pulses could sustain 90 % fidelity in a thermal cloud after $\sim$ 12 applications as augmentation pulses compared to just $\sim$ 3 WALTZ pulses.[3]

[3] Of course, the interferometer contrast would reduce rapidly, roughly as the product of the fidelities for every interferometer pulse, and an extra 4 pulses need to be added per $\hbar k_{\mathrm{eff}}$ of momentum separation between the arms.

In a $\sim 150\,\mu\mathrm{K}$ cloud — much closer to Doppler cooling limit — for which the peak transfer of a rectangular π-pulse is just 54(2) %, the broader spectral profiles of the BAT and WALTZ pulses translate to more efficient state transfer on resonance. This is shown in Figure 9.5(b), with the BAT pulse achieving 89(4) % transfer on resonance compared to the 81(4) % achieved by the WALTZ pulse.

9.3 Optimised atom interferometry pulse sequences

We now consider a basic $\pi/2$-π-$\pi/2$ Mach-Zehnder type interferometer in which inertial effects such as rotation or acceleration of the apparatus relative to the atoms imprint a relative phase Φ between the superposed atomic states. This phase is mapped onto the state probabilities by the final $\pi/2$-pulse, so that the probability of finding an atom in the excited state $|2\rangle$ at the interferometer output is

$$|c_2|^2 = \frac{1}{2}\left[\mathcal{A} - \mathcal{B}\cos\left(\Phi + \Delta\phi\right)\right], \tag{9.5}$$

where $\mathcal{A}$ and $\mathcal{B}$ are the offset and contrast of the interferometer fringes respectively and

$$\begin{aligned}\Delta\phi = &\operatorname{Arg}\langle 2|\,\hat{\Omega}_{\mathrm{a}}\,|1\rangle - \operatorname{Arg}\langle 1|\,\hat{\Omega}_{\mathrm{a}}\,|1\rangle \\ &- 2\operatorname{Arg}\langle 2|\,\hat{\Omega}_{\mathrm{b}}\,|1\rangle \\ &+ \operatorname{Arg}\langle 1|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle + \operatorname{Arg}\langle 2|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle\end{aligned} \tag{9.6}$$

is a phase factor introduced to the output by the atom–light interactions themselves. This should be minimised, or at least fixed, over a range a Raman detunings δ_L and Rabi frequencies Ω_R so that the signal has the same phase for all atoms and the fringes are not washed out when averaged over a thermal ensemble as they are with some pulses [2, 172].

The $\hat{\Omega}_{\mathrm{a,b,c}}$ of Equation (9.6) are propagators for the first, second and third pulses of the interferometer respectively that, with appropriate choices of fidelity functions, can be optimised to maximise the fringe contrast $\mathcal{B}$ and minimise variations in the fringe offset $\mathcal{A}$ and phase $\Delta\phi$ across an atomic ensemble.

9.3.1 Optimisation considerations

Maximising the contrast $\mathcal{B}$ automatically limits any variation in the offset $\mathcal{A}$, and this is achieved by ensuring the following conditions are met:

$$\left|\langle 2|\,\hat{\Omega}_{\mathrm{a,c}}\,|1\rangle\right|^2 = \frac{1}{2}, \text{ and} \tag{9.7}$$

$$\left|\langle 2|\,\hat{\Omega}_{\mathrm{b}}\,|1\rangle\right|^2 = 1. \tag{9.8}$$

This corresponds to the intuitive notion that our beamsplitter pulses $\hat{\Omega}_{\mathrm{a,c}}$ should perform 90° rotations about an axis in the x–y plane of the Bloch sphere and the mirror pulse $\hat{\Omega}_{\mathrm{b}}$ should perform a robust 180° rotation of the quantum state. We thus design our pulses to meet these conditions and also fix the value of $\Delta\phi$ as defined in Equation (9.6).

At the start of the interferometer the atoms are prepared in the ground state $|1\rangle$, and we design the first pulse to transfer the atoms to a state $|\psi_T\rangle = {}^1\!/\!\sqrt{2}\left(|1\rangle + |2\rangle\right)$ with well-defined phase on the equator of the Bloch sphere. This means maximising the fidelity

$$\mathcal{F}_{\mathrm{a}} = \left|\langle\psi_T|\,\hat{\Omega}_{\mathrm{a}}\,|1\rangle\right|^2, \tag{9.9}$$

yielding a PP 90° pulse that satisfies the condition $|\,\langle 2|\,\hat{\Omega}_{\mathrm{a}}\,|1\rangle\,|^2 = {}^1\!/\!2$. Maximising this fidelity over the same range of detuning and Rabi frequency offsets as used for the BAT pulse in Section 9.2, with a pulse duration of slightly more than

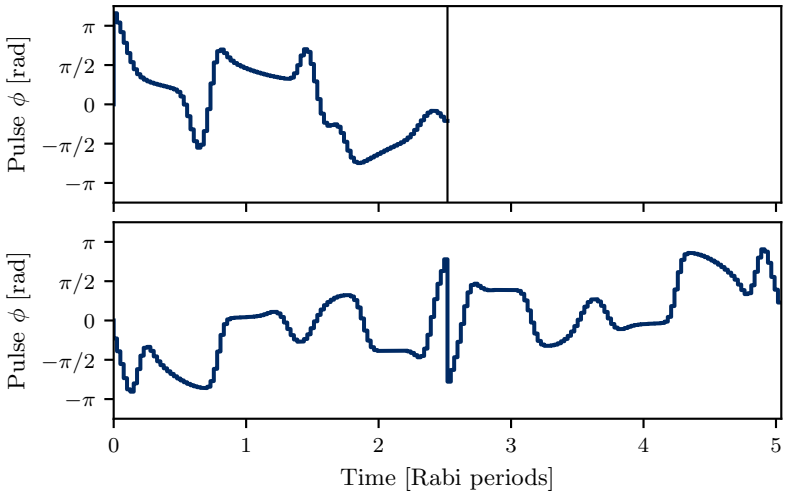

Figure 9.6: Optimisation results for, *above*, a beamsplitter pulse maximising $\mathcal{F}_1$ and, *below*, an antisymmetric mirror pulse constructed following the procedure outlined by Luy *et al.* [173]. The pulse profiles are plotted against time as a fraction of the Rabi period $^{2\pi}/_{\Omega_0}$. Each optimisation was continued until fidelities greater than 0.99 were reached. Each pulse was optimised for an ensemble of atoms with a temperature of 120 μK and a range of coupling strengths of $\pm$ 10% Ω_0.

2.5 Rabi periods, resulted in the phase profile shown in the upper panel of Figure 9.6.

The mirror pulse $\hat{\Omega}_\mathrm{b}$ should swap the internal states and introduce no relative phase between them, satisfying the condition $|\langle 2|\,\hat{\Omega}_\mathrm{b}\,|1\rangle|^2 = 1$ and constraining the third term in Equation (9.6) $\mathrm{Arg}\,\langle 2|\,\hat{\Omega}_\mathrm{b}\,|1\rangle = \text{constant}$. The PP optimisation employed for the BAT pulse is therefore insufficient, as this should enact a 180° rotation about a fixed axis in the x–y plane. It is thus natural to maximise a UR fidelity function that quantifies the overlap of the pulse propagator to that of an ideal 180° rotation $\hat{\Omega}_\pi$,

$$\mathcal{F}_\mathrm{b} = \frac{1}{2}\,\mathrm{Tr}\left(\hat{\Omega}_\pi{}^\dagger \hat{\Omega}_\mathrm{b}\right). \tag{9.10}$$

However, rather than optimising this fidelity directly, we note that if the phase profile of the pulse is antisymmetric about the temporal midpoint then the axis of rotation will be fixed to the x–z plane of the Bloch sphere for all δ_L and Ω_R by virtue of the pulse symmetry — a property known in the NMR literature [53, 174, 175] — automatically ensuring that any modification to the phase $\Delta\phi$ is constant for all resonance offsets.

UR 180° pulses with this symmetry can be constructed by first optimising a PP 90° pulse and then concatenating the phase profile with a time-reversed and phase-inverted version of itself, as

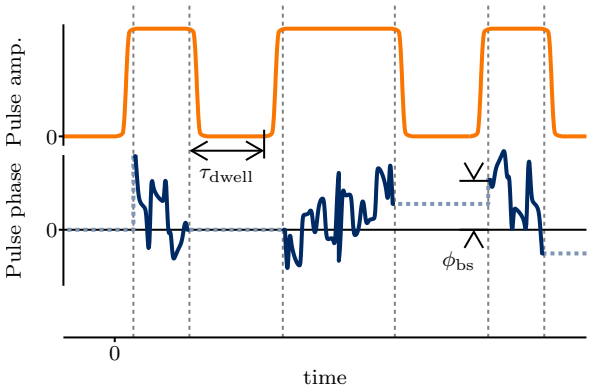

Figure 9.7: Illustration, not to scale, of an optimised interferometer sequence. Three pulses, as depicted in Figure 9.3, are combined, with a time τ_{dwell} between the light of one pulse turning off and the subsequent pulse turning on. Fringes are measured by measuring the fraction of atoms in the excited state at the end of the sequence as a function of a phase offset ϕ_{bs} applied to the phase sequence for the final beamsplitter pulse.

detailed by Luy *et al.* [173]. We thus optimise a UR mirror pulse in this fashion, resulting in the antisymmetric phase profile shown in the bottom panel of Figure 9.6.[4]

[4]Maximising $\mathcal{F}_{\mathrm{b}}$ with a phase profile constructed from antisymmetric basis functions would result in a pulse with the same properties.

The third pulse $\hat{\Omega}_{\mathrm{c}}$ should accurately map the inertial phase Φ imprinted on the superposition phase over the duration of the interferometer onto the probability of finding an atom in the excited state $|2\rangle$. We choose to construct it by taking the phase profile of the first pulse $\phi_{\mathrm{a}}(t)$, reversing it in time and inverting the phase profile to obtain $\phi_{\mathrm{c}}(t) = -\phi_{\mathrm{a}}(\tau_{\mathrm{a}} - t)$ where τ_{a} is the duration of the first pulse.

This *'flip-reverse'* operation results in a pulse that satisfies $|\langle 2|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle|^2 = 1/2$ and also gives the entire interferometer a temporal antisymmetry. Due to the symmetry properties of spin-1/2 propagators [176, 177], this fixes $\mathrm{Arg}\,\langle 2|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle + \mathrm{Arg}\,\langle 1|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle = -\,\mathrm{Arg}\,\langle 2|\,\hat{\Omega}_{\mathrm{a}}\,|1\rangle + \mathrm{Arg}\,\langle 1|\,\hat{\Omega}_{\mathrm{c}}\,|1\rangle + \pi$ and thus ensures a constant $\Delta\phi$.

9.3.2 Preliminary experimental results

An optimal Mach-Zehnder sequence comprised of pulses combined as outlined in the previous section is capable of maintaining high contrast despite significant variations in detuning and Rabi frequency in the atomic cloud. We have started to test these three-pulse interferometer sequences experimentally by performing multiple pulses sequentially, separated by periods of free evolution

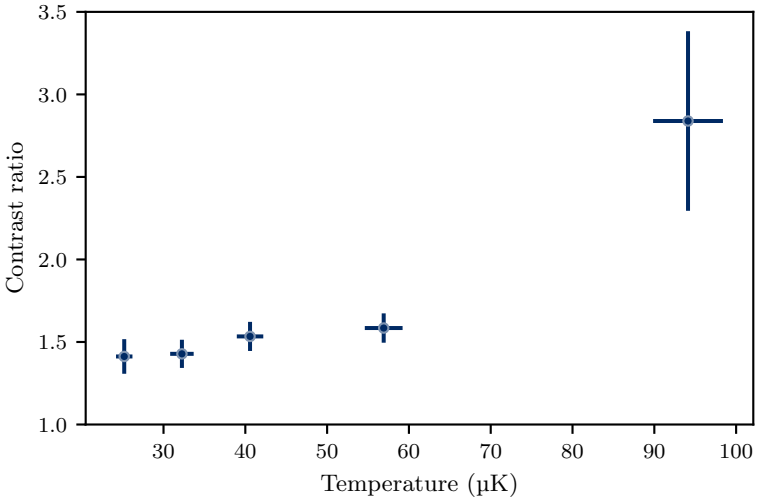

Figure 9.8: Ratio of contrasts obtained by fitting sinusoidal functions to fringes obtained from optimised 'flip-reversed' GRAPE and rectangular interferometer sequences for a range of cloud temperatures. GRAPE consistently improves the interferometer contrast with a significant 2.8(6) times improvement at the highest cloud temperature of 94(4) µK where temperature and not laser phase noise is the dominant source of contrast loss.

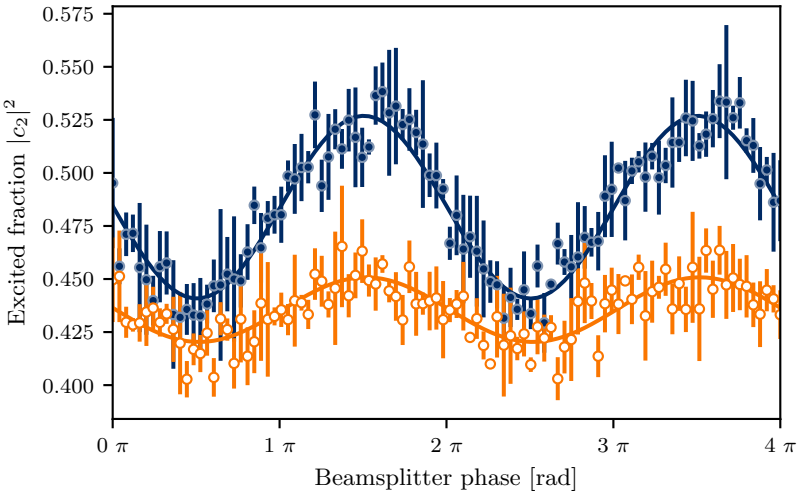

Figure 9.9: Interferometer fringes obtained at 94(4) µK for rectangular pulses (empty circles) and the optimised GRAPE sequence (filled circles). The GRAPE sequence improved the contrast of the fringes by a factor of 2.8(6). The average or, 'effective' Rabi frequency was approximately 420 kHz, and was determined empirically from the optimal duration of a rectangular π-pulse. We attribute the slight deviation of the fringes from a sinusoidal form to a small non-linearity in the response of the I&Q modulator.

τ_{dwell}, as illustrated in Figure 9.7. We measure the interferometer contrast by varying a phase offset ϕ_{bs} applied to the phase sequence of the final beamsplitter pulse between 0 to 4π and fitting a sinusoidal function to the resulting fringes in $|c_2|^2$.

While the dwell time between pulses is limited, at present, to $\tau_{\text{dwell}} \lesssim 100\,\mu\text{s}$ by decoherence that we attribute to transverse motion of the atoms across the irregular wavefronts of the square, top-hat beams manifesting as phase noise of the optical beatnote that differs from atom to atom and hence also limits the overall contrast,[5] initial results are promising. The relative improvement in contrast provided by an optimised sequence as the cloud temperature is varied is shown in Figure 9.8. GRAPE improved the contrast of the fringes at all temperatures investigated and, in a hot 94(4) µK sample where the contrast loss is dominated by atomic temperature and not the laser phase noise, nearly a threefold enhancement is observed, the fringes for which are shown in Figure 9.9.

[5] At these timescales the cloud will have fallen $\sim 1/4$ of a wavelength under gravity so the phase noise is likely dominated by the initial thermal motion of the atoms, complicated by the spatial extent of the cloud meaning that atoms with the same initial velocity in different parts of the cloud will likely experience different phase variations.

Some experimental evidence suggests that flip-reversed optimised sequences are less susceptible to drifts in offset and phase of the interferometer fringes. The variation of the fringe offset and phase from their respective means for a range of temperatures for GRAPE and rectangular interferometers are shown in Figures 9.10(b) and 9.10(c), and we hope to explore this aspect more systematically in future work.

In particular, there appears to be a systematic shift in the interferometer phase as the temperature is increased that we do not fully understand, but which could possibly be caused by an offset of the laser detuning from the centre of the atomic momentum distribution [146]. It is notable that the GRAPE interferometer appears less susceptible to this shift.

Another route of inquiry will be to explain why the contrast increase is quite so significant only when employing a fully optimised pulse sequence. While the mirror pulse, with its increased Doppler sensitivity, should be the dominant source of contrast loss in a Mach-Zehnder interferometer [2], only a slight enhancement was observed when just this pulse was replaced. To see significant improvement from a fully optimised sequence, maintaining the overall anti-symmetry in a flip-reversed configuration proved necessary. This is expected but, when this constraint was met, the contrast improvement far exceeded that of replacing just the mirror, or indeed the beamsplitters, in isolation.

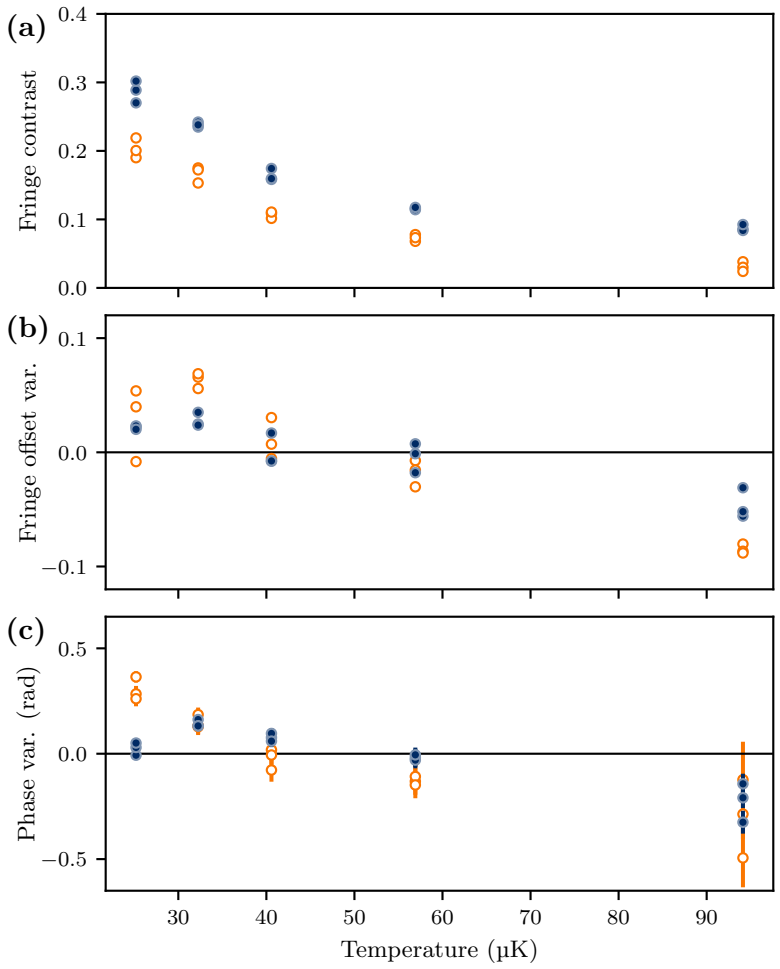

Figure 9.10: Variation in fitting parameters for sinusoidal functions fitted to fringes obtained from GRAPE optimised (filled circles) and rectangular (empty circles) interferometer sequences at a range of cloud temperatures. (a) shows the fringe contrast $\mathcal{B}$, while (b) and (c) show the shifts in the fringe offset $2\mathcal{A}$ and phase from their respective mean values. The GRAPE interferometer has consistently higher contrast than the rectangular counterpart, and exhibits less variation in both fringe offset and phase. Temperature error bars are omitted for clarity, but are the same as in Figure 9.8.

9.4 Conclusion

In this chapter we have used optimal control to design Raman pulses that achieve robust population inversion in the presence of variations in the atom–laser coupling strength and detuning, designed to improve the contrast of LMT interferometers. We have also outlined and demonstrated a design for an optimal pulse sequence that improves the contrast of the three-pulse Mach-Zehnder interferometer.

We expect such optimal pulse sequences to have applications

in improving the sensitivity and robustness of atom interferometric sensors, especially when operating in non-ideal environments, relaxing the requirement for low atomic temperatures and potentially reducing their susceptibility to drifts in the signal phase and offset.

Specifically, we have presented a point-to-point inversion pulse that achieves 99.8(3) % transfer between hyperfine ground states in a thermal cloud of ^{85}Rb atoms. Our BAT pulse has a broad spectral profile making it a good candidate for an augmentation pulse in LMT interferometry where a large velocity acceptance is required. The best results of Raman LMT interferometers to date have been achieved with adiabatic augmentation pulses [67, 85], but optimal control pulses such as this have the potential to realise similarly robust transfer with order of magnitude reductions in pulse duration [61].

Furthermore, our strategy for optimising three-pulse Mach-Zehnder type interferometer sequences, in which optimised beamsplitter and mirror pulses are combined in a 'flip-reversed' configuration to maximise the contrast of interferometer fringes where the phase of atomic superpositions is important, has shown up to a threefold contrast improvement in a proof-of-principle interferometer with hot 94(4) µK atoms, although our current investigations have been limited by experimental phase noise.

It should be noted that there are many open questions. For instance, it is interesting that the BAT pulse in Figure 9.2 has a symmetric form, despite no symmetry being enforced; how and why this symmetry emerges is being investigated, and systematic studies of how the phase profiles arrived at by GRAPE vary with respect to optimisation parameters such as pulse duration, assumed cloud temperature, smoothing penalties etc. will be included in the thesis of Jack Saywell and will hopefully offer insight into the function of different stages within a pulse and how best to utilise the algorithm.

We are also exploring alternative fidelity measures and ways to optimise all pulses concurrently within interferometer sequences, along with considering — both computationally and analytically — how to use the framework of optimal control to engineer robustness to other factors such as laser phase noise in atom interferometers.

Future experimental work will focus on apparatus upgrades that will allow us to extend our study of optimised three-pulse interferometer sequences to test their efficacy and robustness when experimental noise is no longer such a limiting constraint. With extended coherence times, we will also be in a position to experimentally demonstrate LMT interferometry with GRAPE pulses.

Part III

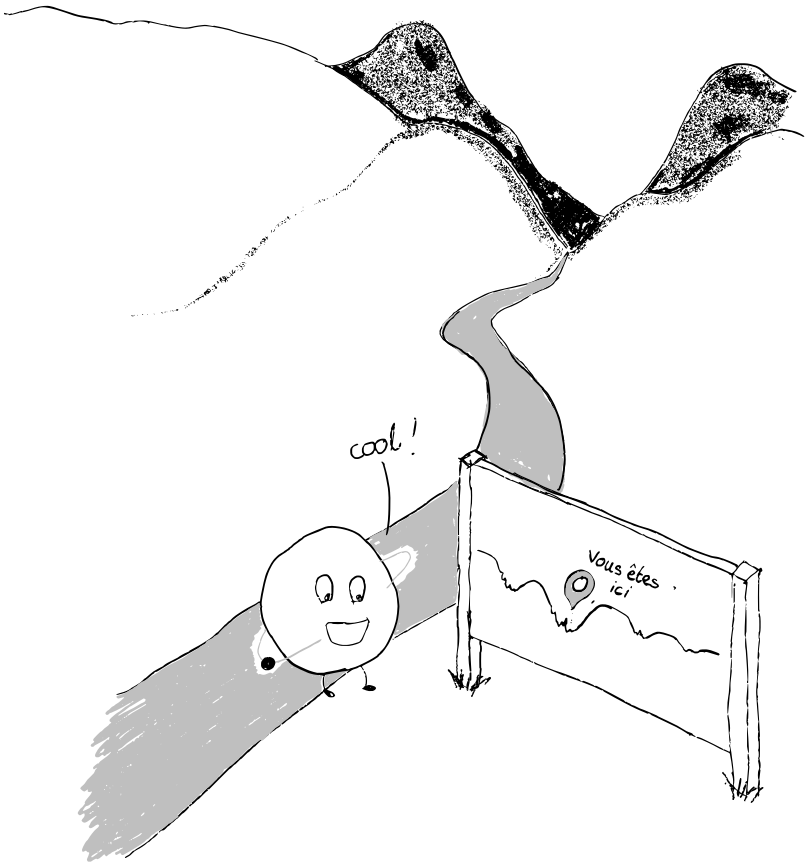

An Englishman in Bordeaux

Experimental setup 10

Everything's a cavity.

— The Front Bottoms, *Ginger*

The third year of my PhD was spent working on the BIARO experiment at the Institut d'Optique in Bordeaux, with BIARO standing for *condensation de Bose-Einstein et interférométrie atomique dans un résonateur optique de grande finesse*.[1] The experiment itself is excellently described and characterised in other theses that have come out of the group [178–180], and the modifications made during my time there will be thoroughly detailed in the PhD thesis of Hodei Eneriz in due course. I will give just a brief precis of the apparatus here; enough to lend context to the results acquired during my time on the experiment that are presented in the following chapter.

[1] Or, in English, Bose-Einstein condensation in an optical, high-finesse resonator. BOHR, perhaps?

Briefly, the experiment consists of a 'bow-tie' optical cavity in the main ultra-high vacuum (UHV) science chamber, comprising four identical mirrors arranged in a 63.6 mm square in the horizontal plane so that the beam crosses itself at the waist formed in the centre, as shown in Figure 10.1. Overlapping this waist, a 3-D MOT of ^{87}Rb atoms is loaded from a 2-D MOT that is separated from the main chamber by a narrow conductance channel, enabling a $\sim 10^3\times$ differential vacuum between the chambers.

In a configuration akin to that described in Section 3.5, a 'push' beam is aligned along the conductance channel in order to increase the atomic flux during loading. A high-resolution absorption imaging system — in which the atoms cast a 'shadow' in a low-intensity, resonant probe beam — is set up on the axis orthogonal to the plane of the cavity and the mirrors of the cavity itself have a dual coating so that, in principle, it can support both 1560 nm and 780 nm light.

The 1560 nm light can be used to create a far off-resonant optical dipole trap (FORT) potential for the ^{87}Rb atoms — operating on the principles established in Section 3.3.1 — with the theoretical finesse of $\mathcal{F}_{1560} \approx 4000$ allowing for high intra-cavity powers, and hence trap depths, to be reached for relatively modest input

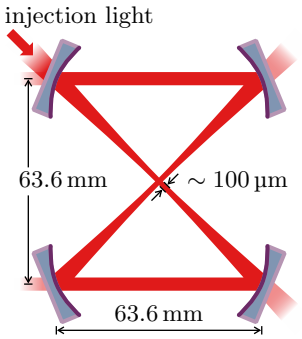

Figure 10.1: Illustration of the fundamental mode in the BIARO cavity. The off-axis reflections from the curved mirrors give rise to an astigmatism of the mode, with the waist being a factor of ~ 1.4 times larger in the axis directed into the page [178].

powers. As the cavity light intersects itself at the $\sim 100\,\mu$m waist, the trap is twice as deep here as anywhere else resulting in a 'dimple' at the centre in which atoms collect and are trapped along all directions. In practice $\sim 200\,$W of intra-cavity power can be obtained with the current experimental setup, resulting in a trap depth of $\sim 1\,$mK in the central dimple.

The cavity is designed to have a much higher theoretical finesse of $\mathcal{F}_{780} \approx 10^5$ for 780 nm light, originally intended for quantum nondemolition measurement-induced spin squeezing experiments [181–183], and making it suited to cavity QED experiments with strong light–atom interactions. This requires active stabilisation of the cavity length to be resonant with — or precisely detuned from — an atomic transition, and one of the mirrors is mounted on a piezoelectric stack for this purpose.

One key feature of the experiment is that, as the FORT laser operates at half the frequency of the rubidium D_2 line, it can be frequency doubled and stabilised to a rubidium reference. Stabilising the length to the 1560 nm light, the cavity resonance at 780 nm can then be maintained at a known detuning from a D_2 line transition without needing to continuously inject (near-)resonant light.[2] A 780 nm probe beam can then be turned on and off in the cavity only when desired, or else the atoms can be driven with interaction lasers orthogonal to the plane of the cavity with the only 780 nm photons to enter the cavity being those that the atoms scatter into it.

[2] In practice, a $\sim 70\,$MHz frequency shift needs to be introduced to the 1560 nm light, relative to the atomic reference, in order to account for differences in the effective optical length of the cavity at the two wavelengths.

The long lifetime of these photons in the cavity, combined with high atomic densities at the crossing region, result in a high likelihood of these photons interacting with the atoms again, mediating long-range interactions. Furthermore, the running wave nature of the bow-tie cavity then offers intriguing prospects for studying phenomena such as self-organisation, with the spontaneous emergence of an optical lattice in the cavity to which the atoms arrange themselves to be in-phase with [184–191].

For the bulk of my time on the experiment, however, the cavity was free-running and the injected light followed it, as described below.

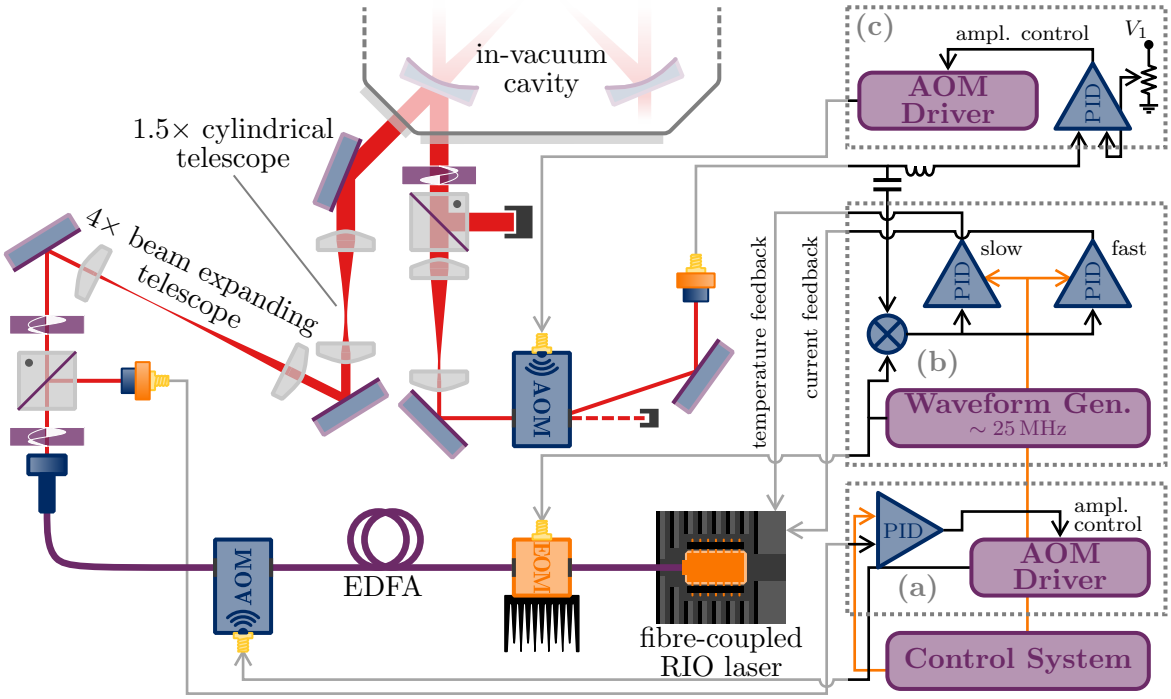

Figure 10.2: Schematic diagram of cavity injection setup, including control loops. Symbols correspond to those defined in Figure 6.3 except where labelled otherwise. The cavity is injected along one diagonal axis, and a power-stabilised fraction of the reflected light is aligned onto a photodiode in order to PDH lock the laser frequency to the cavity resonance. Control loop (a) stabilises and allows variation of the overall power, (b) stabilises the laser frequency and (c) stabilises the power incident on the PDH photodiode to ensure the frequency lock remains stable as the total optical power is varied.

10.1 Injection and locking of FORT light

The cavity is injected with 1560 nm light from a narrow linewidth ($\sim$ 5 kHz) *RIO Planex*™ laser diode amplified to 5 W with an erbium-doped fibre amplifier (EDFA). The 1560 nm light is mode matched to the fundamental cavity mode, first with a 4× telescope that enlarges the beam and then with a cylindrical telescope that further expands it in the horizontal axis to account for the astigmatism described in Figure 10.1. The lengths of the telescopes are adjusted in order to produce a beam that converges at the cavity waist. The laser is then Pound-Drever-Hall (PDH) locked to the cavity, feeding back to the current and temperature of the laser diode. The power of the injected light is controlled by a fibre AOM at the output of the amplifier.

As the power of the FORT light needs to be modulated throughout a typical experimental sequence, and shut off entirely in order to image the atoms in time of flight (TOF) measurements, this locking scheme is more complicated than this makes it sound. A schematic diagram of the injection stage, including the control electronics, is shown in Figure 10.2.

The power of the light emerging from the fibre is monitored by sampling a small fraction onto a photodiode with a PBSC. As shown in Figure 10.2(a), the signal from this photodiode is summed with a control voltage to provide an error signal that is used in a proportional-integral-differential (PID) loop controlling the RF power delivered to the AOM immediately following the EDFA. This both stabilises the optical power and allows it to be varied via the control voltage.

The laser is phase modulated at $\sim 25\,\mathrm{MHz}$ by an EOM prior to amplification in order to produce the sidebands for the PDH lock. In Figure 10.2(b) the 25 MHz signal is mixed with the RF output of a photodiode that monitors a small fraction of the reflected light after the rest is dumped.[3] The mixed signal is fed through low-frequency and high-frequency PID circuits that modulate the temperature and current of the laser diode respectively.

[3]Note that the running-wave cavity geometry means that the reflected light does not overlap the injected light.

Digital control signals are used to turn off the output of the PDH PIDs when the light in the cavity needs to be completely extinguished. The feedback loop needs to operate in a high gain mode in order to maintain a tight lock of the laser frequency to the cavity and so — to avoid immediate saturation when unlocking and relocking — separate controls are used to reduce or raise the gain and to turn the feedback on or off entirely. The staggering of these controls needs to be finely tuned so that the system is stable enough that it will mostly relock upon reengaging the PID circuits. This was later reinforced, however, by using the signal from a photodiode monitoring the cavity transmission in a comparator circuit, the output of which was combined with the digital control signal modulating the gain via a NAND gate in order to ensure that the tight lock was only reengaged when light emerged from the cavity[4]

[4]Prior to introducing the comparator and NAND gate, tens of experimental cycles could be conducted before the system would unlock. Subsequently this was increased to hundreds.

As the optical power can be varied by up to two orders of magnitude during each experimental cycle, a further control loop is required in order to ensure the stability of the PDH lock. This takes the dc signal from the PDH photodiode and stabilises its level to a control voltage by feeding back to a free-space AOM immediately before the photodiode, ensuring that the gain of the PDH feedback remains constant even as the total optical power is varied. This third loop is represented in Figure 10.2(c). Much trial and error was required in order to tune these control loops so that they acted cooperatively and did not interfere with each other.[5]

[5]In practice, the relocking described here is the result of months of work and a large amount of the time spent in Bordeaux. The cavity has a mechanical resonance at $\sim 90\,\mathrm{Hz}$, and is mechanically coupled to the MOT coils within the UHV chamber. When the coil current is quenched, the cavity oscillates dramatically, with the resonant frequency swinging over almost a GHz range. Getting the PDH lock to follow these oscillations and then unlock and relock during them was "non-trivial"; a full account of these technical challenges will form a sizeable portion of the PhD thesis of Hodei Eneriz, and here we will focus just on the results.

10.2 Cooling and imaging light

A 780 nm *New Focus* external-cavity diode laser (ECDL) is locked to the $|5S_{1/2}, F = 1\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition of ^{87}Rb via FM saturated absorption spectroscopy and serves as the 'repump' light for the MOT. A second *Sacher* DFB laser is frequency locked to the repump laser by stabilising the beatnote between them, registered on a fast photodiode, to $\sim$ 6.6 GHz. The error signal for the frequency lock has a linear slope over a $\sim$ 500 MHz range,[6] centred so that the span encompasses both the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ and $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transitions, and this is summed with a control voltage before feedback to the laser current controller so as to provide flexible control of the frequency. This laser provides both the cooling light for the MOT and the probe for the absorption imaging. The frequency of this laser can be controllably shifted across the $\sim$ 500 MHz range in $<$ 1 ms, which is particularly useful for realising the grey molasses scheme described in the following chapter.

[6] The beatnote is mixed with a stable oscillator at a similar frequency in order to translate it to a lower frequency that is further reduced by a factor of 2^9 by a frequency divider. This signal is then fed through an *Analog Devices AD650* frequency to voltage converter in order to generate this error signal.

The cooling and repump light are combined into a single PM fibre for delivery to the 3-D MOT, where their amplitudes are individually controllable via AOMs with an additional mechanical shutter in place to ensure that the MOT light can be completely extinguished when necessary.

They are also combined into a tapered amplifier (TA) for further amplification prior to coupling into another PM fibre for delivery to the 2-D MOT where they are separated to form both the MOT and push beams. The 2-D MOT light shares common AOMs with the 3-D MOT light prior to the TA, with an additional mechanical shutter after the TA and before it is fibre coupled. In practice this means that the 2-D MOT can be extinguished quickly in conjunction with the 3-D MOT, or else slowly in isolation.

Additionally, the cooling laser is separately coupled into a 'probe' fibre after passing through an AOM and mechanical shutter for precise amplitude and frequency control. This light is used for absorption imaging and could also provide a transverse pump for future experiments.

The delivery of the light to the experiment itself is illustrated in Figure 10.3, where a horizontal cross-section of the vacuum chamber is shown. The coils and light for the 2-D MOT are omitted from the diagram, but the cooling and repump light is

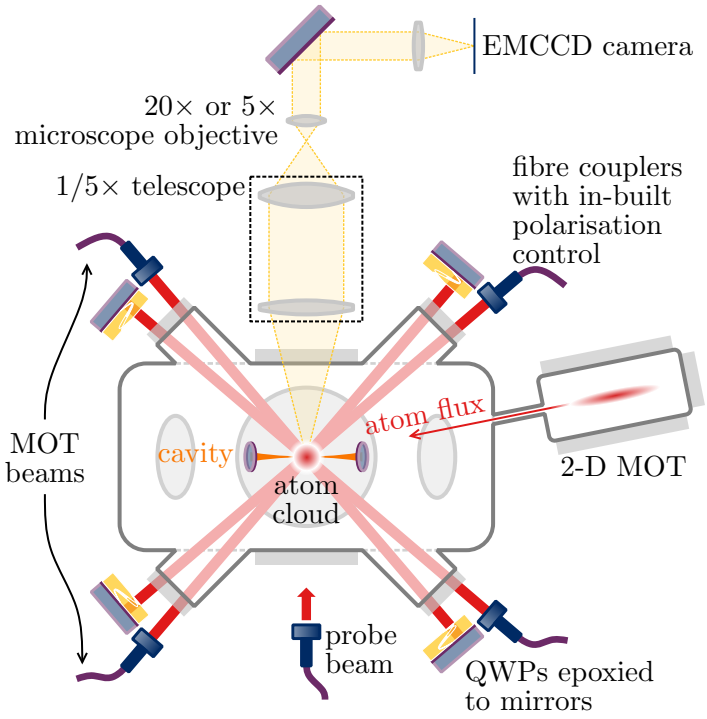

Figure 10.3: Vacuum chamber schematic, with symbols defined as per Figure 6.3 unless otherwise labelled. A 2-D MOT, optics not shown, loads a 3-D MOT in the main science chamber at the centre of the in-vacuum cavity. Additional 3-D MOT beams, along with in-vacuum coils, are aligned along the axis going into the page and are not shown. The remaining MOT beams are aligned in a hybrid retroreflected setup, illustrated here and described in the main text. A high-resolution imaging system images the atoms onto an EMCCD camera above the chamber in either fluorescence or absorption modes. For the latter case, a resonant probe beam is aligned along the vertical axis.

delivered vertically and into the page in a retroreflected configuration. The push beam is aligned to aid atomic flux along the conductance channel to the main science chamber as shown.

The 3-D MOT coils are also not shown, but are inside the vacuum chamber and have their axis aligned into the page. Helmholtz coils along all three axes, external to the chamber, allow fine positioning of the point of zero magnetic field so that the MOT can be positioned at the waist of the in-vacuum cavity.

Two fibres deliver counter-propagating MOT beams into and out of the page, and these are not retroreflected. The remaining four, off-axis, beams are shown, and these have a rather unique arrangement. While they are delivered from 4 separate fibres, they are all aligned slightly off-axis with mirrors to retroreflect each beam. This creates two pairs of nearly-colinear retroreflected setups that cross each other at the location of the MOT. This recycling of the light increases the optical intensity available for cooling, while the

near symmetry of the setup means that the force exerted on the atoms due to the 'shadow' effect — where the retroreflected light is weaker due to absorption by the atom cloud — is effectively cancelled. In practice a small but noticeable force is still present in the leftward direction in the figure, while the optical lattices created by the retroreflections somewhat complicate the theoretical analysis of grey molasses cooling in such a setup.

The high-resolution imaging system is shown above the chamber. This features interchangeable 5× or 20× microscope objectives and forms an image of the atoms on an electron-multiplying charecoupled device (EMCCD) camera. This image can be taken in fluorescence mode, by strobing the MOT beams during the exposure, or else in absorption mode by strobing the probe beam aligned along the imaging axis. This probe is resonant to the cooling transition and so only images atoms in the upper hyperfine manifold. A bright image of all of the atoms can be attained by illuminating with the repump light from the 3-D MOT while the probe is active to continuously repopulate the upper manifold.

Alternatively, a spatially resolved measurement of the fraction of atoms in the upper manifold can be made by imaging once with just the probe beam to get an image of the atoms in the upper manifold, then activating the repump beam for long enough to repump all of the atoms into the upper manifold, before finally taking a second 'normalisation' image with just the probe beam that will now resolve all of the atoms with the same brightness. The ratio of the first to the second image can then be used to determine how the population of the state manifolds varies spatially.

Grey molasses assisted loading of an optical dipole trap

11

"You're really not going to like it," observed Deep Thought.
"Tell us!"
"Alright," said Deep Thought. "The Answer to the Great Question..."
"Yes...!"
"Of Life, the Universe and Everything..." said Deep Thought.
"Yes...!"
"Is..." said Deep Thought, and paused.
"Yes...!"
"Is..."
"Yes...!!!...?"
"Forty-two," said Deep Thought, with infinite majesty and calm.
— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

The focus of this work is to efficiently load the far off-resonant optical dipole trap (FORT) potential formed by the 1560 nm light in the bow-tie cavity with ^{87}Rb atoms from the MOT. While it is attractive to use 1560 nm light for the FORT for the reasons outlined in the previous chapter, along with the readily available telecom lasers around this wavelength, loading such a trap with rubidium atoms poses a significant challenge.

The 1560 nm FORT light results in a light shift of both the $5S_{1/2}$ states and the $5P_{3/2}$ states that form the D_2 line. The $5S_{1/2}$ states receive a negative shift that is desirable for trapping and easy to understand: the FORT light is very far to the red of the 780 nm resonance. However, the $5P_{3/2}$ states also receive a negative shift, as they are coupled to higher-lying states via a 1529 nm transition [192]. The FORT light is much closer to this resonance, and so these excited states actually receive a light shift greater than that of the ground state by a factor of 42.6 [69, 193].

This differential shift means that MOT and optical molasses light that is red detuned from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition in free space will quickly become resonant or blue shifted within a deep FORT and cause heating. This can be seen in Figure

11.1, where the spatial dependence of the D_2 line resonances across the Gaussian profile of the FORT beam are shown. Furthermore, we see that the shift also has an m_F dependence, breaking the degeneracy of the excited state Zeeman sublevels.

Using a red detuned optical molasses to achieve a high atomic density from which to load, and then transferring those atoms into the conservative potential formed by such a FORT, is therefore not an inherently efficient process. Common solutions to this problem include slowly increasing the red detuning of an optical molasses while increasing the FORT power, or else employing additional beams to compensate the light shift [194] or an intermediate trap [195, 196].

Here it is demonstrated that, surprisingly, particularly effective loading can be achieved when MOT atoms are further cooled in an optical molasses that is blue detuned from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition. This cooling mechanism makes use of dark states in a manner similar to that described

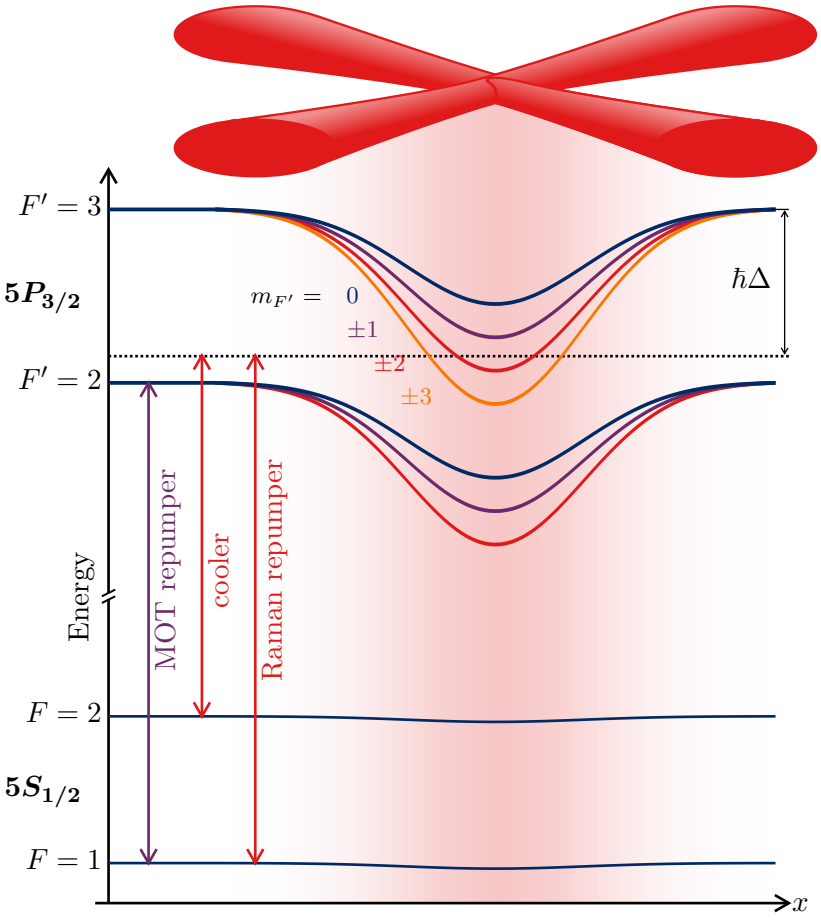

Figure 11.1: Illustration, not to scale, of the m_F dependent light shifts induced by the FORT light on the relevant levels of the ^{87}Rb D_2 line.

in Section 3.6.2, and is shown to be effective even in the presence of a deep FORT, negating the need for an adiabatic ramp between the cooling and trapping phases of the experiment.

11.1 Dark states on the F=2 → F′=2 transition

So called *Zeeman dark states (ZDSs)* exist on the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition that are dark to the cooling light over short timescales. However, these states are of the form [106]

$$|ZDS\rangle = \frac{1}{\sqrt{3}}\left(|-2, \boldsymbol{p} - \hbar\boldsymbol{k}\rangle - |0, \boldsymbol{p}\rangle + |2, \boldsymbol{p} + \hbar\boldsymbol{k}\rangle\right), \tag{11.1}$$

where $|m_F, \boldsymbol{\mathcal{P}}\rangle$ represent Zeeman sublevels m_F of the $|5S_{1/2}, F = 2\rangle$ state and the second index $\boldsymbol{\mathcal{P}}$ denotes the associated momentum, with $\boldsymbol{p}$ being the average atomic momentum and $\boldsymbol{k}$ being the effective wavevector formed by two counter-propagating cooling beams. Such a state is shown in Figure 11.2, where the contributing dipole matrix elements are also illustrated. It is clear that these states are not eigenstates of momentum and thus, even for $\boldsymbol{p} = 0$, they exhibit kinetic energy dephasing on timescales on the order of $^{2m}/_{\hbar k^2}$, associated with the two-photon recoil momentum. They do not, then, remain dark indefinitely, and are more accurately described as *grey states*. Polarisation gradient cooling mechanisms that involve such states are thus referred to as *grey molasses* [110].

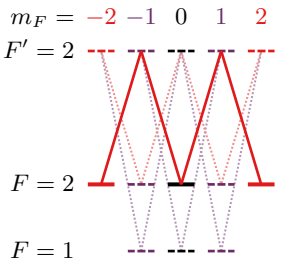

Figure 11.2: Illustration of a Zeeman dark state (ZDS) on the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition. Hyperfine states are denoted F for $5S_{1/2}$ states and F' for $5P_{3/2}$ states. Horizontal lines represent magnetic sublevels of each hyperfine state, and diagonal lines joining them represent the dipole matrix elements of the counter-propagating, circularly polarised laser beams that couple them. The states and dipole matrix elements contributing the the ZDS are represented by solid lines. While dark (over short timescales), a repump laser (whose dipole matrix elements are also shown) is required on the $|5S_{1/2}, F = 1\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition for effective cooling in a polarisation gradient where coupling to the excited level (and associated decay) occurs. This repumper, however, complicates the model and could be detrimental to the darkness of the ZDS [197].

Furthermore, this is an open transition, with dipole coupling from the $|5P_{3/2}, F' = 2\rangle$ state to $|5S_{1/2}, F = 1\rangle$. For polarisation gradient cooling mediated by such states to occur, then, a repump beam is required to return atoms to the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ cooling transition. The MOT repumper can serve this purpose, but the presence of this additional light should then be incorporated into the model and the 'darkness' of the ZDSs becomes less apparent, possibly to the detriment of the cooling [197].

An efficient grey molasses can still be realised, however, if a repump beam is introduced in a Raman condition with the cooling beam in what has been termed a Λ-enhanced configuration [198]. In such cases, a rich set of *hyperfine dark states (HDSs)* can exist

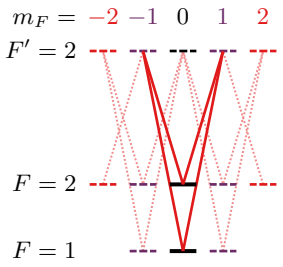

Figure 11.3: Level scheme, as per Figure 11.2, illustrating a hyperfine dark state (HDS) between the $F = 1, m_F = 0$ and $F = 2, m_F = 0$ levels that involves dipole matrix elements to the $F' = 2$ states which form a closed contour loop [199, 200] (phase coherence between the cooling and repump matrix elements is illustrated by employing the same colour). The phase matching condition that ensures such a closed loop varies spatially on the microwave wavelength associated with the ground state hyperfine splitting. Two other such dark states can form, involving the $m_F = \pm 1$ levels; it is worth noting, however, that only one such state would exist if the $F' = 1$ excited state were used.

that are comprised of a mix of Zeeman sublevels from both the $|5S_{1/2}, F = 1\rangle$ and $|5S_{1/2}, F = 2\rangle$ states.

Well studied in the contexts of coherent population trapping and electromagnetically induced transparency [201, 202], the set of potential HDSs becomes particularly rich when the Raman repumper has a fixed phase relationship with the cooling beam, allowing dark states to form in which the dipole transition elements form *closed contour loops* in which the phases sum to an integer multiple of 2π [199, 200]. An example of such a state, again with the contributing dipole matrix elements highlighted, is shown in Figure 11.3.

Other such HDSs — without such stringent requirements on the relative laser phases — are reviewed and categorised in [109], but a description of the precise states produced in our system is beyond the scope of this chapter. It suffices to say that we realise the Raman repumper by employing an EOM to put 6.834 GHz sidebands on the frequency-tunable DFB laser, making it phase-coherent and maximising the available HDSs. The output of the frequency synthesiser driving this EOM is controlled by an RF switch to allow fast switching of the sidebands on and off, so that we can switch between the MOT repumper — locked to the $|5S_{1/2}, F = 1\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition — and our Raman repumper during an experimental cycle.

11.2 HDS molasses optimisation

To realise and see evidence of an HDS molasses, first a small number of atoms are loaded into the 3-D MOT — in ~ 100 ms — so as not to saturate the absorption imaging. A short ~ 10 ms 'compressed MOT' phase, common to all our experiments, then occurs in which the 2-D MOT light and push beam are extinguished, and the cooling beam is detuned $\sim 6\Gamma$ to the red of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ MOT cooling transition, where $\Gamma = 6.066$ MHz is the natural linewidth of the D_2 line. This allows for some sub-Doppler Sisyphus cooling to occur while the magnetic quadrupole field for the MOT remains on. The quadrupole field is then turned off within 3 ms, during which the MOT repumper is turned off, the cooling laser is further detuned $\Delta \approx 42\Gamma$ — or $\sim 3\Gamma$ to the blue of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition — and the coherent repumper is turned on with the RF switch driving the EOM.

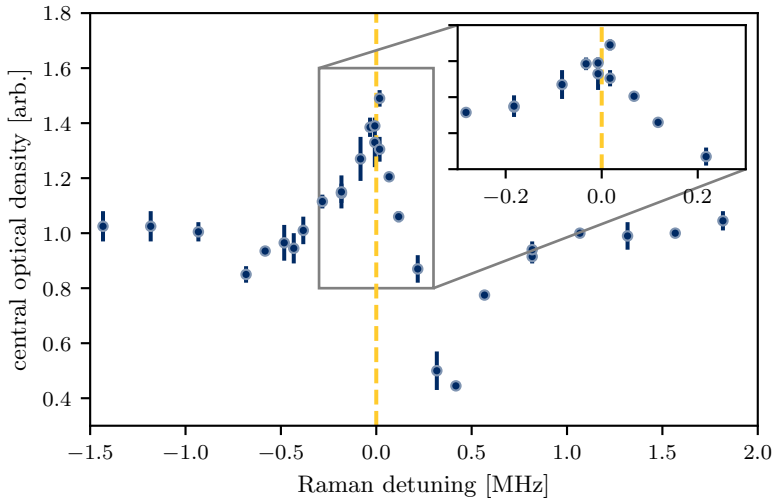

Figure 11.4: Optical density of molasses, imaged in absorption after 10 ms TOF, as a function of Raman detuning δ. The density is highest when $\delta \approx 0$, approaching the hyperfine splitting of the ^{87}Rb ground states. A minimum in density, indicative of heating, is evident for $\delta > 0$, a known characteristic of Λ-enhanced grey molasses.

At this point, the atoms should be cooled in the HDS molasses, where they remain for between 5 and 60 ms.[1] In order to see the effect of this molasses, the cooling light is extinguished and the cloud is imaged in absorption after 10 ms time of flight (TOF). The cooling laser is repurposed for the imaging so, during the TOF, the EOM is turned off and the frequency is shifted to be resonant with the cooling transition. The MOT repumper is illuminated 500 µs prior to imaging and remains so while the imaging probe beam is on to ensure the atoms are pumped into the $|5S_{1/2}, F = 2\rangle$ state and absorb the probe light.

[1]The DFB laser was a recent addition to the experiment, replacing a laser whose frequency could not be moved across the required range in less than 20 ms, setting a lower limit on the molasses time that was overly high. The vast majority of the data presented here were collected with the DFB laser, with molasses times < 10 ms.

We look for evidence of the Raman resonance by scanning the EOM frequency and looking for an effect in these images. The effect is most striking when plotting the resonant optical density (OD) after 10 ms as a function of δ, the detuning of the EOM frequency from the ^{87}Rb ground state hyperfine splitting of 6.834 68 GHz, as shown in Figure 11.4. A clear peak, indicating the coldest sample, is evident about $\delta = 0$. For $\delta > 0$ a heating effect is observed, characteristic of Λ-enhanced grey molasses [198].

Upon finding the Raman peak, the grey molasses proves an excellent magnetometer with which to zero the magnetic field across the atomic sample. The currents through bias coils along all three Cartesian axes are optimised to achieve the highest optical density. Once all three sets of coils are optimised, atoms can remain within the molasses beams for over 2 s after the magnetic trapping potential has been turned off.

11.3 Optimal FORT loading

The optimisation in Section 11.2 was performed without injecting 1560 nm light into the cavity. In practice, to produce the FORT potential, a small amount of 1560 nm light must always be injected in order to keep it locked to the cavity. It is reasonable to expect the differential light shift induced by this light to affect the efficacy of the HDS molasses, even at low power. It is not immediately obvious what the optimum FORT power for loading should be.

Here we seek the optimum loading configuration with a small amount of FORT light, of power P_{low}, maintained in the cavity throughout the loading sequence.[2] When the molasses light is extinguished, the cavity light is increased to a higher power P_{high}, controlled by adjusting the RF power to the AOM at the output of the EDFA, in order to raise the trapping potential for greater confinement.

[2] An even lower power, the minimum required to sustain a lock, is injected while the 3-D MOT is loaded.

The procedure for loading the FORT starts in the manner described in Section 11.2, except that the MOT is loaded for 10 s in order to load from a large sample. Upon extinguishing the cooling light, the trapping light is raised from P_{low} to P_{high} in a sigmoid curve, reaching P_{high} after 2 ms. In these experiments, the atoms are then held in the FORT for 0.5 s before the light is extinguished, by briefly turning off the PDH feedback and the AOM, and they are imaged after 2 ms TOF. The relative number of atoms loaded is then calculated by doing a Gaussian fit to a narrow strip of the absorption image across the centre of the cloud.

The values of different experimental parameters were varied in order to maximise the number of atoms measured in the trap after this period. These include the detuning Δ of the cooling light, measured from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ MOT cooling transition, and that of the EOM frequency δ, measured from the hyperfine splitting $\omega_{\text{hf}} = 2\pi \times 6.834\,68$ GHz. Interestingly, whereas the peak optical density in the absence of the FORT light was found to occur at $\delta = 0$, a small but resolvable ~ 10 kHz shift was found to maximise the number of atoms loaded into the trap, as shown in Figure 11.5. The precise magnitude of this shift exhibited a dependence on the trap depth P_{low} during the loading phase, despite the fact that the associated light shift should not alter the hyperfine splitting. An explanation of this is offered in Section 11.4.

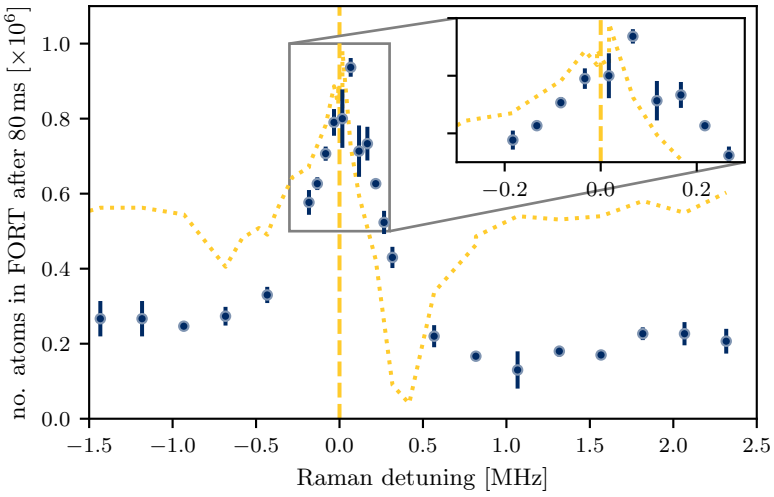

Figure 11.5: Relative number of atoms loaded into the FORT as a function of the Raman detuning δ. The form of the optical density measurements in the absence of the FORT light is shown as a light dotted line for reference. Peak loading appears to occur at a Raman detuning that is shifted from the hyperfine splitting on the order of $\sim 10\,\mathrm{kHz}$, and we verified that this shift exhibits a dependence on the FORT power P_{low}.

The powers P_{low} and P_{high} determining the FORT depth during the loading and confinement phases, and that of the RF driving the EOM to generate the Raman repumper, were also optimised. A surprisingly high value of P_{low}, corresponding to a trap depth $\sim 200\,\mu\mathrm{K}$, was found to optimise the loading. Indeed, while ramping the power up to a P_{high} value corresponding to a trap depth $\sim 1\,\mathrm{mK}$ after the cooling light is extinguished results in a greater lifetime of atoms in the trap, it is not crucial to the actual loading.

In contrast to other loading schemes, in which the conservative potential is raised adiabatically in order to properly phase match the loading between the cooling and trapping phases of the experiment, we observe optimal loading with a deep FORT present throughout the cooling process. Atoms lose kinetic energy by scattering photons, enabling them to collect in the conservative potential; this conserves atom number, and constitutes a significant advantage over loading schemes that rely on three-body collisions to dissipate energy.

It is conventional to load an optical dipole trap directly from Sisyphus cooled atoms, with a cooling beam red detuned from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition. For large red detuning Δ, however, it is more accurate to describe the laser as blue detuned from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition and we would expect grey molasses cooling processes to

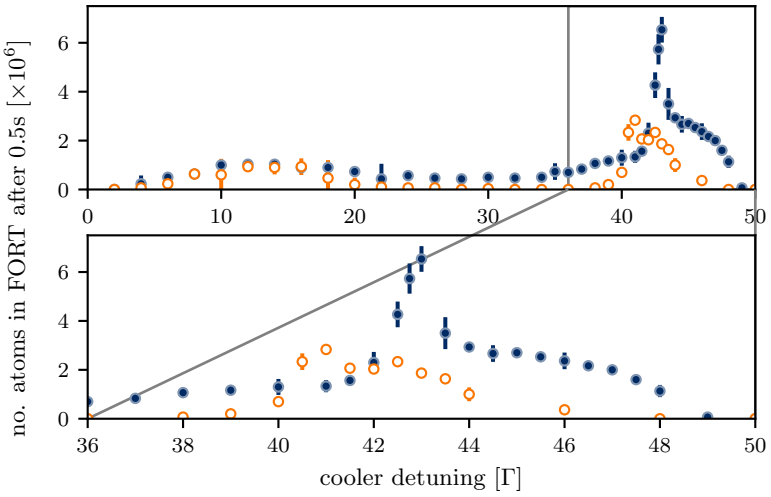

Figure 11.6: Number of atoms in the FORT after 0.5 s as a function of the cooling laser detuning Δ from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition in optimal conditions. Blue (filled) and orange (empty) data points correspond to Raman and MOT repumpers respectively.

dominate. Having identified the optimal loading conditions for the HDS molasses, we thus measure the number of confined atoms over a large range of Δ that encompasses both of these regimes. We perform these measurements using first the MOT repumper — locked on resonance to the $|5S_{1/2}, F = 1\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition — with which ZDS cooling can occur, and then with the Raman repumper that enables the formation of HDSs.

Data for these measurements are shown in Figure 11.6, with blue (open) circles corresponding to the Raman repumper and orange (filled) circles corresponding to the MOT repumper. In both cases a broad peak in loading is observed around $\Delta \approx 14\Gamma$; the atoms are being Sisyphus cooled on the cycling transition, the conventional condition in which one would choose to load a dipole trap. For both repumpers, a reduction in loading is observed for $\Delta > 18\Gamma$, with no loading observed with the MOT repumper for $\Delta > 22\Gamma$.

However, as the detuning approaches the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ resonance — equivalent to $\Delta \approx 44.5\Gamma$ — from the blue side, a revival in loading is observed with both repumpers.[3] Both are the result of grey molasses cooling: for the MOT repumper, cooling in a ZDS molasses results in a broad loading peak around $\Delta \approx 42\Gamma$ that diminishes as the detuning is increased to the red of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ resonance and results in heating; for the Raman repumper, a much more intense loading peak is observed at $\Delta \approx 43\Gamma$, associated with much more efficient cooling mediated by truly dark HDSs.[4]

[3] This revival was observed with the MOT repumper in the early days of the BIARO experiment and is where atoms were historically loaded from, but it was never fully explained.

[4] There is an intrinsic uncertainty in absolute values of Δ as the repump laser, the frequency reference for the adjustable laser, drifts slightly from day to day on the order of 0.5Γ.

The difference in detuning between the ZDS and HDS loading peaks is in agreement with previous work, conducted by the authors of [198], in which peaks were observed in the central atomic density of a similarly cooled sample of ^{87}Rb atoms in the absence of any additional trapping potential. The truly dark HDSs, centred about $\boldsymbol{p} = 0$, allow atoms to accumulate in low kinetic energy states while tuning much closer to the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition, resulting in higher atomic densities and much more efficient loading. The HDS loading peak corresponds to a factor of 7 increase in the number of confined atoms when compared to loading from a Sisyphus molasses. What is initially surprising is that this peak sits atop a broader background feature that does not immediately diminish for detunings $\Delta > 44.5\Gamma$ to the red of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition where we would expect heating to occur.

11.4 Effect of FORT-induced light shifts on the formation of HDSs

We would expect the differential shift to the $|5P_{3/2}, F' = 2\rangle$ state m_F levels induced by the FORT light, illustrated in Figure 11.1, to cause dark states to rapidly dephase within the FORT beam. However, at least some of the HDSs appear resilient to these differential shifts, and we can attribute the broader feature in Figure 11.6 to atoms being loaded from HDS-mediated cooling occurring within the trap volume. While for $\Delta > 44.5\Gamma$ atoms outside of the trapping region are heated, atoms subject to the differential light shifts within the trapping beam are still cooled.

To demonstrate this explicitly, we retake the HDS loading data about the peak at $\Delta \approx 43\Gamma$ for a range of trap depths P_{low}. These data are plotted in Figure 11.7(a), where we observe a shift and a broadening of the background feature with increasing power of the FORT light, while the position of the sharp peak at $\Delta \approx 43\Gamma$ remains constant. Our attribution of the broader peak to atoms being loaded from within the trapping volume of the FORT appears confirmed; increasing the FORT depth increases the range of optimal detunings Δ_{opt} for cooling/loading atoms across the spatial profile of the trapping beam.

We conclude that, by changing the detuning, we can load atoms from different parts of the spatial atomic distribution overlapping

with the FORT. The shape of the broader loading curve is a convolution of the atomic density distribution and the spatial profile of the FORT: the deeper the FORT, the broader the loading curve. The sharpness of the peak at $\Delta \approx 43\Gamma$ occurs due to a competition between the enhanced atomic densities outside the FORT and the effect of these light shifts on atoms approaching the FORT frontier. Its position does not change with the FORT power as it ultimately arises from a volumetric effect associated with the peak in atomic density.

That HDS-mediated cooling can occur while subject to a breaking of the m_F degeneracy of the excited state can be explained by the types of HDSs that can form [109, 199]. Some form that only rely on dipole matrix elements to symmetric pairs of excited Zeeman sublevels $m_{F'} = \pm 1$ and $m_{F'} = \pm 2$ that experience the same light shift and so do not dephase; the closed loop dark state illustrated in Figure 11.3 is just one example of these.[5] Those formed from other combinations will experience differential shifts, introducing extra phase factors in the Rabi frequencies of the contributing transitions so that they no longer cancel. This effect will be more severe the deeper an atom is in the trap, but can be partially counteracted by varying the Raman detuning δ. This explains why, when optimising the HDS molasses loading, the optimum Raman detuning appeared to depend on the FORT depth.

[5] It should be noted that this is also true of the ZDSs formed with the MOT repumper. As this repumper is kept at a fixed frequency, however, it becomes increasingly ineffective within the trap where atoms then decay to the $|5S_{1/2}, F = 1\rangle$ level.

This rather speculative explanation is backed up by further data, soon to be published [203]. Evidence of the types of HDSs that are formed at different trap depths is presented in the form of data for the fraction of atoms measured in the $|5S_{1/2}, F = 2\rangle$ state after the cooling. In addition, loading was performed from a HDS molasses blue detuned from the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 1\rangle$ transition, and the dependence of the optimum Raman detuning on trap depth was shown to vanish.

This latter point could be an initial indication that the HDSs that prove resilient to the FORT-induced light shifts are of the close contour loop nature: when just the $|5P_{3/2}, F' = 1\rangle$ excited level is considered, only one such dark state can occur and it involves a symmetric pair of excited Zeeman levels $m_{F'} = \pm 1$. As a counterpoint to this, however, it should be noted that the likelihood of the phase matching condition for such a loop being realised in our experiment by chance is relatively low and, in reality, considering dark states to involve only one excited level in isolation is a questionable assumption, even near resonance. In practice the neighbouring excited levels will always have a role to play [204].

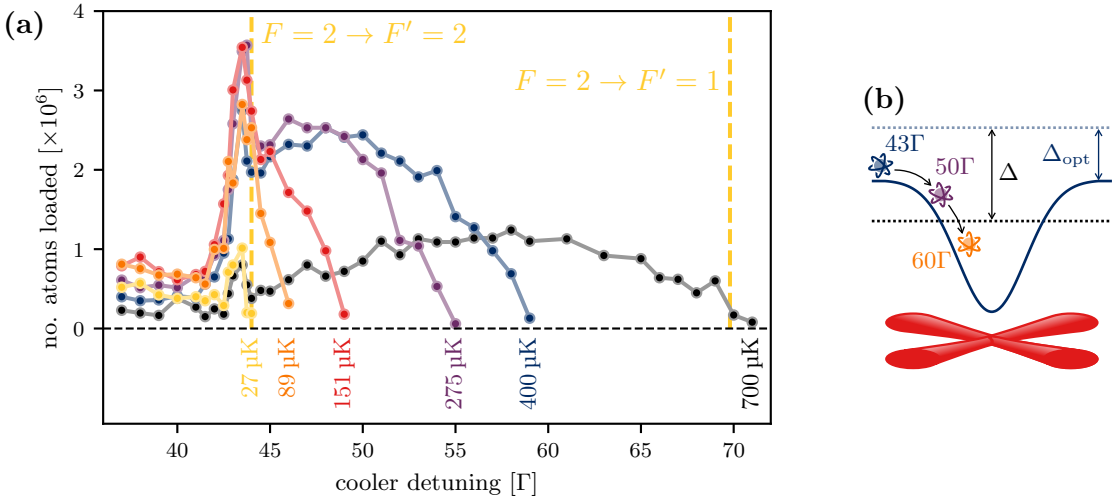

Figure 11.7: (a) Number of atoms loaded into the optical dipole trap as a function of cooler detuning for a variety of FORT trap depths. A loading peak to the blue of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ resonance remains fixed, on top of a background feature that broadens with increasing trap depth. This feature can be attributed to the loading of atoms from within the trap volume, with the broadening occurring due to the light shift induced selectivity illustrated in (b). The optimum cooler detuning for the HDS-mediated cooling Δ_{opt} varies across the spatial profile of the trapping beam — so that only atoms at particular locations within the trap are optimally cooled — and this selectivity increases with increasing trap depth.

These data, however, were collected after my time on the experiment and — while investigations into the precise nature of the HDS that mediate the effective loading are ongoing — they are beyond the scope of this chapter. Instead, we explore the possibilities for further cooling atoms within the conservative potential, subsequent to loading them, that the existence of these resilient HDSs naturally suggests.

11.5 Velocity-tuned HDS cooling

The ability to selectively cool atoms at different locations within the trap corresponds to an ability to address different momentum classes once atoms are loaded. In a conservative trap, the average displacement of an atom depends on its momentum, with higher momentum states spending more time near the edges — where the light shifts are small — and lower momentum states accumulating near the centre — where the light shifts are large. As illustrated in Figure 11.7(b), the optimum detuning Δ_{opt} for HDS-mediated cooling will depend upon the atomic momentum; the deeper the trap, the greater the selectivity.

This suggests the possibility of cooling atoms once they have been loaded into the trap: the detuning of the cooling light can be

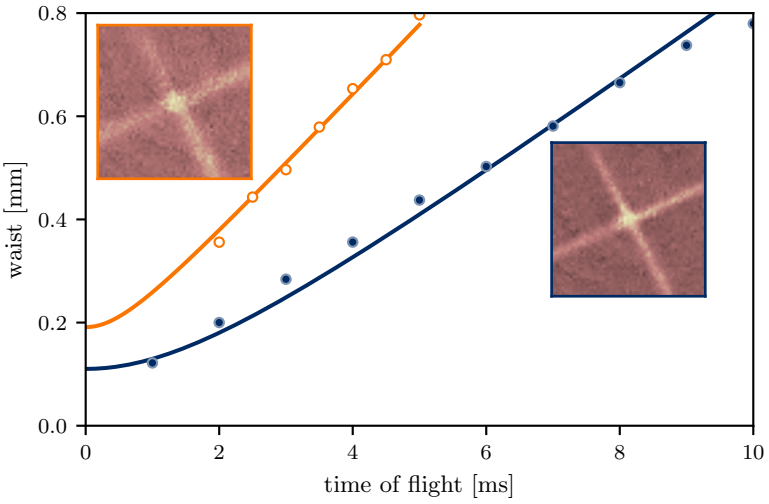

Figure 11.8: Expansion of atoms released from FORT over 10 ms TOF with and without a cooling sweep, represented by blue (filled) and orange (empty) data points respectively. The cooling sweep results in a factor of 4 reduction in the sample temperature, from 198 µK to 48 µK, as inferred from best fit theory curves (solid lines). The inset images show absorption images of the corresponding samples after 700 µs TOF; the central crossing region of both images contains $\sim 6 \times 10^6$ atoms yet, after cooling, the atoms are more tightly confined to the centre of the trap, both in the crossing region and in the wings.

ramped from a detuning that is optimal for atoms near the edges to a detuning that cools atoms at the centre of the trap. If this ramp is much slower than the photon scattering rate then — at each point of the ramp — all of the atoms being addressed will be cooled and fall deeper into the trap; atoms, then, will accumulate in momentum states below the level determined by the laser detuning. There is an analogy to be made here with the *RF knife* method employed in magnetic traps [70, 71] except that, while such a 'knife' eliminates atoms above a certain momentum from the trap and then allows the sample to rethermalise, our detuning ramp pushes atoms deeper into the trap and conserves atom number.

A preliminary demonstration of this in-trap cooling is presented in Figure 11.8. After the initial loading phase, the power of the FORT beam is adiabatically increased to P_{high} corresponding to a trap depth ~ 1 mK. This increases the momentum selectivity of the HDS-mediated cooling and results in an effective temperature of 198 µK for the trapped atoms, as inferred from the TOF data represented in Figure 11.8 by orange (empty) circles. Turning on the Raman cooling beams once the trap is raised, and sweeping their frequency from $\Delta = 55\Gamma$ to $\Delta = 68\Gamma$ in 6 ms, we observe the TOF data represented by blue (filled) circles that demonstrate

a cooling of the trapped sample to 48 µK. The atom number is conserved, with 6×10^6 atoms remaining in the central crossing region of the trap until the laser is detuned beyond $\Delta = 68\Gamma$; this takes the laser across the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ resonance at the centre of the trap and the atoms are all quickly lost due to the associated heating.

This initial demonstration shows a fourfold reduction in temperature of the trapped atoms. Ideally, the trap would be deepened in order to increase the momentum selectivity further, but here we are limited by the nature of the D_2 line currently used for the cooling. As shown in Figure 11.1, when the Raman lasers are tuned just to the blue of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 2\rangle$ transition in order to optimally cool atoms at the edge of the trap then — in a deep trap — the spatially dependent light shift means that low momentum atoms near the centre are close to resonance with the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ transition. If the trap is shallow enough then they will remain on the red side of this transition and, while the light shift may impede the formation of dark states [106], it is not overly problematic; however, if the trap is deep enough that the light shift results in the central atoms crossing to the blue of the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ resonance then they will be heated while the atoms nearer the edges are cooled.

This competition between cooling and heating mechanisms fundamentally limits the effectiveness of the cooling in our current experiment, but could be overcome by switching to the D_1 line and performing HDS-mediated cooling on the $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{1/2}, F' = 2\rangle$ transition. In this case the heating associated with higher energy excited states is no longer a factor and it should be possible to cool to much lower temperatures in millisecond timescales without losing atoms.

11.6 Conclusions

This work demonstrates a new type of all-optical cooling within a conservative optical trap, making use of hyperfine dark states present in a Raman configuration [109, 198]. The method is not reliant on closed, cycling transitions, but rather a diffusion in momentum space that drives atoms towards dark states that exhibit minimal excitation and scattering.

This elimination of the excited state, in conjunction with the fast rate of the cooling and conservation of atom number, could open

new avenues in the production of ultracold atomic and molecular gases; as the cooling and loading occur simultaneously, the issue of mode matching a cold sample to the conservative potential is bypassed with no lengths intermediate steps required. The technique is thus particularly suited to the rapid production of ultracold, trapped, samples in unusual geometries [205] and environments [30, 116].

Discussion and outlook

12

Just wanted to let you know
Before I say
Au revoir...

— The Front Bottoms, *Au revoir (adios)*

To bring this work to a close, we here reflect on the results of this and the previous Parts in the context of other work and consider their potential impact.

12.1 Velocity-sensitive atom interferometers

The Ramsey-type atom interferometers discussed in Chapter 7 are often overlooked, or else combined into more complex interferometry sequences, due to their 'unclosed' nature. However — just as applications have been found for unclosed optical interferometers in, for example, shear plate in interferometry [206] — with these velocimetry results we have demonstrated one application of the velocity-dependent phase inherent to the paths not closing.

The decay of coherence with increasing wavepacket separation at the output of a matter-wave interferometer is well understood [122, 207]; measuring it is a good way to quantify the coherence length of the atomic wavepacket and thus derive an effective ensemble temperature [147–149]. We are the first to to make a full measurement of the atomic velocity distribution by varying the wavepacket separation and employing Fourier transform spectroscopy techniques [142, 143].

This decoherence arises from the dephasing of velocity-dependent fringes in the time domain. The corollary of these fringes is a set of time-dependent fringes in velocity space, related by a Fourier transform and on which our velocimetry technique ultimately relies. Weitz and Hänsch identified that the pulse sequence of a Ramsey interferometer can be tuned to vary the periodicity and phase of these fringes, that ultimately map onto the probability of finding an atom with a particular velocity in the excited state at the output of the interferometer [43]. As this excitation

has an associated change in atomic momentum, the interferometer thus effects a tunable and velocity-dependent impulse.

Weitz and Hänsch stressed the applicability of this impulse to the cooling of atoms and molecules, but this was only recently verified experimentally by Dunning *et al.* [44]. Dunning's results and the velocimetry of Chapter 7 can be considered as two sides of the same coin, as explored in a poster presented at ICOLS 2019, Queenstown, and reproduced in Figure C.2. These closely related works both represent promising and under explored applications of velocity-sensitive matter-wave interferometers.

An interesting recent study by Kritsotakis *et al.* suggests that a velocity-sensitive interferometer may, in fact, constitute the optimal way to perform an atom interferometric gravimetry measurement [45]. Quantum Fisher information is maximised by inferring the gravitational acceleration from a momentum resolving measurement at the output of a Ramsey interferometer, rather an electronic state or position resolving measurement at the output of a Mach-Zehnder interferometer. Experimentally, the challenge then becomes how to best perform this momentum measurement in order to take advantage of this benefit.[1]

[1] It should not escape our notice that, having demonstrated the momentum resolving powers of the Ramsey interferometer, we now consider the quandary of how best to measure fringes in momentum space at the output of a Ramsey interferometer in order to make a gravimetry measurement. We verge perilously close to reinventing the Mach-Zehnder interferometer, but with significantly more questions than when we started.

Velocity-sensitive interferometers have already formed the basis of another new gravimetry method. Published even more recently by D'Amico *et al.* [46] — subsequent to the publication of our velocimetry technique — two Ramsey interferometers are used to measure and compare the velocities of atom clouds travelling in opposite directions in order to determine their acceleration due to gravity. The difference in free-fall time is large with respect to the actual interrogation time of each interferometer, allowing a sensitive measurement that is resilient to seismic vibration of the optics.

When it is studied, then, the velocity-sensitive interferometer has a tendency to yield new applications — or else new insights and approaches to old problems — and the work in this thesis is just one of several recent demonstrations of this [45, 46].

12.2 Optimal control in atom interferometry

The work in Chapter 9 is an initial foray into the realm of what can be achieved by applying optimal control to atom interferometry.

In the immediate future we are looking to experimentally demonstrate a LMT interferometer with GRAPE optimised Raman augmentation pulses.

Such an interferometer warrants comparison to interferometers that already employ Bragg diffraction [25, 65, 170] and Bloch oscillations [169] to achieve large wavepacket separation. Broadband optimised Raman pulses place much less stringent requirements on the initial momentum distribution of the atomic sample for a potential signal enhancement. One application that is of particular interest is their potential to enhance the sensitivity of point source interferometers such as the one described in [41, 208].

A similar interferometer is currently being constructed in Southampton, and the concept relies on performing interferometry with relatively hot atoms that expand over a large area in the transverse plane of the interferometry beams in order to spatially resolve fringes across a ballistically expanding MOT cloud. The periodicity of these fringes scales with the rotation rate of the apparatus, and such an interferometer lends itself ideally to optimised Raman pulses, enabling point source LMT interferometry that could potentially realise a rotation sensitivity that is of interest for inertial navigation.

We have many ideas about how to extend our use of optimal control in atom interferometry and get a better understanding of the power of the algorithms. A potential control parameter that we have not yet experimented with varying is the pulse amplitude. It is well known, and can be shown by Fourier analysis, that pulses with smoothly varying amplitude profiles are more robust to high-frequency phase noise of the interferometry beams [172, 209] and can be used to realise sharper resonances [210]. As we are able to smoothly vary the pulse amplitude in our experiment, it is a logical step to include this as a control parameter in the GRAPE algorithm.

It is also worth noting that — rather than dividing each pulse into time slices and treating the phase and amplitude of each slice as control parameters — pulses could be constructed from a basis of sine and cosine functions and the control parameters could be the relative amplitudes of the various frequency components. It would be interesting to see what insight can be gleaned from studying how the frequency response of optimal pulses varies as this basis set is restricted. A similar analysis could be performed in the time

domain, studying how the frequency response varies as the temporal pulse profiles are truncated. These are the type of questions that will be studied in the thesis of Jack Saywell.

There are a few aspects of atom interferometry experiments that differentiate them from other coherent control experiments to which optimal control techniques have previously been applied. It makes sense to tailor how we employ the algorithms in order to take these aspects into effect. For instance, in a LMT interferometer the resonance condition for each 'arm' changes after each augmentation pulse, as they receive Doppler shifts in opposite directions. The augmentation pulses then need to be resonant with two distinct peaks in frequency space and could be optimised to account for this.

In addition, many atom interferometers employ retroreflected beams so that, in total, four frequency components are present during each pulse. In such interferometers the cloud is usually free falling along the beam axis, and the associated Doppler shift can be employed to ensure that only one pair of frequency components is resonant. There is, however, always a small probability of double diffraction, and this probability is increased when employing broadband pulses. Employing optimal control in such interferometers, the fidelity measure should therefore be chosen so as to suppress these unwanted resonances or — perhaps more interestingly — make use of them [66].[2]

[2]Clearly this latter option is of particular interest for atom interferometers in micro gravity environments where the apparatus is also in free fall [30, 207, 211].

This represents just a flavour of the problems in atom interferometry that optimal control could ultimately be applied to. In the long term, one could envisage an atom interferometric sensor that is free from such notions as beamsplitter and mirror pulses; an optimal control algorithm could be given control of the phase and amplitude of the laser beam throughout the interferometer sequence in order to maximise the variation of some measurable atomic property with respect to some quantity that one wishes to determine.

Even the interferometer duration could be flexible, and the optimisation need not be static. Active feedback could be implemented, such that the laser interactions are tailored to respond to changes in, for example, the motion of a vehicle in the case of an inertial sensor. Such a sensor could be multiplexed, with several atom clouds being juggled at any one time, and the optimal manner in which to address these clouds could be algorithmically determined.

Such a device is clearly a long way from being realised, but is a motivating vision towards which — in beginning to understand the application of optimal control to atom interferometry — we are taking an important first step.

12.3 Grey molasses loading and cooling of a far off-resonant optical dipole trap

The work in Chapter 11 demonstrates how grey molasses formed with hyperfine dark states (HDSs) can be used to efficiently load a far off-resonant optical dipole trap (FORT), even in the presence of large and uncompensated light shifts. The adaptations required to existing experiments that presently load a FORT from Sisyphus-cooled atoms are modest and, as both the cooling and loading happen simultaneously, may in fact simplify loading schemes that depend upon an intermediate — often magnetic — trap to overcome difficulties in mode-matching the cooling and loading phases. Indeed, the scheme has already proved pivotal in the recent all-optical production of BECs in microgravity [212].

Furthermore, the fourfold reduction in temperature while maintaining atom number demonstrates the potential for velocity-selective cooling of atoms within the trap volume. This could be significantly improved upon by switching to a grey molasses scheme on the ^{87}Rb D_1 line, avoiding the heating caused by unwanted resonances with the higher-lying excited states of the D_2 line. Such a solution should permit trapped atoms to be cooled to recoil-limited temperatures in millisecond timescales without loss of atoms. Work is ongoing to demonstrate this experimentally, and the technique could prove an effective complement to existing techniques for all-optical production of degenerate quantum gases [213, 214] that is particularly suited to rapid production in unusual geometries [205] and environments [116, 196].

Part IV

Appendices

Incorporation of spontaneous emission in the OBEs

A

"Cheshire-Puss," began Alice, rather timidly, "would you please tell me which way I ought to go from here?"
"In that direction," the Cat said, waving the right paw 'round, "lives a Hatter; and in that direction," waving the other paw, "lives a March Hare. Visit either you like; they're both mad."
"But I don't want to go among mad people," Alice remarked.
"Oh, you can't help that," said the Cat; "we're all mad here. . . "
— Lewis Carroll, *Alice in Wonderland*

The form of the OBEs presented in Chapter 3, with damping terms included to account for the spontaneous emission associated with an excited state decay rate of $\gamma = 1/\tau$, are not entirely intuitive. In this appendix we offer a form of justification.

It should be said that this is a somewhat non-standard derivation, and more rigorous arguments can be found in [81, 215], for example, or in many graduate courses such as [216]. The standard treatment employs the density matrix formalism, a reformulation of quantum mechanics that naturally lends itself to the treatment of open quantum systems and statistical mixtures. We will not formally introduce the density matrix here, but rather think of the state vector $\tilde{\mathbf{\Psi}} = (\tilde{u}, \tilde{v}, \tilde{w})$ as an ensemble average, thus relaxing the constraint that it must be of unit length. This is largely what the density matrix accomplishes, and it will be sufficient to describe all of the physics in the thesis.

A.1 Spontaneous emission: a naïve treatment

On timescales much shorter than τ, the probability of any spontaneous emission occurring is negligible and the dynamics of a two-level atom interacting with near-resonant light are well described by the form of the OBEs in Equations (2.34). The Rabi oscillations in Figure 2.3 then provide an accurate representation of how the probability $|c_2|^2$ of finding a single atom in the excited state evolves with time when the Rabi frequency $\Omega_R \gg \gamma$.

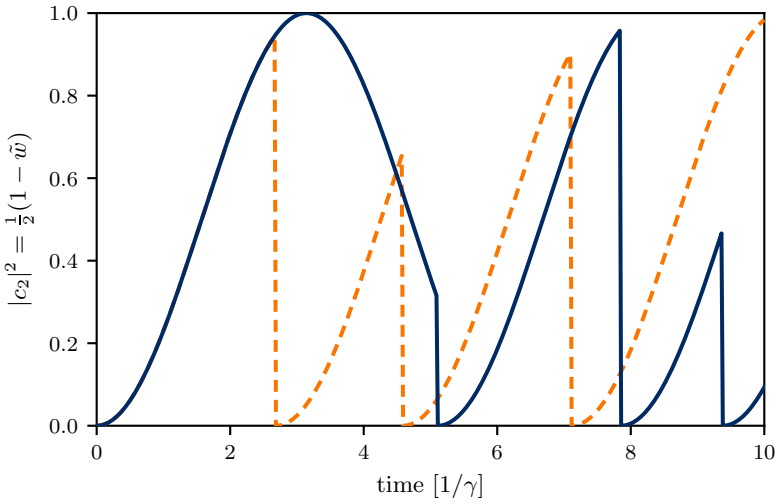

Figure A.1: Possible single-atom Rabi oscillations with $\Omega_R = \gamma$ and $\delta = 0$, assuming that the atom evolves unitarily but undergoes probabilistic decay to the ground state with probability $P_1 = \gamma|c_2|^2\delta t$ for a time interval δt.

If Ω_R is on the order of γ then, driven on resonance, one might expect the excited state probability for a single atom to evolve something like the curves in Figure A.1. Here the OBEs in Equations (2.34) are numerically integrated in small timesteps of fixed duration δt, but after each step we consider the possibility that the the atom could decay to the ground state with probability $P_1 = \gamma|c_2|^2\delta t$ given by the known decay rate of the excited state. This results in stochastic, discontinuous, 'jumps' in the state evolution as the atom is reset to the ground state by a spontaneous emission event. Logically, the higher the excited state probability $|c_2|^2$, the more likely it is that such a jump will occur.

A typical experiment to observe Rabi oscillations might involve preparing an ensemble of atoms in the ground state and inferring the evolution of the excited state probability by measuring the fraction of atoms $N_2/(N_2 + N_1)$ in the excited state after they have interacted with the laser for a series of times t.

We can make an educated guess as to what these data would look like as t is increased. At the start of the data set, where $t < \tau$, the probability that any of the atoms will spontaneously emit while the laser is on is small and so our measurement of the fraction of atoms in the excited state should closely follow the Rabi oscillations that we would expect for a single atom. For $\Omega_R \gg \gamma$ we would observe many oscillations, much like those shown in Figure 2.3. However, as t approaches τ, we reach the point where it is is likely that quantum jumps will have occurred for a significant fraction of the atoms. While the atoms all started in phase, each quantum jump

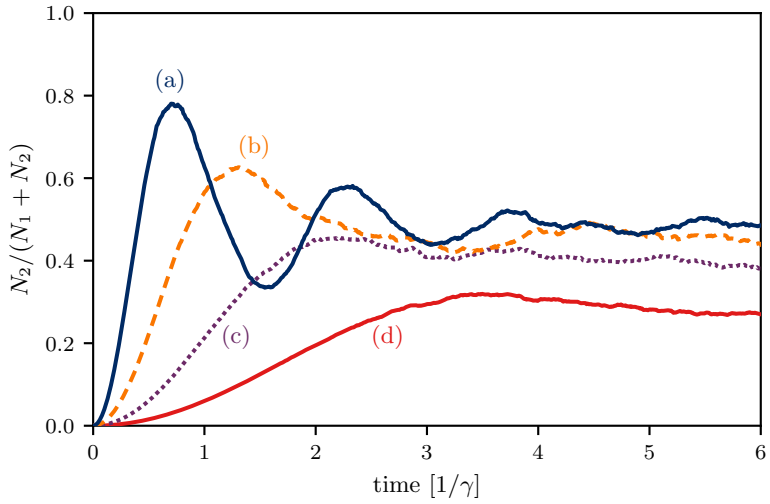

Figure A.2: Sum of 1000 stochastic trajectories such as those shown in Figure A.1 for $\delta = 0$ and (a) $\Omega_R = 4\gamma$ (b) $\Omega_R = 2\gamma$ (c) $\Omega_R = \gamma$ (d) $\Omega_R = \gamma/2$. Again, unitary evolution as per Equations (2.34) is assumed between the stochastic jumps.

randomises the phase of the Rabi oscillations for the atom that jumped relative to all of the others.

For $t \gg \tau$, then, we would expect there to be no coherent phase relationship between the oscillations and so, in a large enough sample of atoms, when we measure the number in the excited state then we would just observe a steady value, plus or minus a few atoms here or there from $1/\sqrt{N}$ fluctuations. The magnitude of this steady value should depend on the ratio of Ω_R to γ. If $\Omega_R \gg \gamma$ then a single atom will undergo many oscillations before jumping and so the average excited state probability $\langle |c_2|^2 \rangle = 0.5$ is just the cycle average and we would expect to find half of the atoms in the excited state. If $\Omega_R < \gamma$ then it is likely that a single atom will not complete a full cycle before experiencing a jump, and so the jumps have the effect of 'clamping' $\langle |c_2|^2 \rangle$ closer to zero and we might expect to find most of the atoms in the ground state.

In brief, for large Ω_R we would expect to see many coherent oscillations of the fraction of atoms $N_2/(N_1 + N_2)$ in the excited state, but they would eventually decay to a steady value of 0.5. As Ω_R decreases we would expect to see fewer oscillations before the ensemble settles to a steady state with $N_2/(N_1 + N_2) < 0.5$. This is illustrated in Figure A.2, where resonant stochastic Rabi oscillations such as those in Figure A.1 are summed over $N = 1000$ atoms for a range of Rabi frequencies between $\Omega_R = 4\gamma$ and $\Omega_R = \gamma/2$.

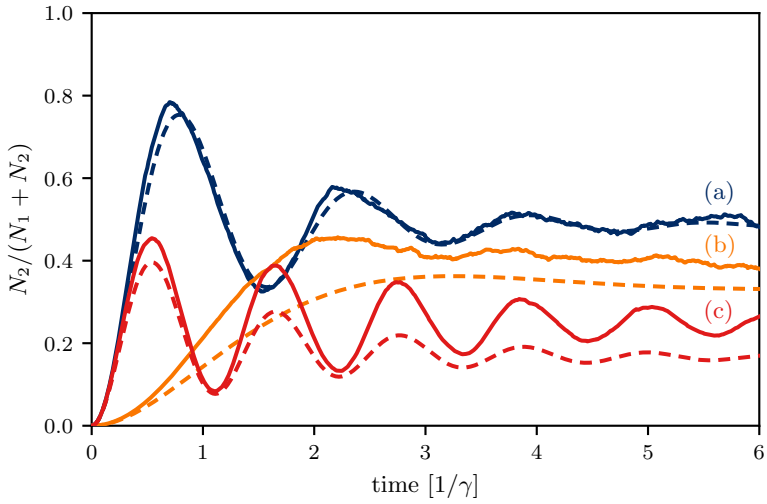

Figure A.3: Monte Carlo simulations conducted as per Figure A.2 for N=1000 atoms (solid lines) compared to the physically accurate result (dashed lines) found in many texts for (a) $\Omega_R = 4\gamma$, $\delta = 0$, (b) $\Omega_R = 1\gamma$, $\delta = 0$, (c) $\Omega_R = \delta = 4\gamma$. We revise our assumptions to account for the differences and derive the correct result in the next section.

These curves agree qualitatively with experiments, but show quantitative discrepancies that are most apparent at low Rabi frequencies or when introducing a detuning. This is demonstrated in Figure A.3, where Monte Carlo results incorporating our assumptions thus far (solid lines) are compared to a physically accurate theory (dashed lines) presented in many texts. Good agreement is observed on resonance with a sufficiently high Rabi frequency, but relaxing either of these conditions results in considerable differences emerging. We shall reexamine our assumptions in the next section to account for these discrepancies, and derive the physically accurate result in a manner that, while less rigorous than many treatments, hopefully offers deeper physical insight. This one example has a lot to tell us about the nature of quantum mechanics.

A.2 Spontaneous emission: what's really going on?

To square spontaneous emission with the unitary evolution of a quantum state dictated by the Schrödinger equation we must expand our basis to consider the surrounding environment. Then, rather than considering energy as being lost from the system upon spontaneous emission, we instead account for it by the appearance of a photon in a *vacuum mode*.[1] In this basis, we might write an arbitrary state of a two level atom as

[1] Due warning: we are about to expose a flagrant disregard for mathematical rigour that can only be partly excused by my being an 'experimentalist'.

$$\begin{aligned} |\Psi(t)\rangle &= |\Psi_{\text{atom}}\rangle \otimes |\Psi_{\text{environ}}\rangle \\ &= \left(\tilde{c}_1 |1\rangle + \tilde{c}_2 |2\rangle\right) \otimes |n=0\rangle \,, \end{aligned} \tag{A.1}$$

where we employ our now familiar trick of entering a frame rotating at the transition frequency ω_0 and employing the associated state amplitudes $\tilde{c}_{g,e} = e^{\mp\omega_0 t/2} c_{g,e}$ to eliminate their time-dependence. The right-hand ket represents 0 photons in the surrounding vacuum; more accurately, we could sum (or integrate) over all possible wavevectors and polarisations [97], but we shall keep things simple. We know empirically that, in a small time interval $\delta t \ll \tau$, the probability of the atom spontaneously decaying and thus emitting a photon into the vacuum is $P_1 = |\tilde{c}_2|^2 \gamma \delta t$. Thus it seems sensible to conclude that, at a time $t + \delta t$, the state of the system will have evolved into

$$\begin{aligned} |\Psi(t+\delta t)\rangle =& \tilde{c}_1 |1\rangle \otimes |n=0\rangle + \tilde{c}_2 \Big(\sqrt{\gamma \delta t}\, |1\rangle \otimes |n=1\rangle \\ &+ \sqrt{1-\gamma\delta t}\, |2\rangle \otimes |n=0\rangle \Big) \\ =& \left(\tilde{c}_1 |1\rangle + \tilde{c}_2 \sqrt{1-\gamma\delta t}\, |2\rangle\right) \otimes |n=0\rangle \\ &+ \left(\tilde{c}_2 \sqrt{\gamma\delta t}\, |1\rangle\right) \otimes |n=1\rangle \,. \end{aligned} \tag{A.2}$$

Note that this encapsulates the physics of the probabilistic decay that we just described, while preserving the norm of the wavefunction $\langle \Psi(t+\delta t) | \Psi(t+\delta t) \rangle = 1$.

Now, let us model the environment as a unit quantum efficiency photon detector around the atom, encompassing a 4π solid angle so that all photons emitted are detected, but without any which-way information so that our notation of either no vacuum photon $|n=0\rangle$ or one vacuum photon $|n=1\rangle$ is somewhat justified. How far away from the atom is this gedanken detector? Well in a sense it is defined by how far a photon travels in our time interval δt. It should be close enough so that the time it takes for the emitted photon to reach it is much smaller than τ and there is no chance of multiple spontaneous emission events taking place in our time interval, but it should be far enough away that the photon has 'separated' from the atom before the detection, which can then be taken to be instantaneous and independent of anything that happened during the previous time interval in what is referred to as the Markov approximation, more formally defined in [81].

In this picture, we think of the environment as making repeated projective measurements on the system, namely measuring whether or not a photon has been emitted into the vacuum. Employing the projection operator $\hat{M}_1 = \mathbb{I} \otimes |n=1\rangle \langle n=1|$, the probability of detecting a vacuum photon is

$$\begin{aligned} P_1 =& \langle \Psi(t+\delta t)| \hat{M}_1^\dagger \hat{M}_1 |\Psi(t+\delta t)\rangle \\ =& |\tilde{c}_2|^2 \gamma \delta t, \end{aligned} \tag{A.3}$$

corresponding to the probability of spontaneous emission and, as further justification for our argument, the measurement projects the system into the state

$$\frac{1}{\sqrt{P1}} \hat{M}_1 |\Psi(t+\delta t)\rangle = \frac{1}{\sqrt{P_1}} \left(\tilde{c}_2 \sqrt{\gamma \delta t} \right) |1\rangle \otimes |n=1\rangle \,, \tag{A.4}$$

corresponding to the atom resetting to the ground state as anticipated. For the next time interval, we can set the number of vacuum photons back to zero again and proceed iteratively.[2]

[2] It is this step that relies on our gedanken 'detector' being far enough away from the atom so that the atom does not 'see' the photon once it has arrived at the detector and the time taken for the detector to absorb the photon and 'click' can be considered instantaneous on the timescale of δt. From the atom's perspective the electromagnetic mode is empty again at the start of each timestep. It seems logical that this critical distance should be on the order of the optical wavelength. In practise this means that the photon propagates far enough into the environment to get lost and carry away information. We are using our hypothetical detector to temporarily keep track of this information and demonstrate how its loss from the system affects its evolution.

So far so good, but what is most interesting — and what we omitted in the previous section — is what happens when the detector does not click. This, again, corresponds to a projective measurement; this time of no vacuum photons, and represented by the projection operator $\hat{M}_0 = \mathbb{I} \otimes |n=0\rangle \langle n=0|$. The probability of this outcome is

$$\begin{aligned} P_0 =& \langle \Psi(t+\delta t)| \hat{M}_0^\dagger \hat{M}_0 |\Psi(t+\delta t)\rangle \\ =& |\tilde{c}_1|^2 + |\tilde{c}_2|^2 (1 - \gamma \delta t) \,, \end{aligned} \tag{A.5}$$

satisfying $P_1 + P_0 = 1$. This collapses the system into the state

$$\frac{1}{\sqrt{P_0}} \hat{M}_0 |\Psi(t+\delta t)\rangle = \frac{1}{\sqrt{P_0}} \left(\tilde{c}_1 |1\rangle + \tilde{c}_2 \sqrt{1-\gamma \delta t} |2\rangle \right) \otimes |n=0\rangle \,. \tag{A.6}$$

It appears that the atomic part of the state has changed by virtue of the fact that no photon was detected. This may seem counter-intuitive, but a few things are obvious from inspection that help clarify the issue. Firstly, if the atom started in the ground state then observing no photon has no effect. We have $P_0 = |\tilde{c}_1|^2 = 1$

and so $|\Psi_{\text{atom}}(t+\delta t)\rangle = |\Psi_{\text{atom}}(t)\rangle = \tilde{c}_1\,|1\rangle$. As we would expect, the atom remains in the ground state indefinitely.

Similarly, if the atom starts out in the excited state with $\tilde{c}_2 = 1$ and $\tilde{c}_1 = 0$ then measuring no photon leaves the state unaltered. The atom will remain in the excited state until, with probability P_1, spontaneous emission occurs, a photon is detected and it is projected into the ground state.

However, if the atom is in a superposition then a null measurement alters the excited state amplitude $\tilde{c}_2$ relative to $\tilde{c}_1$. A Taylor expansion of $\sqrt{1-\gamma\delta t}$ shows that, for suitably small δt, the collapse induced by the measurement is equivalent to reducing $\tilde{c}_2 \to \tilde{c}_2 - \tilde{c}_2\frac{\gamma}{2}\delta t$ and renormalising. This looks like the excited state amplitude decays exponentially

$$\frac{\partial \tilde{c}_2}{\partial t} = -\left(\frac{\gamma}{2}\right)\tilde{c}_2, \tag{A.7}$$

but the renormalisation step — encompassed in the projective measurement — is important. Note that the decay does not occur if $\tilde{c}_2 = 1$.

A.2.1 Philosophically, what's really going on?

Nonetheless, in coupling the atom to our hypothetical detector we have provided a route for an atomic superposition to return to the ground state *without the emission of a photon* [215]. If we sit and monitor our detector for a long time then we may see a photon emitted from the atom or we may not. Either way, if we look at the atom at $t = \infty$ then we will find it in the ground state.

This may seem strange, but is it more strange than expecting the atom to remain in a superposition state for a duration that is thousands of times greater than its intrinsic lifetime? Quantum mechanics is fundamentally a way of quantifying ignorance, and observing no photon emitted from the atom tells us something about its state. The longer we don't observe a photon for, the more likely it is that the atom is in the ground state. If we sit and watch the box containing Schrödinger's cat for days and don't hear it meowing in distress or clawing at the walls then we can be pretty certain that the poor thing is dead, regardless of whether it died by poisoning or from natural causes.

Now, our hypothetical detector is not a necessary ingredient in this decay. It serves as a nice way to simplify the mathematics and keep track of the information leaving the system, but if we took it away then the spontaneously emitted photons would not remain 'unmeasured' for long. They would still interact with, and play a role in, the surrounding environment. Whether or not some experimentalist contrives that environment so that they hear a click when the photon is emitted — and then choose to *call* that a measurement — is of no consequence. Coupling to a macroscopic universe is enough; atomic physics experiments are conducted in UHV chambers located in dark laboratories for a reason.

A.2.2 Geometrically, what's really going on?

As is so often the case, the Bloch sphere offers some rich geometric insights into how these decays, with and without the emission of a photon, affect the state of a single atom and of an ensemble.

Let us start with the decay in the absence of emitting a photon. We can use Equation (A.6) to determine how the Bloch vector coordinates $\tilde{u} = \tilde{c}_1^*\tilde{c}_2 + \tilde{c}_1\tilde{c}_2^*$, $v = i(\tilde{c}_1\tilde{c}_2^* - \tilde{c}_1^*\tilde{c}_2)$, and $w = \tilde{c}_1^*\tilde{c}_1 - \tilde{c}_2^*\tilde{c}_2$ evolve in the small time interval δt by identifying

$$\tilde{c}_1(t+\delta t) = \frac{1}{\sqrt{P_0}}\tilde{c}_1(t), \text{ and} \tag{A.8a}$$

$$\tilde{c}_2(t+\delta t) = \frac{\sqrt{1-\gamma\delta t}}{\sqrt{P_0}}\tilde{c}_2(t) \tag{A.8b}$$

to yield[3]

$$\begin{aligned}
\tilde{u}(t+\delta t) &= \frac{\sqrt{1-\gamma\delta t}}{P_0}\tilde{u}(t), \\
\tilde{v}(t+\delta t) &= \frac{\sqrt{1-\gamma\delta t}}{P_0}\tilde{v}(t), \text{ and} \\
\tilde{w}(t+\delta t) &= \frac{1+\tilde{w}(t)-\big(1-\tilde{w}(t)\big)\,(1-\gamma\delta t)}{2P_0}.
\end{aligned} \tag{A.9}$$

[3] Here we assume no laser interactions, $\Omega_R = \delta = 0$, and just focus on the dynamics driven by the detection, or not, of spontaneously emitted photons.

Taking Taylor expansions to first order about $\gamma\delta t = 0$ we obtain

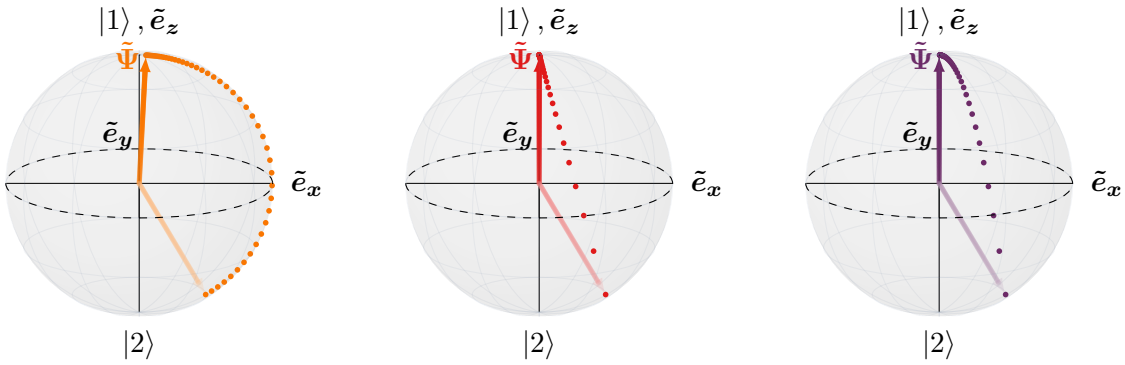

(a) Pure measurement decay. **(b)** Pure stochastic decay. **(c)** Combined decay.

Figure A.4: Different decay processes of an atomic superposition to the ground state depicted on the Bloch sphere. (a) A single atom decaying without emitting a photon via a measurement-induced 'pull' of the state vector along the surface of the Bloch sphere towards the ground state. This pull is strongest at the equator and vanishes at the poles. (b) The average of many atoms undergoing immediate, stochastic jumps to the ground state with probability $P_0 = |c_2|^2\gamma\delta t$. (c) A combination of the two processes that corresponds to what is observed in a physical ensemble. Each dot corresponds to a $\tau/5$ timestep.

$$\begin{aligned}
\tilde{u}(t+\delta t) - \tilde{u}(t) &\approx -\frac{\gamma}{2}\tilde{w}(t)\tilde{u}(t)\delta t, \\
\tilde{v}(t+\delta t) - \tilde{v}(t) &\approx -\frac{\gamma}{2}\tilde{w}(t)\tilde{v}(t)\delta t, \text{ and} \\
\tilde{w}(t+\delta t) - \tilde{w}(t) &\approx \frac{\gamma}{2}\left(1-\tilde{w}(t)^2\right)\delta t,
\end{aligned} \tag{A.10}$$

inferring that, for as long as the atom does not spontaneously emit,

$$\frac{\partial \tilde{u}}{\partial t} = -\frac{\gamma}{2}\tilde{w}\tilde{u} \tag{A.11a}$$

$$\frac{\partial \tilde{v}}{\partial t} = -\frac{\gamma}{2}\tilde{w}\tilde{v} \tag{A.11b}$$

$$\frac{\partial \tilde{w}}{\partial t} = \frac{\gamma}{2}\left(1-\tilde{w}^2\right). \tag{A.11c}$$

Making use of $\tilde{u}^2 + \tilde{v}^2 + \tilde{w}^2 = 1$, it is pleasing to verify that this is equivalent to a rotation $\frac{\partial \tilde{\boldsymbol{\Psi}}}{\partial t} = \boldsymbol{R} \times \tilde{\boldsymbol{\Psi}}$, where

$$\boldsymbol{R} \equiv \frac{\gamma}{2}\begin{pmatrix} \tilde{v} \\ -\tilde{u} \\ 0 \end{pmatrix}. \tag{A.12}$$

The state vector rotates at a rate $R = |\boldsymbol{R}| = \gamma/2\sqrt{\tilde{u}^2 + \tilde{v}^2}$ about an axis in the equatorial plane that is orthogonal to it. It is always pulled along the surface of the Bloch sphere to the north pole, corresponding to the ground state, and the pull is greatest at the

equator. It vanishes at the poles, where $\tilde{u} = \tilde{v} = 0$, in agreement with our assertion that the measurement of no photon has no effect on an atom in the ground or excited state. The nature of this rotation for a superposition state is represented on the Bloch sphere in Figure A.4(a), where each dot represents a fixed timestep so that the rate of the rotation can be inferred from their spacing. We see that the state asymptotically approaches the north pole.

It is this rotation that we omitted from our naïve treatment in the previous section, and we can now see why that treatment was in quantitative agreement for the resonant case with large Ω_R. On resonance $\boldsymbol{R}$ is aligned with the field vector $\tilde{\boldsymbol{\Omega}}$ that drives the Rabi oscillations. Therefore, when the field vector is driving the state vector towards the ground state then the measurement-induced rotation will slightly increase the rotation rate, and when the field vector is driving the state vector towards the excited state then $\boldsymbol{R}$ will act to oppose it. Over a single Rabi cycle the net effect will be zero and will not be observed. We would only expect to see an effect either off resonance, where $\boldsymbol{R}$ is no longer parallel to $\tilde{\boldsymbol{\Omega}}$, or else for small Rabi frequencies $\Omega_R < \gamma$ where the state vector rarely completes a full cycle before a spontaneous decay and therefore the change to the rotation rate does not completely cancel. This is exactly the explanation for the discrepancies observed in Figure A.3.

The effect that we did include in Section A.1, namely the stochastic and instantaneous jumps to the ground state corresponding to the spontaneous emission of a photon, is represented for the same state on the Bloch sphere in Figure A.4(b). Clearly this is not a trajectory for a single atom, which will either have jumped or not jumped. Rather, each time step is the normalised average of the state vectors for many atoms, each decaying at random times but with equal instantaneous probability. At any point in time, every atom is either in the initial state or the ground state so that the average of all of the state vectors will always lie somewhere on the line between the two.[4] The longer we wait, the more atoms will have jumped and the average tends towards the ground state.

[4]This line passes through the inside of the Bloch sphere and is not restricted to the surface. This is a feature, not a bug; by relaxing the constraint that a state vector should have unit length we can use the Bloch sphere to represent statistical mixtures such as the one we are describing. We shall explore this in more detail shortly.

It remains to combine both of these effects into a complete physical description. We can start by amending our Monte Carlo simulation, and this has been done to produce Figure A.5, in which the fraction of $N = 1000$ in the excited state is shown as a function of time, first for an ensemble of atoms starting in an equal superposition of the ground and excited states and then for an ensemble of atoms starting in the excited state. Obviously at $t = 0$ the curves

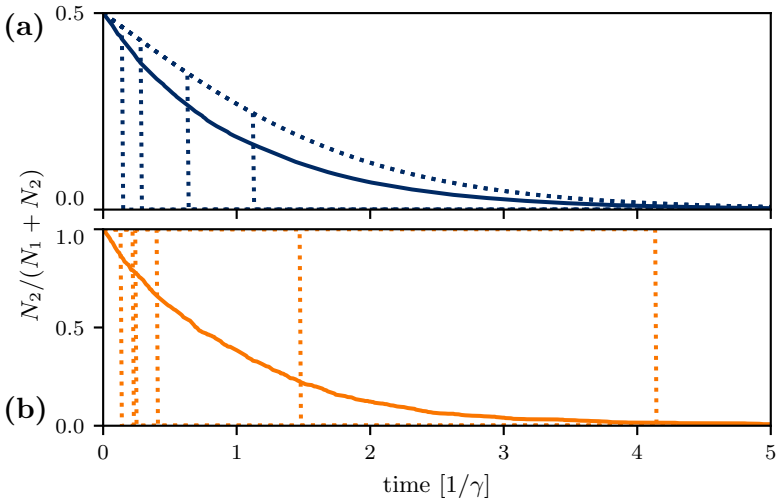

Figure A.5: Monte Carlo simulations of the fraction of N=1000 atoms in the excited state (solid lines) as a function of time, incorporating the decay mechanisms introduced in this section, (a) starting from an equal superposition state and (b) starting from the excited state. A sample of single-atom trajectories are shown as dashed lines for each case.

start at $N_2/(N_1 + N_2) = 1/2$ and 1 respectively, but they both decay to zero at the same rate.

However, a sample of single-atom trajectories are represented by dashed lines demonstrating that, as we have shown, the ensemble behaves very differently in each case. When the atoms all start in a superposition then they begin a smooth, measurement-induced decay to the ground state, but some also spontaneously emit and jump straight to their destination, expediting the decay of the ensemble as a whole. Conversely, when the atoms start in the excited state then the only way they can decay is via emission of a photon, and so the individual trajectories are step functions. However, by virtue of the fact that they are in the excited state, these jumps are much more likely to occur. The result is a decay of the ensemble at the same rate.

We can incorporate an analytical expression for the decay effect that the stochastic jumps have on the ensemble averaged state vector into differential Equations (A.11). The rate at which the jumps occur is $\Gamma = \gamma|\tilde{c}_2|^2 = \gamma(1-\tilde{w})/2$. Each jump returns $\tilde{u}$ and $\tilde{v}$ to the origin and so, averaged over an ensemble, the rates of change of these components acquire $-\Gamma\tilde{u}$ and $-\Gamma\tilde{v}$ terms respectively. The $\tilde{w}$ component however, is reset to $\tilde{w} = 1$ and so the associated rate of change is $\Gamma(1 - \tilde{w})$. Including these terms, Equations (A.11) become

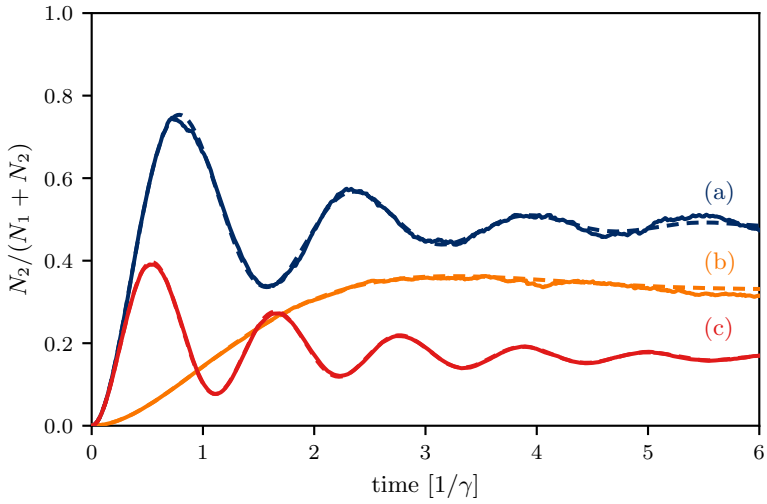

Figure A.6: Revised Monte Carlo simulations of the fraction of N=1000 atoms in the excited state during a laser excitation (solid lines) compared to a numerical integration of Equations (A.14) (dashed lines) for (a) $\Omega_R = 4\gamma$, $\delta = 0$, (b) $\Omega_R = 1\gamma$, $\delta = 0$, (c) $\Omega_R = \delta = 4\gamma$. Quantitative agreement is now observed in all cases.

$$\frac{\partial \tilde{u}}{\partial t} = -\frac{\gamma}{2}\tilde{w}\tilde{u} - \frac{\gamma}{2}(1-\tilde{w})\tilde{u} = -\frac{\gamma}{2}\tilde{u}, \tag{A.13a}$$

$$\frac{\partial \tilde{v}}{\partial t} = -\frac{\gamma}{2}\tilde{w}\tilde{v} - \frac{\gamma}{2}(1-\tilde{w})\tilde{v} = -\frac{\gamma}{2}\tilde{v}, \text{ and} \tag{A.13b}$$

$$\frac{\partial \tilde{w}}{\partial t} = \frac{\gamma}{2}\left(1-\tilde{w}^2\right) + \frac{\gamma}{2}\left(1-\tilde{w}\right)^2 = \gamma\left(1-\tilde{w}\right). \tag{A.13c}$$

These equations are used to show the total decay of a superposition on the Bloch sphere in Figure A.4(c). Note that again the state vector passes through the interior of the sphere, and contributions from the individual decay methods shown alongside it are evident. From geometric considerations we would expect the decay to be dominated by spontaneous emission in the southern hemisphere, with the measurement-induced decay dominating when the state vector is near the north pole, and this is indeed evident.

Tacking these decay terms onto the coherent OBEs from Equations (2.34), we arrive at them in their complete form

$$\frac{\mathrm{d}\tilde{u}}{\mathrm{d}t} = -\delta\tilde{v} + \Omega_R \sin(\phi)\tilde{w} - \frac{\gamma}{2}\tilde{u} \tag{A.14a}$$

$$\frac{\mathrm{d}\tilde{v}}{\mathrm{d}t} = \delta\tilde{u} - \Omega_R \cos(\phi)\tilde{w} - \frac{\gamma}{2}\tilde{v} \tag{A.14b}$$

$$\frac{\mathrm{d}\tilde{w}}{\mathrm{d}t} = -\Omega_R \sin(\phi)\tilde{u} + \Omega_R \cos(\phi)\tilde{v} + \gamma\left(1-\tilde{w}\right). \tag{A.14c}$$

These are the equations used the generate the theory curves in Figure A.3 and variations of them are found in most quantum optics texts. Having incorporated the measurement-induced decay into our Monte Carlo simulation, it is now in quantitative agreement with them, as shown in Figure A.6. Monte Carlo methods have been applied to this particular problem before, and mathematically our simulation is now equivalent to those described by [217] in which the differential equations for the complex state amplitudes, rather than the Bloch vector components, are integrated directly.

Optimal channel crossing in a time-varying tide

B

The boat zipped and skipped across the sea, the sea that lay between the main islands of the only archipelago of any useful size on the whole planet. Zaphod Beeblebrox was on his way from the tiny spaceport on Easter Island (the name was an entirely meaningless coincidence — in Galacticspeke, easter means small flat and light brown) to the Heart of Gold island, which by another meaningless coincidence was called France.

— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

I want to cross the English channel by boat in the shortest time possible. What course should I steer?

Let's simplify the problem. I suppose that the maximum speed I can make through the water is the same in all directions and that I will always travel at full speed. Furthermore we'll assume that the channel is infinitely long in the east-west x direction and does not change width in the north-south y direction, justifying a tidal model that assumes the tide only flows along the x axis with a strength $c(y, t)$ that depends only on the time of day and my y coordinate.

We will employ the model

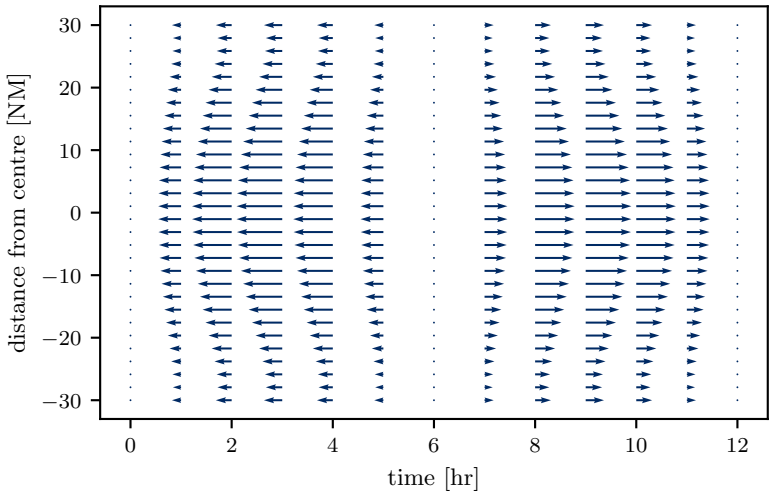

Figure B.1: Simplified tidal model described by Equation (B.1).

$$c(y,t) = c_0 \left(\left(\frac{y}{30} \right)^2 - 1 - \frac{2}{5} \left(\frac{y}{30} \right)^8 \right) \sin \left(2\pi \frac{t}{12} + \phi \right), \quad \text{(B.1)}$$

with distance measured in nautical miles (NM) and time measured in hours, assuming the channel to be 60 NM wide with the English coast at $y = 30$ NM and the French coast at $y = -30$ NM. The tide oscillates over a 12 hr period and, for convenience, we include a parameter ϕ to specify the phase of the tide at the start of the voyage at $t = 0$ although in practice this is just equivalent to selecting a different start time.

This function is graphically portrayed in Figure B.1 and corresponds to the tide flowing most strongly in the deepest waters at the centre of the channel with some revival of flow inshore because I am used to sailing in the stretch of water where the Cherbourg peninsula, Portland and the Isle of Wight create a pinch point where such things can happen. So, you see, I model edge effects in a system that I just said was infinite because it's my system and I can do what I like.

My boat's position (x, y) is thus governed by the equations of motion

$$\begin{aligned} \frac{\partial x}{\partial t} &= v \sin\big(\theta(t)\big) + c(y,t), \\ \frac{\partial y}{\partial t} &= v \cos\big(\theta(t)\big), \end{aligned} \quad \text{(B.2)}$$

where v is my boatspeed through the water and $\theta(t)$ is my course to steer (CtS), the instantaneous direction in which I point the bow of my boat that is not, in general, equal to my actual *course over ground (COG)* due to the additional tidal vector.

Suppose that, at its strongest, the tide flows at $c_0 = 6$ kt and my modest vessel can make $v = 8$ kt through the water. I want to go to Alderney, just a touch west of due south. Knowing how quickly I am likely to get there then, by looking at tidal predictions, I can quite trivially estimate how far the tide is likely to push me east or west when integrated over the voyage and pick a constant CtS in order to compensate for it and arrive safely at my destination.

Mathematically I calculate a CtS for which $\frac{\mathrm{d}\theta}{\mathrm{d}t} = 0$ and that gets me to where I want to go. The Royal Yachting Association (RYA) offers good guidance on how to approximate such a solution with

very little effort, and it is employed by sailors globally.[1] But this is a problem that is crying out to be treated by employing optimal control techniques; I seek an optimal CtS that minimises the total time required to travel from an initial location (x_i, y_i) to a final destination (x_f, y_f).

[1] A good sailor will of course build *leeway* into their passage plan to account for possible delays, and will be continually revising their plan en route. They will also consider the wind forecasts, as this will affect their potential boatspeed that will, in general, be a function of their steering angle and the local wind speed and direction. We will neglect all of these complications.

Reading [218, §2.7], I discover that Pontryagin's minimum principle says that there will exist an *adjoint* system (λ_x, λ_y) that obeys the equations of motion

$$\begin{aligned} \frac{\partial \lambda_x}{\partial t} &= 0, \\ \frac{\partial \lambda_y}{\partial t} &= \frac{\partial c}{\partial y}\,\lambda_x(t), \end{aligned} \tag{B.3}$$

such that the optimal CtS $\bar{\theta}(t)$ maximises the *effective Hamiltonian*

$$H = \lambda_x(t)\left[v\sin\big(\theta(t)\big) + c(y,t)\right] + \lambda_y(t)\left[v\cos\big(\theta(t)\big)\right] + 1, \tag{B.4}$$

for all times $0 < t < t_f$, where t_f is a free parameter — corresponding to the time of arrival at the target destination — that I seek to minimise [158].[2] Hence

$$\frac{\partial H}{\partial \theta} = \lambda_x v\cos(\theta) - \lambda_y v\sin(\theta) = 0, \Rightarrow \tan(\theta) = \frac{\lambda_x}{\lambda_y}, \tag{B.5}$$

[2] Note that, while this is a necessary condition for a path to be optimal, it does not strictly guarantee that it will be. Proving that a derived path is truly optimal would technically require more effort.

omitting the dependence on t for brevity.

Equations (B.2), (B.3) and (B.5), in conjunction with the boundary conditions $x(0) = x_i$, $y(0) = y_i$, $x(t_f) = x_f$, $y(t_f) = y_f$, and the transversality condition (again detailed in [218])

$$\begin{aligned} \lambda_x(t_f)\left[v\sin\big(\theta(t_f)\big) + c(y_f, t_f)\right] \\ + \lambda_y(t_f)\left[v\cos\big(\theta(t_f)\big)\right] = -1, \end{aligned} \tag{B.6}$$

constitute a boundary-value problem that can be solved numerically. So I do so, employing a simple shooting method, hoping to find an optimal CtS that significantly outperforms the simple method prescribed by the RYA... and it turns out that I can't.

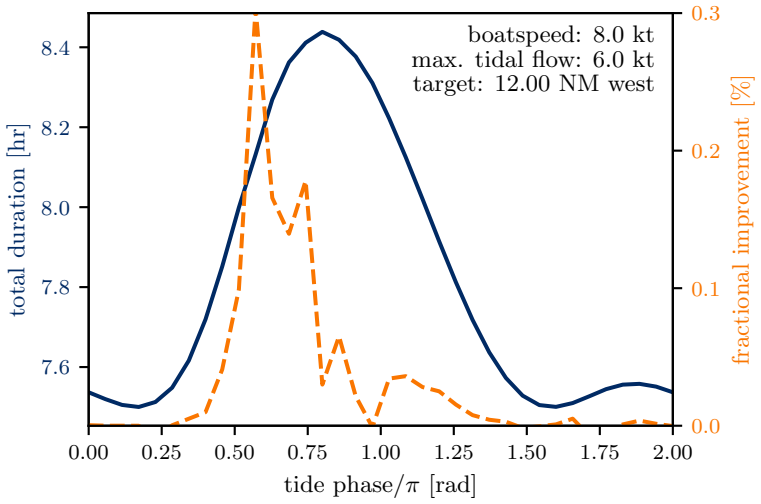

Figure B.2: Total passage time in hours and amount of time saved, in minutes, when steering an optimal course as opposed to a constant one as a function of the tidal phase upon leaving port to a destination just west of due south.

Figure B.2 shows the total passage time and the fractional improvement that an optimised CtS can yield compared to a constant course, as a function of the phase of the tide when I depart. The best improvement I could possibly hope to see is just under 0.3 %, with a saving of 58 s over around 8 hours and 5 minutes when $\phi = 0.55\pi$. A comparison of the trajectory for this CtS and an equivalent constant course are shown in Figure B.4.

So, annoyed that I've gone to all this effort for nothing, I start scouting out alternative destinations that will require me to put in

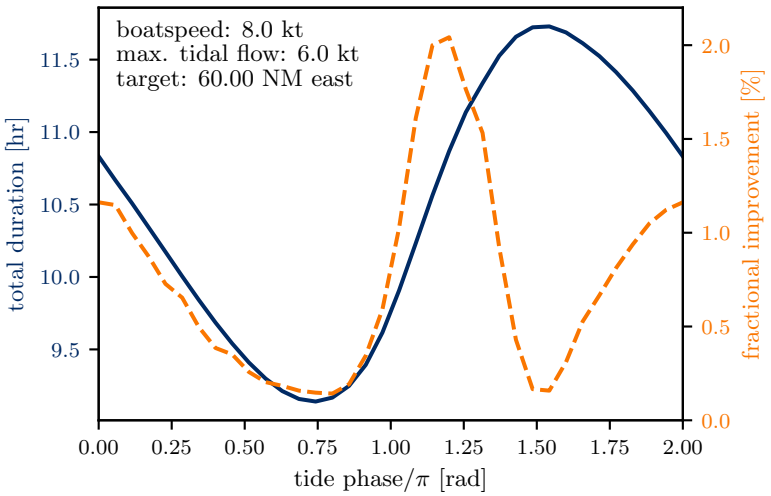

Figure B.3: Total passage time in hours and amount of time saved, in minutes, when steering an optimal course as opposed to a constant one as a function of the tidal phase upon leaving port to a destination directly southeast of me.

considerable miles towards or against the direction of tidal flow, figuring that this will be where my newfound mastery of optimal control will really come into its own. I choose to head directly southeast, perhaps to Dieppe. Sure enough, about a 1–2 % improvement can be made for certain departure times, as shown in Figure B.3.

Looking more closely at the departure time $\phi = 1.2\pi$ where the most improvement is made, 13 minutes and 18 seconds over about 11 hours, we see that it corresponds to the tide flowing towards my destination when I depart. The computer then arrives at the course, shown in Figure B.5, that I might well have sailed instinctively: getting out into the fast flowing tide quickly while it can still carry me east towards my destination and thus doing less work to fight it once it turns. In the case where the most benefit from the tide is felt in the middle of the passage (Figure B.6) such a strategy is no longer favourable and the optimal CtS is again only marginally faster than a constant one.

So the take-home message would appear to be that the course that a sailor will be taught to steer is very close to optimal. In retrospect, this should not come as much of a surprise. As discussed in Section 8.3.2, the nautical community has effectively acted as a global feedback loop, approaching an optimal solution. It operates at a much lower frequency than anything I might be used to in the lab, but then it has had hundreds of years to operate, culminating in the current RYA guidelines.

As technology evolves, of course, the optimal approach might change. The rise of hydrofoils, for instance, represents a paradigm shift in how we think about our sport. The advantage of an active feedback loop is that it surely will not be long before the guidelines are updated to reflect the new status quo.

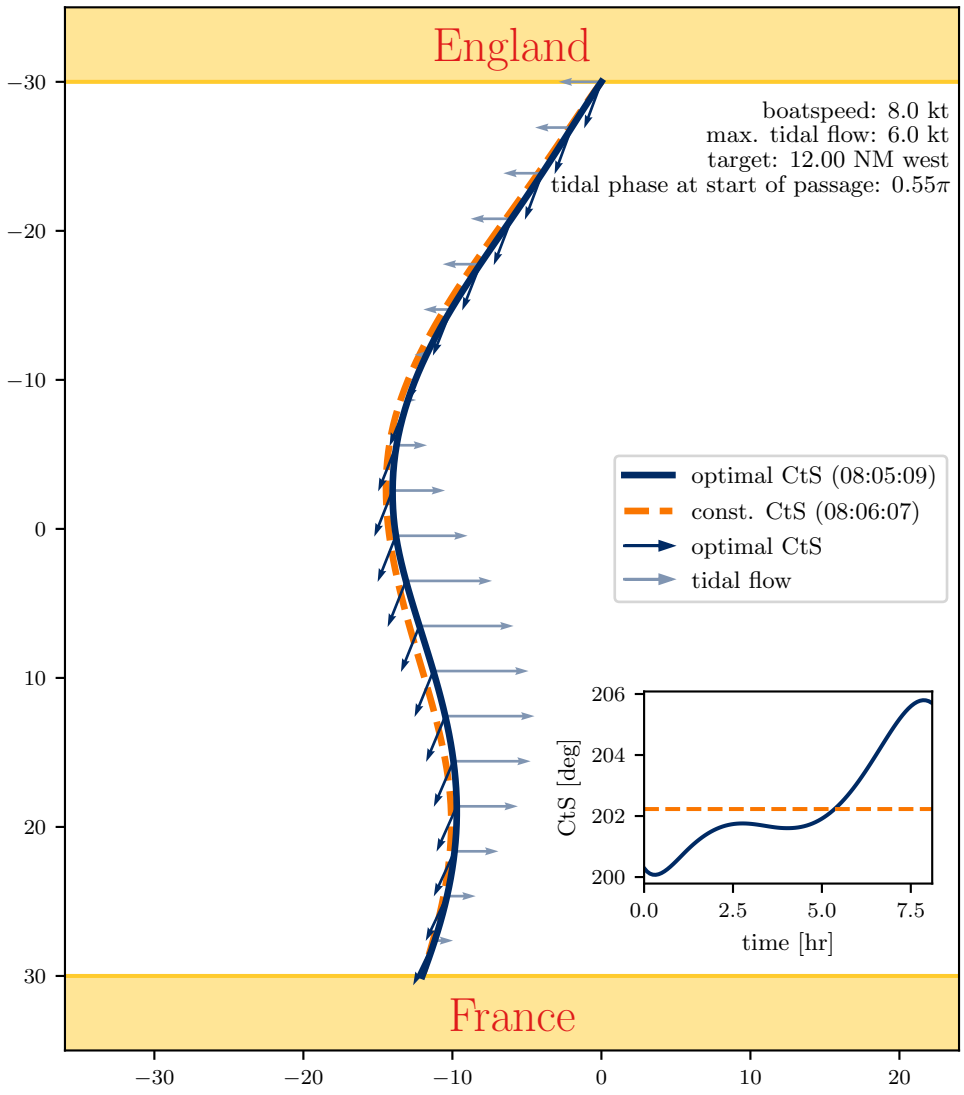

Figure B.4: Trajectories for an optimal course, with arrows representing the course to steer and tidal vectors along it, compared to a constant course for the best case improvement on a passage to "Alderney".

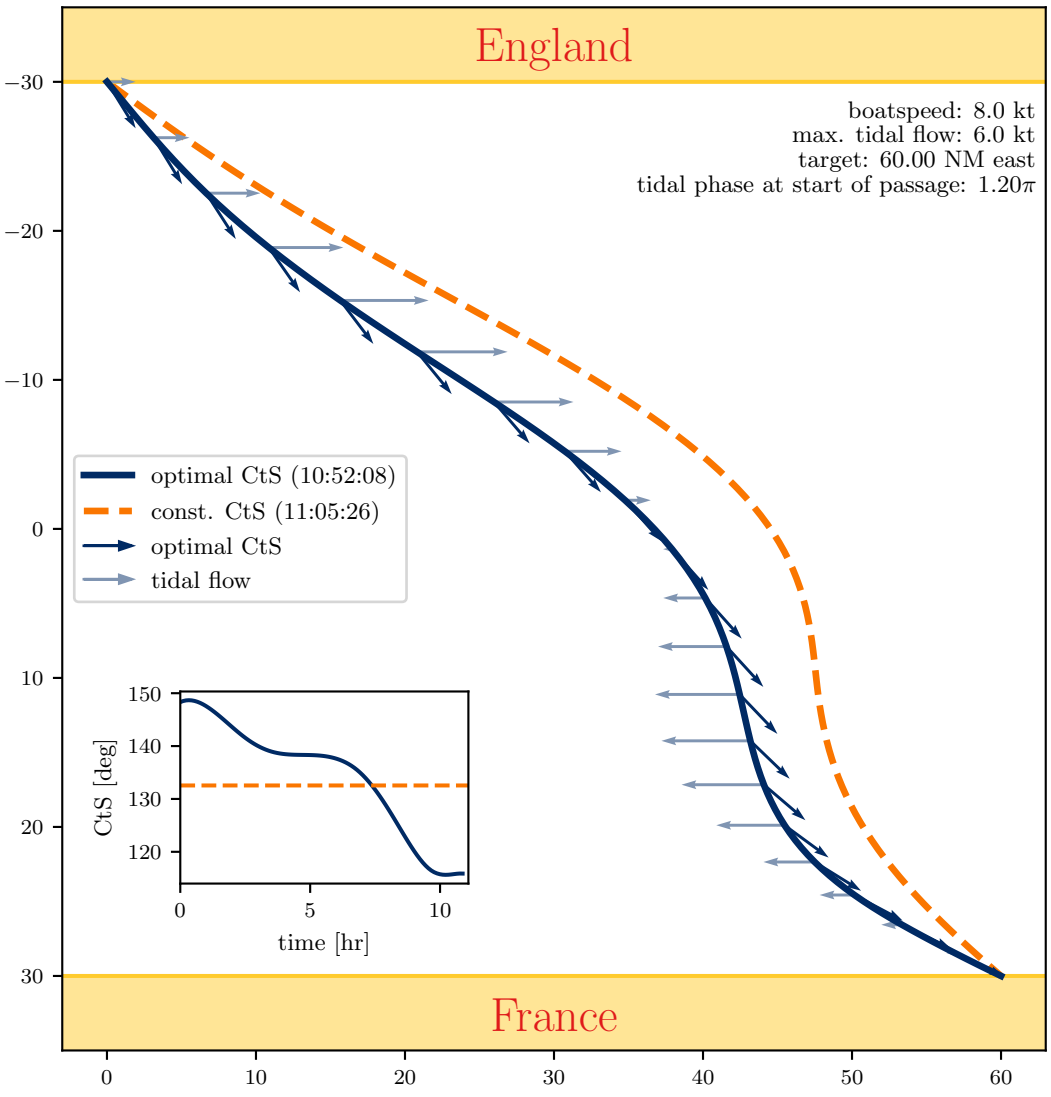

Figure B.5: Trajectories for an optimal course, with arrows representing the course to steer and tidal vectors along it, compared to a constant course for the best case improvement on a passage to "Dieppe".

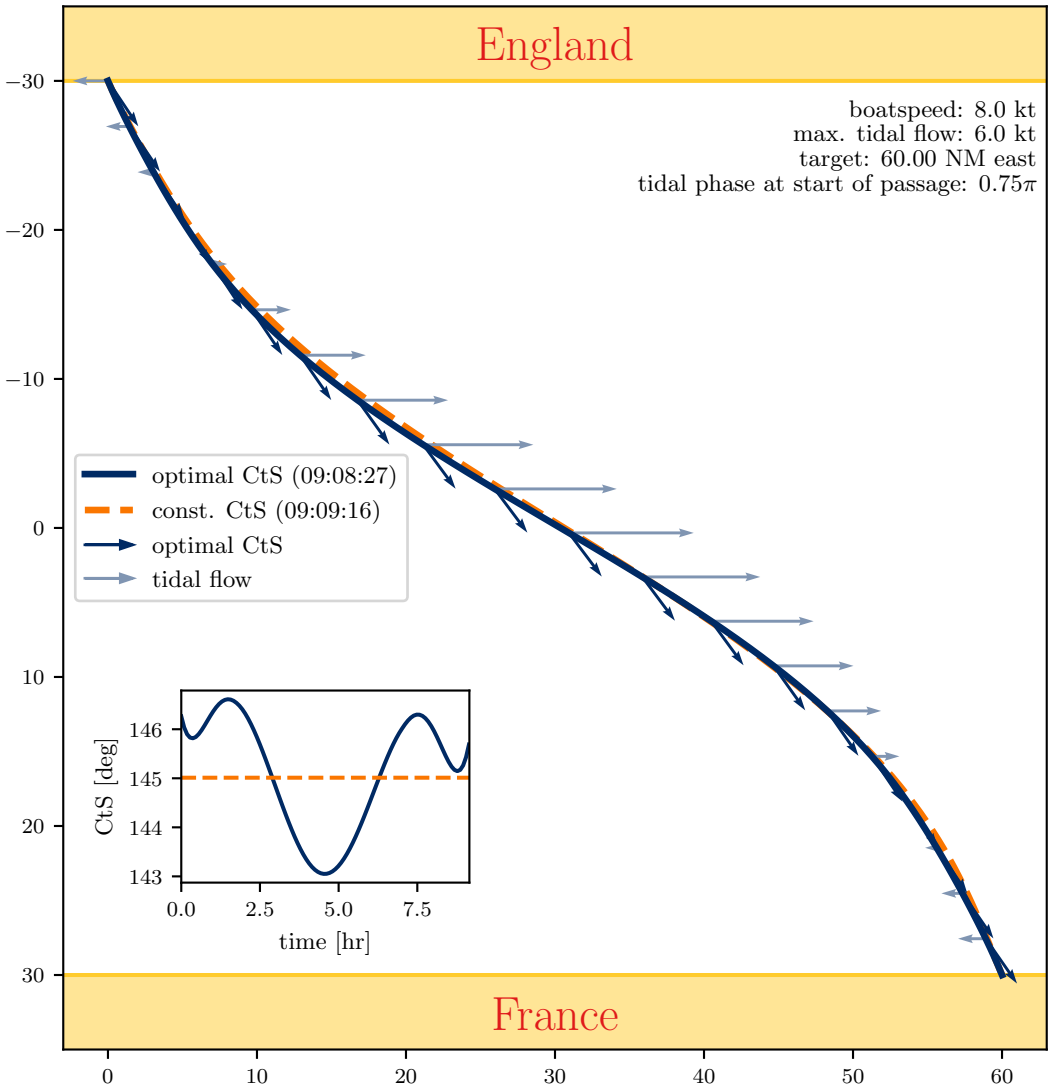

Figure B.6: Trajectories for an optimal course, with arrows representing the course to steer and tidal vectors along it, compared to a constant course on another passage to "Dieppe", this time leaving at a phase of the tide when not much can be done to improve on a constant course.

Conference posters

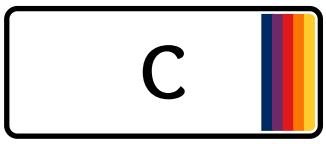

Did I ever tell you the story of how they arranged the names on the poster for 'The Towering Inferno' to give Steve McQueen and Paul Newman equal billing...

— Mark Kermode, *repeatedly.*

A number of posters, presented at academic conferences during the course of writing this thesis, will assault your retinas on the following pages.

Measurement of ultracold atom velocity distributions by matterwave interferometry

M Carey, J Saywell, D Elcock, M Belal & T Freegarde
School of Physics & Astronomy, University of Southampton, SO17 1BJ, UK

We present a method for the velocimetry of cold atoms by matterwave interferometry, which we have used to measure the velocity distributions within freely-expanding clouds of ^{85}Rb atoms at temperatures of 17 and 33 μK. Quadrature measurement using an asymmetric three-pulse technique allows determination of the full analytic interferometer signal which yields the 1-D velocity distribution with excellent fidelity. The technique is particularly suited to ultracold samples, and avoids some of the artefacts of conventional Doppler and time-of-flight methods.

Concept

The Fourier transform of data acquired by scanning the pulse separation of a Ramsey interferometer is the velocity distribution of the atomic sample

Matterwave interferometry is performed by pulsed interactions between two-level atoms and a resonant control field.

An initial pulse prepares atoms into an equal **superposition** of two states.

The effect of subsequent pulses depends upon the **phase accrued** between the atoms and the control field during periods of **free evolution** [1].

When the control field is a laser beam, whose phase varies spatially with an optical wavelength, we can think of it as a **ruler** with **sub-micron** spacing along which we can measure the motion of very slow-moving atoms using interferometry.

This picture suggests a method for measuring the speed of atoms along the ruler by measuring their distance travelled along it during a period T.

Consider a **Ramsey interferometer** [2], as in fig 1a, of two "π/2-pulses" separated by a period τ of free propagation. The first pulse prepares atoms in an equal superposition.

The phase accrued during the free propagation period τ for an atom with velocity v will be

$$\Phi = \mathbf{k}\cdot\mathbf{v}\tau = kv_z\tau,$$

where k is the laser wavenumber and vz is the projection of the atom's velocity along its axis.

The second pulse **maps this phase onto the atomic state probabilities** so that the fraction of atoms in the excited state at the end of a sequence in a statistical ensemble is

$$\mathcal{S}(\tau,\phi) = \int_{-\infty}^{\infty} P(v_z)\frac{1}{2}[1+\cos(kv_z\tau-\phi)]dv_z.$$

The parameter ϕ represents an advance in laser phase that can be introduced between the two pulses, allowing us to make a **quadrature measurement**. The velocity distribution P(v$_z$) can then be reconstructed by taking the **Fourier transform** of the quantity [3]

$$\mathcal{S}_I + i\mathcal{S}_Q \equiv \mathcal{S}(\tau,0) + \mathcal{S}(\tau,\pi/2).$$

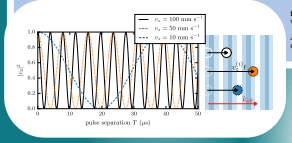

fig. 2: The interferometer output for atoms of different velocities is sinusoidal with a velocity-dependent frequency.

A general output will be a sum of these components, which can be separated by means of a Fourier transform.

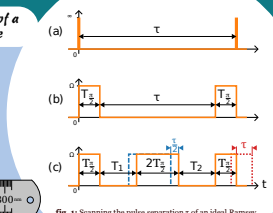

fig. 1: Scanning the pulse separation τ of an ideal Ramsey interferometer (a) reveals the atomic velocities. Phase shifts in real pulses (b) cancel in a Mach-Zehnder interferometer with a temporal asymmetry τ (c).

Implementation

We use an asymmetric Mach-Zehnder geometry to compensate for Doppler sensitive interferometer pulses

In a two pulse interferometer, the pulse separation τ can only be positive. This **truncation** in the time domain **broadens** the profile in the frequency (velocity) domain.

With perfect, instantaneous, interactions (fig 1a) that prepare all atoms in equal superpositions with the same phase, this broadening can be **deconvolved** [3].

In practice, pulses have a finite time and so have Doppler sensitivity themselves (fig 1b). **The atoms travel along the ruler during the pulses!**

An elegant solution is to employ a Mach-Zehnder sequence, with a "mirror" π-pulse between two "beamsplitter" π/2-pulses (fig 1c).

In its **time-symmetric** form (T1 = T2), atoms spend equal amounts of time in both states and velocity-dependent phases accrued in the first period of free evolution T1 are reversed and cancelled in the second T2.

If the apparatus undergoes acceleration, the phases no longer cancel. The output is a **differential** measurement of velocity: the basis of matterwave **inertial sensing**.

By introducing a **temporal asymmetry** τ, we reestablish the **velocity dependence** of the free evolution phase while maintaining **cancellation** of the phases accrued during the pulses.

Furthermore, as the sign of the asymmetry can vary, we are no longer limited to acquiring data for positive time values [4].

Experimental setup

Our interferometry uses ^{85}**Rb** atoms from a magneto-optical trap.

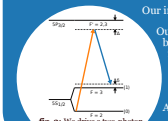

fig. 3: We drive a two-photon ground-state hyperfine transition .

Our "ruler" comprises two **counter-propagating** laser beams, detuned by ~10s of GHz from the 780nm D2 transition and separated in frequency by the ground state hyperfine splitting.

These beams drive 2-photon **Raman** transitions between the two ground states.

A frequency detuning δ of one of the beams tunes the interaction.

We benefit from the **stability** and **microwave frequency** of the ground state transition, while the counter-propagating arrangement means the "ruler" retains the sensitivity of the optical wavelength.

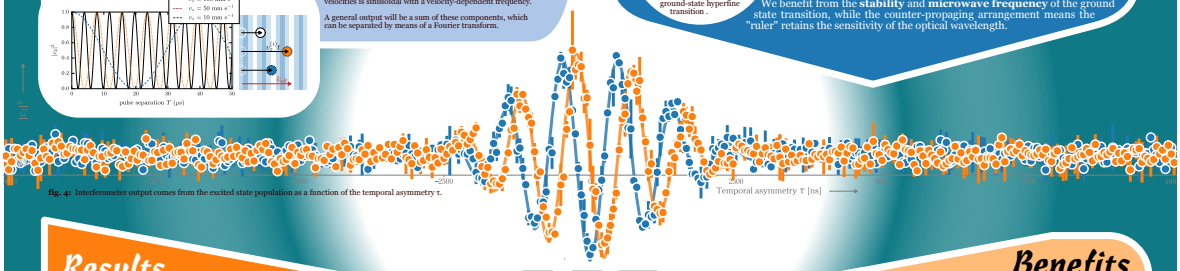

fig. 4: Interferometer output comes from the excited state population as a function of the temporal asymmetry τ.

FFT

Results

We take the **Fourier transform** of the background-subtracted interferometer output from **fig 4** and overlay the result upon a velocity profile taken by Doppler Raman spectroscopy in **fig 5a**.

Our lasers are detuned from the "undressed" two-photon resonance to compensate for an AC stark shift, and this detuning has been subtracted.

Gaussian fits to the Doppler and interferometric profiles are shown, with corresponding temperatures of 33μK and 30μK respectively.

This discrepancy disappears when the interferometric data are corrected to account for residual **Doppler sensitivity** within the pulses in **fig 5b**, where a dotted line shows the theoretical spectral profile of our pulses and the data in the highlighted central region have been multiplied by its reciprocal.

The interferometric data exhibit a noticeably improved **signal-to-noise ratio** compared to their spectroscopic counterparts [4].

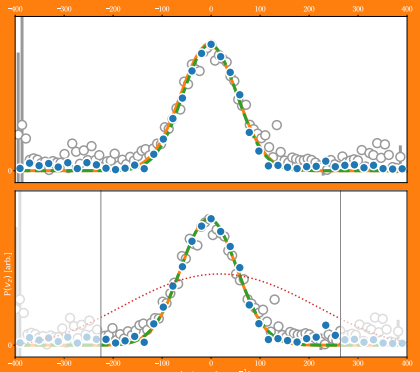

fig. 5: FFT of interferometer output reveals a velocity distribution in good agreement with Doppler spectroscopy.

Benefits

Other cold atom velocimetry methods (e.g. Doppler spectroscopy [5] or time-of-flight imaging [6]) rely on interacting with atoms in a **narrow spread of velocities** for each data point. Increasing the resolution reduces the **signal-to-noise ratio**.

In contrast, our method utilises strong interactions with **all atoms at once**; the resolution is limited by the **interferometer duration**.

Fig 6 demonstrates an interferometric measurement **free of the artefacts** present in a Doppler spectroscopy measurement:

Interferometric data are well fitted by a Gaussian for a temperature of 17μK, while spectroscopic data exhibit a second, broader component.

This broad background depends upon probe intensity, suggesting it could be due to **off-resonant excitation** [7]. Other researchers have noted similar phenomena [8-10]. If the velocity distribution of trapped atoms really extended this far, it should show in the interferometric measurement.

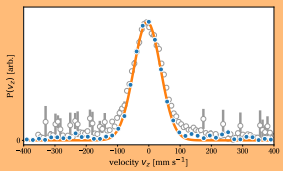

fig. 6: When used to measure a narrower 17μK distribution, our interferometric method proves to be free of broadening artefacts present in our spectroscopic measurement.

References

[1] T. L. Gustavson, et al. Class Quantum Gra. 17, 2385 (2000)
[2] N. F. Ramsey, Phys Rev 78, 695 (1950)
[3] M. Carey, et al. J Mod Opt 65, 657 (2018)
[4] M. Carey, et al. arXiv:1802.02190
[5] M. Kasevich, et al. Phys Rev Lett 66, 2297 (1991)
[6] P. D. Lett, et al. Phys Rev Lett 61, 169 (1988)
[7] I. G. Hughes, J Mod Opt 65, 640 (2018)
[8] P. D. Featonby et al. Phys Rev Lett 81, 495 (1998)
[9] L. P. Parazzoli et al. Phys Rev Lett 109, 230401 (2012)
[10] E. Marechal et al. Phys Rev A 62, 053603 (2000)

UNIVERSITY OF Southampton [dstl] EPSRC Pioneering research and skills

Figure C.1: 50$^{\text{th}}$ Annual Meeting of the APS Division of Atomic, Molecular and Optical Physics (DAMOP), Milwaukee, 27–31 May 2019.

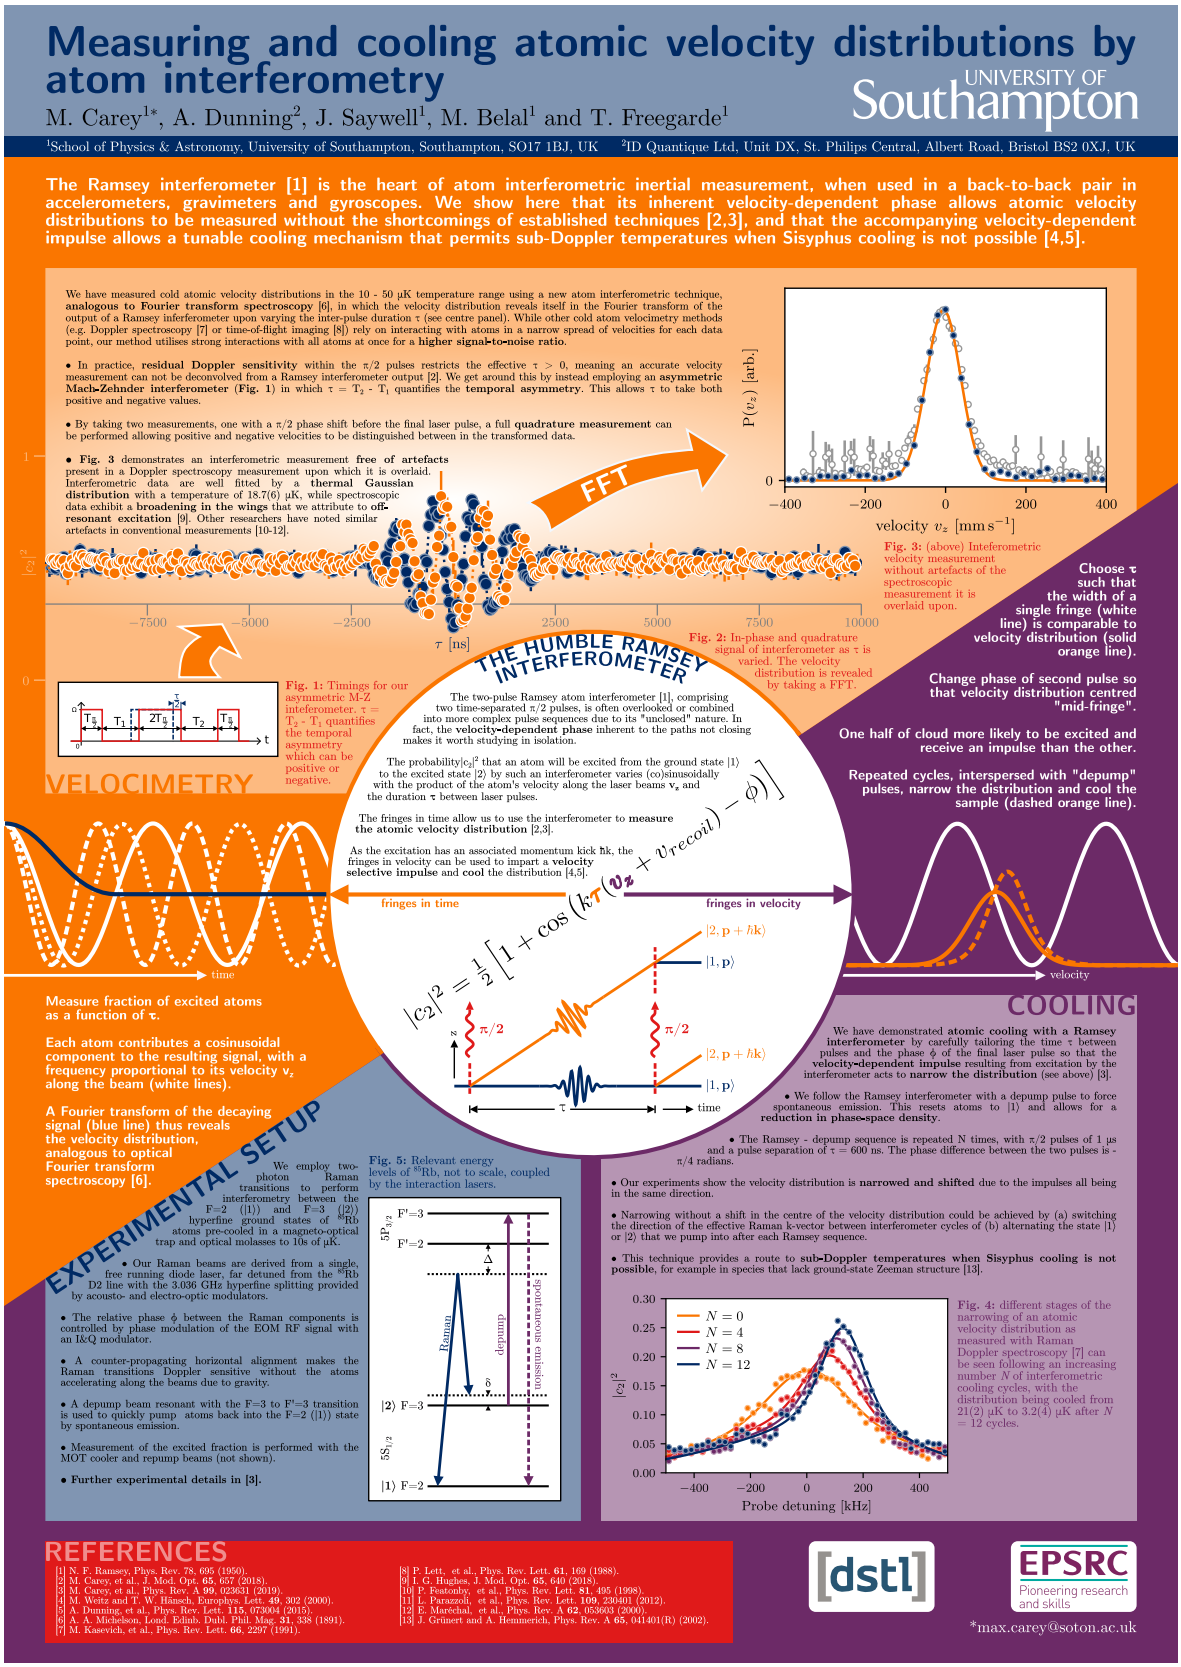

Figure C.2: International Conference of Laser Spectroscopy (ICOLS) 2019, Queenstown, 8–12 July 2019.

Optimised Raman pulses for atom interferometry

Jack Saywell*[1], **Max Carey***[1], Mohammad Belal[2], Ilya Kuprov[3], and Tim Freegarde[1]

*max.carey@soton.ac.uk; j.c.scaywell@soton.ac.uk

[1]School of Physics & Astronomy, [3]Department of Chemistry, University of Southampton, Highfield, Southampton, SO17 1BJ, UK

[2]National Oceanography Centre, European Way, Southampton, SO14 3ZH, UK

- The contrast and sensitivity of atom interferometers are limited by the fidelity of the atom-light interactions comprising the beamsplitter and mirror operations.
- We have used **optimal control theory** to design broadband Raman pulses that achieve **99.8(3)% state transfer** in a 35 μK ensemble of ^{85}Rb and, with a 100 μK sample, realise a **threefold improvement** in fringe visibility.
- We have also designed **augmentation pulses** that track the velocities of the two arms of a **large momentum transfer interferometer** and promise significant improvements over previous pulses including adiabatic rapid passage.

Atom interferometry

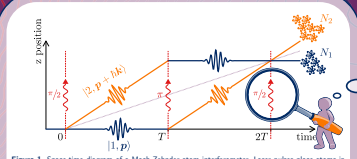

Figure 1. Space-time diagram of a Mach-Zehnder atom interferometer. Laser pulses place atoms in a superposition of two spatial paths. The phase accrued between the two paths is sensitive to external forces, and is mapped onto internal state probabilities so that it can be measured by counting atoms in each state at the output, provided the atom-laser interactions are perfect.

Atom interferometers (Figure 1) use laser pulses ("**atom optics**") to spatially separate and recombine atomic wavepackets, mapping the relative phase accrued — due to e.g. acceleration or rotation of the apparatus — onto the internal state probabilities.

This simple picture is complicated by imperfect light-atom interactions comprising the atom optics. In real experiments these are **fractional Rabi oscillations**, with π/2- and π-pulses acting as "beamsplitters" and "mirrors".

Experimental inhomogeneities mean these Rabi oscillations differ from atom to atom (Figure 2). No pulse duration is perfect for all atoms, leading to **contrast loss**.

Optimal control: gradient ascent pulse engineering

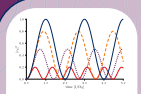

Figure 2. Atom optics comprise fractional Rabi oscillations. A perfect π/2- or π-pulse for one atom is imperfect for others which oscillate at different rates/ amplitudes due to experimental inhomogeneities.

NMR literature divides inhomogeneities that affect a **single atom-light interaction** into 2 categories:

- **Off-resonance errors** → non-zero detuning from resonance δ (finite temperature sample)
- **Pulse-length errors** → variations in coupling strength $\Omega_R/\Omega_{\text{eff}}$ (laser intensity variations)

Fringe contrast is lost when these variations are too large, but minimising them typically involves wasting atoms and/or laser power, **limiting signal-to-noise improvements**.

Inhomogeneous Rabi oscillations can be **rephased** with laser phase variations (Figure 3) [4], enabling the design of pulses **effective over a larger range of variations**. Contrast is improved at the expense of pulse duration, as in composite [5] and adiabatic rapid passage [6-8] pulses.

We design such pulses with **optimal control** [1], identifying laser phase $\phi(t)$ and amplitude $\Omega_R(t)$ as **control parameters** in the Hamiltonian

$$\hat{H}(t) = \frac{\hbar}{2}\begin{pmatrix} \delta & \Omega_R(t)e^{i\phi(t)} \\ \Omega_R(t)e^{-i\phi(t)} & -\delta \end{pmatrix}.$$

Expressing them as piecewise constant waveforms $\phi(t) = \phi_1, \phi_2, \ldots, \phi_N$ and $\Omega_R(t) = \Omega_1, \Omega_2, \ldots, \Omega_N$ with a fixed duration and number of steps N results in 2N discrete control parameters, for which the result of a pulse $|\psi_{\text{final}}\rangle$ can be computed and we can define a **fidelity** such as

$$\mathcal{F} = |\langle\psi_{\text{target}}|\psi_{\text{final}}\rangle|^2,$$

that we average over a range of errors characteristic of an experimental ensemble

$$\mathcal{F} \to \sum_{\delta,\Omega_R} \mathcal{F}(\delta, \Omega_R).$$

We compute the derivatives of this fidelity w.r.t. the control parameters and use the **gradient ascent pulse engineering (GRAPE)** algorithm [9,10] to **efficiently climb the fidelity landscape**, iteratively updating the pulse profile until a target value is attained.

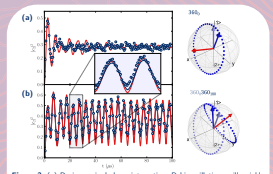

Figure 3. (a) During a single laser interaction, Rabi oscillations will quickly dephase. However, each atom still undergoes a coherent interaction with the laser. (b) Introducing a π phase jump to the laser every Rabi cycle brings the oscillations back into phase, revealing long-term coherences that can be used to design more robust atom optics.

Large momentum transfer

Large momentum transfer (LMT) atom interferometers gain a sensitivity advantage by increasing the enclosed spatial area with **additional atom optics** [8,11,12].

These should have high fidelity to avoid **reductions in fringe contrast cancelling out potential sensitivity gains**.

Increased momentum separation detunes the resonance condition for each arm (Figure 4). Pulses must address two regions of frequency space, maintaining **high fidelity over a large range of off-resonant errors**.

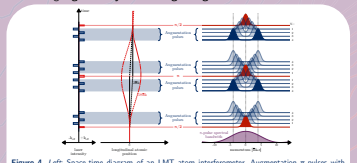

Figure 4. *Left:* Space-time diagram of an LMT atom interferometer. Augmentation π-pulses with alternating wavevectors increase momentum splitting between the arms for an inherent sensitivity increase. *Right:* Typical momentum distribution for 1 μK ensemble of rubidium atoms during such a sequence. The spectral bandwidth of a 200 kHz π-pulse (*bottom*) no longer covers both arms as they diverge in momentum space, limiting augmentation pulse fidelity and thus fringe contrast.

Apparatus

Figure 5. *Left:* We drive two-photon Raman transitions between ^{85}Rb ground states. *Right:* Phase profiles applied to phase of RF driving EOM that provides one Raman frequency component.

We perform **2-photon Raman transitions** between hyperfine ground states of **^{85}Rb atoms**. Read-out measures the fraction of atoms excited from the lower to upper level.

Raman frequency components realised with acousto- and electro-optical modulators impinge upon the atoms from opposing horizontal directions.

Optimised phase sequences are applied to the **phase of the RF driving the EOM**, shifting the beatnote phase.

Optimised 3-pulse Mach-Zehnder interferometer

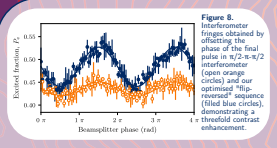

Figure 8. Interferometer fringes obtained by offsetting the phase of the final pulse in π/2-π-π/2 interferometer (open orange circles) and our optimised "flip-reverse" sequence (filled blue circles), demonstrating a 4 π threefold contrast enhancement.

We optimised a **3-pulse Mach-Zehnder** interferometer sequence [2] that yields a **threefold increase in fringe contrast** in a proof-of-principle interferometer with a 95 μK ensemble of rubidium atoms (Figure 8).

While many pulses designed for efficient state transfer introduce phase variations that make them ill-suited to atom interferometry [1], our **"flip-reverse"** approach (Figure 9) gives the interferometer a temporal anti-symmetry that ensures **inhomogeneities do not affect the interferometric phase**.

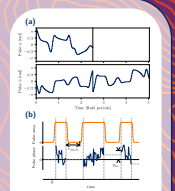

Figure 9. (a) Optimised beamsplitter (*top*) and mirror pulse (*bottom*). The pulses are concatenated as shown in (b), with the final beamsplitter pulse flipped in phase and reversed in time to give the entire sequence a temporal anti-symmetry. In this proof of principle, τ_{pulse}=100 μs and a phase offset is applied to the final pulse to resolve fringes.

Broadband Atomic Transfer pulse

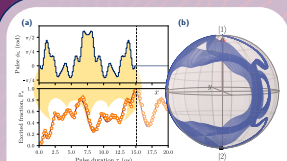

Figure 6. (a) *Above:* BAT pulse Raman phase as a function of time and *below:* fraction of atoms in a 35 μK ensemble measured in the excited state throughout the pulse, peaking at the end of the phase modulation. (b) Simulated Bloch sphere trajectories for a range of atoms during the pulse.

We have optimised a **broadband atomic transfer (BAT)** pulse [2] with a duration of ~3.5 Rabi periods by averaging the fidelity

$$\mathcal{F}_{\text{eq.}} = |\langle e|\psi_{\text{final}}\rangle|^2$$

over ±10% Rabi frequency variations and a range of detunings corresponding to a 120 μK thermal ensemble of rubidium atoms ($|\delta| \lesssim 1.5\Omega_R$ for our experimental Rabi frequency $\Omega_R \approx 2\pi \times 300$ kHz).

Figure 6(a) shows the fraction of a 35 μK ensemble excited throughout the pulse, **peaking at 99.8(3)%** when the modulations ends. This compares with 75(3)% after a π-pulse in the same experimental conditions.

Figure 6(b) shows simulated Bloch sphere trajectories for a range of off-resonance errors; they diverge during the pulse but **reconverge upon a single point at the end**. A range of such simulations, reflecting experimental inhomogeneities, yields the solid purple theory line in Figure 6(a).

The BAT pulse exhibits **>90% transfer over a four times broader detuning range** than the WALTZ composite pulse, similarly designed for efficient state transfer.

This large detuning acceptance makes the BAT pulse **attractive for LMT** (see panel above). WALTZ [11] and ARP [8] state transfer pulses have already proved effective as augmentation pulses which are typically applied in pairs so that the **phase factors introduced by one pulse can be cancelled out by a subsequent one**.

The fidelity gains for optimal control pulses are particularly striking at **short pulse durations** (Figure 7(b)). This is useful for **mitigating spontaneous emission**, another source of contrast loss in interferometers comprised of many concatenated pulses.

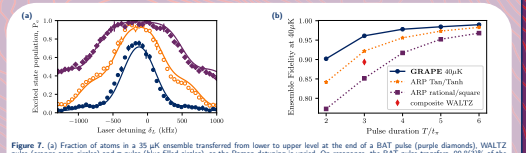

Figure 7. (a) Fraction of atoms in a 35 μK ensemble transferred from lower to upper level at the end of a BAT pulse (purple diamonds), WALTZ pulse (orange open circles) and π-pulse (blue filled circles), as the Raman detuning is varied. On resonance, the BAT pulse transfers 99.8(3)% of the ensemble — compared with 96(2)% for WALTZ and 75(3)% for a π-pulse — and maintains >90% transfer over a 400 kHz range.
(b) Simulated fidelities of gradient ascent optimised (GRAPE) pulses in a 40 μK ensemble for a range of pulse durations, compared to common ARP pulse shapes and the WALTZ composite pulse. The relative fidelity gains are most apparent at short pulse durations, favourable for minimising spontaneous emission.

Bi-selective optimisation for LMT

Broadband state transfer pulses make good LMT augmentation pulses, but their finite bandwidth still **limits the achievable momentum splitting.**

We propose an alternative [3], using a different optimised pulse for each augmentation pulse, tailored to the expected momentum distribution.

These **bi-selective** pulses remain resonant with both arms, **"tracking"** their separation **without an increase in pulse duration** (Figure 10).

Figure 11(a) shows simulated LMT contrast, comparing π augmentation pulses to our bi-selective pulses and ARP pulses of the same duration. Bi-selective pulses retain **>60% contrast** in a **42hk** interferometer where **ARP contrast is <5%**.

The phase cancellation that makes state transfer pulses good for LMT is perfect if the Rabi frequency is constant throughout the interferometer, but **variations in Rabi frequency between pulses cause a steep degradation in contrast**, evident in Figure 11(b). This can be overcome by optimising pulses with a **phase-sensitive fidelity measure**

$$\mathcal{F}_{\text{ps.}} = \text{Re}\left(\langle e|\psi_{\text{final}}\rangle\right),$$

yielding efficient state transfer pulses of well-defined phase over a range of Rabi frequencies. Shown by red points in Figure 11(b), this results in a retention of high simulated contrast, **resilient to the Rabi frequency variations that degrade other schemes.**

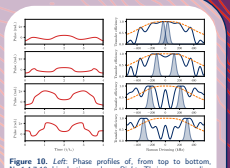

Figure 10. *Left:* Phase profiles of, from top to bottom, N=1,4,7,10 bi-selective pulses. *Right:* Their corresponding frequency responses (solid blue lines). The high-fidelity peaks in frequency space track the divergence of a 1 μK atomic momentum distribution (filled Gaussians) throughout the interferometer. The response of an equal duration ARP pulse is shown (dashed orange lines) for comparison.

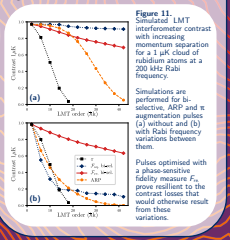

Figure 11. Simulated LMT interferometer contrast with increasing momentum separation for a 1 μK cloud of rubidium atoms at a 200 kHz Rabi frequency.

Simulations are performed for bi-selective, ARP and π augmentation pulses (a) without and (b) with Rabi frequency variations between them.

Pulses optimised with a phase-sensitive fidelity measure $\mathcal{F}_{\text{ps}}$ prove resilient to the contrast losses that would otherwise result from these variations.

References

[1] J. Saywell *et al.*, Phys. Rev. A **98**, 023625 (2018)
[2] J. Saywell *et al.*, J. Phys. B, *accepted* (2020)
[3] J. Saywell *et al.*, *in preparation*
[4] A. Dunning *et al.*, Phys. Rev. A **90**, 033608 (2014)
[5] A. J. Shaka et al., J. Magn. Reson. 1969 52, 335 (1983)
[6] T. L. Hwang *et al.*, J. Magn. Reson. **133**, 200, (1998)
[7] C. J. Hardy *et al.*, J. Magn. Reson. 1969 **66**, 470 (1986)
[8] K. Kotru et al., Phys. Rev. Lett. 155, 103001 (2015)
[9] N. Khaneja et al., J. Magn. Reson. 172, 296 (2005)
[10] H. J. Hogben et al., J. Magn. Reson. 208, 179 (2011)
[11] D. Butts *et al.*, JOSA B **30**, 922 (2013)
[12] J. M. McGuirk *et al.*, Phys. Rev. Lett. **85**, 4498 (2000)

NCM dstl EPSRC Pioneering research and skills UNIVERSITY OF Southampton

Figure C.3: Quantum Optics 2020, Obergurgl, 23–29 February 2020.

Bibliography

[1] M. Carey, M. Belal, M. Himsworth, J. Bateman, and T. Freegarde, "Matterwave interferometric velocimetry of cold Rb atoms," Journal of Modern Optics **65**, 657 (2018).

[2] J. C. Saywell, I. Kuprov, D. Goodwin, M. Carey, and T. Freegarde, "Optimal control of mirror pulses for cold-atom interferometry," Physical Review A **98**, 023625 (2018).

[3] M. Carey, J. Saywell, D. Elcock, M. Belal, and T. Freegarde, "Velocimetry of cold atoms by matter-wave interferometry," Physical Review A **99**, 023631 (2019).

[4] J. Dalton, *A New System of Chemical Philosophy* (Cambridge University Press, Cambridge, 2010).

[5] A. Einstein, "Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen," Annalen der Physik **17**, 549 (1905).

[6] J. J. Thomson, "Cathode rays," The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **44**, 293 (1897).

[7] N. Robotti and F. Pastorino, "Zeeman's discovery and the mass of the electron," Annals of Science **55**, 161 (1998).

[8] E. Rutherford, "The scattering of α and β particles by matter and the structure of the atom," The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **21**, 669 (1911).

[9] P. H. Lee and M. L. Skolnick, "Saturated neon absorption inside a 6238-Å laser," Applied Physics Letters **10**, 303 (1967).

[10] S. Haroche and F. Hartmann, "Theory of saturated-absorption line shapes," Physical Review A **6**, 1280 (1972).

[11] D. J. Wineland and W. M. Itano, "Laser cooling of atoms," Physical Review A **20**, 1521 (1979).

[12] A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, "First observation of magnetically trapped neutral atoms," Physical Review Letters **54**, 2596 (1985).

[13] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," Physical Review Letters **59**, 2631 (1987).

[14] N. F. Ramsey, "A molecular beam resonance method with separated oscillating fields," Physical Review **78**, 695 (1950).

[15] N. F. Ramsey, "Radiofrequency spectra of H_2 and D_2 by a new molecular beam resonance method," Physica **17**, 388 (1951).

[16] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye, "A Fermi-degenerate three-dimensional optical lattice clock," Science **358**, 90 (2017).

[17] S. Weyers, V. Gerginov, M. Kazda, J. Rahm, B. Lipphardt, G. Dobrev, and K. Gibble, "Advances in the accuracy, stability, and reliability of the PTB primary fountain clocks," Metrologia **55**, 789 (2018).

[18] R. J. Hendricks, F. Ozimek, K. Szymaniec, B. Nagorny, P. Dunst, J. Nawrocki, S. Beattie, B. Jian, and K. Gibble, "Cs fountain clocks for commercial realizations — an improved and robust design," IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **66**, 624 (2019).

[19] J. L. Hall, C. J. Bordé, and K. Uehara, "Direct optical resolution of the recoil effect using saturated absorption spectroscopy," Physical Review Letters **37**, 1339 (1976).

[20] T. W. Mossberg and S. R. Hartmann, "Diagrammatic representation of photon echoes and other laser-induced ordering processes in gases," Physical Review A **23**, 1271 (1981).

[21] C. J. Bordé, "Atomic interferometry with internal state labelling," Physics Letters A **140**, 10 (1989).

[22] M. Kasevich and S. Chu, "Atomic interferometry using stimulated Raman transitions," Physical Review Letters **67**, 181 (1991).

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

[24] B. Barrett and A. Kumarakrishnan, "Atom interferometric techniques for measuring gravitational acceleration and constant magnetic field gradients," Physics **3**, 16 (2011).

[25] P. A. Altin, M. T. Johnsson, V. Negnevitsky, G. R. Dennis, R. P. Anderson, J. E. Debs, S. S. Szigeti, K. S. Hardman, S. Bennetts, G. D. McDonald, L. D. Turner, J. D. Close, and N. P. Robins, "Precision atomic gravimeter based on Bragg diffraction," New Journal of Physics **15**, 023009 (2013).

[26] T. Farah, C. Guerlin, A. Landragin, P. Bouyer, S. Gaffet, F. Pereira Dos Santos, and S. Merlet, "Underground operation at best sensitivity of the mobile LNE-SYRTE cold atom gravimeter," Gyroscopy and Navigation **5**, 266 (2014).

[27] T. Kovachy, P. Asenbaum, C. Overstreet, C. A. Donnelly, S. M. Dickerson, A. Sugarbaker, J. M. Hogan, and M. A. Kasevich, "Quantum superposition at the half-metre scale," Nature **528**, 530 (2015).

[28] S. Dimopoulos, P. W. Graham, J. M. Hogan, and M. A. Kasevich, "Testing general relativity with atom interferometry," Physical Review Letters **98**, 111102 (2007).

[29] L. Zhou, S. Long, B. Tang, X. Chen, F. Gao, W. Peng, W. Duan, J. Zhong, Z. Xiong, J. Wang, Y. Zhang, and M. Zhan, "Test of equivalence principle at 10^{-8} level by a dual-species double-diffraction Raman atom interferometer," Physical Review Letters **115**, 013004 (2015).

[30] B. Altschul, Q. G. Bailey, L. Blanchet, K. Bongs, P. Bouyer, L. Cacciapuoti, S. Capozziello, N. Gaaloul, D. Giulini, J. Hartwig, L. Iess, P. Jetzer, A. Landragin, E. Rasel, S. Reynaud, S. Schiller, C. Schubert, F. Sorrentino, U. Sterr, J. D. Tasson, G. M. Tino, P. Tuckey, and P. Wolf, "Quantum tests of the Einstein Equivalence Principle with the STE–QUEST space mission," Advances in Space Research **55**, 501 (2015).

[31] B. Barrett, L. Antoni-Micollier, L. Chichet, B. Battelier, P.-A. Gominet, A. Bertoldi, P. Bouyer, and A. Landragin, "Correlative methods for dual-species quantum tests of the weak equivalence principle," New Journal of Physics **17**, 085010 (2015).

[32] J. Williams, S.-w. Chiow, N. Yu, and H. Müller, "Quantum test of the equivalence principle and space-time aboard the International Space Station," New Journal of Physics **18**, 025018 (2016).

[33] D. S. Weiss, B. C. Young, and S. Chu, "Precision measurement of $\hbar/m_{Cs}$ based on photon recoil using laser-cooled atoms and atomic interferometry," Applied Physics B **59**, 217 (1994).

[34] D. Hanneke, S. Fogwell, and G. Gabrielse, "New measurement of the electron magnetic moment and the fine structure constant," Physical Review Letters **100**, 120801 (2008).

[35] 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," Physical Review Letters **106**, 080801 (2011).

[36] R. H. Parker, C. Yu, W. Zhong, B. Estey, and H. Müller, "Measurement of the fine-structure constant as a test of the Standard Model," Science **360**, 191 (2018).

[37] T. L. Gustavson, *Precision Rotation Sensing Using Atom Interferometry*, Ph.D. thesis, Stanford University (2000).

[38] J. Fang and J. Qin, "Advances in atomic gyroscopes: a view from inertial navigation applications," Sensors **12**, 6331 (2012).

[39] 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," Physical Review Letters **111**, 083001 (2013).

[40] B. Barrett, A. Bertoldi, and P. Bouyer, "Inertial quantum sensors using light and matter," Physica Scripta **91**, 053006 (2016).

[41] S. Riedl, G. W. Hoth, B. Pelle, J. Kitching, and E. A. Donley, "Compact atom-interferometer gyroscope based on an expanding ball of atoms," Journal of Physics: Conference Series **723**, 012058 (2016).

[42] P. Cheiney, L. Fouché, S. Templier, F. Napolitano, B. Battelier, P. Bouyer, and B. Barrett, "Navigation-compatible hybrid quantum accelerometer using a Kalman filter," Physical Review Applied **10**, 034030 (2018).

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

[44] A. Dunning, R. Gregory, J. Bateman, M. Himsworth, and T. Freegarde, "Interferometric laser cooling of atomic rubidium," Physical Review Letters **115**, 073004 (2015).

[45] M. Kritsotakis, S. S. Szigeti, J. A. Dunningham, and S. A. Haine, "Optimal matter-wave gravimetry," Physical Review A **98**, 023629 (2018).

[46] G. D'Amico, L. Cacciapuoti, M. Jain, S. Zhan, and G. Rosi, "Measuring the gravitational acceleration with matter-wave velocimetry," The European Physical Journal D **73**, 98 (2019).

[47] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, "Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms," Journal of Magnetic Resonance **172**, 296 (2005).

[48] H. Hogben, M. Krzystyniak, G. Charnock, P. Hore, and I. Kuprov, "Spinach – A software library for simulation of spin dynamics in large spin systems," Journal of Magnetic Resonance **208**, 179 (2011).

[49] P. de Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, "Second order gradient ascent pulse engineering," Journal of Magnetic Resonance **212**, 412 (2011).

[50] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, "Training Schrödinger's cat: quantum optimal control," The European Physical Journal D **69**, 279 (2015).

[51] J. Saywell, M. Carey, M. Belal, I. Kuprov, and T. Freegarde, "Optimal control of Raman pulse sequences for atom interferometry," Journal of Physics B: Atomic, Molecular and Optical Physics **53**, 085006 (2020).

[52] M. H. Levitt and R. Freeman, "Compensation for pulse imperfections in NMR spin-echo experiments," Journal of Magnetic Resonance **43**, 65 (1981).

[53] S. Odedra and S. Wimperis, "Use of composite refocusing pulses to form spin echoes," Journal of Magnetic Resonance **214**, 68 (2012).

[54] J. M. McGuirk, G. T. Foster, J. B. Fixler, M. J. Snadden, and M. A. Kasevich, "Sensitive absolute-gravity gradiometry using atom interferometry," Physical Review A **65**, 1 (2002).

[55] 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," Journal of the Optical Society of America B **30**, 922 (2013).

[56] A. Dunning, R. Gregory, J. Bateman, N. Cooper, M. Himsworth, J. A. Jones, and T. Freegarde, "Composite pulses for interferometry in a thermal cold atom cloud," Physical Review A **90**, 033608 (2014).

[57] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, "Fast machine-learning online optimization of ultra-cold-atom experiments," Scientific Reports **6**, 25890 (2016).

[58] R. Heck, O. Vuculescu, J. J. Sørensen, J. Zoller, M. G. Andreasen, M. G. Bason, P. Ejlertsen, O. Elíasson, P. Haikka, J. S. Laustsen, L. L. Nielsen, A. Mao, R. Müller, M. Napolitano, M. K. Pedersen, A. R. Thorsen, C. Bergenholtz, T. Calarco, S. Montangero, and J. F. Sherson, "Remote optimization of an ultracold atoms experiment by experts and citizen scientists," Proceedings of the National Academy of Sciences **115**, E11231 (2018).

[59] G. Jäger, D. M. Reich, M. H. Goerz, C. P. Koch, and U. Hohenester, "Optimal quantum control of Bose-Einstein condensates in magnetic microtraps: Comparison of gradient-ascent-pulse-engineering and Krotov optimization schemes," Physical Review A **90**, 033628 (2014).

[60] S. van Frank, A. Negretti, T. Berrada, R. Bücker, S. Montangero, J.-F. Schaff, T. Schumm, T. Calarco, and J. Schmiedmayer, "Interferometry with non-classical motional states of a Bose–Einstein condensate," Nature Communications **5**, 4009 (2014).

[61] M. H. Goerz, E. J. Halperin, J. M. Aytac, C. P. Koch, and K. B. Whaley, "Robustness of high-fidelity Rydberg gates with single-site addressability," Physical Review A **90**, 032329 (2014).

[62] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, "Parallel implementation of high-fidelity multi-qubit gates with neutral atoms," (2019), arXiv:1908.06101 .

[63] T. Keating, C. H. Baldwin, Y.-Y. Jau, J. Lee, G. W. Biedermann, and I. H. Deutsch, "Arbitrary Dicke-state control of symmetric Rydberg ensembles," Physical Review Letters **117**, 213601 (2016).

[64] J. M. McGuirk, M. J. Snadden, and M. A. Kasevich, "Large area light-pulse atom interferometry," Physical Review Letters **85**, 4498 (2000).

[65] H. Müller, S. W. Chiow, Q. Long, S. Herrmann, and S. Chu, "Atom interferometry with up to 24-photon-momentum-transfer beam splitters," Physical Review Letters **100**, 180405 (2008).

[66] T. Lévèque, A. Gauguet, F. Michaud, F. Pereira Dos Santos, and A. Landragin, "Enhancing the area of a Raman atom interferometer using a versatile double-diffraction technique," Physical Review Letters **103**, 080405 (2009).

[67] M. Jaffe, V. Xu, P. Haslinger, H. Müller, and P. Hamilton, "Efficient adiabatic spin-dependent kicks in an atom interferometer," Physical Review Letters **121**, 040402 (2018).

[68] M. D. Barrett, J. A. Sauer, and M. S. Chapman, "All-optical formation of an atomic Bose-Einstein condensate," Physical Review Letters **87**, 010404 (2001).

[69] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer, "All-optical runaway evaporation to Bose-Einstein condensation," Physical Review A **79**, 061406 (2009).

[70] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews, and W. Ketterle, "Evaporative cooling of sodium atoms," Physical Review Letters **74**, 5202 (1995).

[71] J. L. Roberts, N. R. Claussen, J. P. Burke, C. H. Greene, E. A. Cornell, and C. E. Wieman, "Resonant magnetic field control of elastic scattering in cold ^{85}Rb," Physical Review Letters **81**, 5109 (1998).

[72] D. A. Steck, "Rubidium 85 D line data," (revision 2.1.6, 20 September 2013), available online at `https://steck.us/alkalidata/`.

[73] D. A. Steck, "Rubidium 87 D line data," (revision 2.1.5, 13 January 2015), available online at `https://steck.us/alkalidata/`.

[74] F. Bloch, "Nuclear induction," Physica **17**, 272 (1951).

[75] G. G. Stokes, "On the composition and resolution of streams of polarized light from different sources," Transactions of the Cambridge Philosophical Society **9**, 399 (1851), (also published by Cambridge University Press, DOI:10.1017/CBO9780511702266.010).

[76] R. P. Feynman, F. L. Vernon, and R. W. Hellwarth, "Geometrical representation of the Schrödinger equation for solving maser problems," Journal of Applied Physics **28**, 49 (1957).

[77] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, New York, 1995).

[78] I. I. Rabi, "Space quantization in a gyrating magnetic field," Physical Review **51**, 652 (1937).

[79] C. J. Foot, *Atomic Physics*, Oxford Master Series in Physics (Oxford University Press, Oxford, 2004).

[80] B. R. Mollow, "Power spectrum of light scattered by two-level systems," Physical Review **188**, 1969 (1969).

[81] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom–Photon Interactions* (Wiley-VCH, Weinheim, 2004).

[82] M. Jaffe, *Atom Interferometry in an Optical Cavity*, Ph.D. thesis, University of California, Berkeley (2018).

[83] M. M. T. Loy, "Observation of population inversion by optical adiabatic rapid passage," Physical Review Letters **32**, 814 (1974).

[84] V. Malinovsky and J. Krause, "General theory of population transfer by adiabatic rapid passage with intense, chirped laser pulses," The European Physical Journal D **14**, 147 (2001).

[85] K. Kotru, J. M. Brown, D. L. Butts, J. M. Kinast, and R. E. Stoner, "Robust Ramsey sequences with Raman adiabatic rapid passage," Physical Review A **90**, 053611 (2014).

[86] P. Tommasini, E. Timmermans, and A. F. R. de Toledo Piza, "The hydrogen atom as an entangled electron–proton system," American Journal of Physics **66**, 881 (1998).

[87] N. L. Harshman and K. S. Ranade, "Observables can be tailored to change the entanglement of any pure state," Physical Review A **84**, 12303 (2011).

[88] B. Young, M. Kasevich, and S. Chu, "Precision atom interferometry with light pulses," in *Atom Interferometry*, edited by P. R. Berman (Academic Press, San Diego, 1997) pp. 363–406.

[89] A. J. Dunning, *Coherent atomic manipulation and cooling using composite optical pulse sequences*, Ph.D. thesis, University of Southampton (2014), (also published by Springer, DOI:10.1007/978-3-319-21738-3).

[90] L. Allen and C. Stroud, "Broadening and saturation in n-photon absorption," Physics Reports **91**, 1 (1982).

[91] B. W. Shore, "Coherent manipulations of atoms using laser light," Acta Physica Slovaca **58**, 243 (2008).

[92] E. Brion, L. H. Pedersen, and K. Mølmer, "Adiabatic elimination in a lambda system," Journal of Physics A **40**, 1033 (2007).

[93] V. Paulisch, H. Rui, H. K. Ng, and B.-G. Englert, "Beyond adiabatic elimination: A hierarchy of approximations for multi-photon processes," The European Physical Journal Plus **129**, 12 (2014).

[94] R. Han, H. Khoon Ng, and B.-G. Englert, "Raman transitions without adiabatic elimination: a simple and accurate treatment," Journal of Modern Optics **60**, 255 (2013).

[95] R. J. Roy, *Ytterbium and Lithium Quantum Gases: Heteronuclear Molecules and Bose-Fermi Superfluid Mixtures*, Ph.D. thesis, University of Washington (2017).

[96] J. Dalibard and C. Cohen-Tannoudji, "Dressed-atom approach to atomic motion in laser light: the dipole force revisited," Journal of the Optical Society of America B **2**, 1707 (1985).

[97] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer, New York, 1999).

[98] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook, "Optical molasses," Journal of the Optical Society of America B **6**, 2084 (1989).

[99] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven, "Two-dimensional magneto-optical trap as a source of slow atoms," Physical Review A **58**, 3891 (1998).

[100] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. B. Ovchinnikov, and T. Pfau, "Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap," Physical Review A **66**, 023410 (2002).

[101] J. Catani, P. Maioli, L. De Sarlo, F. Minardi, and M. Inguscio, "Intense slow beams of bosonic potassium isotopes," Physical Review A **73**, 033415 (2006).

[102] J. Dalibard and C. Cohen-Tannoudji, "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models," Journal of the Optical Society of America B **6**, 2023 (1989).

[103] D. Rio Fernandes, *Trapping and Cooling of Fermionic Alkali Atoms to Quantum Degeneracy. Sub-Doppler Cooling of Potassium-40 and Lithium-6 in Gray Molasses*, Ph.D. thesis, Université Pierre et Marie Curie (2014).

[104] S. M. Barnett, "On the recoil and Doppler shifts," Journal of Modern Optics **57**, 1445 (2010).

[105] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji, "Laser cooling below the one-photon recoil energy by velocity-selective coherent population trapping," Physical Review Letters **61**, 826 (1988).

[106] F. Papoff, F. Mauri, and E. Arimondo, "Transient velocity-selective coherent population trapping in one dimension," Journal of the Optical Society of America B **9**, 321 (1992).

[107] M. S. Shahriar, P. R. Hemmer, M. G. Prentiss, P. Marte, J. Mervis, D. P. Katz, N. P. Bigelow, and T. Cai, "Continuous polarization-gradient precooling-assisted velocity-selective coherent population trapping," Physical Review A **48**, R4035 (1993).

[108] M. Weidemüller, T. Esslinger, M. A. Ol'shanii, A. Hemmerich, and T. W. Hänsch, "A novel scheme for efficient cooling below the photon recoil limit," Europhysics Letters **27**, 109 (1994).

[109] D. Finkelstein-Shapiro, S. Felicetti, T. Hansen, T. Pullerits, and A. Keller, "Classification of dark states in multilevel dissipative systems," Physical Review A **99**, 053829 (2019).

[110] D. Boiron, C. Triché, D. R. Meacher, P. Verkerk, and G. Grynberg, "Three-dimensional cooling of cesium atoms in four-beam gray optical molasses," Physical Review A **52**, R3425 (1995).

[111] D. Rio Fernandes, F. Sievers, N. Kretzschmar, S. Wu, C. Salomon, and F. Chevy, "Sub-Doppler laser cooling of fermionic ^{40}K atoms in three-dimensional gray optical molasses," Europhysics Letters **100**, 63001 (2012).

[112] A. T. Grier, I. Ferrier-Barbut, B. S. Rem, M. Delehaye, L. Khaykovich, F. Chevy, and C. Salomon, "Λ-enhanced sub-Doppler cooling of lithium atoms in D_1 gray molasses," Physical Review A **87**, 063411 (2013).

[113] D. Nath, R. K. Easwaran, G. Rajalakshmi, and C. S. Unnikrishnan, "Quantum-interference-enhanced deep sub-Doppler cooling of ^{39}K atoms in gray molasses," Physical Review A **88**, 053407 (2013).

[114] F. Sievers, N. Kretzschmar, D. R. Fernandes, D. Suchet, M. Rabinovic, S. Wu, C. V. Parker, L. Khaykovich, C. Salomon, and F. Chevy, "Simultaneous sub-Doppler laser cooling of fermionic ^{6}Li and ^{40}K on the D_1 line: Theory and experiment," Physical Review A **91**, 023426 (2015).

[115] K. N. Jarvis, J. A. Devlin, T. E. Wall, B. E. Sauer, and M. R. Tarbutt, "Blue-detuned magneto-optical trap," Physical Review Letters **120**, 083201 (2018).

[116] D. Becker, M. D. Lachmann, S. T. Seidel, H. Ahlers, A. N. Dinkelaker, J. Grosse, O. Hellmig, H. Müntinga, V. Schkolnik, T. Wendrich, A. Wenzlawski, B. Weps, R. Corgier, T. Franz, N. Gaaloul, W. Herr, D. Lüdtke, M. Popp, S. Amri, H. Duncker, M. Erbe, A. Kohfeldt, A. Kubelka-Lange, C. Braxmaier, E. Charron, W. Ertmer, M. Krutzik, C. Lämmerzahl, A. Peters, W. P. Schleich, K. Sengstock, R. Walser, A. Wicht, P. Windpassinger, and E. M. Rasel, "Space-borne Bose-Einstein condensation for precision interferometry," Nature **562**, 391 (2018).

[117] L. De Marco, G. Valtolina, K. Matsuda, W. G. Tobias, J. P. Covey, and J. Ye, "A degenerate Fermi gas of polar molecules," Science **363**, 853 (2019).

[118] P. Asenbaum, C. Overstreet, T. Kovachy, D. D. Brown, J. M. Hogan, and M. A. Kasevich, "Phase shift in an atom interferometer due to spacetime curvature across its wave function," Physical Review Letters **118**, 183602 (2017).

[119] 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" , 493 (2013).

[120] R. Beach, S. R. Hartmann, and R. Friedberg, "Billiard-ball echo model," Physical Review A **25**, 2658 (1982).

[121] P. Storey and C. Cohen-Tannoudji, "The Feynman path integral approach to atomic interferometry. A tutorial," Journal de Physique II **4**, 1999 (1994).

[122] R. Friedberg and S. R. Hartmann, "Billiard balls and matter-wave interferometry," Physical Review A **48**, 1446 (1993).

[123] A. Peters, K. Y. Chung, B. Young, J. Hensley, and S. Chu, "Precision atom interferometry," Philosophical Transactions of the Royal Society A **355**, 2223 (1997).

[124] C. Antoine and C. J. Bordé, "Quantum theory of atomic clocks and gravito-inertial sensors: an update," Journal of Optics B **5**, S199 (2003).

[125] K. Bongs, R. Launay, and M. Kasevich, "High-order inertial phase shifts for time-domain atom interferometers," Applied Physics B **84**, 599 (2006).

[126] B. Barrett, R. Geiger, I. Dutta, M. Meunier, B. Canuel, A. Gauguet, P. Bouyer, and A. Landragin, "The Sagnac effect: 20 years of development in matter-wave interferometry," Comptes Rendus Physique **15**, 875 (2014).

[127] C. J. Bordé, "Atomic interferometry and laser spectroscopy," in *Laser Spectroscopy*, Vol. X, edited by M. Ducloy, E. Giacobino, and G. Camy (World Scientific, Singapore, 1991) pp. 239–245.

[128] A. Landragin, "Conference presentation," International Conference of Laser Spectroscopy, Arcachon (2017).

[129] V. Xu, M. Jaffe, C. D. Panda, S. L. Kristensen, L. W. Clark, and H. Müller, "Probing gravity by holding atoms for 20 seconds," Science **366**, 745 (2019).

[130] D. Sobel, *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time* (Walker Publishing Co., New York, 1995).

[131] J. C. Hafele and R. E. Keating, "Around-the-world atomic clocks: predicted relativistic time gains," Science **177**, 166 (1972).

[132] J. C. Hafele and R. E. Keating, "Around-the-world atomic clocks: observed relativistic time gains," Science **177**, 168 (1972).

[133] R. Stevenson, M. R. Hush, T. Bishop, I. Lesanovsky, and T. Fernholz, "Sagnac interferometry with a single atomic clock," Physical Review Letters **115**, 163001 (2015).

[134] N. Cooper, J. Woods, J. Bateman, A. Dunning, and T. Freegarde, "Stabilized fiber-optic Mach–Zehnder interferometer for carrier-frequency rejection," Applied Optics **52**, 5713 (2013).

[135] F. K. Fatemi, M. L. Terraciano, M. Bashkansky, and Z. Dutton, "Cold atom Raman spectrography using velocity-selective resonances," Optics Express **17**, 12971 (2009).

[136] M. Kasevich, D. S. Weiss, E. Riis, K. Moler, S. Kasapi, and S. Chu, "Atomic velocity selection using stimulated Raman transitions," Physical Review Letters **66**, 2297 (1991).

[137] K. Moler, D. S. Weiss, M. Kasevich, and S. Chu, "Theoretical analysis of velocity-selective Raman transitions," Physical Review A **45**, 342 (1992).

[138] S. Dimopoulos, P. W. Graham, J. M. Hogan, M. A. Kasevich, and S. Rajendran, "Gravitational wave detection with atom interferometry," Physics Letters B **678**, 37 (2009).

[139] D.-F. Gao, J. Wang, and M.-S. Zhan, "Atomic Interferometric Gravitational-Wave Space Observatory (AIGSO)," Communications in Theoretical Physics **69**, 37 (2018).

[140] R. Stoner, D. Butts, J. Kinast, and B. Timmons, "Analytical framework for dynamic light pulse atom interferometry at short interrogation times," Journal of the Optical Society of America B **28**, 2418 (2011).

[141] R. P. Feynman, "Space-time approach to non-relativistic quantum mechanics," Reviews of Modern Physics **20**, 367 (1948).

[142] A. A. Michelson, "On the application of interference-methods to spectroscopic measurements," The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **31**, 338 (1891).

[143] R. J. Bell, *Introductory Fourier Transform Spectroscopy* (Academic Press, 1972).

[144] 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," Physical Review A **52**, 1423 (1995).

[145] I. G. Hughes, "Velocity selection in a Doppler-broadened ensemble of atoms interacting with a monochromatic laser beam," Journal of Modern Optics **65**, 640 (2018).

[146] P. Gillot, B. Cheng, S. Merlet, and F. Pereira Dos Santos, "Limits to the symmetry of a Mach-Zehnder-type atom interferometer," Physical Review A **93**, 013609 (2016).

[147] L. P. Parazzoli, A. M. Hankin, and G. W. Biedermann, "Observation of free-space single-atom matter wave interference," Physical Review Letters **109**, 230401 (2012).

[148] B. Saubaméa, T. W. Hijmans, S. Kulin, E. Rasel, E. Peik, M. Leduc, and C. Cohen-Tannoudji, "Direct measurement of the spatial correlation function of ultracold atoms," Physical Review Letters **79**, 3146 (1997).

[149] P. D. Featonby, G. S. Summy, C. L. Webb, R. M. Godun, M. K. Oberthaler, A. C. Wilson, C. J. Foot, and K. Burnett, "Separated-path Ramsey atom interferometer," Physical Review Letters **81**, 495 (1998).

[150] E. Maréchal, R. Long, T. Miossec, J.-L. Bossennec, R. Barbé, J.-C. Keller, and O. Gorceix, "Atomic spatial coherence monitoring and engineering with magnetic fields," Physical Review A **62**, 053603 (2000).

[151] 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**, 198 (1995).

[152] T. M. Brzozowski, M. Maczynska, M. Zawada, J. Zachorowski, and W. Gawlik, "Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances," Journal of Optics B **4**, 62 (2002).

[153] G. Afek, J. Coslovsky, A. Courvoisier, O. Livneh, and N. Davidson, "Observing power-law dynamics of position-velocity correlation in anomalous diffusion," Physical Review Letters **119**, 060602 (2017).

[154] B. Deh, C. Marzok, S. Slama, C. Zimmermann, and P. W. Courteille, "Bragg spectroscopy and Ramsey interferometry with an ultracold Fermi gas," Applied Physics B **97**, 387 (2009).

[155] E. Zermelo, "Über die Navigation in der Luft als Problem der Variationsrechnung," Jahresbericht der Deutschen Mathematiker-Vereinigung **39**, 44 (1930).

[156] E. Zermelo, "Über das Navigationsproblem bei ruhender oder veränderlicher Windverteilung," Zeitschrift für Angewandte Mathematik und Mechanik **11**, 114 (1931).

[157] C. G. Fraser, "Introductory note to 1930c and 1931a," in *Ernst Zermelo - Collected Works/Gesammelte Werke II* (Springer, Berlin, Heidelberg, 2013) pp. 672–721.

[158] M. Levi, *Classical Mechanics with Calculus of Variations and Optimal Control: An Intuitive Introduction* (American Mathematical Society, Providence, 2014).

[159] R. P. Feynman, *QED: The Strange Theory of Light and Matter* (Penguin, London, 1990).

[160] S. G. Karshenboim, "Precision physics of simple atoms: QED tests, nuclear structure and fundamental constants," Physics Reports **422**, 1 (2005).

[161] T. Aoyama, T. Kinoshita, and M. Nio, "Revised and improved value of the QED tenth-order electron anomalous magnetic moment," Physical Review D **97**, 036001 (2018).

[162] R. von Mises, "Zum Navigationsproblem der Luftfahrt," Zeitschrift für Angewandte Mathematik und Mechanik **11**, 373 (1931).

[163] P. Frank, "Die schnellste flugverbindung zwischen zwei punkten," Zeitschrift für Angewandte Mathematik und Mechanik **13**, 88 (1933).

[164] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, "The quantum technologies roadmap: a European community view," New Journal of Physics **20**, 080201 (2018).

[165] K. Bongs, M. Holynski, J. Vovrosh, P. Bouyer, G. Condon, E. Rasel, C. Schubert, W. P. Schleich, and A. Roura, "Taking atom interferometric quantum sensors from the laboratory to real-world applications," Nature Reviews Physics **1**, 731 (2019).

[166] T. Nöbauer, A. Angerer, B. Bartels, M. Trupke, S. Rotter, J. Schmiedmayer, F. Mintert, and J. Majer, "Smooth optimal quantum control for robust solid-state spin magnetometry," Physical Review Letters **115**, 190801 (2015).

[167] C. P. Koch, J. P. Palao, R. Kosloff, and F. Masnou-Seeuws, "Stabilization of ultracold molecules using optimal control theory," Physical Review A **70**, 013402 (2004).

[168] K. Kotru, D. L. Butts, J. M. Kinast, and R. E. Stoner, "Large-area atom interferometry with frequency-swept Raman adiabatic passage," Physical Review Letters **115**, 103001 (2015).

[169] G. D. McDonald, C. C. N. Kuhn, S. Bennetts, J. E. Debs, K. S. Hardman, M. Johnsson, J. D. Close, and N. P. Robins, "80$\hbar k$ momentum separation with Bloch oscillations in an optically guided atom interferometer," Physical Review A **88**, 053620 (2013).

[170] S. S. Szigeti, J. E. Debs, J. J. Hope, N. P. Robins, and J. D. Close, "Why momentum width matters for atom interferometry with Bragg pulses," New Journal of Physics **14**, 023009 (2012).

[171] J. Bateman and T. Freegarde, "Fractional adiabatic passage in two-level systems: Mirrors and beam splitters for atomic interferometry," Physical Review A **76**, 013416 (2007).

[172] Y. Luo, S. Yan, Q. Hu, A. Jia, C. Wei, and J. Yang, "Contrast enhancement via shaped Raman pulses for thermal cold atom cloud interferometry," The European Physical Journal D **70**, 262 (2016).

[173] B. Luy, K. Kobzar, T. E. Skinner, N. Khaneja, and S. J. Glaser, "Construction of universal rotations from point-to-point transformations," Journal of Magnetic Resonance **176**, 179 (2005).

[174] K. Kobzar, S. Ehni, T. E. Skinner, S. J. Glaser, and B. Luy, "Exploring the limits of broadband 90° and 180° universal rotation pulses," Journal of Magnetic Resonance **225**, 142 (2012).

[175] R. Tycko, A. Pines, and J. Guckenheimer, "Fixed point theory of iterative excitation schemes in NMR," The Journal of Chemical Physics **83**, 2775 (1985).

[176] M. H. Levitt, "Symmetry in the design of NMR multiple-pulse sequences," The Journal of Chemical Physics **128**, 052205 (2008).

[177] M. Braun and S. J. Glaser, "Concurrently optimized cooperative pulses in robust quantum control: application to broadband Ramsey-type pulse sequence elements," New Journal of Physics **16**, 115002 (2014).

[178] S. Bernon, *Trapping and Nondemolition Measurement of Cold Atoms in a High-Finesse Ring Cavity*, Ph.D. thesis, Institut d'Optique Graduate School (2011).

[179] R. Kohlhaas, *Feedback Control of Collective Spin States for Atom Interferometry*, Ph.D. thesis, Institut d'Optique Graduate School (2014).

[180] G. Kuyumjyan, *Multiple Bose-Einstein Condensation in Higher Order Modes of a Dual Frequency Optical Cavity*, Ph.D. thesis, Université de Bordeaux (2017).

[181] T. Vanderbruggen, S. Bernon, A. Bertoldi, A. Landragin, and P. Bouyer, "Spin-squeezing and Dicke-state preparation by heterodyne measurement," Physical Review A **83**, 013821 (2011).

[182] O. Hosten, N. J. Engelsen, R. Krishnakumar, and M. A. Kasevich, "Measurement noise 100 times lower than the quantum-projection limit using entangled atoms," Nature **529**, 505 (2016).

[183] K. C. Cox, G. P. Greve, J. M. Weiner, and J. K. Thompson, "Deterministic squeezed states with collective measurements and feedback," Physical Review Letters **116**, 093602 (2016).

[184] P. Domokos and H. Ritsch, "Collective cooling and self-organization of atoms in a cavity," Physical Review Letters **89**, 253003 (2002).

[185] B. Nagorny, T. Elsässer, and A. Hemmerich, "Collective atomic motion in an optical lattice formed inside a high finesse cavity," Physical Review Letters **91**, 153003 (2003).

[186] S. Gopalakrishnan, B. L. Lev, and P. M. Goldbart, "Atom-light crystallization of Bose-Einstein condensates in multimode cavities: nonequilibrium classical and quantum phase transitions, emergent lattices, supersolidity, and frustration," Physical Review A **82**, 043612 (2010).

[187] A. Xuereb, P. Horak, and T. Freegarde, "Amplified optomechanics in a unidirectional ring cavity," Journal of Modern Optics **58**, 1342 (2011).

[188] S. Bux, C. Gnahm, R. A. W. Maier, C. Zimmermann, and P. W. Courteille, "Cavity-controlled collective scattering at the recoil limit," Physical Review Letters **106**, 203601 (2011).

[189] D. Schmidt, H. Tomczyk, S. Slama, and C. Zimmermann, "Dynamical instability of a Bose-Einstein condensate in an optical ring resonator," Physical Review Letters **112**, 115302 (2014).

[190] R. Landig, L. Hruby, N. Dogra, M. Landini, R. Mottl, T. Donner, and T. Esslinger, "Quantum phases from competing short- and long-range interactions in an optical lattice," Nature **532**, 476 (2016).

[191] J. Léonard, A. Morales, P. Zupancic, T. Esslinger, and T. Donner, "Supersolid formation in a quantum gas breaking a continuous translational symmetry," Nature **543**, 87 (2017).

[192] H.-R. Noh and H. S. Moon, "Calculation of line shapes in double-resonance optical pumping," Physical Review A **80**, 022509 (2009).

[193] A. Bertoldi, S. Bernon, T. Vanderbruggen, A. Landragin, and P. Bouyer, "In situ characterization of an optical cavity using atomic light shift," Optics Letters **35**, 3769 (2010).

[194] S. Palacios, S. Coop, P. Gomez, T. Vanderbruggen, Y. N. Martinez de Escobar, M. Jasperse, and M. W. Mitchell, "Multi-second magnetic coherence in a single domain spinor Bose–Einstein condensate," New Journal of Physics **20**, 053008 (2018).

[195] M. P. Baden, K. J. Arnold, A. L. Grimsmo, S. Parkins, and M. D. Barrett, "Realization of the Dicke model using cavity-assisted Raman transitions," Physical Review Letters **113**, 020408 (2014).

[196] D. N. Aguilera, H. Ahlers, B. Battelier, A. Bawamia, A. Bertoldi, R. Bondarescu, K. Bongs, P. Bouyer, C. Braxmaier, L. Cacciapuoti, C. Chaloner, M. Chwalla, W. Ertmer, M. Franz, N. Gaaloul, M. Gehler, D. Gerardi, L. Gesa, N. Gürlebeck, J. Hartwig, M. Hauth, O. Hellmig, W. Herr, S. Herrmann, A. Heske, A. Hinton, P. Ireland, P. Jetzer, U. Johann, M. Krutzik, A. Kubelka, C. Lämmerzahl, A. Landragin, I. Lloro, D. Massonnet, I. Mateos, A. Milke, M. Nofrarias, M. Oswald, A. Peters, K. Posso-Trujillo, E. Rasel, E. Rocco, A. Roura, J. Rudolph, W. Schleich, C. Schubert, T. Schuldt, S. Seidel, K. Sengstock, C. F. Sopuerta, F. Sorrentino, D. Summers, G. M. Tino, C. Trenkel, N. Uzunoglu, W. von Klitzing, R. Walser, T. Wendrich, A. Wenzlawski, P. Weßels,

A. Wicht, E. Wille, M. Williams, P. Windpassinger, and N. Zahzam, "STE-QUEST—test of the universality of free fall using cold atom interferometry," Classical and Quantum Gravity **31**, 115010 (2014).

[197] H. Y. Ling, Y.-Q. Li, and M. Xiao, "Coherent population trapping and electromagnetically induced transparency in multi-Zeeman-sublevel atoms," Physical Review A **53**, 1014 (1996).

[198] S. Rosi, A. Burchianti, S. Conclave, D. S. Naik, G. Roati, C. Fort, and F. Minardi, "Λ-enhanced grey molasses on the D2 transition of Rubidium-87 atoms," Scientific Reports **8**, 1301 (2018).

[199] D. V. Kosachiov, B. G. Matisov, and Y. V. Rozhdestvensky, "Coherent phenomena in multilevel systems with closed interaction contour," Journal of Physics B: Atomic, Molecular and Optical Physics **25**, 2473 (1992).

[200] D. Kosachiov, B. Matisov, and Y. Rozhdestvensky, "Phase-sensitive laser cooling of double-Λ atoms," Europhysics Letters **22**, 11 (1993).

[201] F. T. Hioe and J. H. Eberly, "N-level coherence vector and higher conservation laws in quantum optics and quantum mechanics," Physical Review Letters **47**, 838 (1981).

[202] J. R. Morris and B. W. Shore, "Reduction of degenerate two-level excitation to independent two-state systems," Physical Review A **27**, 906 (1983).

[203] D. S. Naik, H. Eneriz-Imaz, M. Carey, T. Freegarde, F. Minardi, B. Battelier, P. Bouyer, and A. Bertoldi, "Loading and cooling in an optical trap via hyperfine dark states," Physical Review Research **2**, 013212 (2020).

[204] D. Lucas, P. Horak, and G. Grynberg, "Sisyphus cooling of rubidium atoms on the D_2 line: the role of the neighbouring transitions," The European Physical Journal D **7**, 261 (1999).

[205] M. Xin, W. S. Leong, Z. Chen, and S.-Y. Lan, "An atom interferometer inside a hollow-core photonic crystal fiber," Science Advances **4**, e1701723 (2018).

[206] W. J. Bates, "A wavefront shearing interferometer," Proceedings of the Physical Society **59**, 940 (1947).

[207] B. Barrett, L. Antoni-Micollier, L. Chichet, B. Battelier, T. Lévèque, A. Landragin, and P. Bouyer, "Dual matter-wave inertial sensors in weightlessness," Nature Communications **7**, 13786 (2016).

[208] G. W. Hoth, B. Pelle, S. Riedl, J. Kitching, and E. A. Donley, "Point source atom interferometry with a cloud of finite size," Applied Physics Letters **109**, 071113 (2016).

[209] B. Fang, N. Mielec, D. Savoie, M. Altorio, A. Landragin, and R. Geiger, "Improving the phase response of an atom interferometer by means of temporal pulse shaping," New Journal of Physics **20**, 023020 (2018).

[210] R. Freeman, "Shaped radiofrequency pulses in high resolution NMR," Progress in Nuclear Magnetic Resonance Spectroscopy **32**, 59 (1998).

[211] E. Giese, A. Roura, G. Tackmann, E. M. Rasel, and W. P. Schleich, "Double Bragg diffraction: a tool for atom optics," Physical Review A **88**, 053608 (2013).

[212] G. Condon, M. Rabault, B. Barrett, L. Chichet, R. Arguel, H. Eneriz-Imaz, D. Naik, A. Bertoldi, B. Battelier, P. Bouyer, and A. Landragin, "All-optical Bose-Einstein condensates in microgravity," Physical Review Letters **123**, 240402 (2019).

[213] S. Stellmer, B. Pasquiou, R. Grimm, and F. Schreck, "Laser cooling to quantum degeneracy," Physical Review Letters **110**, 263003 (2013).

[214] A. Urvoy, Z. Vendeiro, J. Ramette, A. Adiyatullin, and V. Vuletić, "Direct laser cooling to Bose-Einstein condensation in a dipole trap," Physical Review Letters **122**, 203202 (2019).

[215] S. Haroche and J.-M. Raimond, *Exploring the Quantum* (Oxford University Press, Oxford, 2006).

[216] I. H. Deutsch, "Physics 566, University of New Mexico," (2017).

[217] K. Mølmer, Y. Castin, and J. Dalibard, "Monte Carlo wave-function method in quantum optics," Journal of the Optical Society of America B **10**, 524 (1993).

[218] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control: Optimization, Estimation, and Control* (Halsted Press, New York, 1975).