

University of Southampton

Faculty of Engineering and Physical Sciences
School of Physics & Astronomy

**Optimal control of cold atoms for
ultra-precise quantum sensors**

by

Jack Cameron Saywell

ORCID: 0000-0002-7867-2128

*A thesis for the degree of
Doctor of Philosophy*

December 2020

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences

School of Physics & Astronomy

Doctor of Philosophy

Optimal control of cold atoms for ultra-precise quantum sensors

by Jack Cameron Saywell

Atom interferometric sensors can enable extremely precise measurements of inertial motion and external fields by manipulating and interfering atomic states using pulses of laser light. However, like many experiments that require the coherent control of a quantum system, the interaction fidelity is limited by inhomogeneities in the control fields. Variations in atomic velocity and laser intensity lead different atoms to experience different interactions under the same pulse, reducing the interference fringe contrast, introducing bias, and limiting the sensitivity.

We present the theoretical design and experimental demonstration of pulses for atom interferometry which compensate inhomogeneities in atomic velocity and laser intensity. By varying the laser phase throughout a pulse and choosing an appropriate fidelity measure to be maximised, pulses are optimised by adapting optimal control techniques originally designed for nuclear magnetic resonance applications. We show using simulations that optimised pulses significantly improve the fidelity of interferometer operations and verify this experimentally using Raman transitions within a cold sample of ^{85}Rb atoms. We demonstrate a robust state-transfer pulse that achieves a fidelity of 99.8(3)% in a $\sim 35\ \mu\text{K}$ sample and obtain a threefold increase in the fringe contrast using a full sequence of optimised pulses. Many of the pulse shapes found by optimal control are simple and symmetrical, and we show that certain symmetries are integral to error compensation. By systematically exploring the dependence of these solutions on the model and optimisation parameters, we demonstrate a stability which underlines the general applicability of optimised pulses to a range of interferometer configurations.

Finally, we introduce and computationally analyse a novel theoretical approach to improve the sensitivity of large-momentum-transfer (LMT) interferometers, whereby “bi-selective” pulses are optimised to track the changing resonance conditions encountered in extended pulse sequences that are designed to increase the measurement sensitivity. When conventional pulses of steady phase are used, the interference contrast decays rapidly as extra pulses are added because of the change in resonance. Using numerical simulations, we show that bi-selective pulses maintain interaction fidelity throughout extended pulse sequences, allowing significant increases in the sensitivity that may be obtained using LMT.

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. 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;
5. I have acknowledged all main sources of help;
6. 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;
7. Parts of this work have been published as: [1], [2], [3], and [4].

Signed:.....

Date:.....

Contents

Declaration of Authorship	v
List of Figures	xi
List of Tables	xvii
Acknowledgements	xix
1 Introduction	1
1.1 Thesis outline	4
I Theoretical background	7
2 Coherent manipulation of atoms with light pulses	9
2.1 Raman transitions	10
2.1.1 Interaction of light with a three-level atom	10
2.1.2 Momentum transfer	12
2.1.3 Evaluating the 3-level Hamiltonian	13
2.1.4 The effective two-level atom	16
2.2 Rubidium-85	18
2.2.1 Beam polarisations and multiple Raman routes	19
2.3 Rabi oscillations, coherent pulses, and the Bloch sphere picture	21
2.3.1 Rotations on the Bloch sphere	22
2.3.2 $\pi/2$ and π pulses	24
2.4 Pulse imperfections: off-resonance and pulse-length errors	25
3 Atom interferometry	29
3.1 The Mach-Zehnder interferometer	31
3.1.1 Phase shift due to acceleration	33
3.2 The effect of pulse imperfections	35
3.3 Large-momentum-transfer (LMT) interferometers	39
3.4 Summary	42
II Composite pulses, adiabatic rapid passage, and optimal control	43
4 Composite pulses and adiabatic rapid passage	45
4.1 Composite Raman pulses	45

4.2	Adiabatic rapid passage	49
4.2.1	Dynamic phase	51
5	Optimal control methods	53
5.1	Approaches to quantum optimal control	55
5.2	Gradient ascent pulse engineering	56
5.2.1	Gradient ascent	56
5.2.2	Setting up the control problem	57
5.2.3	Propagator derivatives and matrix exponentials	60
5.3	Second-order gradient ascent	61
5.3.1	The Limited-memory Broyden-Fletcher-Goldfarb-Shanno method	62
5.3.2	Completing the line-search	63
5.4	Optimised Raman pulse construction	65
5.5	Robustness	66
5.5.1	Robust Raman pulses	66
5.6	Penalties	67
III	Experimental setup and procedures	71
6	Experimental procedures	73
6.1	Magneto-optical trapping of atomic rubidium	73
6.2	Raman laser system	77
6.2.1	Switching the effective wavevector direction	79
6.2.2	Implementing optimal pulses	80
6.3	Experimental procedure	80
6.4	Raman velocimetry	82
IV	Results	85
7	Robust mirror pulses for atom interferometry	87
7.1	Mirror pulse fidelities and design objectives	88
7.2	Mirror pulse performance	91
7.2.1	Robust state-transfer	93
7.2.2	Interferometric phase variation and contrast	96
7.3	Dependence on optimisation parameters	100
7.3.1	Number of time-steps	100
7.3.2	Ensemble size and specification	102
7.3.2.1	Ensemble density	102
7.3.2.2	Ensemble distribution	104
7.4	Amplitude modulation	106
7.5	Experimental results	109
7.5.1	Temporal scans	111
7.5.2	Robustness	112
7.6	Conclusions	114
8	Optimised interferometer sequences	117

8.1	Beamsplitter pulses	117
8.1.1	“Flip-reverse” operation	121
8.2	Antisymmetric mirror pulses	125
8.3	Interferometric phase noise	129
8.3.1	Phase noise and GRAPE sequences	132
8.4	Preliminary experimental results	134
8.5	Conclusions	137
9	Bi-selective pulses for large-area atom interferometry	141
9.1	LMT using robust pulses	142
9.2	Bi-selective pulses	145
9.3	Characterisation of bi-selective pulses	148
9.3.1	Parabolic approximation of $\mathcal{F}_{\text{square}}$ bi-selective pulses	152
9.4	LMT contrast	154
9.5	Suppression of off-resonant excitation	157
9.6	Conclusions	160
10	Conclusions and outlook	161
10.1	Physical interpretation of optimised pulses	163
V	Appendices	165
A	Mach-Zehnder interferometer output	167
B	Numerical Simulations	171
B.1	Modelling individual Raman pulses	171
B.2	Modelling Mach-Zehnder sequences	172
B.3	Modelling LMT sequences	172
C	Conference posters	175
	Bibliography	179

List of publications

- **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**, M. Belal, and T. Freegarde, “Velocimetry of cold atoms by matter-wave interferometry,” *Physical Review A* **99**, 023631 (2019).
- **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). **equal contribution*
- **J. Saywell**, M. Carey, I. Kuprov, and T. Freegarde, “Biselective pulses for large-area atom interferometry,” *Physical Review A* **101**, 063625 (2020).

List of Figures

2.1	Energy levels in a 3-level Raman system.	10
2.2	Energies of atomic states as a function of atomic momentum in different Raman arrangements.	13
2.3	Hyperfine energy levels in ^{85}Rb	18
2.4	The Bloch sphere representation of a two-level system.	23
2.5	Representation of a Raman pulse on the Bloch sphere.	24
2.6	Representation of $\pi/2$ and π pulses on the Bloch sphere.	25
2.7	Off-resonance and pulse-length errors on the Bloch sphere.	26
2.8	Contour plot showing the effect of control errors on the π pulse efficiency.	27
2.9	Rabi flopping in the presence of control errors.	28
3.1	Atomic state trajectories in the Mach-Zehnder interferometer sequence.	31
3.2	Representation of the Mach-Zehnder interferometer on the Bloch sphere.	32
3.3	Illustration of different interferometer phase errors arising from finite-duration beamsplitter pulses.	37
3.4	The effect of control errors on the interferometer fringe contrast.	38
3.5	Energies of atomic states as a function of atomic momentum and the effect of switching the effective wavevector in large-momentum-transfer.	39
3.6	Atomic state trajectories in a LMT interferometer.	41
4.1	Representation of the Levitt composite pulse on the Bloch sphere.	47
4.2	Representation of the WALTZ composite pulse on the Bloch sphere.	48
4.3	Representation of ARP on the Bloch sphere.	50
6.1	Cooling and repump transitions used in a ^{85}Rb MOT.	75
6.2	Arrangement of the MOT lasers and magnetic field.	77
6.3	Illustration of the apparatus used to generate Raman beams.	79
6.4	Schematic of a typical experimental sequence.	81
6.5	Illustration of Raman velocimetry.	82
7.1	Histograms showing the distribution of terminal fidelities for mirror pulse optimisations.	93
7.2	Waveforms of the optimised mirror pulses with highest fidelity.	94
7.3	Illustration of an optimised mirror pulse on the Bloch sphere.	95
7.4	Fraction of atoms excited by mirror pulses of different durations within a $40\text{ }\mu\text{K}$ atom cloud.	95
7.5	Contour plots showing the fidelity of different mirror pulses as a function of off-resonance and pulse-length errors.	97

7.6	Interferometric phase variation following sequences with different mirror pulses.	98
7.7	Simulated interferometer contrast for different pulse sequences as a function of the atomic cloud temperature.	99
7.8	Dependence of the terminal fidelity on the number of time-steps in a pulse.	101
7.9	Colour-map showing how optimised pulse shapes depend on the number of time-steps in a pulse when using the $\mathcal{F}_{\text{square}}$ fidelity measure with a parabolic initial guess.	102
7.10	Colour-map showing how optimised pulse shapes depend on the number of time-steps in a pulse when using the $\mathcal{F}_{\text{real}}$ fidelity measure with a random initial guess.	103
7.11	Colour-map showing how optimised pulse shapes depend on the number of time-steps in a pulse when using the $\mathcal{F}_{\text{square}}$ fidelity measure with a random initial guess.	104
7.12	Illustration showing how the number of detunings used in the optimisation ensemble affects the pulse performance.	105
7.13	Illustration showing how the shape of detuning distribution used in the optimisation ensemble affects the pulse performance.	106
7.14	Depiction of a phase and amplitude-modulated pulse optimised with no constraint on the maximum Rabi frequency.	107
7.15	Depiction of a phase and amplitude-modulated pulse optimised with a constraint on the maximum Rabi frequency.	108
7.16	Waveform of a robust phase-modulated mirror pulse optimised for a temperature of $120 \mu\text{K}$	110
7.17	Illustration of timing parameters used to realise optimised Raman pulses.	111
7.18	Experimental data showing evolution of the excited state population throughout an optimised mirror pulse.	112
7.19	Experimental data showing how the fraction of atoms transferred to the excited state following different mirror pulses depends on the laser detuning.	113
8.1	A comparison of two beamsplitter pulses found using optimal control.	120
8.2	Contour plots showing the beamsplitter pulse fidelity as a function of off-resonance and pulse-length errors.	122
8.3	The effect of a variation in Rabi frequency between beamsplitter pulses on the interferometer phase.	126
8.4	Contour plots depicting the Mach-Zehnder interferometer contrast following different beamsplitter pulses as a function of the variation in Rabi frequency between pulses.	127
8.5	Illustration of the differences in performance between mirror pulses with no symmetry constraint and antisymmetric pulses.	129
8.6	Interferometer sensitivity to phase noise following a Mach-Zehnder sequence using rectangular and Gaussian pulses.	131
8.7	Interferometer sensitivity to phase noise following a Mach-Zehnder sequence using optimised pulses.	133
8.8	Interferometer sensitivity to phase noise following a Mach-Zehnder sequence using optimised pulses with a smooth amplitude profile.	135
8.9	Waveforms of phase-modulated beamsplitter and mirror pulses optimised for a temperature of $120 \mu\text{K}$	136

8.10	Illustration of the experimental timings used to realise an optimised pulse sequence.	137
8.11	Ratio of the experimental interference fringe contrasts obtained using optimised pulse sequences applied at different temperatures.	138
8.12	Data showing interferometer fringes following rectangular pulses and optimised pulses.	138
8.13	Variation in the fitting parameters for fringes obtained using rectangular pulses and optimised pulses applied at different temperatures.	139
9.1	Atomic state trajectories and changing resonance conditions within a LMT interferometer sequence.	143
9.2	Bi-selective pulse phase profiles found optimising the $\mathcal{F}_{\text{real}}$ fidelity measure.	148
9.3	Bi-selective pulse phase profiles found optimising the $\mathcal{F}_{\text{square}}$ fidelity measure	149
9.4	A comparison of the state-transfer efficiency between a bi-selective pulse, the composite WALTZ pulse, and the tanh/tan ARP pulse.	151
9.5	The phase imparted to atoms by different augmentation pulses as a function of off-resonance and pulse-length errors.	152
9.6	Comparison between an optimised $\mathcal{F}_{\text{square}}$ pulse and a parabolic approximation to the pulse.	153
9.7	Illustration of a bi-selective augmentation pulse on the Bloch sphere. . . .	154
9.8	Interferometer contrast following LMT sequences composed of different augmentation pulses.	155
9.9	Illustration of the suppression of off-resonant excitation using a modified bi-selective fidelity measure.	158
9.10	Illustration of the suppression of off-resonant excitation using bi-selective pulses with different constraints.	159
B.1	Diagram showing how LMT interferometer sequences are simulated. . . .	173
C.1	Poster presented at the 2018 Frontiers of Matterwave Optics (FOMO) conference 17-21 September, Orthodox Academy of Crete, Greece.	176
C.2	Poster presented at the 2019 International Conference on Quantum Error Correction (QEC) 29 July - 2 August, Senate House, London, UK.	177

List of Tables

7.1	Simulated Mach-Zehnder contrast values for different mirror pulses. . . .	100
9.1	Simulated interferometer contrast values following LMT sequences composed of different augmentation pulses.	157

Acknowledgements

I would like to express my immense gratitude to my supervisors, Tim Freegarde and Ilya Kuprov, for their support and guidance throughout my time in Southampton.

It has been a pleasure to work alongside Max Carey, Mohammad Belal, David Goodwin, Andrei Dragomir, Matt Aldous, Chester Camm, Leila Mashhadi, and Nikolaos Dedes. In particular, I must thank Max Carey: our joint exploration of the atomic world has been greatly rewarding and I have benefited from your advice throughout. I am also grateful to Matt Himsworth, Jonathan Woods, and David Goodwin for providing many invaluable insights on how to survive the trials of Ph.D. research.

I am also thankful for the funding provided by the Centre for Doctoral Training in Next Generation Computational Modelling.

Finally, I must thank my parents, Brenda and Mark. I could not have completed my Ph.D. without their support.

To my grandfather, Ernest Saywell.

Chapter 1

Introduction

The ability to prepare and manipulate quantum states with high fidelity in the presence of systematic inhomogeneities and noise is essential for many experiments investigating coherent quantum phenomena. As a result, a wealth of techniques has been developed to mitigate these errors and increase operation fidelity. Nowhere have these techniques reached a similar degree of sophistication as in the field of nuclear magnetic resonance (NMR) spectroscopy. In NMR, single pulsed interactions with radio-frequency control fields are replaced by those in which the phase, frequency, and amplitude of the control fields are varied as a function of time to replicate the action of the original pulse but compensate for inhomogeneities such as variations in field strength and offsets from resonance. Examples of these techniques include composite pulses [5], where typically a small number of pulses are concatenated with variable phases and durations; and adiabatic rapid passage [6], which relies on slow continuous variation of the field frequency, phase, and amplitude.

NMR spectroscopy has also been a significant catalyst for the development and demonstration of quantum optimal control techniques [7–11]. Optimal control theory tries to find the best way to steer a system from a specific initial state to a desired target state subject to experimental constraints on available energy and time. In quantum mechanics, this typically means finding the optimal sequence of the available control fields that can realise some desired transformation of the system. Upon carefully choosing a measure of fidelity that defines how well a given pulse works, a pulse may be numerically optimised using an optimal control algorithm.

Well-established applications of quantum optimal control include the control of chemical reactions using shaped laser pulses [12–15], and the control of nuclear spins using radio-frequency fields [7, 16–18]. Since these early demonstrations, a range of sophisticated quantum control algorithms have been developed [9, 14, 19–26] and the techniques have found other applications including: nitrogen-vacancy (NV) centre magnetometry [27, 28], gate-design in quantum information processors [29–32], and the production

and control of Bose-Einstein condensates (BECs) [33–35]. The advantage of optimal control is that it can be used to optimise pulses that are tailored to specific problems and systems when provided with the correct fidelity measures and constraints.

Atom interferometric quantum sensors [36–39] use sequences of optical pulses to manipulate and interfere atomic states and obtain measurements of forces and fields with remarkable precision and stability. Atom interferometers work by exploiting the wavelike behaviour of matter, a manifestation of wave-particle duality, just as conventional optical interferometers exploit the wavelike behaviour of light. The light pulses can therefore be considered as the optical equivalents of the physical mirrors and beamsplitters used in an optical interferometer. However, one of the central problems in atom interferometry is that the efficiency of the conventional $\pi/2$ (beamsplitter) and π (mirror) pulses of constant laser intensity and phase is limited by inhomogeneities in the atom-light interaction. Atomic samples contain a distribution of atomic velocities, states, and locations; and such effects lead different atoms to undergo different evolution under the same pulse, thus reducing the visibility, or contrast, of the interference patterns and causing errors in the measurement.

It is sometimes possible to solve these problems by selecting atoms with a narrow range of velocities, or by discarding atoms in unwanted locations or unwanted states. However, by reducing the number of atoms that can participate in a measurement, all of these processes reduce the signal-to-noise ratio (SNR). Furthermore, there has been a recent effort to miniaturise atom interferometric sensors in an attempt to push the technology outside laboratory environments and make them commercially viable [39]. In these cases, the importance of designing interferometers that operate in noisy and inhomogeneous environments is paramount. Imperfect pulses are also a serious limitation when large-momentum-transfer (LMT) techniques are employed to increase an interferometer’s intrinsic sensitivity [40]. LMT interferometers rely on the momentum transferred by extended sequences of light-pulses - known as augmentation pulses - to increase the space-time area enclosed by the interfering paths, which is proportional to the measurement scale-factor. However, any improvement in sensitivity can be lost because the visibility of the interference fringes decays following the repeated application of pulses with imperfect fidelity.

Despite the similarities between the description of nuclear spin-1/2 systems in NMR and the two-level atomic systems used in atom interferometry, techniques such as ARP and composite pulses have found limited application in atom interferometers. In 2002, McGuirk *et al.* [41] used the very first composite pulse developed in NMR - the Levitt pulse [5] - to increase the efficiency of state preparation in a cold atom interferometer. Later, in 2013, Butts *et al.* [42] investigated the application of two composite mirror pulses - the Levitt [5] and WALTZ [43] pulses - as augmentation pulses within a LMT interferometer using a laser-cooled and velocity-selected sample of ^{133}Cs at 0.5 μK . Of the pulses tested, only the WALTZ pulse led to an improvement in the contrast,

increasing the intrinsic sensitivity by a factor of 9 before visibility was lost. However, when WALTZ was again used in a similar experiment with a hotter atomic sample, there was no improvement in contrast when compared with conventional π pulses [44]. More recently, Dunning *et al.* [45] compared the state-transfer efficiency of a selection of NMR composite mirror pulses within an inhomogeneous cloud of ^{85}Rb at a temperature of $80\text{ }\mu\text{K}$. They found that all the composite pulses tested improved the fidelity and robustness compared with the standard π pulse. However, there was no consideration of the effect of each pulse on the interferometer phase, and the applicability of composite beamsplitter pulses was not explored.

Frequency-swept adiabatic pulses, while typically longer and more technically difficult to implement than composite pulses [41], have also been applied in atom interferometry [46–48]. Bateman and Freearge [49] designed an adiabatic mirror and beamsplitter pulse using analytical methods, and a similar design was implemented by Jaffe *et al.* [48, 50] in a cavity-based interferometer. Kotru *et al.* [46] employed a pulse with a tangential frequency sweep, known as the tanh/tan pulse [51], to improve the contrast within a LMT interferometer with ^{133}Cs atoms at a temperature of $9\text{ }\mu\text{K}$, leading to an increase of the intrinsic sensitivity by a factor of 15. However, using adiabatic pulses in LMT relies on a precise cancellation of velocity and intensity-dependent phases between the pulses [50], and rapid dephasing can occur if there is variation in these quantities, making the interferometer susceptible to variations in field strength and detuning throughout the sequence.

The use of NMR pulse techniques in atom interferometry has therefore relied on using existing pulse shapes and schemes originally developed to tackle the problems encountered in the control of nuclear spins, and there is a clear need for pulses that are tailored to tackle the distinct problems within atom interferometry. In this thesis, we present the first systematic application of optimal control to the design of pulses for atom interferometers where the atom-light interactions are formed using stimulated Raman transitions [52]. Unlike conventional $\pi/2$ and π pulses where the laser phase and intensity are constant throughout the interaction, we treat the optical phase and intensity as controllable parameters and optimise pulses where these parameters are varied in time by adapting NMR optimal control methods. We demonstrate that these optimised pulses yield significant increases in the fidelity of individual beamsplitter and mirror pulses using a combination of simulations and experiments.

We first explore optimisation of the interferometric mirror operation and compare the resulting optimised mirror pulses with the best composite and adiabatic alternatives that have been employed in interferometry to-date. We find significant improvements in error compensation and fidelity with our optimised pulses, and experimentally verify our designs using a cold sample of ^{85}Rb . We also investigate the effect that different composite, adiabatic, and optimised pulses have on the interferometer phase within a Mach-Zehnder interferometer, and demonstrate that pulses designed using a carefully

chosen phase-sensitive fidelity measure minimise unwanted variation in the measurement scale-factor across an inhomogeneous atom cloud, thus leading to improvements in fringe contrast.

We then extend our optimisation to the beamsplitter operation, and propose a symmetry-based construction procedure for these pulses that enables the cancellation of imperfections between the first and final pulse in the sequence. We use simulations to demonstrate an improvement in error compensation over composite and adiabatic approaches, and deploy symmetry arguments to show that mirror pulses designed using antisymmetry are immune to velocity-dependent phase shifts that limit the application of many composite and adiabatic pulses. Furthermore, we experimentally demonstrate a threefold improvement in interferometer contrast with our optimised sequence over that obtained using conventional π and $\pi/2$ pulses in an inhomogeneous cloud of ^{85}Rb at a temperature of $\sim 100\ \mu\text{K}$.

LMT interferometry is a promising approach to increase the intrinsic sensitivity of an atom interferometer, and Raman LMT interferometers suffer from a decay of contrast as extra pulses are added due to the change in atomic velocity and resonance conditions throughout the interferometer. Even when composite pulses are used, a significant decay of contrast is still observed [42, 44], and when adiabatic pulses are employed, the results rely on a delicate cancellation of phase errors between pulses [46, 50]. We therefore introduce a “bi-selective” LMT pulse design to increase the sensitivity of atom interferometers by maintaining the interaction fidelity as the relative momentum difference between atomic wave-packets is increased. Each bi-selective pulse is individually optimised using optimal control methods to track the changing bi-modal momentum distribution throughout the interferometer, while providing robustness to variations in laser intensity and detuning. This technique allows for a greater measurement sensitivity without requiring higher laser powers. Using an experimentally validated model, we find our designs yield considerable improvements in interferometer contrast and the suppression of parasitic phase shifts compared with the best-established composite pulses and adiabatic rapid passage techniques.

1.1 Thesis outline

Section 1 presents a theoretical treatment of atom interferometry and the problem of fidelity loss caused by systematic inhomogeneities in the atom-light interaction. In Chapter 2 we introduce stimulated Raman transitions - used to form the “atom-optics” in many atom interferometric sensors - and the $\pi/2$ (beamsplitter) and π (mirror) pulses formed from fractional Rabi oscillations. We explain how the fidelity of these ‘rectangular’ pulses is decreased by variations in the atom-light interaction. In Chapter 3, we present an introduction to atom interferometry using Raman pulses, including the

technique of LMT interferometry, and explain how pulse errors limit interferometer performance and sensitivity.

Section 2 reviews prior and current approaches to improving the fidelity and robustness of atomic state manipulation using NMR pulse shaping techniques. We also introduce the state-of-the-art optimal control methods adapted in this thesis to design tailored pulses for interferometric applications. Chapter 4 reviews the techniques of composite pulses and adiabatic rapid passage (ARP). In Chapter 5, we review optimal control methods as an alternative approach to pulse design. In particular, we describe the gradient ascent pulse engineering (GRAPE) algorithm, which we use to optimise tailored pulses for atom interferometry.

Section 3 details the experimental apparatus used to test the pulse schemes in this thesis. Chapter 6 explains how we produce a cold sample of ^{85}Rb atoms using laser cooling and trapping in a three-dimensional magneto-optical trap (MOT). We also explain how we can control and measure the temperature of our atomic sample, and how we can induce Raman transitions between the hyperfine ground states in ^{85}Rb . In this chapter, we also explain how our optimised pulses are implemented in our experiment.

Section 4 presents the main results of this thesis. In Chapter 7, we address the problem of designing robust individual mirror pulses which mitigate experimental imperfections in detuning and coupling strength. We discuss the topic of mirror pulse fidelity and present and characterise pulses optimised using different fidelity measures. We show using simulations that mirror pulse designs found using GRAPE significantly improve upon the pulse fidelity and interferometer contrast compared with composite pulses and adiabatic rapid passage. These fidelity improvements are also verified experimentally in a cold sample of ^{85}Rb . In this chapter we also explore how the shape of the optimised pulse waveforms depends on the parameters of the optimisation.

In Chapter 8, we extend our optimisation to each component operation within a three-pulse “Mach-Zehnder” interferometer sequence. We optimise pairs of beamsplitter pulses using a “flip-reverse” symmetry construction procedure that are robust to variations in detuning and coupling strength. We also explore the role of pulse symmetry in the design of individual pulses, and the effect of phase noise on interferometric sensitivity when employing optimised pulses. Experimental results are presented which show our computational designs improve the interferometer contrast compared with conventional pulses.

In Chapter 9, we turn our attention to LMT interferometry, and introduce a scheme whereby individual pulses are tailored to track the evolving momentum distribution throughout an LMT interferometer. We demonstrate using simulations that these “bi-selective” pulses lead to considerable increases in interferometer area and sensitivity compared with conventional composite and ARP pulses. We also explain how, by using

a careful choice of fidelity measure, we can engineer robustness to variations in Rabi frequency between individual pulses. Finally, we show how these pulses may be adapted to suppress unwanted off-resonant excitation present in some interferometer configurations.

Finally, in Chapter 10 we draw conclusions and discuss questions raised throughout this thesis, thus providing a framework for future investigations.

Part I

Theoretical background

Chapter 2

Coherent manipulation of atoms with light pulses

In many atom interferometers, stimulated Raman transitions between ground hyperfine states in alkali atoms form the beamsplitter and mirror operations [36, 37, 39] used to diffract, redirect, and interfere the atomic matter-waves. Raman transitions enable an atom to transition between two long-lived ground states from which spontaneous emission is negligible via a two-photon process, namely stimulated absorption to and stimulated emission from a short-lived upper intermediate state. This requires two lasers that are both far detuned from the upper state with a frequency difference that closely matches the energy splitting of the ground states. Under these conditions the upper intermediate state can be eliminated from the problem, and the atom is reduced to a simple and stable two-state system with a negligible decay of coherence.

In this chapter we present a theoretical treatment of Raman transitions and introduce the building blocks of conventional interferometer sequences: the $\pi/2$ (beamsplitter) and π (mirror) pulses formed from fractional Rabi oscillations between the states of a two-level atom. This leads us to introduce the Bloch sphere picture of atomic state evolution, which provides an intuitive geometric representation of the action of interferometer pulses on atomic states. We conclude with a brief summary of the atomic structure of the alkali metal atom ^{85}Rb , the atomic species used in the simulations and experiments throughout this thesis.




FIGURE 2.1: Energy levels in a 3-level Raman system (not to scale). Diagonal arrows indicate the interaction of each level with lasers 1 and 2; laser 1 couples $|g\rangle$ and $|i\rangle$ only, and laser 2 couples $|e\rangle$ and $|i\rangle$ only. Δ is the single-photon detuning of both lasers from the intermediate level $|i\rangle$ (Equation 2.24), and δ is the two-photon Raman detuning defined by Equation 2.26. If the laser beams are co-propagating, $\mathbf{k}_1 \approx \mathbf{k}_2$ and the momentum imparted to an atom initially in the state $|g\rangle$ is $\hbar|\mathbf{k}_1 - \mathbf{k}_2| \approx 0$. If the beams are counter-propagating, $\mathbf{k}_1 \approx -\mathbf{k}_2$ and therefore the momentum imparted is $\hbar|\mathbf{k}_1 - \mathbf{k}_2| \approx 2|\mathbf{k}_1|$.

2.1 Raman transitions

2.1.1 Interaction of light with a three-level atom

We begin by considering a three-level atom interacting with two counter-propagating laser beams with frequencies ω_1 and ω_2 , as depicted in Figure 2.1. The two lower levels, $|g\rangle$ and $|e\rangle$ represent ground hyperfine states, and $|i\rangle$ represents an upper intermediate level. Spontaneous decay is prohibited from $|g\rangle$ and $|e\rangle$ but possible from the intermediate level $|i\rangle$. Although there is no permitted single-photon electric dipole transition or coupling between $|g\rangle$ and $|e\rangle$, it is possible by a *two-photon* Raman transition via the intermediate level. Furthermore, we will show that if both lasers are far detuned from single-photon resonance with $|i\rangle$, and if their frequency difference matches the ground hyperfine splitting ω_{eg} between $|g\rangle$ and $|e\rangle$, then $|i\rangle$ is never occupied and transitions may occur in the effective two-level system formed by $|g\rangle$ and $|e\rangle$.

In our experiments with ^{85}Rb , the states $|g\rangle$ and $|e\rangle$ are separated by a microwave frequency difference of approximately 3 GHz. In contrast, the intermediate state $|i\rangle$ is separated from both $|g\rangle$ and $|e\rangle$ by a much larger optical frequency splitting (equivalent to a wavelength of approximately 780 nm). In practice both laser beams are generated

from a source at 780 nm. The light from this source is split spatially into two components, and frequency shifts are applied such that the frequency difference between the beams $\omega_1 - \omega_2$ matches the GHz splitting between $|g\rangle$ and $|e\rangle$. We discuss the atomic structure of ^{85}Rb and the specific energy levels used to form the Raman transition in Section 2.2 of this chapter; a description of how this is realised experimentally is given in Chapter 6, Section 6.2.

In general, we may describe the quantum state of our three-level atom as the following superposition of basis states:

$$|\Psi(t)\rangle = a_g(t) |g, \mathbf{p}_g\rangle + a_e(t) |e, \mathbf{p}_e\rangle + a_i(t) |i, \mathbf{p}_i\rangle. \quad (2.1)$$

$a_g(t)$, $a_e(t)$, and $a_i(t)$ represent the (time-dependent) quantum amplitudes associated with each basis state of our atom, and $\mathbf{p}_{g,e,i}$ are the atomic momenta associated with each state. Under the Born rule, the probability of finding the atom in a given state, say $|g, \mathbf{p}_g\rangle$, at time t is given by the squared modulus of the corresponding complex amplitude $a_g(t)$. Our aim in this section is to calculate how the quantum state evolves in time due to the interaction with our two laser beams.

The time-evolution of our three-level atom is governed by the time-dependent Schrödinger equation (TDSE)

$$i\hbar \frac{\partial}{\partial t} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle, \quad (2.2)$$

where $\hat{H}$ is the Hamiltonian, given by the sum of the atomic Hamiltonian $\hat{H}_A$ and the interaction Hamiltonian $\hat{H}_I$.

Our choice to include the atomic momentum in our basis is motivated by the fact that atoms possess both an internal electronic structure and an external degree of freedom due to their centre-of-mass momentum, meaning the atomic Hamiltonian may be expressed as the sum of these two contributions $\hat{H}_A = \hat{H}_{\text{int}} + \hat{H}_{\text{ext}}$, where $\hat{H}_{\text{int}} = \sum_n \hbar\omega_n |n\rangle \langle n|$ and $\hat{H}_{\text{ext}} = \hat{\mathbf{p}}^2/2m$. Here, ω_n represents the frequency of the electronic level $|n\rangle$, $\hat{\mathbf{p}}$ is the momentum operator, and m is the atomic mass. The basis states of our atom may therefore be expressed as tensor products $|n, \mathbf{p}_n\rangle = |n\rangle \otimes |\mathbf{p}_n\rangle$ of momentum plane-waves and the internal energy states [37].

The interaction Hamiltonian $\hat{H}_I = -\hat{\mathbf{d}} \cdot \mathbf{E} = e\hat{\mathbf{r}} \cdot \mathbf{E}$ is given by the energy of an electric dipole $\hat{\mathbf{d}} = -e\hat{\mathbf{r}}$ in the presence of a classical electric field, where $\hat{\mathbf{r}}$ is the operator for the position of a single electron relative to the atomic centre of mass and $\mathbf{E}$ is the electric field due to the laser beams.¹

¹This dipole approximation is valid if the wavelength of the electromagnetic radiation is much larger than the size of the atom [53].

The electric field due to both lasers with frequencies ω_1 and ω_2 , which are counter-propagating along the z -axis, is given by

$$\mathbf{E} = \frac{1}{2}\mathbf{E}_1 e^{i(k_1 z - \omega_1 t + \phi_1)} + \frac{1}{2}\mathbf{E}_2 e^{i(-k_2 z - \omega_2 t + \phi_2)} + c.c., \quad (2.3)$$

where $\mathbf{k}_{1,2}$ are the wave vectors of each beam, $\mathbf{E}_{1,2}$ are the (vector) amplitudes of each field, $\phi_{1,2}$ are phases of each light field, and $c.c.$ represents the complex conjugate of the first terms. Since the lasers are counter-propagating, $\mathbf{k}_1 \approx -\mathbf{k}_2$.

2.1.2 Momentum transfer

Terms such as $e^{\pm i\mathbf{k}_{1,2}z}$ appear in the interaction Hamiltonian matrix elements. The effect of these parts of the interaction is to act on the momentum part of the atomic states and change the momentum of the atom by $\pm\hbar\mathbf{k}_{1,2}$ in the z direction - a physical consequence of the absorption and/or emission of a photon and the conservation of momentum. This can be seen explicitly by rewriting $e^{i\mathbf{k}_1 z}$ using the identity operator $\mathbf{I} = \int dp_z |p_z\rangle \langle p_z|$, where $|p_z\rangle = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} dz e^{ip_z z/\hbar} |z\rangle$ represents a momentum plane-wave eigenstate:

$$\mathbf{I} \times e^{i\mathbf{k}_1 z} = \int_{-\infty}^{\infty} dp_z e^{i\mathbf{k}_1 z} |p_z\rangle \langle p_z| \quad (2.4)$$

$$= \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} dp_z \int_{-\infty}^{\infty} dz e^{i\mathbf{k}_1 z} e^{ip_z z/\hbar} |z\rangle \langle p_z| \quad (2.5)$$

$$= \int_{-\infty}^{\infty} dp_z |p_z + \hbar\mathbf{k}_1\rangle \langle p_z|. \quad (2.6)$$

We can therefore rewrite our momentum-inclusive basis states as $|g, \mathbf{p}\rangle$, $|i, \mathbf{p} + \hbar\mathbf{k}_1\rangle$, and $|e, \mathbf{p} + \hbar\mathbf{k}_1 - \hbar\mathbf{k}_2\rangle$ such that the momentum part forms a “closed momentum family” [52]. Physically, this means that an atom in the state $|g, \mathbf{p}\rangle$ can absorb a photon with momentum $\hbar\mathbf{k}_1$ from laser 1 and end up in state $|i, \mathbf{p} + \hbar\mathbf{k}_1\rangle$, and undergo stimulated emission of a photon with momentum $\hbar\mathbf{k}_2$ to end up in the state $|e, \mathbf{p} + \hbar\mathbf{k}_1 - \hbar\mathbf{k}_2\rangle$.

We will show in Section 2.1.4 that if both lasers are far detuned from single-photon resonance with the intermediate state $|i, \mathbf{p} + \hbar\mathbf{k}_1\rangle$ then Raman transitions occur between $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_1 - \hbar\mathbf{k}_2\rangle$, forming an effective two-level atom. We define the *effective wavevector* such that $\hbar\mathbf{k}_{\text{eff}} \equiv \hbar\mathbf{k}_1 - \hbar\mathbf{k}_2$ is the momentum transferred by the two-photon transition from $|g, \mathbf{p}\rangle$ to $|e, \mathbf{p} + \hbar\mathbf{k}_1 - \hbar\mathbf{k}_2\rangle$. If the Raman beams are *co-propagating*, $\mathbf{k}_1 \approx \mathbf{k}_2$ and therefore $\hbar|\mathbf{k}_{\text{eff}}| \approx 0$. However, in the case of *counter-propagating* beams $\mathbf{k}_1 \approx -\mathbf{k}_2$ and $\hbar|\mathbf{k}_{\text{eff}}| \approx 2|\mathbf{k}_1|$ and the momentum transferred is approximately twice the single-photon impulse. The difference between these two arrangements is shown in Figure 2.2, and for additional detail we refer the reader to Section 9.4.3 of Cohen-Tannoudji and Guéry-Odelin [54]. The transfer of momentum is essential to atom interferometers that are sensitive to external forces; this is a point we will return to in Chapter 3.




FIGURE 2.2: Energies (not to scale) of atomic states $|g\rangle$, $|e\rangle$, and $|i\rangle$ are depicted as a function of atomic momentum along the z -axis. The solid (blue) and dashed (red) diagonal arrows correspond to photons from each laser beam with frequencies ω_1 and ω_2 , respectively. On the left we show the co-propagating Raman arrangement, where $\mathbf{k}_1 \approx \mathbf{k}_2$, and on the right we show the counter-propagating arrangement, where $\mathbf{k}_1 \approx -\mathbf{k}_2$. We can see that photons in the co-propagating Raman arrangement couple the states $|g\rangle$ and $|e\rangle$ with a negligible momentum difference. However, in the counter-propagating arrangement, the momentum difference is much larger - approximately twice the momentum carried by photons in each beam.

Although it is possible to directly induce a microwave transition between the states $|g\rangle$ and $|e\rangle$, and thereby take advantage of the long lifetime of these states, the momentum transfer would be small. Using a Raman transition with counter-propagating beams allows us to retain the advantage of coupling two states with a long lifetime and achieve a considerably larger momentum transfer due to the optical frequency splitting between the two lower states and the upper intermediate level $|i\rangle$.

2.1.3 Evaluating the 3-level Hamiltonian

We now proceed to evaluate the Hamiltonian and solve the TDSE for the time-evolution of the internal state amplitudes. We begin by defining the following single-photon Rabi

frequencies:

$$\Omega_{1g} \equiv -\frac{1}{\hbar} \langle g | \hat{\mathbf{d}} \cdot \mathbf{E}_1 | i \rangle \quad (2.7)$$

$$\Omega_{2e} \equiv -\frac{1}{\hbar} \langle e | \hat{\mathbf{d}} \cdot \mathbf{E}_2 | i \rangle \quad (2.8)$$

$$\Omega_{1e} \equiv -\frac{1}{\hbar} \langle e | \hat{\mathbf{d}} \cdot \mathbf{E}_1 | i \rangle \quad (2.9)$$

$$\Omega_{2g} \equiv -\frac{1}{\hbar} \langle g | \hat{\mathbf{d}} \cdot \mathbf{E}_2 | i \rangle. \quad (2.10)$$

Next, we write the total Hamiltonian for the three-level atom as

$$\hat{H} = \begin{pmatrix} E_g & 0 & \hat{H}_{gi} \\ 0 & E_e & \hat{H}_{ei} \\ \hat{H}_{gi}^* & \hat{H}_{ei}^* & E_i \end{pmatrix}, \quad (2.11)$$

where $E_n \equiv \hbar\omega_n + \mathbf{p}_n^2/2m$ are the energy eigenvalues of the atomic Hamiltonian $\hat{H}_A$ ($n = g, e, i$). We assume that laser 1 couples only states $|g\rangle$ and $|i\rangle$, and that laser 2 couples only $|e\rangle$ and $|i\rangle$. This means we can ignore Ω_{1e} and Ω_{2g} , such that elements $\hat{H}_{gi}$ and $\hat{H}_{ei}$ are given by:

$$\hat{H}_{gi} = \hbar \frac{\Omega_{1g}}{2} (e^{-i\omega_1 t - i\phi_1} + e^{i\omega_1 t + i\phi_1}) \quad (2.12)$$

$$\hat{H}_{ei} = \hbar \frac{\Omega_{2e}}{2} (e^{-i\omega_2 t - i\phi_2} + e^{i\omega_2 t + i\phi_2}). \quad (2.13)$$

Now, using the TDSE we may write the following coupled differential equations for the internal state amplitudes:

$$i\hbar \frac{\partial}{\partial t} a_g(t) = E_g a_g(t) + \hat{H}_{gi} a_i(t) \quad (2.14)$$

$$i\hbar \frac{\partial}{\partial t} a_e(t) = E_e a_e(t) + \hat{H}_{ei} a_i(t) \quad (2.15)$$

$$i\hbar \frac{\partial}{\partial t} a_i(t) = \hat{H}_{gi}^* a_g(t) + \hat{H}_{ei}^* a_e(t) + E_i a_i(t). \quad (2.16)$$

The above equations may be simplified by making the unitary transformation $|\tilde{\Psi}(t)\rangle = \hat{T}(t) |\Psi(t)\rangle$, where $\hat{T}(t) \equiv e^{iE_n t/\hbar} \mathbf{I}$ and $\langle n | \tilde{\Psi}(t) \rangle \equiv b_n(t) = a_n(t) e^{iE_n t/\hbar}$. This reduces the above set of coupled differential equations to:

$$i\hbar \frac{\partial}{\partial t} b_g(t) = \hat{H}_{gi} b_i(t) e^{\frac{it}{\hbar}(E_g - E_i)} \quad (2.17)$$

$$i\hbar \frac{\partial}{\partial t} b_e(t) = \hat{H}_{ei} b_i(t) e^{\frac{it}{\hbar}(E_e - E_i)} \quad (2.18)$$

$$i\hbar \frac{\partial}{\partial t} b_i(t) = \hat{H}_{gi}^* b_g(t) e^{\frac{it}{\hbar}(E_i - E_g)} + \hat{H}_{ei}^* b_e(t) e^{\frac{it}{\hbar}(E_i - E_e)}. \quad (2.19)$$

Under such a unitary transformation, the transformed Hamiltonian $\hat{\hat{H}}$ is given by

$$\hat{\hat{H}} = \hat{T} \hat{H} \hat{T}^\dagger - i\hbar \hat{T} \frac{\partial \hat{T}^\dagger}{\partial t}, \quad (2.20)$$

which can be shown by substituting $|\Psi(t)\rangle = \hat{T}^\dagger(t) |\tilde{\Psi}(t)\rangle$ into the Schrödinger equation for the untransformed state $|\Psi(t)\rangle$.

For this particular unitary transformation, $\hat{\hat{H}}$ is therefore given by

$$\hat{\hat{H}} = \begin{pmatrix} 0 & 0 & \hat{\hat{H}}_{gi} \\ 0 & 0 & \hat{\hat{H}}_{ei} \\ \hat{\hat{H}}_{gi}^* & \hat{\hat{H}}_{ei}^* & 0 \end{pmatrix}, \quad (2.21)$$

where the elements $\hat{\hat{H}}_{gi}$ and $\hat{\hat{H}}_{ei}$ are

$$\hat{\hat{H}}_{gi} = \hbar \frac{\Omega_{1g}}{2} \left[e^{\frac{it}{\hbar} (E_g - E_i) - i\omega_1 t + \phi_1} + e^{\frac{it}{\hbar} (E_g - E_i) + i\omega_1 t - \phi_1} \right] \quad (2.22)$$

$$\hat{\hat{H}}_{ei} = \hbar \frac{\Omega_{2e}}{2} \left[e^{\frac{it}{\hbar} (E_e - E_i) - i\omega_2 t + \phi_2} + e^{\frac{it}{\hbar} (E_e - E_i) + i\omega_2 t - \phi_2} \right]. \quad (2.23)$$

We now define the single photon detuning Δ of laser 1 from the $|g\rangle \leftrightarrow |i\rangle$ transition to be

$$\begin{aligned} \Delta &\equiv \omega_1 - (E_i - E_g)/\hbar \\ &= \omega_1 + \left(\omega_g + \frac{\mathbf{p}^2}{2m\hbar} \right) - \left(\omega_i + \frac{(\mathbf{p} + \hbar\mathbf{k}_1)^2}{2m\hbar} \right), \end{aligned} \quad (2.24)$$

and the much smaller *velocity-sensitive* two-photon detuning δ_{12} of both lasers from the forbidden single-photon transition $|g\rangle \leftrightarrow |e\rangle$ as

$$\begin{aligned} \delta_{12} &\equiv (\omega_1 - \omega_2) - (E_e - E_g)/\hbar \\ &= (\omega_1 - \omega_2) + \left(\omega_g + \frac{\mathbf{p}^2}{2m\hbar} \right) - \left(\omega_e + \frac{(\mathbf{p} + \hbar\mathbf{k}_{\text{eff}})^2}{2m\hbar} \right) \end{aligned} \quad (2.25)$$

$$= (\omega_1 - \omega_2) - \left((\omega_e - \omega_g) + \mathbf{k}_{\text{eff}} \cdot \mathbf{v} + \frac{\hbar\mathbf{k}_{\text{eff}}^2}{2m} \right) \quad (2.26)$$

$$= +\delta_L - \delta_{\text{Doppler}} - \delta_{\text{recoil}}. \quad (2.27)$$

We have defined the laser detuning by $\delta_L \equiv (\omega_1 - \omega_2) - (\omega_e - \omega_g)$. $\delta_{\text{Doppler}} \equiv \mathbf{k}_{\text{eff}} \cdot \mathbf{v}$ is the Doppler shift due to the initial velocity of the atom relative to the rest-frame of the lasers. $\delta_{\text{recoil}} \equiv \hbar\mathbf{k}_{\text{eff}}^2/2m$ is the two-photon recoil shift². Using these definitions of

²The recoil shift vanishes if we replace the initial velocity in the Doppler shift with the mean of initial and final velocities [55].

Δ and δ_{12} , Equations 2.22 and 2.23 become

$$\hat{H}_{\text{gi}} = \hbar \frac{\Omega_{1\text{g}}}{2} [e^{it(\Delta-2\omega_1)+\phi_1} + e^{it(\Delta)-\phi_1}] \quad (2.28)$$

$$\hat{H}_{\text{ei}} = \hbar \frac{\Omega_{2\text{e}}}{2} [e^{it(\Delta-\delta_{12}-2\omega_2)+\phi_2} + e^{it(\Delta-\delta_{12})-\phi_2}], \quad (2.29)$$

and we can see that the “non-resonant” terms with factors $e^{it(\Delta-2\omega_1)}$ and $e^{it(\Delta-\delta_{12}-2\omega_2)}$ oscillate much faster than the “resonant” terms with factors $e^{it(\Delta)}$ and $e^{it(\Delta-\delta_{12})}$. The rapidly oscillating non-resonant terms therefore average out and may be ignored. This is known as the rotating-wave approximation (RWA). Finally, we arrive at the following Hamiltonian

$$\hat{H} = \frac{\hbar}{2} \begin{pmatrix} 0 & 0 & \Omega_{1\text{g}} e^{i\Delta t - i\phi_1} \\ 0 & 0 & \Omega_{2\text{e}} e^{i(\Delta - \delta_{12})t - i\phi_2} \\ \Omega_{1\text{g}}^* e^{-i\Delta t + i\phi_1} & \Omega_{2\text{e}}^* e^{-i(\Delta - \delta_{12})t + i\phi_2} & 0 \end{pmatrix}. \quad (2.30)$$

2.1.4 The effective two-level atom

Suppose $|\Delta| \gg |\Omega_{1\text{g}}|, |\Omega_{2\text{e}}|, |\delta_{12}|$. Under these conditions the intermediate level $|i\rangle$ may be adiabatically eliminated from the system, leading to an effective two-level atom where transitions occur only between the long-lived states $|g\rangle$ and $|e\rangle$.

We begin by explicitly writing out the time-dependence of the amplitudes $b_n(t)$ using the transformed Hamiltonian $\hat{H}$ (Equation 2.30) and the TDSE,

$$i\hbar \frac{\partial}{\partial t} b_{\text{g}}(t) = \hbar \frac{\Omega_{1\text{g}}}{2} e^{i\Delta t - i\phi_1} b_{\text{i}}(t) \quad (2.31)$$

$$i\hbar \frac{\partial}{\partial t} b_{\text{e}}(t) = \hbar \frac{\Omega_{2\text{e}}}{2} e^{i(\Delta - \delta_{12})t - i\phi_2} b_{\text{i}}(t) \quad (2.32)$$

$$i\hbar \frac{\partial}{\partial t} b_{\text{i}}(t) = \hbar \frac{\Omega_{1\text{g}}^*}{2} e^{-i\Delta t + i\phi_1} b_{\text{g}}(t) + \hbar \frac{\Omega_{2\text{e}}^*}{2} e^{-i(\Delta - \delta_{12})t + i\phi_2} b_{\text{e}}(t), \quad (2.33)$$

and assume that amplitudes $b_{\text{g}}(t)$ and $b_{\text{e}}(t)$ oscillate slowly compared with much faster oscillations of $b_{\text{i}}(t)$. This allows us to integrate Equation 2.33 with $b_{\text{g}}(t)$ and $b_{\text{e}}(t)$ assumed to be constant [52]:

$$b_{\text{i}}(t) = \left(\frac{\Omega_{1\text{g}}^*}{2\Delta} e^{-i\Delta t + i\phi_1} b_{\text{g}}(t) + \frac{\Omega_{2\text{e}}^*}{2\Delta} e^{-i(\Delta - \delta_{12})t + i\phi_2} b_{\text{e}}(t) \right). \quad (2.34)$$

Substituting this expression for $b_i(t)$ back into Equations 2.31 and 2.32 we obtain

$$i\hbar \frac{\partial}{\partial t} b_g(t) = \frac{\hbar}{2} \left(\frac{|\Omega_{1g}|^2}{2\delta} b_g(t) + \frac{\Omega_{1g}\Omega_{2e}^*}{2\Delta} e^{i\delta t + i\phi_2 - i\phi_1} b_e(t) \right) \quad (2.35)$$

$$i\hbar \frac{\partial}{\partial t} b_e(t) = \frac{\hbar}{2} \left(\frac{|\Omega_{2e}|^2}{2\Delta} b_e(t) + \frac{\Omega_{2e}\Omega_{1g}^*}{2\Delta} e^{-i\delta t - i\phi_2 + i\phi_1} b_g(t) \right) \quad (2.36)$$

$$i\hbar \frac{\partial}{\partial t} b_i(t) = \hbar \frac{\Omega_{1g}^*}{2} e^{-i\Delta t + i\phi_1} b_g(t) + \hbar \frac{\Omega_{2e}^*}{2} e^{-i(\Delta - \delta_{12})t + i\phi_2} b_e(t). \quad (2.37)$$

We have therefore uncoupled the evolution of the states $|g\rangle$ and $|e\rangle$ from the evolution of state $|i\rangle$. On the time-scale of the oscillations of $b_g(t)$ and $b_e(t)$, $b_i(t)$ undergoes rapid oscillations which average out. Therefore, providing the intermediate level is initially unoccupied it remains unoccupied and we have an effective two-level system with the following Hamiltonian

$$\hat{H} = \frac{\hbar}{2} \begin{pmatrix} 2\Omega_g^{\text{AC}} & \Omega_R e^{+i\delta_{12}t - i\phi_L} \\ \Omega_R^* e^{-i\delta_{12}t + i\phi_L} & 2\Omega_e^{\text{AC}} \end{pmatrix}, \quad (2.38)$$

where we have defined the two-photon Rabi frequency as $\Omega_R \equiv \Omega_{1g}\Omega_{2e}^*/(2\Delta)$ and the relative laser phase as $\phi_L \equiv \phi_1 - \phi_2$. Ω_R represents the coupling strength of the two-photon Raman transition between $|g\rangle$ and $|e\rangle$. The diagonal elements represent the AC Stark or light shifts to the energy levels $|g\rangle$ and $|e\rangle$ of the atom, which we have defined as $\Omega_{g,e}^{\text{AC}} \equiv |\Omega_{1g,2e}|^2/(4\Delta)$. In reality, however, the Stark shifts are modified by the couplings we previously ignored between laser 2 and transition $|g\rangle \leftrightarrow |i\rangle$ and between laser 1 and transition $|e\rangle \leftrightarrow |i\rangle$. If these are included in the calculation the Stark shifts are modified to become [56]:

$$\Omega_{g,e}^{\text{AC}} = \frac{|\Omega_{1g,2e}|^2}{4\Delta} + \frac{|\Omega_{2g,1e}|^2}{4\Delta}. \quad (2.39)$$

We may simplify $\hat{H}$ further still and apply a uniform shift to the energy scale. This is equivalent to applying the transformation $e^{i(\Omega_g^{\text{AC}} + \Omega_e^{\text{AC}})t/2} \mathbf{I}$ to the state vector $(b_g(t) \ b_e(t))^T$ and turns the diagonal elements of the Hamiltonian into $\mp\delta_{\text{AC}}$, where the relative Stark shift is given by $\delta_{\text{AC}} \equiv (\Omega_e^{\text{AC}} - \Omega_g^{\text{AC}})$ [37].

Finally, the time-dependence may be completely removed from the Hamiltonian by transforming into a frame rotating at a rate δ about the z -axis. This transformation is represented by the rotation operator $\hat{R} = e^{-i\delta_{12}t/2\hat{\sigma}_z}$. Using the identity $e^{i\alpha\hat{\sigma}_z} \equiv \cos(\alpha)\mathbf{I} + \sin(\alpha)\hat{\sigma}_z$, this transformation becomes

$$\hat{R} = \begin{pmatrix} e^{-i\delta_{12}t/2} & 0 \\ 0 & e^{i\delta_{12}t/2} \end{pmatrix}. \quad (2.40)$$




FIGURE 2.3: Hyperfine structure within the ^{85}Rb D_2 line ($|5^2P_{3/2}\rangle \rightarrow |5^2S_{1/2}\rangle$ transition) with data taken from [57]. Raman transitions act to couple the two ground hyperfine levels $|5^2S_{1/2}, F = 2\rangle$ and $|5^2S_{1/2}, F = 3\rangle$ via upper hyperfine levels $|5^2P_{3/2}, F' = 2\rangle$ and $|5^2P_{3/2}, F' = 3\rangle$. Each hyperfine level F is further split into $2F + 1$ sub-levels (not shown) in the presence of a weak external magnetic field.

Applying Equation 2.20, we find the rotating frame Hamiltonian

$$\hat{H}_R = \frac{\hbar}{2} \begin{pmatrix} \delta & \Omega_R e^{-i\phi_L} \\ \Omega_R^* e^{i\phi_L} & -\delta \end{pmatrix}, \quad (2.41)$$

where we have absorbed the relative Stark shift into the Raman detuning and defined $\delta \equiv \delta_{12} - \delta_{AC}$. We now have a manageable time-independent Hamiltonian for an effective two-level atom undergoing Raman transitions between two ground hyperfine levels $|g\rangle$ and $|e\rangle$. In Section 2.3 we solve the TDSE using this Hamiltonian in the rotating frame and find the resulting time-dependence of the state amplitudes.

2.2 Rubidium-85

^{85}Rb is our element of choice in the atom interferometry experiments and simulations in this thesis³. It is relatively convenient to cool and trap using the techniques outlined in Chapter 6 due to the existence of a cyclic cooling transition. It is also technically convenient because the D_2 optical transition at 780 nm is easily accessed by existing

³The other isotope widely used in atom interferometry experiments is ^{87}Rb . Although naturally less abundant than ^{85}Rb and with a higher radio-frequency splitting between ground hyperfine levels, ^{87}Rb has a simpler hyperfine and Zeeman structure.

diode lasers such as those originally designed for CD writers, or by frequency-doubling telecommunications laser sources at 1560 nm. Furthermore, as an alkali metal, rubidium has a single electron in its outer shell, making the energy level structure ‘hydrogenic’ and consequently simple to model.

Our interferometry pulses use stimulated Raman transitions between the ground hyperfine levels $|5^2S_{1/2}, F=2\rangle \equiv |g\rangle$ and $|5^2S_{1/2}, F=3\rangle \equiv |e\rangle$, which form part of the hyperfine structure within the $|5^2P_{3/2}\rangle \rightarrow |5^2S_{1/2}\rangle$ transition known as the D_2 line, shown in Figure 2.3 [57]. The interaction between the nuclear spin magnetic moment and the magnetic field generated by the electrons gives rise to what is known as hyperfine structure with energy levels described by the quantum number F [53]. F is the magnitude of the total angular momentum $\mathbf{F}$, i.e. the sum of the total angular momentum of the electrons $\mathbf{J}$ and the nuclear spin $\mathbf{I}$. Since the nuclear spin I is $5/2$ for ^{85}Rb , the level $|5^2S_{1/2}\rangle$ is split into two hyperfine states ($F=2, 3$) and the level $|5^2P_{3/2}\rangle$ into four ($F=1, 2, 3, 4$), following the rules for the addition of angular momenta. As per the dipole selection rules, single-photon transitions must obey $\Delta F = 0, \pm 1$, and therefore our Raman transitions must occur via the intermediate levels $F' = 2, 3$.⁴

In the presence of a weak magnetic field of magnitude B aligned along the positive z direction ($\mathbf{B} = B\hat{z}$), each hyperfine level F splits into $2F + 1$ Zeeman sub-levels described by the quantum number m_F , where $m_F = -F, -F + 1, \dots, F$, representing the projection of the total angular momentum along $\hat{z}$. These sub-levels are shifted in energy by

$$\Delta E_{\text{Zeeman}} = g_F \mu_B B m_F. \quad (2.42)$$

g_F is the Landé g-factor, which may be calculated following the treatment given by Foot [53], and μ_B is the Bohr magneton. This Zeeman shift also leads to a modification of the Raman detuning.

2.2.1 Beam polarisations and multiple Raman routes

The direction of the external magnetic field $\mathbf{B} = B\hat{z}$ provides a quantisation axis for the atom, and both the polarisation and direction of propagation of an incident photon determine which dipole transitions can occur between m_F sub-levels. For single-photon transitions $m_F \rightarrow m'_F$, the conservation of angular momentum implies that $\Delta m_F = m'_F - m_F = 0, \pm 1$, dependent on the specific polarisation of the light relative to $\hat{z}$. When a photon is absorbed (emitted) by an atom its angular momentum is added to (subtracted from) the atomic angular momentum.

⁴Raman transitions cannot occur via the intermediate levels $F' = 1$ or $F' = 4$ because the effective Rabi frequency - given by the product of the two single-photon Rabi frequencies - is zero for these Raman routes: single-photon transitions between $F = 2$ and $F' = 4$ and between $F = 3$ and $F' = 1$ are dipole-forbidden.

A polarised photon whose electric field oscillates in a direction parallel to the quantisation axis causes a π^0 transition to occur where $\Delta m_F = 0$. For a circularly polarised photon propagating along the quantisation axis $\hat{z}$, the polarisation direction rotates in a plane perpendicular to $\hat{z}$ and $\Delta m_F = \pm 1$. If the photon is left-circularly polarised (LCP) and propagating in $\hat{z}$ direction (the polarisation rotates counter-clockwise when viewed towards the source) then σ^+ transitions are driven with $\Delta m_F = +1$. As expected therefore, σ^- transitions are driven by right-circularly polarised (RCP) light propagating along $\hat{z}$ and the result is that $\Delta m_F = -1$.

Linearly polarised light propagating along the quantisation axis with an electric field perpendicular to $\hat{z}$ can be expressed as an equal superposition of left- and right- circularly polarised light, resulting in π^+ (polarisation parallel to $\hat{y}$) and π^- (polarisation parallel to $\hat{x}$) transitions where $\Delta m_F = \pm 1$.

Raman transitions occur between different m_F sub-levels depending on the specific polarisation arrangement of both beams. Here we limit our discussion to two particular arrangements used in the experiments in this thesis: the $\sigma^+ - \sigma^+$ arrangement and the $\pi^+ - \pi^-$ or ‘lin-perp-lin’ arrangement.

In the $\sigma^+ - \sigma^+$ arrangement both Raman beams carry photons with circular polarisations such that the Raman transitions occur between levels $|F = 2, m_F\rangle$ and $|F = 3, m_F\rangle$ via levels $|F' = 2, 3; m_F + 1\rangle$. The ‘lin-perp-lin’ or $\pi^+ - \pi^-$ arrangement occurs when the photons in each Raman beam have orthogonal linear polarisations perpendicular to the quantisation axis. In this situation, since π^+ and π^- polarised light is an equal superposition of left- and right- circularly polarised light, Raman transitions can occur between the levels $|F = 2, m_F\rangle$ and $|F = 3, m_F\rangle$ ⁵ via four possible routes: $|F' = 2, 3; m_F \pm 1\rangle$.

Since our Raman transitions occur via more than one intermediate level our treatment of Raman transitions must be modified to account for multiple Raman routes. We first generalise our single-photon couplings of each laser to the possible intermediate states by defining the vectors

$$\mathbf{\Omega}_{1g} \equiv (\Omega_{|2,m_F\rangle|1,m'_F\rangle}^1, \Omega_{|2,m_F\rangle|2,m'_F\rangle}^1, \Omega_{|2,m_F\rangle|3,m'_F\rangle}^1, 0)^T \quad (2.43)$$

$$\mathbf{\Omega}_{2e} \equiv (0, \Omega_{|3,m_F\rangle|2,m'_F\rangle}^2, \Omega_{|3,m_F\rangle|3,m'_F\rangle}^2, \Omega_{|3,m_F\rangle|4,m'_F\rangle}^2)^T, \quad (2.44)$$

where $\Omega_{|F,m_F\rangle|F',m'_F\rangle}^a$ is the single-photon Rabi frequency for the coupling of states $|F, m_F\rangle$ and $|F', m'_F\rangle$ by laser a . As before we have ignored couplings between the wrong laser and the wrong ground state e.g. $\mathbf{\Omega}_{1e}$ and $\mathbf{\Omega}_{2g}$ although these will affect the light shift. Furthermore we assume that the single-photon detuning from all upper intermediate states is large and approximately equal, taken to be Δ .

⁵Transitions between $|F = 2, m_F\rangle$ and $|F = 3, m_F \pm 2\rangle$ do not occur in the ‘lin-perp-lin’ arrangement because the dipole matrix elements cancel.

In this vector formalism for multiple Raman routes, the two-photon Rabi frequency Ω_R and the Stark shift δ_{AC} become [58, 59]:

$$\Omega_R = \frac{|\mathbf{\Omega}_{1g} \cdot \mathbf{\Omega}_{2e}|}{2\Delta} \quad (2.45)$$

$$\begin{aligned} &= \frac{1}{2\Delta} \sum_{F'=2,3} |\Omega_{|2,m_F\rangle|F',m'_F\rangle}^1 \Omega_{|3,m_F\rangle|F',m'_F\rangle}^2| \\ \delta_{AC} &= \frac{||\mathbf{\Omega}_{2e}||^2 + ||\mathbf{\Omega}_{1e}||^2}{4\Delta} - \frac{||\mathbf{\Omega}_{1g}||^2 + ||\mathbf{\Omega}_{2g}||^2}{4\Delta}. \end{aligned} \quad (2.46)$$

We can therefore see that the two-photon Rabi frequency and Stark shifts of the ground hyperfine levels will depend not only on the polarisation of the light but also on the occupation of the various m_F sub-levels by the atoms, leading to a potential source of Rabi frequency inhomogeneity unless atoms are pumped into a single sub-level.

2.3 Rabi oscillations, coherent pulses, and the Bloch sphere picture

The TDSE tells us how quantum states evolve in time under a given Hamiltonian. Consider the ansatz

$$|\Psi(t)\rangle = \exp\left(-\frac{i}{\hbar}\hat{H} \times (t - t_0)\right) |\Psi(t_0)\rangle. \quad (2.47)$$

By substituting this expression into the TDSE and assuming $\hat{H}$ for our Raman interaction is time-independent (constant laser intensity, frequency, and phase), we see that Equation 2.47 is a solution to the TDSE, and gives us the time-evolution of $|\Psi(t)\rangle = c_g(t)|g\rangle + c_e(t)|e\rangle$. We therefore define the *propagator*

$$\hat{U}(t_0, t) \equiv \exp\left(-\frac{i}{\hbar}\hat{H} \times (t - t_0)\right) \quad (2.48)$$

as the unitary transformation that evolves states from time t_0 to t under the action of time-independent Hamiltonian $\hat{H}$.

For Raman transitions where the laser intensity, phase and frequency are constant for the duration of the interaction, $\hat{U}$ may be evaluated with relative ease by first rewriting $\hat{H}_R$ using the Pauli spin matrices

$$\hat{H}_R = \frac{\hbar}{2} \begin{pmatrix} \delta & \Omega_R e^{-i\phi_L} \\ \Omega_R^* e^{i\phi_L} & -\delta \end{pmatrix} = \frac{\hbar}{2} \mathbf{\Omega}(t) \cdot \hat{\boldsymbol{\sigma}}, \quad (2.49)$$

where $\hat{\boldsymbol{\sigma}} = (\hat{\sigma}_x \hat{\sigma}_y \hat{\sigma}_z)^T$ is the vector of Pauli matrices and $\mathbf{\Omega}$ is the “field vector”, defined by

$$\mathbf{\Omega} = \Omega_R \cos(\phi) \mathbf{x} + \Omega_R \sin(\phi) \mathbf{y} + (\delta) \mathbf{z}. \quad (2.50)$$

The propagator $\hat{U}$ is then given by the exponential

$$\hat{U}(t_0, t) = \exp \left(- \frac{i\mathbf{\Omega}(t) \cdot \hat{\boldsymbol{\sigma}}}{2} (t - t_0) \right), \quad (2.51)$$

which may be evaluated explicitly using the identity $\exp(i\alpha\hat{\mathbf{n}} \cdot \hat{\boldsymbol{\sigma}}) \equiv \mathbf{I} \cos(\alpha) + i\hat{\mathbf{n}} \cdot \hat{\boldsymbol{\sigma}} \sin(\alpha)$, yielding the expression [60]

$$\hat{U} = \begin{pmatrix} C^* & -iS^* \\ -iS & C \end{pmatrix}, \quad (2.52)$$

where C and S are defined as:

$$\begin{aligned} C &\equiv \cos \left(\frac{\sqrt{\Omega_R^2 + \delta^2}(t - t_0)}{2} \right) \\ &+ i \left(\frac{\delta}{\sqrt{\Omega_R^2 + \delta^2}} \right) \sin \left(\frac{\sqrt{\Omega_R^2 + \delta^2}(t - t_0)}{2} \right) \end{aligned} \quad (2.53)$$

$$S \equiv e^{i\phi} \left(\frac{\Omega_R}{\sqrt{\Omega_R^2 + \delta^2}} \right) \sin \left(\frac{\sqrt{\Omega_R^2 + \delta^2}(t - t_0)}{2} \right). \quad (2.54)$$

These equations may be used in accordance with Equation 2.47 to find the final state amplitudes in the rotating frame $c_g(t)$ and $c_e(t)$ following a Raman “pulse” of fixed duration, coupling strength Ω_R , detuning δ , and relative laser phase ϕ_L . If these quantities change during the interaction, then the Hamiltonian becomes time-dependent and the solution in Equation 2.47 is no longer valid.

The Raman Hamiltonian describing the interaction between atomic states and optical radiation in the rotating frame studied in this chapter (Equation 2.49) is isomorphic to the Hamiltonian that describes a nuclear spin interacting with an applied radio-frequency (rf) field studied extensively in nuclear magnetic resonance (NMR) spectroscopy [61]. We make use of this similarity in later chapters by investigating how NMR techniques designed to improve the control of nuclear spins in the presence of systematic errors may be used in the laser control of atoms for interferometry.

2.3.1 Rotations on the Bloch sphere

The state vector of any two-level system may, neglecting a global phase factor, be written as

$$|\Psi\rangle = \cos \left(\frac{\vartheta}{2} \right) |g\rangle + e^{i\varphi} \sin \left(\frac{\vartheta}{2} \right) |e\rangle, \quad (2.55)$$

where ϑ and φ are the polar and azimuthal coordinates of a point on the surface of a unit sphere, known as the Bloch sphere [62] (shown in Figure 2.4). We define the Bloch vector $\boldsymbol{\lambda} = (u \ v \ w)^T$ as the vector representation of the quantum state $|\Psi\rangle$ in




FIGURE 2.4: Bloch sphere representation of the quantum state, $|\psi\rangle$, for a single two-level atom. The north pole corresponds to the lower state $|g\rangle$ and the south pole corresponds to the upper state $|e\rangle$.

this picture. λ has unit length and real elements given by the expectation values of the Pauli spin matrices:

$$u = \langle \Psi | \hat{\sigma}_x | \Psi \rangle = c_g^* c_e + c_g c_e^* \quad (2.56)$$

$$v = \langle \Psi | \hat{\sigma}_y | \Psi \rangle = i(c_g c_e^* - c_g^* c_e) \quad (2.57)$$

$$w = \langle \Psi | \hat{\sigma}_z | \Psi \rangle = |c_g|^2 - |c_e|^2. \quad (2.58)$$

The Bloch sphere provides an intuitive geometric way to analyse the dynamics of our two-level atomic system interacting with the laser beams. All possible states in our two-level system are mapped onto the surface of the Bloch sphere, where the north and south poles correspond to the states $|g\rangle$ and $|e\rangle$ respectively. If the atom is in an equal superposition of the two states, then its state may be represented on the Bloch sphere by a point on the equator.

The three Pauli spin matrices, when exponentiated⁶, lead to operators that describe rotations of the state $|\Psi\rangle$ on the Bloch sphere. Rotations by an angle α about the x , y , and z -axes are given by $\exp(-i\alpha\hat{\sigma}_x/2)$, $\exp(-i\alpha\hat{\sigma}_y/2)$, and $\exp(-i\alpha\hat{\sigma}_z/2)$ respectively. This can be shown by applying these transformations to an arbitrary state and evaluating the components of the resulting Bloch vector. Furthermore, a rotation by α about an arbitrary axis, described by unit vector $\hat{n}$, is given by $\exp(-i\alpha\hat{n} \cdot \boldsymbol{\sigma}/2)$.

By examining the form of our propagator (Equation 2.51), we can therefore see that its action is equivalent to a rotation of the state vector $|\Psi(t)\rangle$ on the surface of the Bloch sphere by an angle given by $|\boldsymbol{\Omega}| \times (t - t_0) = \sqrt{\Omega_R^2 + \delta^2} \times (t - t_0)$ about an axis aligned with the field vector, $\boldsymbol{\Omega}$. The projection of this axis in the xy plane is determined by the relative laser phase ϕ_L , and the inclination from this plane determined by the detuning

⁶The exponential of an $m \times m$ square matrix $\hat{A}$ is defined through the power series $\exp(\hat{A}) \equiv \sum_{n=0}^{\infty} \frac{1}{n!} \hat{A}^n$, where $\hat{A}^0$ is the $m \times m$ identity matrix.




FIGURE 2.5: The action of the propagator (Equation 2.51) rotates the quantum state from $|\Psi(0)\rangle$ to $|\Psi(t_f)\rangle$ on the surface of the Bloch sphere by an angle $\sqrt{\Omega_R^2 + \delta^2} \times (t_f)$ about an axis defined by the field vector $\mathbf{\Omega}$. The field vector's orientation is specified by the detuning from resonance and the laser phase. Note that in this picture we have depicted the normalised field vector $\hat{\mathbf{\Omega}} \equiv \mathbf{\Omega}/\sqrt{\Omega_R^2 + \delta^2}$.

from resonance δ , as illustrated in Figure 2.5. Resonant pulses therefore lead to rotations about axes in the equatorial plane and for an atom initially in the internal state $|g\rangle$, continuous resonant interaction will cause the state of the atom to “flop” between the levels $|g\rangle$ and $|e\rangle$ - known as Rabi flopping.

2.3.2 $\pi/2$ and π pulses

The building blocks of conventional interferometer sequences are rectangular π (mirror) and $\pi/2$ (beam splitter) pulses, which result from tuning the duration of a pulse of constant intensity and laser phase to be either (π/Ω_R) or $(\pi/2\Omega_R)$ respectively. When the interaction is resonant ($\delta = 0$), the π pulse rotates the state vector by π about an axis in the equatorial plane of the Bloch sphere defined by the laser phase ϕ_L and the $\pi/2$ pulse yields a $\pi/2$ rotation about the same axis. These pulses are said to be rectangular in the sense that the Rabi frequency Ω_R is fixed for the duration of the interaction. The evolution of the atomic state vector under the action of resonant rectangular $\pi/2$ and π pulses is shown in Figures 2.6a and 2.6b respectively.

Interferometer sequences also consist of periods of time for which the laser beams are extinguished, known as periods of “free evolution”. In these circumstances Ω_R is zero, and the field vector becomes parallel or anti-parallel to the z -axis of the Bloch sphere depending on the sign of the detuning. Consequently, the quantum state rotates about the z -axis of the Bloch sphere during free evolution, at a rate given by the detuning δ .




FIGURE 2.6: Bloch sphere trajectories for a resonant $\pi/2$ pulse (a) and a π pulse (b). The field vector is represented by the red vector and the initial state in each case is $|g\rangle$. The $\pi/2$ pulse creates an equal superposition of the two basis states, and the π pulse completely inverts the atomic state from pole to pole.

2.4 Pulse imperfections: off-resonance and pulse-length errors

In this section we discuss two main classes of systematic control error that affect Raman pulses: *off-resonance* and *pulse-length* or -strength errors [45]. This categorisation of control errors comes from the NMR literature [61, 63, 64]. Since the Hamiltonian for a spin-1/2 nucleus interacting with an applied rf field is isomorphic to the two-level Raman Hamiltonian of Equation 2.49, NMR terminology may be applied to Raman transitions. However, the sources and possible correlations between these control errors are quite different in atom interferometry.

An off-resonance error arises when the Raman detuning δ is non-zero and consequently the field vector does not lie perfectly in the equatorial plane of the Bloch sphere. As a result, the trajectory followed by the quantum state on the Bloch sphere under a rectangular pulse (e.g. a π or $\pi/2$ pulse) is deflected from its resonant trajectory and the final state is different from that obtained following a resonant pulse. The effect of an off-resonance error may be reduced by increasing the Rabi frequency for a given atom. This is because increasing Ω_R for a given atom relative to fixed δ will decrease the angle between the rotation axis and the equatorial plane of the Bloch sphere. Ω_R may be increased by a) increasing the laser intensity or b) decreasing the single-photon detuning from the intermediate upper level Δ . This means that short high-power pulses are favourable but in practice one is limited experimentally by the maximum available laser power.




FIGURE 2.7: Bloch sphere trajectories for a π pulse acting on the state $|g\rangle$ under the influence of (a) an off-resonance error and (b) a pulse-length error. The off-resonance error deflects the trajectory on the sphere and the pulse-length error leads to an undershooting of the “target” state $|e\rangle$. The rotation axis, or field vector, is shown in each case by the red arrow.

A pulse-length or amplitude error occurs when the desired total rotation angle around the field vector is incorrect, due to either a miscalibration in the pulse duration or an error in the coupling strength (Rabi frequency) for a given atom Ω_R leading to a deviation from its nominal or average value, which we call the effective Rabi frequency Ω_{eff} . The ratio $\Omega_R/\Omega_{\text{eff}}$ provides a natural way to quantify pulse-length errors. Figure 2.7 illustrates both classes of error on the Bloch sphere for a π pulse acting on the initial state $|g\rangle$. The off-resonance error caused by a non-zero Raman detuning leads to a deflection of the trajectory and the pulse-length errors caused by an offset in the Rabi frequency leads to an undershooting of the excited state $|e\rangle$. In other words, both errors lead to a reduction in the fidelity or efficiency of the pulse - meaning that it no longer performs the intended operation.

We can quantify the fidelity of a pulse in many different ways. One option, supposing our atom is initially in the state $|\psi_i\rangle$, is to measure how close the final state is to a desired target state $|\psi_T\rangle$, e.g. $|\langle\psi_T|\hat{U}|\psi_i\rangle|^2$ where $\hat{U}$ is the pulse propagator. For the π pulse we could choose the initial state $|g\rangle$ and a target state $|e\rangle$, and measure the state-transfer efficiency by $|\langle e|\hat{U}|g\rangle|^2$ - which is also the probability for the atom to be in the excited state after the pulse P_e . We can simulate this state-transfer fidelity for the π pulse in the presence of pulse-length and off-resonance errors, and produce a fidelity contour plot such as the one in Figure 2.8 showing how the fidelity decreases as the magnitude of each error is increased.

Although the state-transfer fidelity measure is useful, it only makes sense when considering one particular pairing of initial and target states and tells us nothing about how different initial states are transformed by the pulse. Another more general measure




FIGURE 2.8: Simulated state-transfer fidelity $|\langle e|\hat{U}|g\rangle|^2$ for the rectangular π pulse as a function of off-resonance errors $\delta/\Omega_{\text{eff}}$ and pulse-length errors $\Omega_R/\Omega_{\text{eff}}$. Contours correspond to state-transfer fidelities of 0.5, 0.7, and 0.9.

involves quantifying how close the error-prone pulse propagator $\hat{U}$ in the presence of control errors is to the ideal target propagator $\hat{U}_T$ in the absence of errors, e.g. $\text{Tr}(\hat{U}_T^\dagger \hat{U})$. This measure is different to the point-to-point fidelity measure because it quantifies how the operation of a given pulse varies from the ideal transformation, and thereby removes the need to specify pairs of initial and target states. We return to the topic of fidelity in later chapters when we design and implement replacements to the conventional rectangular Rabi π and $\pi/2$ pulses that are robust to off-resonance and pulse-length errors.

Examining our expressions for the Raman detuning (Equation 2.26) and the two-photon Rabi frequency (Equation 2.45), we can gain insight into the physical origin of off-resonance and pulse-length errors. The Raman detuning depends on the velocity of an atom relative to the fixed laser beams, and therefore for a cloud of atoms with non-zero temperature, the spread in atomic velocities parallel to the Raman beams will lead to a distribution of off-resonance errors: the hotter the cloud, the larger the range of off-resonance errors. Other sources of off-resonance errors include the Zeeman shift (Equation 2.42) which is caused by an uncompensated small magnetic field and also depends on which Zeeman sub-levels are populated within an atomic sample, and the light-shift (Equation 2.46). Examining the two-photon Rabi frequency, we see that any variation in the intensity of the laser beams either spatially or temporally will lead to pulse-length errors, and often this occurs due to the thermal expansion of an atomic cloud across a non-uniform laser intensity profile.

Off-resonance and pulse-length errors limit the ability to accurately control atomic states when there is a variation of these errors within an ensemble of atoms. Any variation in pulse-length or off-resonance errors in an ensemble leads to what is known as *dephasing*




FIGURE 2.9: Simulated Rabi flopping with off-resonance and pulse-length errors. (a) shows how the excited state probability $P_e = |c_e(t)|^2$ of a single atom varies as a function of the time of interaction with the Raman beams for three different cases: zero detuning and no amplitude error (blue solid curve), a detuning of $0.5\Omega_{\text{eff}}$ and no amplitude error (dashed orange curve), and zero detuning but an amplitude error of $\Omega_R/\Omega_{\text{eff}} = 0.5$. The effective Rabi frequency was 200 kHz. (b) shows how the fraction of atoms in the excited state $\langle P_e \rangle$ varies with time for a simulated cloud of ^{85}Rb atoms with a temperature of 50 μK . The distribution of atomic velocities results in a spread of detunings and leads to dephasing within the ensemble - the amplitude of the oscillations decays and precise control of the ensemble is lost.

as different atoms follow different trajectories under the same pulse, limiting one's ability to control the atomic states. In other words, a π pulse will not be a π pulse for all atoms. This is illustrated in Figure 2.9 using the example of Rabi flopping in an atomic ensemble with pulse-length and off-resonance errors. All atoms begin in the same state but, since different atoms have different detunings in a sample with non-zero temperature, their trajectories rapidly diverge on the Bloch sphere and the expected fraction of atoms in the excited state $\langle P_e \rangle$ exhibits Rabi oscillations with decaying amplitude.

Chapter 3

Atom interferometry

One of the extraordinary discoveries of quantum mechanics is the phenomenon of wave-particle duality whereby massive particles can exhibit wavelike properties such as diffraction and interference. Atom interferometers [37] can exploit the interference of these “matter-waves” to enable precise measurements of forces and fields. Essential to the operation of an atom interferometer, are the “atom-optics” that are used to split, redirect, and interfere the atomic matter-waves. After the matter-waves are divided to follow separate spatial paths by a beamsplitter operation, atoms in each path accumulate a different phase due to their interaction with an external field. The resulting phase difference upon recombination of the matter-waves can be used to generate an interference pattern and enable a measurement of the external field.

Neutral atoms make excellent identical test masses and, unlike the neutrons or electrons used in early matter-wave interferometers [65, 66], their internal energy levels allow for precise control of their momentum using resonant laser light. Light-pulse atom interferometry has also been made possible in part by the development of methods to cool and trap atoms using laser light [67–71], which allow a cold ($\sim \mu\text{K}$) sample of many millions of atoms to be prepared and localised in space.

Since the first demonstration of atom interferometry with light pulses in the 1990s [36], atom interferometric sensors have been used to measure gravitational accelerations and Newton’s gravitational constant [41, 72, 73], probe the validity of Einstein’s equivalence principle [74], measure the fine-structure constant [75–78], and test dark energy models [79–81]. Furthermore, they have been proposed as a means to detect gravitational waves [82, 83] and the development of mobile quantum gyroscopes and accelerometers [38, 84–88] has the potential to realise quantum inertial navigation systems with a performance surpassing the best classical devices.

Although many of these atom interferometric sensors have already achieved impressive sensitivities, the performance of all atom interferometers is limited by the fidelity of the

beamsplitter and mirror operations. All the atoms in an interferometer are in principle identical test masses, but their interaction with the light is subject to inhomogeneities due to variations in atomic velocity, magnetic fields, laser intensity, and atomic sub-state. This limits the efficiency of the atomic diffraction: the beamsplitter and mirror pulses do not split and redirect all atoms in the same way, introducing phase errors and reducing the visibility or contrast of the resulting interference fringes. Consequently, in many lab-based interferometers, the atomic sample is filtered to reduce the inhomogeneities present [40, 89], but this inevitably results in a loss of signal [90]. Furthermore, the development of mobile interferometric sensors requires interferometer pulse sequences that operate with high-fidelity in noisy and inhomogeneous environments.

The sensitivity of an atom interferometer scales with the area enclosed by the interfering paths. In an atom interferometer used for inertial sensing the enclosed area, and hence sensitivity, is proportional to the momentum transferred by the light to the atoms. This has motivated the development of large-area or large-momentum-transfer (LMT) atom-optics [40, 89, 91], where additional light pulses may be used to increase the interferometer area and sensitivity. LMT interferometers, described in Section 3.3, are particularly susceptible to pulse infidelity, since in multi-pulse sequences the errors can accumulate, resulting in a loss of interference visibility which also limits any potential sensitivity increase.

Atom interferometers typically use either two-photon Raman transitions [36, 92] or multi-photon Bragg pulses [89] to diffract atomic wave-packets. In a Raman transition, two counter-propagating laser beams induce a transition between two ground hyperfine states in an alkali atom - thus changing the internal (electronic) and external (momentum) state of the atom. Bragg diffraction, however, does not lead to a change in electronic state; the counter-propagating laser beams instead couple two states that differ only in their momentum. Since Raman transitions lead to a change in electronic and momentum state, they enable the use of electronic state readout employed in many current inertial sensor designs [39]. Furthermore, efficient Bragg diffraction requires an atomic source with a very narrow momentum width [90] and attaining these lower temperatures in practice means filtering the atomic sample and hence incurring a loss in the signal-to-noise ratio of the interferometer.

In this chapter we provide a brief overview of atom interferometry where the atom-optics are formed using stimulated Raman transitions between the hyperfine ground states in alkali atoms. Furthermore, we show how such an interferometer may be used to measure acceleration by deriving the resulting interferometric phase shift that is imprinted on the atomic wave-functions. We introduce the important class of large-momentum-transfer (LMT) interferometers, which use extra light-pulses to increase the enclosed area of the interferometer and hence increase the intrinsic sensitivity. In the final section of this chapter we discuss how inhomogeneities present in the atom-light interaction in




FIGURE 3.1: Positions of atomic trajectories as a function of time within a three-pulse Mach Zehnder interferometer sequence composed of ideal rectangular $\pi/2$ and π pulses that are assumed to act instantaneously. The initial beamsplitter pulse prepares a superposition of the two internal states, which separate spatially due to the momentum imparted by the atom-optics ($\hbar\mathbf{k}_{\text{eff}}$). The central mirror pulse acts to swap population between the two states in each interferometer “arm”. The two wave-packets then converge before a final “recombiner” pulse converts the accumulated phase difference between the wave-packets into a difference in atomic population which is then detected at the output. The initial atomic momentum $\mathbf{p}$ is zero in this picture and T is the “dwell-time” between the beamsplitter and mirror pulses.

interferometers reduce the efficiency of the beamsplitter and mirror pulses. We characterise these errors and explain how they ultimately limit interferometer contrast and sensitivity.

3.1 The Mach-Zehnder interferometer

By applying beamsplitter and mirror pulses to separate, redirect, and recombine atomic wave-packets as depicted in Figure 3.1, phase shifts due to inertial effects such as rotations and accelerations are imprinted on the difference between atomic state populations at the output of the interferometer. The most common sequence of pulses used in interferometric sensing applications is the $\pi/2 - \pi - \pi/2$ or “Mach-Zehnder” sequence, illustrated using the Bloch sphere picture in Figure 3.2. With atoms initially in the ground state $|g, \mathbf{p}\rangle$, a Raman $\pi/2$ (beamsplitter) pulse forms an equal superposition of the two states $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}\rangle$. The laser beams are switched off and the two states in the superposition separate spatially due to the momentum imparted by the




FIGURE 3.2: The Mach-Zehnder interferometer sequence depicted for a single atom on the Bloch sphere ((a) to (e)). At the end of the interferometer, the accumulated phase difference between the two internal states is mapped onto a difference in the excited state population, and may be used to generate a fringe pattern. ϕ_{T1} and ϕ_{T2} are the phases accumulated during the first and second dwell-times respectively. The azimuthal angle for the quantum state on the Bloch sphere is shown throughout the sequence. (f) shows a top-down view of the Bloch sphere during the entire sequence. The beamsplitter and mirror pulses are depicted as ideal $\pi/2$ and π pulses respectively.

Raman transition $\hbar \mathbf{k}_{\text{eff}}$, accumulating a relative phase difference. A π (mirror) pulse is used to redirect the two “arms” of the interferometer (by swapping both the internal state and momentum in each part of the superposition) and, after a second period of free evolution with the lasers turned off, a final $\pi/2$ or ‘recombiner’ pulse maps the resulting superposition onto a difference in the atomic population. This may then be measured by detecting the number of atoms in each internal state after the final pulse.

For a Mach-Zehnder interferometer composed of ideal resonant $\pi/2$ and π pulses, we observe interference fringes where the final excited state probability oscillates according to

$$P_e = \frac{1}{2} - \frac{1}{2} \cos(\Phi), \quad (3.1)$$

where Φ is the difference between the phase accrued during the first and second periods of free evolution, which depends upon the acceleration and/or rotation of the atoms relative to the lasers. These fringes allow one to determine the inertial phase and hence make a measurement with the interferometer. For a succinct overview of atom interferometry and the matter-wave analogue of the optical Mach-Zehnder interferometer introduced in this section we refer the reader to Section 18.3 of Cohen-Tannoudji and Guéry-Odelin [54].

3.1.1 Phase shift due to acceleration

As an example, we evaluate the interferometer phase Φ for the case where atoms experience a constant acceleration parallel to the propagation direction of one of the Raman beams. This arrangement is used in experiments with vertically orientated beams to measure the local gravitational acceleration g [36].

We begin by rewriting our *time-dependent* Raman detuning (Equation 2.26) for a single atom as

$$\delta(t) = (\mathbf{v} + \mathbf{a}t) \cdot \mathbf{k}_{\text{eff}}, \quad (3.2)$$

where we have ignored the Stark and recoil shifts and assumed the laser detuning δ_L is zero. $\mathbf{v}$ represents the initial velocity due, for example, to the temperature distribution of the atoms, and $\mathbf{a}$ is a constant acceleration common to all atoms (e.g. due to gravity). We can calculate the relative phase accrued during the two periods of free evolution (with durations T_1 and T_2) between the pulses by solving the Schrödinger equation using the Hamiltonian (Equation 2.49) in the rotating frame with $\Omega_R = 0$:

$$\phi_{T1} = \int_{t_i}^{t_i+T_1} \delta(t) dt = \mathbf{k}_{\text{eff}} \cdot \left(\mathbf{v}t + \frac{1}{2} \mathbf{a}t^2 \right) \Big|_{t_i}^{t_i+T_1} \quad (3.3)$$

$$\phi_{T2} = \int_{t_i+T_1}^{t_i+T_1+T_2} \delta(t) dt = \mathbf{k}_{\text{eff}} \cdot \left(\mathbf{v}t + \frac{1}{2} \mathbf{a}t^2 \right) \Big|_{t_i+T_1}^{t_i+T_1+T_2}. \quad (3.4)$$

t_i represents an arbitrary time at which the pulse sequence commences. For equal periods of free evolution of time $T \equiv T_1 = T_2$, the interferometric phase Φ becomes

$$\Phi = \phi_{T_2} - \phi_{T_1} \quad (3.5)$$

$$= \mathbf{k}_{\text{eff}} \cdot \left(\mathbf{v}(T_2 - T_1) + \frac{1}{2} \mathbf{a}[(t_i + T_1 + T_2)^2 - 2(t_i + T_1)^2 + t_i^2] \right) \quad (3.6)$$

$$= \mathbf{k}_{\text{eff}} \cdot \left(\frac{1}{2} \mathbf{a}[(t_i + 2T)^2 - 2(t_i + T)^2 + t_i^2] \right) \quad (3.7)$$

$$= \mathbf{k}_{\text{eff}} a T^2 \quad (3.8)$$

$$= \mathcal{S} a, \quad (3.9)$$

where we have assumed that the acceleration is parallel with the effective wavevector. $\mathcal{S} \equiv \mathbf{k}_{\text{eff}} T^2$ is the measurement scale-factor of the interferometer - defined as the ratio between the interferometer phase shift and the quantity we wish to measure. The scale-factor is also proportional to the space-time area enclosed by the diffracting wave-packets (see Figure 3.1). We see that for equal periods ($T_1 = T_2 = T$) of free evolution the Doppler contribution to the interferometer phase vanishes. This effect is analogous to the “spin echo” familiar to NMR spectroscopists [93], and relies on the perfect 180° rotation of the mirror pulse in the interferometer sequence.

The phase shift of the interferometer therefore provides a way to measure the acceleration experienced by the atoms relative to the fixed laser beams during the interferometer. By detecting the fraction of the atoms in the excited state following the Mach-Zehnder sequence, this phase shift can be determined. This technique has been used to measure rotations, gravity, and gravity gradients and forms the basis for state-of-the-art ultra-precise atom-interferometric quantum sensors.

The measurement sensitivity of the atom interferometer is defined as the smallest acceleration that results in a detectable change in the measured fraction of atoms in the excited state. When operating at mid-fringe ($\Phi \approx \pi/2$), a change in the excited state probability is directly proportional to a change in the inertial phase shift [94] so the scale-factor can be considered a measure of the interferometer’s intrinsic sensitivity. In practice, however, the actual measurement sensitivity will depend both on the level of detection noise (which determines the smallest detectable change in the excited state population) and the contrast of the interferometer fringes. In line with previous work [42, 46], we refer to the product of the scale-factor and contrast as the measurement sensitivity. In Section 3.3 we discuss how the intrinsic sensitivity (scale-factor) may be increased by extended sequences of Raman pulses, but stress that this does not always guarantee an increase in measurement sensitivity due to an accompanying loss in visibility.

In deriving Equation 3.5 we have glossed over many important details and made several assumptions about the light-atom interactions. For example, we have assumed that

the beamsplitter and mirror pulses occur instantaneously and are perfectly resonant. In reality, in interferometers with vertically orientated Raman beams, the laser frequency difference must be continually chirped to compensate for the Doppler shift caused by the gravitational acceleration and to keep the pulses approximately resonant [36]. Accounting for the finite pulse durations leads to a modification of the interferometric phase [60, 95–98] (known as the pulse duration effect). However, in the limit where the free evolution period is long compared to the pulse durations, the phase shift is still given by $k_{\text{eff}}aT^2$. We have also assumed that all atoms experience the same two-photon Rabi frequency and that this does not vary from pulse to pulse.

There are many different pulse sequences beyond the three-pulse ‘Mach-Zehnder’ that may be employed in atom interferometry and each one is suited to a different purpose [37]. One example is the ‘Ramsey’ 2-pulse $\pi/2 - \pi/2$ sequence widely used in atomic clocks [99–101], which may also be adapted to measure the velocity distribution of an atomic ensemble [102, 103] and implement interferometric cooling schemes [104, 105]. Another pulse sequence is the Ramsey-Bordé $\pi/2 - \pi/2 - \pi/2 - \pi/2$ sequence used to measure the recoil velocity and hence determine the fine-structure constant [78, 106]. More complex multi-pulse sequences are used to increase the enclosed area and hence sensitivity of atom interferometers through large-momentum-transfer (LMT) [40], and for algorithmic cooling applications [107].

3.2 The effect of pulse imperfections

What happens to the interferometer when the beamsplitter and mirror pulses are no longer ideal? We can evaluate the final excited state probability for a given atom following the Mach-Zehnder interferometer by chaining together the propagators for each pulse (including the dwell-times between pulses) and applying the result to the initial state of the atom, which we take to be $|g\rangle$:

$$P_e = |\langle e | \hat{U}_3 \hat{U}_{T_2} \hat{U}_2 \hat{U}_{T_1} \hat{U}_1 | g \rangle|^2. \quad (3.10)$$

$\hat{U}_1$, $\hat{U}_2$, and $\hat{U}_3$ represent the propagators for the initial beamsplitter, central mirror, and final beamsplitter pulses respectively, which we take to be rectangular $\pi/2$ and π pulses. $\hat{U}_{T_1}$ and $\hat{U}_{T_2}$ represent the propagators for the two periods of free-evolution: the first of duration T_1 and the second of duration T_2 .

The resulting excited state probability for a given atom may be written as [60]¹

$$P_e = \frac{1}{2}(\mathcal{A} - \mathcal{B} \cos(\Phi + \Delta\phi)), \quad (3.11)$$

¹A derivation of this expression for the interferometer output may be found in Appendix A.

where $\mathcal{A}$ and $\mathcal{B}$ are the offset and contrast of the interferometer fringes respectively. Φ is the inertial phase term defined by Equation 3.5 and $\Delta\phi$ represents a fixed shift to the phase that may be used to scan the fringe pattern. The detected fraction of atoms in the excited state following an interferometer sequence is given by the integral of Equation 3.11 over the distribution of atoms in the sample. This integral may be computed analytically [108, 109], or else using a Monte Carlo approach [44].

We also find the following expressions for $\mathcal{A}$, $\mathcal{B}$, and $\Delta\phi$ assuming the Rabi rate, laser frequency, phase, and detuning do not vary much on the time-scale of individual pulses:

$$\begin{aligned}\mathcal{A} = & 2(|\langle e|\hat{U}_1|g\rangle\langle e|\hat{U}_2|g\rangle\langle e|\hat{U}_3|g\rangle|^2 \\ & + |\langle g|\hat{U}_1|g\rangle\langle e|\hat{U}_2|g\rangle\langle g|\hat{U}_3|g\rangle|^2 \\ & + |\langle e|\hat{U}_1|g\rangle\langle g|\hat{U}_2|g\rangle\langle g|\hat{U}_3|g\rangle|^2 \\ & + |\langle g|\hat{U}_1|g\rangle\langle g|\hat{U}_2|g\rangle\langle e|\hat{U}_3|g\rangle|^2)\end{aligned}\quad (3.12)$$

$$\mathcal{B} = 4|\langle g|\hat{U}_1|g\rangle\langle e|\hat{U}_1|g\rangle||\langle e|\hat{U}_2|g\rangle|^2|\langle g|\hat{U}_3|g\rangle\langle e|\hat{U}_3|g\rangle|\quad (3.13)$$

$$\begin{aligned}\Delta\phi = & \phi(\langle e|\hat{U}_1|g\rangle) - \phi(\langle g|\hat{U}_1|g\rangle) \\ & - 2\phi(\langle e|\hat{U}_2|g\rangle) \\ & + \phi(\langle g|\hat{U}_3|g\rangle) + \phi(\langle e|\hat{U}_3|g\rangle).\end{aligned}\quad (3.14)$$

We use the notation $\phi(\langle a|b\rangle)$ to indicate the argument of the overlap $\langle a|b\rangle$. We see that the optimal contrast and offset are achieved for all atoms only if the following conditions on the 3 pulses are met:

$$|\langle e|\hat{U}_{i=1,3}|g\rangle|^2 = \frac{1}{2}\quad (3.15a)$$

$$|\langle e|\hat{U}_2|g\rangle|^2 = 1.\quad (3.15b)$$

The contrast of the fringes is therefore reduced if the π pulse can no longer invert the atomic population between the two internal states and/or if the beamsplitters cannot produce an equal superposition of the two internal states. We interpret these conditions as requiring our beamsplitter pulses (pulses 1 and 3) to perform a 90° rotation of the quantum state about an axis in the xy plane of the Bloch sphere. Similarly, the optimal mirror pulse (pulse 2) must perform a 180° rotation of the quantum state. By examining the form of the propagator for each pulse (Equation 2.52) we see that these conditions are met for $\pi/2$ and π pulses in the absence of a detuning or pulse-length error. However, as shown in Chapter 2, in the presence of off-resonance or pulse-length errors, these rotations are no longer performed around the intended axes or by the desired angles and so the contrast of the interferometer fringes is reduced. This can happen in one of two ways: the Rabi rate or Raman detuning may vary for a given atom from pulse to pulse;




FIGURE 3.3: Illustration of different interferometer phase errors arising from a finite-duration beamsplitter pulse. Different atomic velocities δ_1 (red dot-dashed line) and δ_2 (orange dashed line) and acceleration (solid blue line) during the pulse lead different atoms to reach the Bloch sphere equator with different superposition phases. We can associate an effective “pulse-origin” for each atom by tracing back the phases $\phi(\delta_1)$ and $\phi(\delta_2)$ to the phase ϕ_0 that is accumulated without any detuning, finding that $T_1 \neq T_2$. Without mitigation or cancellation by a subsequent pulse, this variation in pulse-origin will modify the measurement scale-factor of the interferometer sequence.

or else remain the same for each atom but vary across the atomic sample. In both cases the ideal pulse conditions (Equations 3.15a and 3.15b) will not be met.

It is also crucial for interferometry that the inertial phase Φ is not modified by a different amount for each atom as a result of the pulse sequence. This is equivalent to requiring the combined phase shift due to the pulse sequence, $\Delta\phi$, to be fixed or cancelled by the pulses for every atom. If $\Delta\phi$ varies from atom to atom, the resulting contrast after thermal averaging will be washed out due to the contribution of fringes with different phases [108] and a phase error may be introduced into the interferometric measurement.

Figure 3.3 illustrates different contributions to the interferometer phase that can arise during a realistic (finite duration) beamsplitter pulse. We see that atoms with different constant detunings, e.g. δ_1 and δ_2 , are brought down to the xy plane with different phases $\phi(\delta_1)$ and $\phi(\delta_2)$. A changing detuning due to acceleration within a pulse also leads to a further modification of this phase. We can associate the spread in phase due to these effects with different effective “pulse-origins” for different atoms by tracing back $\phi(\delta_1)$ and $\phi(\delta_2)$ to the phase ϕ_0 that results from zero detuning and acceleration. We see that the pulse origins, which we refer to by times T_1 and T_2 , are different. Without cancellation or compensation of this effect, the measurement scale-factor will be modified




FIGURE 3.4: (a) Simulated interference fringes for a single atom with no control errors (solid blue line), an off-resonance error (dashed orange line), and a pulse-length error (purple dotted line). (b) The thermally integrated fringe computed numerically for a $50 \mu\text{K}$ cloud of ^{85}Rb with a uniform distribution of pulse-length errors in the range $\pm 30\% \Omega_{\text{eff}}$. The nominal Rabi rate was 200 kHz.

from the value obtained in Equation 3.5. For the $\pi/2 - \pi - \pi/2$ sequence consisting of fractional Rabi oscillations, the phase spread due to different velocities during the first $\pi/2$ pulse is exactly cancelled by the final $\pi/2$ pulse, but only when there is no change in detuning (acceleration) or Rabi frequency between the pulses [109].

To illustrate the effect of imperfect pulses on the contrast of a Mach-Zehnder interferometer, we have simulated the interferometer fringe (Equation 3.11) following the $\pi/2 - \pi - \pi/2$ sequence for two atoms: one with a large pulse-length error and the other with an off-resonance error. The results are shown in Figure 3.4a, where the x -axis depicts a fixed phase shift applied to the final beamsplitter pulse. We see that in both cases the contrast of the fringe is lowered. Figure 3.4b shows the simulated interferometer fringe when the excited state probability is averaged over the velocity distribution due to a cloud of ^{85}Rb at a temperature of $50 \mu\text{K}$, which we assume is a Maxwell-Boltzmann distribution, and a uniform distribution of Rabi frequencies in the range $\pm 30\% \Omega_{\text{eff}}$. The velocity distribution leads to a spread in the Raman detuning, and this coupled with the spread in pulse-length errors leads to a reduction in the contrast of the resulting fringe.

The effect of inhomogeneities in detuning may be reduced by filtering the atomic sample to lower the temperature but this reduces the SNR because fewer atoms can participate in the interferometer. The effect of pulse-length errors due to expansion across inhomogeneous beam-fronts can also be reduced by using beam-shaping optics, though




FIGURE 3.5: Energies (not to scale) of the states $|g\rangle$, $|e\rangle$, and $|i\rangle$ as a function of atomic momentum along the z -axis. The shorter diagonal (red) arrows represent photons in laser 2 that couple the states $|e\rangle$ and $|i\rangle$; the longer diagonal (blue) arrows represent photons in laser 1 that couple the states $|g\rangle$ and $|i\rangle$. The solid (dotted) diagonal arrows correspond to photons whose propagation direction is parallel (anti-parallel) with the z -axis. By reversing the propagation direction of each laser, an atom initially at rest in the state $|g\rangle$ may be given a two-photon recoil momentum shift $\hbar k_{\text{eff}} \approx 2\hbar k_1$ in either direction along the z -axis.

it remains difficult to reduce intensity variations to below $\sim 10\%$ [56, 110]. Another mitigation strategy involves increasing the waist of the (typically Gaussian) beam at the expense of a loss in power and hence Rabi frequency.

3.3 Large-momentum-transfer (LMT) interferometers

It is clear from the interferometer scale-factor $k_{\text{eff}} T^2$ that the intrinsic sensitivity of the interferometer may be increased by extending the time between the pulses T and/or by increasing the momentum transferred to the atoms by the light $\hbar k_{\text{eff}}$. This second method is the basis for large-momentum-transfer (LMT) interferometers [40]. In a Raman LMT interferometer, the beamsplitter and mirror pulses are extended by a sequence of additional π pulses, referred to as augmentation pulses. The direction of the effective wavevector is reversed from pulse to pulse, allowing the augmentation pulses following the initial $\pi/2$ pulse to impart additional two-photon momentum kicks ($\hbar k_{\text{eff}} \approx 2\hbar k$) to atoms in each interferometer arm.

Reversing the direction of the effective wavevector $\mathbf{k}_{\text{eff}} \equiv \mathbf{k}_1 - \mathbf{k}_2$ can be achieved experimentally by reversing the propagation direction of each laser, as described by McGuirk *et al.* [40] and in Chapter 6, Section 6.2.1. In the counter-propagating Raman arrangement, laser 1, which couples the states $|g\rangle$ and $|i\rangle$, has a wavevector $\mathbf{k}_1$ while laser 2, which couples the states $|e\rangle$ and $|i\rangle$, has a wavevector $\mathbf{k}_2 \approx -\mathbf{k}_1$. This means the effective wavevector is given by $\mathbf{k}_{\text{eff}} \equiv \mathbf{k}_1 - \mathbf{k}_2 \approx 2\mathbf{k}_1$: an atom initially in state $|g, \mathbf{p}\rangle$ is coupled to $|e, \mathbf{p} + 2\hbar\mathbf{k}_1\rangle$. However, if we now reverse the propagation directions of each laser, $\mathbf{k}_1 \rightarrow -\mathbf{k}_1$ and $\mathbf{k}_2 \rightarrow -\mathbf{k}_2$, and the effective wavevector becomes $\mathbf{k}_{\text{eff}} \equiv -\mathbf{k}_1 + \mathbf{k}_2 \approx -2\mathbf{k}_1$. The momentum transfer is therefore in the opposite direction: an atom initially in state $|g, \mathbf{p}\rangle$ will now be coupled to $|e, \mathbf{p} - 2\hbar\mathbf{k}_1\rangle$. Figure 3.5 illustrates the effect of this effective wavevector reversal on the states that are coupled in a Raman transition.

By increasing the momentum separation between the diffracting wave-packets, the space-time area enclosed by the interferometer and hence the scale-factor is increased. Figure 3.6 illustrates the Raman LMT process by depicting the atomic state trajectories and pulse sequence within a three-pulse Mach-Zehnder interferometer that has been extended by a total of 8 additional augmentation pulses, leading to an increased momentum splitting of $10 \hbar k$ between the arms.

In a general LMT interferometer, where the beamsplitter operation is extended by a sequence of N augmentation pulses and the mirror pulse is extended by $2N$ such pulses, the momentum separation between the interferometer arms is given by $(2N + 1)\hbar k_{\text{eff}}$, and the scale-factor therefore increases by the same factor compared with the nominal three-pulse sequence². Thus LMT interferometers provide a method to increase inertial sensitivity without significantly increasing the entire duration of the interferometer.

However, one of the main practical limitations of Raman pulse LMT is the decrease in pulse fidelity caused by the limited velocity acceptance of conventional π augmentation pulses. As the arms separate in momentum, their resonance conditions vary. The origin of this detuning shift can be seen by examining how the basis states change throughout an LMT sequence. The basis states for each pulse in an LMT sequence vary when the direction of the effective wavevector is switched. For example, during the initial $\pi/2$ pulse the basis states are $|g, \mathbf{p}\rangle$, and $|e, \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}\rangle$, where $\mathbf{p}$ represents the initial atomic momentum. For the first augmentation pulse (where the wavevector is reversed) the upper arm has a basis described by the states $|g, \mathbf{p} + 2\hbar\mathbf{k}_{\text{eff}}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}\rangle$, and the lower arm has a basis described by the states $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} - \hbar\mathbf{k}_{\text{eff}}\rangle$. The two arms therefore have resonance conditions that are separated by $4\delta_{\text{recoil}} = 4\hbar k_{\text{eff}}^2/2m$.

By extension, the resonance conditions between the upper and lower arms for the n th augmentation pulse ($n = 1, 2, 3, \dots, N$) are separated by $4n\delta_{\text{recoil}}$ [42, 44, 50], meaning the

²This assumes that the time between the augmentation pulses is significantly smaller than the time between the extended beamsplitter sequence and the extended mirror sequence [40].




FIGURE 3.6: Positions of atomic trajectories as a function of time for a three-pulse Mach-Zehnder interferometer and a LMT interferometer, where the beamsplitter and mirror operations are each extended by a sequence of 2 augmentation pulses ($N = 2$), which we label A in the diagram. By increasing the enclosed area of the interferometer the scale-factor and hence intrinsic sensitivity is increased. Solid and dotted lines represent atoms in the internal states $|g\rangle$ and $|e\rangle$ respectively.

resonance conditions continue to change as the sequence is extended³. This causes the fidelity of each augmentation pulse to decrease, and the contrast of the interferometer rapidly decays as more pulses are added [40, 42, 46]. This limits the efficiency of the momentum transfer, and the sensitivity improvements resulting from a larger enclosed area are limited by a reduction in the fringe visibility resulting from the accumulated effect of many imperfect pulses.

³The resonance conditions for the initial $\pi/2$, central π , and final $\pi/2$ pulses remain identical to the nominal three-pulse sequence.

3.4 Summary

Atom interferometers using light pulses to form the beamsplitter and mirror operations have many applications in obtaining ultra-precise measurements of forces and fields. In an atom interferometer where Raman pulses are used as the ‘atom-optics’, we observe interference fringes where the fraction of atoms found in a particular electronic state following the pulse sequence depends on the phase difference accrued by the atoms in each interferometer arm. This phase difference depends, for example, on the acceleration of the atoms relative to the fixed laser beams, and may therefore form the basis of a measurement. The magnitude of the phase shift for a given acceleration (known as the scale-factor) may be increased by increasing the time between the pulses when the light is turned off, or else by increasing the momentum transferred to the atoms by the light - the basis of a LMT interferometer.

Inhomogeneities in the atom-light interaction reduce the sensitivity and performance of atom interferometers. We desire atom-optics that are robust to these effects, thereby enabling high-contrast interferometry with warmer atom samples without the need to filter and reduce SNR; and LMT interferometers capable of maintaining high contrast as the enclosed area is increased. In the remainder of this thesis, we review existing techniques and develop new schemes to tackle this problem by designing robust tailored pulses for atom interferometers using optimal control methods.

Part II

Composite pulses, adiabatic rapid passage, and optimal control

Chapter 4

Composite pulses and adiabatic rapid passage

We have seen how the standard π and $\pi/2$ Raman pulses used to build interferometer sequences are susceptible to variations in the field strength (giving rise to pulse-length errors) and detuning (yielding off-resonance errors). Over many years techniques have been developed in the field of nuclear magnetic resonance (NMR) spectroscopy to mitigate similar problems in the control of nuclear spins. Due to the isomorphism between the nuclear spin- $\frac{1}{2}$ systems studied in NMR and the two-level atom picture employed in interferometry, many of these techniques can be applied to improve the fidelity of operations in atom interferometry.

In this chapter we introduce two NMR techniques that may be used to enhance the light-pulse fidelity in atom interferometers: composite pulses [5, 111] and adiabatic rapid passage (ARP) [6]. In both of these techniques, naive single pulses such as π and $\pi/2$ pulses are replaced by alternatives in which the amplitude, phase, and/or frequency of the driving field are varied in time to replicate the original pulse but, crucially, achieve suppression of common experimental imperfections. Composite pulses consist generally of a small number of rectangular pulses of variable durations, concatenated with discrete phase shifts, while ARP involves slow and smooth modulation of the field amplitude and frequency in order to achieve efficient manipulations of the quantum state.

4.1 Composite Raman pulses

Composite pulses [5], replace the single pulse operations in NMR with sequences of concatenated constant-amplitude pulses with variable phases and durations. The overall action of the composite pulse is designed to replicate the original single pulse, while suppressing off-resonance and/or pulse length errors. There exist a great many composite

pulses developed by the NMR community [63], designed to achieve different nuclear spin transformations but we restrict this introduction to a discussion of inversion or 180° pulses designed to replace the single rectangular π pulse.

The Bloch sphere provides an invaluable tool for visualising composite pulses. Since an individual pulse of constant amplitude and phase may be represented as a rotation about a fixed axis of the Bloch sphere (the field vector), composite pulses may be visualised as sequential rotations on the surface of the sphere. We therefore denote a general composite pulse consisting of n steps with durations t_i , by the sequence

$$\theta_{\phi_1}^1 \theta_{\phi_2}^2 \dots \theta_{\phi_n}^n, \quad (4.1)$$

where time increases from left to right and $\theta^i \equiv \Omega_{\text{eff}} \times t_i$ is the angle the quantum state is rotated by during the i th step under resonant interaction. ϕ_i is the phase of the control field during that specific element of the pulse (the combined Raman laser phase in our system) and specifies the rotation axis of that element, which on resonance lies in the equatorial plane of the Bloch sphere.

Since the Hamiltonian during each step is constant, the overall evolution of the quantum state may be computed by chaining together the propagators for each time-step

$$\hat{U}_{\text{pulse}} = \hat{U}_n \hat{U}_{n-1} \dots \hat{U}_2 \hat{U}_1, \quad (4.2)$$

where in this case time increases from right to left and $\hat{U}$ for a given step is given by Equation 2.52.

A careful choice of θ_i and ϕ_i can result in a pulse where the overall rotation is robust to variations in the detuning δ and the Rabi frequency Ω_R ¹. As an example, we consider the very first composite pulse, the $90_{90}180_090_{90}$ ‘Levitt’ sequence [5, 111], which was designed to invert a spin from pole to pole on the Bloch sphere. We can see that this pulse consists of a π pulse sandwiched between two $\pi/2$ pulses with a 90° phase flip. The action of this pulse is depicted in Figure 4.1, where atomic trajectories have been simulated for four different Rabi frequencies in the range $0.75\Omega_{\text{eff}}$ to $0.9\Omega_{\text{eff}}$ (yielding pulse-length errors). Figure 4.1 shows that the sequence effectively compensates for the undershooting caused by the first incorrect $\pi/2$ rotation with the middle phase-shifted π pulse, and all trajectories end up closer to the target $|e\rangle$ state than under a single π pulse with the same range of pulse-length errors. We can also see from the structure of the Levitt pulse that in the case of no off-resonance or pulse-length error the state of an individual spin is perfectly inverted from $|g\rangle$ to $|e\rangle$. The Levitt pulse is one of the very few composite pulses where the method of error compensation is easily understood pictorially; for longer and more complex sequences this becomes impossible.

¹Composite pulses were not originally designed to compensate variations in detuning and field strength that occur during individual pulses [112]; or variations in the timings or phases of individual segments [63].




FIGURE 4.1: State trajectories on the Bloch sphere under the Levitt composite pulse. Each trajectory corresponds to a different pulse-length error in the range 0.75-0.9 $\Omega_R/\Omega_{\text{eff}}$. We have chosen this one-sided range of pulse-length errors for clarity.

A distinction is drawn in the NMR literature between *point-to-point* (PP) transformations and general- or *universal-rotation* (UR) transformations [8, 10, 45, 63, 64]. The PP pulse is designed to achieve efficient transformation of a specific initial state $|\psi_i\rangle$ to some target state $|\psi_T\rangle$, with some tolerance of off-resonance and/or pulse-length errors. However, under a PP pulse, a different initial state will not necessarily be rotated by the same angle and about the same axis in the presence of errors. A UR pulse, by contrast, is designed to effect a rotation by a fixed angle about a fixed axis for all conceivable initial states with some tolerance of control errors. UR pulses are in general longer and harder to find than PP pulses, as the sequence must provide a robust rotation for a range of initial states [63].

The WALTZ pulse [43] is an example of a PP inversion pulse designed to transfer the quantum state of a two-level system from pole to pole on the Bloch sphere. WALTZ consists of three sub-pulses, and is described by the sequence $90_0 180_{180} 270_0$. Figure 4.2 illustrates the PP behaviour of WALTZ. We have simulated the atomic state trajectories following WALTZ for two different starting states $|\psi_i\rangle$: one on the equator of the Bloch sphere, and one at the pole. When the initial state is at the pole, WALTZ rotates the Bloch vector by 180° about an axis in the equatorial plane with high fidelity in the presence of large off-resonance errors. For large positive and negative detunings, the final state ends up close to the opposite pole. However, an initial state lying on the Bloch sphere's equator undergoes a rotation that is highly dependent upon the offset from resonance.




FIGURE 4.2: Bloch sphere trajectories for the state of a two-level atom subjected to the WALTZ PP inversion pulse for two initial states: $|g\rangle$ and $\frac{1}{\sqrt{2}}(|g\rangle + |e\rangle)$. (a) and (b) show the trajectories for each starting state under resonant interaction. (c) and (d) show the trajectories under a large positive detuning $\delta = \Omega_{\text{eff}}$, and (e) and (f) show the trajectories for each starting state under the opposite negative detuning $\delta = -\Omega_{\text{eff}}$.

4.2 Adiabatic rapid passage

Frequency-swept adiabatic rapid passage (ARP) [6] is a fundamentally different technique to composite pulses. In place of short sequences of concatenated pulses with constant amplitude and phase, ARP works by slowly sweeping the driving field frequency through resonance. If this is done slowly enough, a quantum state initially coincident with the field vector can be made to follow the field with high accuracy, allowing it to be transported to any desired target state [49]. In this introduction we limit our discussion to ARP pulses designed to invert atomic spins.

Consider an atom initially in the state $|g\rangle$ with the Bloch vector aligned along $\hat{\mathbf{z}}$. If the driving field starts far off-resonance then the field vector, $\mathbf{\Omega}$, will also be aligned along $\hat{\mathbf{z}}$ and parallel to the quantum state. If the frequency is then swept to a large negative detuning at a rate much slower than the rate at which the state precesses about the field vector, the angle between the state and field vectors will remain constant and the quantum state will be inverted with high fidelity.

We therefore have two requirements for an adiabatic inversion pulse of duration T with a frequency sweep $\delta(t)$ and amplitude $\Omega_R(t)$. Firstly, the initial and final field vectors must be parallel and anti-parallel with the z -axis respectively. This means that

$$|\delta_0 + \delta(0, T)| \gg \Omega_R(0, T), \quad (4.3)$$

where δ_0 represents a fixed detuning or frequency offset for an individual atom. The second requirement for an efficient adiabatic inversion pulse is that the sweep rate of the field must be much slower than the precession rate of the state about the field. We quantify this condition using the adiabaticity parameter $Q(t)$ [49] defined as

$$Q(t) \equiv \sqrt{\Omega_R(t)^2 + \delta(t)^2} / |\dot{\theta}(t)|. \quad (4.4)$$

$\sqrt{\Omega_R(t)^2 + \delta(t)^2}$ is the rate at which the quantum state rotates about the field vector and $\dot{\theta}(t)$ is the rate of change of the angle between the field vector and the z -axis. The larger $Q(t)$, the more adiabatic the pulse.

As an example, we consider the case of the simplest ARP pulse: the linear frequency chirp [113]. For this pulse, the amplitude Ω_R remains fixed and the Raman detuning or field frequency $\delta(t)$ is swept linearly between $\pm\delta_{\max}$ according to

$$\delta(t) = \delta_{\max}(1 - 2t/T), \quad (4.5)$$

where T is the total duration of the sweep. Figure 4.3 shows the simulated trajectory of the quantum state followed by this ARP sweep on the Bloch sphere; the quantum state undergoes many rapid oscillations about the field as it is swept from a large positive to a large negative detuning. In order to keep the initial and final field vectors aligned




FIGURE 4.3: Bloch sphere trajectory for a resonant atom undergoing a linear frequency sweep. The pulse duration was $50t_\pi$ and δ_0 was $5\Omega_{\text{eff}}$. The adiabaticity $Q(t)$ is at its lowest in the centre of the sweep on resonance.

parallel and anti-parallel with respect to the z -axis respectively, we see from Equation 4.3 that the pulse must sweep over a range of detuning that is considerably larger than the spread in detuning in the atomic ensemble. Furthermore, in order to maintain a high Q throughout, the pulse must be considerably longer than the simple π pulse.

We can simulate ARP pulses by dividing the sweep into many time-steps such that the frequency and amplitude are approximately constant for a given time-step. The evolution of the quantum state for a given step may be approximated using the propagator of the form introduced in Chapter 2 (Equation 2.52) and by chaining together the propagators for each step we can approximate the overall propagator for the pulse.

If we allow both the pulse amplitude and frequency to vary, then ARP pulses can be designed with reductions in the time needed for efficient state inversion [49]. This can be seen by examining Equation 4.3. If Ω_R starts and ends at 0, the field vector will be parallel with the z -axis without the need for a very large initial detuning, thus reducing the total duration of the frequency sweep. Furthermore, the frequency sweep itself need not remain linear. The minimum of $Q(t)$ occurs when the field is resonant and hence the rate of precession of the state about the field vector is at its minimum. This means that although $|\dot{\theta}(t)|$ must remain small on resonance, it may become larger towards the start and end of the pulse since at these times $\sqrt{\Omega_R(t)^2 + \delta(t)^2}$ will be larger [113].

Many different frequency sweeps have been developed in NMR that are more efficient than the linear frequency sweep and achieve faster rates of inversion [6, 44, 51, 114–116]. One example is the tanh/tan pulse [51], where the frequency sweep follows a tangential function of time and the amplitude modulation follows a hyperbolic tangential function

of time:

$$\delta(t) = \Omega_{\text{eff}} \tan \left[\alpha(2t/T - 1) \right] \quad (4.6)$$

$$\Omega_R(t) = \Omega_{\text{eff}} \tanh \left[\beta(1 - |2t/T - 1|) \right]. \quad (4.7)$$

Ω_{eff} is the maximum amplitude and β is a number which quantifies the smoothness of the amplitude variation. $\alpha \equiv \arctan(\delta_{\text{max}}/\Omega_{\text{eff}})$ and the maximum detuning of the frequency sweep is δ_{max} .

4.2.1 Dynamic phase

Although ARP can provide efficient state transfer from pole to pole on the Bloch sphere, the action of the adiabatic frequency sweep on a superposition state results in the accumulation of a large *dynamic phase* [46, 49, 50]. This dynamic phase is a problem when using ARP to improve the fidelity of atom interferometers as it has the potential to lead to significant dephasing of an atomic ensemble. The origin of this dynamic phase can be made clear by moving into the “dressed state” picture of the interaction where the basis states are parallel and anti-parallel to the field vector. In this picture, the “dressed” basis states ($|+\rangle, |-\rangle$) are related to the “bare” basis states ($|g\rangle, |e\rangle$) by the following rotation:

$$|+\rangle = e^{-i\theta(t)\hat{\sigma}_y/2} |g\rangle \quad (4.8)$$

$$|-\rangle = e^{-i\theta(t)\hat{\sigma}_y/2} |e\rangle. \quad (4.9)$$

$\theta(t)$ is the angle between the field vector $\mathbf{\Omega}(t)$ and the z -axis of the Bloch sphere. We have assumed that $\phi_L = 0$ so that the field vector remains in the xz plane. The amplitudes d_+, d_- of a state in the dressed basis are related to those in the bare basis c_g, c_e by the transformation

$$\begin{pmatrix} d_+ \\ d_- \end{pmatrix} = \begin{pmatrix} \cos(\frac{\theta(t)}{2}) & \sin(\frac{\theta(t)}{2}) \\ -\sin(\frac{\theta(t)}{2}) & \cos(\frac{\theta(t)}{2}) \end{pmatrix} \begin{pmatrix} c_g \\ c_e \end{pmatrix}, \quad (4.10)$$

and the equation governing the time-dependence of the dressed-state amplitudes is

$$i\hbar \begin{pmatrix} \dot{d}_+ \\ \dot{d}_- \end{pmatrix} = \frac{\hbar}{2} \begin{pmatrix} \sqrt{\Omega_R^2 + \delta^2} & i\dot{\theta} \\ -i\dot{\theta} & -\sqrt{\Omega_R^2 + \delta^2} \end{pmatrix} \begin{pmatrix} d_+ \\ d_- \end{pmatrix}, \quad (4.11)$$

where we have used Equation 2.20 to transform the bare-state Hamiltonian into the dressed-state basis. In the adiabatic limit, we have that $\sqrt{\Omega_R^2 + \delta^2} \gg |\dot{\theta}|$ and therefore Equation 4.11 simplifies to

$$i\hbar \begin{pmatrix} \dot{d}_+ \\ \dot{d}_- \end{pmatrix} = \frac{\hbar}{2} \begin{pmatrix} \sqrt{\Omega_R^2 + \delta^2} & 0 \\ 0 & -\sqrt{\Omega_R^2 + \delta^2} \end{pmatrix} \begin{pmatrix} d_+ \\ d_- \end{pmatrix}. \quad (4.12)$$

This means that evolution of dressed state amplitudes in the adiabatic limit is very simple:

$$d_{\pm}(t) = \exp(\mp i\gamma) d_{\pm}(0). \quad (4.13)$$

The dressed state amplitudes acquire a *dynamic phase* γ , given by

$$\gamma \equiv \int_0^t \sqrt{\Omega_R(t')^2 + \delta(t')^2} dt' / 2. \quad (4.14)$$

Under an ARP inversion pulse of duration T the field vector is initially parallel with the z -axis, before undergoing a frequency sweep and ending up anti-parallel to z : $\theta(t=0) = 0$ and $\theta(t=T) = \pi$. Therefore, at the beginning and end of the sweep, the dressed states $|+\rangle, |-\rangle$ are equivalent to the bare states $|g\rangle, |e\rangle$. When $\theta = 0$, $|+\rangle = |g\rangle$ and $|-\rangle = |e\rangle$. When $\theta = \pi$, $|+\rangle = |e\rangle$ and $|-\rangle = |g\rangle$. Suppose the atom starts in $|g\rangle$. By considering the evolution in the dressed state picture, we see that following an ARP inversion pulse it will end up in $e^{-i\gamma} |e\rangle$. Similarly for an atom initially in $|g\rangle$ the final state will be $e^{i\gamma} |e\rangle$. The dynamic phase accumulated by the atom does not affect the inversion efficiency but if an atom starts in a superposition state which is not aligned with the field vector then it will accumulate a large relative phase under the ARP inversion pulse ($a|g\rangle + b|e\rangle \rightarrow ae^{-i\gamma} |e\rangle + be^{i\gamma} |g\rangle$). Since this phase depends on the Rabi frequency and detuning during the sweep, this can lead to significant dephasing if these parameters vary from atom to atom or from pulse to pulse [46, 50].

Chapter 5

Optimal control methods

The search for accurate and precise control of a physical system lies at the heart of the development of all advanced technology. Quantum optimal control theory [11] tries to answer the following general question: how should we vary the experimental control parameters at our disposal to realise a desired transformation of a quantum system in the ‘best’ possible way? The word ‘best’ here can mean several different things: we often want to accomplish the desired transformation in the shortest possible time and with the least expenditure of energy, for example. However, we may also seek controls that are robust to experimental inhomogeneities and that still work when conditions vary from the ideal. In practical terms the control parameters are the external fields that are used to interact with a quantum system, and the control sequence is the way in which these fields are varied as a function of time by the experimentalist - typically the pulse shape or sequence - that leads to the optimal transformation of the system.

Optimal control problems may be formulated and solved using ideas from the calculus of variations, the branch of mathematics that deals with finding extrema of functionals, i.e. functions that map functions to scalars. The idea is to find a control sequence - in general a continuous function of time - that maximises an objective functional quantifying how well a system is steered from an initial state to a desired target state by the control sequence. In addition to finding a control sequence that accurately transforms the system from initial to target state, the solution must obey the equations of motion of the system and satisfy any additional constraints imposed on the dynamics. Pontryagin’s Maximum Principle [117] can be used to derive a set of equations that may be solved to find an optimal control sequence. For problems in quantum control, however, it is often impossible to solve such equations analytically and one must employ numerical approaches that re-formulate the control problem in the language of numerical optimisation [11].

It is important to distinguish between closed- and open-loop optimal control algorithms: closed-loop algorithms - also known as feedback control methods - incorporate real-time feedback from the experiment to inform the search for a control sequence, while open-loop control relies solely on a mathematical model of the system. Open-loop quantum control can use gradient-based optimisation algorithms such as gradient ascent [9], quasi-Newton [19] or Newton-type [23] methods, which all use derivatives of the objective function to inform the search. An alternative option that is useful when gradients are difficult to compute is to employ a gradient-free method such as a genetic or evolutionary algorithm [118], or a simplex-type method such as the Nelder-Mead algorithm [20, 119]. In this thesis, our focus is on the application of open-loop control pulse design to problems in atom interferometry. We do not consider feedback control in this thesis because a) we have an experimentally validated model of our experiment [45, 56] and b) we want to design pulses that can improve the performance of atom interferometers with a range of different experimental parameters and arrangements.

Although the composite pulses and adiabatic schemes discussed in the previous chapter can improve the fidelity of given operations, many of these were developed using a combination of intuition [5], analytical calculation [6], and iterative expansion methods [120], with no way to tailor the result to a specific system automatically. Optimal control theory provides a way to design control pulses that are tailored for specific systems and problems, with applications including: the control of chemical reactions using shaped laser pulses [12–15], the design of radio-frequency pulses to control nuclear spins in NMR [7, 8, 10], the control of Bose-Einstein condensates [33–35], the stabilisation of ultra-cold molecules [121], the optimisation of magneto-optical traps [122], nitrogen-vacancy centre magnetometry [27, 28], and the design of high-fidelity gates for quantum information processors using trapped ions [29, 30] and superconducting qubits [31, 32].

The purpose of this chapter is to introduce and explain the computational techniques that we have adapted to design pulses for atom interferometry. In this chapter, we provide an introduction to quantum optimal control theory in the context of the NMR gradient ascent pulse engineering (GRAPE) algorithm [9] and explain how it can be used to optimise shaped Raman pulses for atom interferometry. We discuss how a general control problem is set up and how we can search for solutions using numerical optimisation algorithms. Specifically, we explain how GRAPE may be combined with either first-order gradient ascent, or second-order Newton or quasi-Newton optimisation algorithms [123], and we introduce the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton method [124] - the algorithm used to optimise pulses in this thesis. Throughout this chapter, we explain how the GRAPE method is implemented in the NMR spin dynamics software suite for MATLAB, *Spinach* [125], which we have adapted for atom interferometry pulse design. Further detail on the computational implementation of the optimal control algorithms in *Spinach* may be found in the thesis of David Goodwin [126]. We conclude by introducing the concept of robustness and

present two useful penalties that are used to constrain the smoothness of the solutions and limit the maximum Rabi frequency during a pulse.

We have chosen to use the L-BFGS GRAPE method to design pulses for our purposes because it is computationally efficient and has a proven track record in pulse design in NMR applications. However, in the context of this thesis, optimal control is a tool used to solve particular problems in atom interferometry. The application of this tool through the design of the control problem and choice of fidelity measure is therefore more important than the particular algorithm used - providing it can find an optimal solution.

5.1 Approaches to quantum optimal control

The invention of the NMR gradient ascent pulse engineering (GRAPE) algorithm by Khaneja *et al.* [9] provided an efficient way to compute the derivatives required for gradient ascent optimisation of piece-wise constant control pulses divided into many small time-slices. GRAPE makes the optimisation of arbitrarily shaped pulses characterised by many parameters feasible and quick. Furthermore, improvements to the original GRAPE method [19] enabled the use of quasi-Newton algorithms such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [123, 127–131], resulting in super-linear convergence. GRAPE was also extended, by the development of methods to compute analytical second derivatives of the objective functions, to enable Newton-type optimisation with quadratic convergence properties [23, 126]. Another algorithm adapted for quantum control problems that uses piece-wise constant pulses is Krotov’s method [14, 132, 133]. Whereas GRAPE updates each time-slice in a control pulse concurrently, Krotov uses a sequential update of every time-slice [134]. Krotov’s method has also been extended to second-order Newton and quasi-Newton type optimisation [24, 135].

An alternative to using piece-wise constant control pulses is to reduce the number of parameters by using a Fourier series representation of the pulse shapes. Not only does this approach guarantee smooth solutions, but it reduces the dimension of the pulse optimisation problem by reducing the task to finding a small number of basis coefficients, and was used in early NMR pulse design [118, 136, 137]. Methods that take such an approach include the (gradient-free) chopped random basis (CRAB) algorithm [20, 21], the gradient optimisation of analytical controls (GOAT) method [25], and the algorithms outlined in [22, 26].

5.2 Gradient ascent pulse engineering

Gradient ascent pulse engineering, known as GRAPE, is a well-known optimal control algorithm originally developed by Khaneja *et al.* [9] for the control of spin systems in NMR. As the name suggests, GRAPE combines the numerical optimisation algorithm gradient ascent (often known by its complementary name gradient descent) with control pulse design. In this section we begin with a description of the numerical optimisation method known as gradient ascent, before explaining how a quantum control problem may be set up and solved using GRAPE.

5.2.1 Gradient ascent

Gradient ascent is an algorithm used to optimise a multi-variable function by repeatedly computing the gradient of the function with respect to the variables. The gradient gives the direction of the greatest ascent of the objective function, and by making a step in this direction or against it, one moves towards a local maximum or minimum in the space of variables.

Let $\Phi(\mathbf{x})$ be a general objective function we want to maximise, where $\mathbf{x}$ is a vector of real numbers. Let $\mathbf{x}_0$ be our initial guess for the values of our variables which maximises Φ . We assume that Φ is a smooth function and that it can therefore be continuously differentiated at any point. This property allows us to approximate the value of our objective function at a new position $\mathbf{x}_1 = \mathbf{x}_0 + \alpha \mathbf{h}$ displaced from $\mathbf{x}_0$ by α in a direction $\mathbf{h}$ using a Taylor expansion

$$\Phi(\mathbf{x}_0 + \alpha \mathbf{h}) \approx \Phi(\mathbf{x}_0) + \alpha \nabla \Phi(\mathbf{x}_0)^T \mathbf{h}, \quad (5.1)$$

where $\nabla \Phi(\mathbf{x}_0)$ is the gradient of the function Φ evaluated at position $\mathbf{x}_0$. If we therefore choose our step direction as the gradient of the function at our point $\mathbf{x}_0$, and assume that the linear approximation of the Taylor series is valid i.e. that the function is smooth and the step size is small, the objective function at the new position $\mathbf{x}_1$ will be higher than at $\mathbf{x}_0$: $\Phi(\mathbf{x}_1) \geq \Phi(\mathbf{x}_0)$. We can therefore summarise gradient ascent by the following update rule at each iteration j :

$$\mathbf{x}_{j+1} = \mathbf{x}_j + \alpha \nabla \Phi(\mathbf{x}_j). \quad (5.2)$$

Using a constant step length α for each iteration of gradient ascent is clearly sub-optimal: one could end up repeatedly overshooting the target or end up taking needless additional steps to reach it. A line-search method is used to adapt the step size at each iteration by choosing the length of step that maximises the objective function along the step direction (provided by the gradient). We return to this problem and discuss line-search methods in more detail in Section 5.3.2.

Gradient ascent needs objective functions that are smooth functions of their variables to be able to compute the gradient, and often slows down when nearing maxima or minima as the gradient becomes small. There is also no guarantee of finding a global maximum because the algorithm can become stuck if the landscape has local maxima.

5.2.2 Setting up the control problem

We begin by writing the total time-dependent Hamiltonian for our system

$$\hat{H}(t) = \hat{H}^{(0)} + \sum_{n=1}^M c^{(n)}(t) \hat{H}^{(n)}, \quad (5.3)$$

where we have split the Hamiltonian into controllable and uncontrollable parts. The part that is beyond experimental control is known as the drift Hamiltonian $\hat{H}^{(0)}$ and represents the free evolution of the system. The Hermitian control operators $\hat{H}^{(n)}$ correspond to the M available experimental control fields and $c^{(n)}(t)$ are the corresponding real-valued control amplitudes, where $n = 1, \dots, M$.

Our aim is to find the particular values of the control amplitudes - a control sequence or ‘pulse’ - that maximises a chosen measure of performance or fidelity. Typically, this fidelity is the accuracy with which an initial state or set of states is transformed into a particular target state or set of target states by a specific modulation of the control fields. In such a case, we may define our fidelity $\mathcal{F}$ as

$$\mathcal{F} \equiv f(\langle \sigma | \psi(t_f) \rangle), \quad (5.4)$$

where $|\sigma\rangle$ is the target state, $|\psi(t_f)\rangle$ is the final state of the system after evolution from time 0 to final time t_f , and f is a differentiable function. The choice of fidelity used in the optimisation is crucial, determining both what the pulse does and affecting the convergence properties of the pulse optimisation.

In general, we will have experimental limits on the control amplitudes, such as the total available energy or maximum available frequencies. These costs must be minimised and are therefore subtracted from the fidelity in the resulting optimisation. In practice, however, applying penalties in this way requires careful consideration of the balance between the possibly competing goals of maximising the fidelity and minimising the penalties, and therefore each contribution to the total objective function must be weighted appropriately. We discuss some relevant penalties in Section 5.6, but here we focus on the more challenging problem of maximising the fidelity.

Upon solving the time-dependent Schrödinger equation, the final state of a system evolving from initial state $|\psi(0)\rangle$ under the action of a time-dependent Hamiltonian is given

by

$$|\psi(t_f)\rangle = \underbrace{\exp_{(o)} \left(\int_0^{t_f} -\frac{i}{\hbar} \hat{H}(t) dt \right)}_{\hat{U}(0; t_f)} |\psi(0)\rangle, \quad (5.5)$$

where $\exp_{(o)}$ is the time-ordered exponential [138] and $\hat{U}$ is the propagator for time-evolution.

Using our expression for the final state (Equation 5.5) and our Hamiltonian (Equation 5.3), our expression for the fidelity becomes

$$\mathcal{F} = f(\langle \sigma | \exp_{(o)} \left[\int_0^{t_f} -\frac{i}{\hbar} \left(\hat{H}^{(0)} + \sum_{n=1}^M c^{(n)}(t) \hat{H}^{(n)} \right) dt \right] | \psi(0) \rangle) \quad (5.6)$$

and is therefore a functional of the control sequences $c^{(n)}(t)$, and we have a variational calculus problem where we require the M continuous functions of time $c^{(n)}(t)$ that maximise $\mathcal{F}$.

GRAPE begins by instead assuming piece-wise constant controls. We therefore discretise the M control sequences $c^{(n)}(t)$ into finite-dimensional vectors each with N slices $c_k^{(n)}$ of equal duration Δt , where during each time-slice the control amplitude $c_k^{(n)}$ is a constant and $k = 1, \dots, N$. With this simplification, the propagator for time-evolution under a specific control sequence becomes a time-ordered product of N propagators for each slice

$$\hat{U}(0; t_f) = \overrightarrow{\prod}_k \exp \left(-\frac{i}{\hbar} \underbrace{\left(\hat{H}^{(0)} + \sum_{n=1}^M c_k^{(n)} \hat{H}^{(n)} \right)}_{\hat{H}_k} \Delta t \right) \equiv \overrightarrow{\prod}_k \hat{U}_k, \quad (5.7)$$

where $\hat{H}_k$ is the total Hamiltonian and $\hat{U}_k$ is the propagator, for the time-slice k . Using piece-wise constant controls is not an assumption in our case (or indeed in many NMR situations [19]) because the arbitrary waveform generator used to program our pulses outputs a piece-wise constant waveform, which is then applied to the Raman beams as described in Chapter 6.

Our example fidelity, Equation 5.4, can now be written as

$$\mathcal{F} = f(\langle \sigma | \hat{U}_N \hat{U}_{N-1} \dots \hat{U}_k \dots \hat{U}_2 \hat{U}_1 | \psi(0) \rangle), \quad (5.8)$$

making the fidelity a function of $M \times N$ control amplitudes $c_k^{(n)}$ instead of a functional of M functions $c^{(n)}(t)$ and allowing us to employ gradient ascent. In *Spinach*, the M vectors, each with N control amplitudes for each time-step, are combined into a single control vector $\mathbf{c}$ [126]

$$\mathbf{c} \equiv (c_1^{(1)} \dots c_N^{(1)} | c_1^{(2)} \dots c_N^{(2)} | \dots | c_1^{(M)} \dots c_N^{(M)})^T, \quad (5.9)$$

and the fidelity becomes a scalar function of $\mathbf{c}$: $\mathcal{F}(\mathbf{c})$.

To use the gradient ascent algorithm summarised in Section 5.2.1, we require the derivatives of the fidelity with respect to each control amplitude during each time-slice, namely the derivatives $\partial\mathcal{F}/\partial c_k^{(n)}$ for every n and every k . Once we know the gradient, we can iteratively update the sequence of $M \times N$ control amplitudes $c_k^{(n)}$ from some initial guess and climb up the fidelity landscape. GRAPE computes the gradient of $\mathcal{F}$ in an efficient way, which can be seen by making the observation that the only element within the fidelity (Equation 5.8) that depends on $c_k^{(n)}$ is the propagator $\hat{U}_k$ ¹:

$$\begin{aligned} \frac{\partial}{\partial c_k^{(n)}} \left(\langle \sigma | \hat{U}_N \hat{U}_{N-1} \dots \hat{U}_k \dots \hat{U}_2 \hat{U}_1 | \psi(0) \rangle \right) = \\ \underbrace{\langle \sigma | \hat{U}_N \hat{U}_{N-1} \dots \hat{U}_{k+1}}_{\text{backward propagation}} \frac{\partial \hat{U}_k}{\partial c_k^{(n)}} \overbrace{\hat{U}_{k-1} \dots \hat{U}_2 \hat{U}_1 | \psi(0) \rangle}^{\text{forward propagation}}. \end{aligned} \quad (5.10)$$

Therefore, computing the entire gradient requires just *two* complete simulations: one complete backwards propagation from the target state $|\sigma\rangle$, and one complete forwards propagation from the initial state $|\psi(0)\rangle$. It is important to store each intermediate state in the forwards and backwards propagation:

$$|\psi_k\rangle \equiv \hat{U}_k \hat{U}_{k-1} \dots \hat{U}_2 \hat{U}_1 |\psi(0)\rangle \quad (5.11)$$

$$\langle \sigma_k | \equiv \langle \sigma | \hat{U}_N \hat{U}_{N-1} \dots \hat{U}_{k+1} \hat{U}_k. \quad (5.12)$$

The only remaining calculations required to obtain the full gradient are the $M \times N$ derivatives of the individual propagators for each slice, $\partial\hat{U}_k/\partial c_k^{(n)}$. Once these derivatives are computed, the entire gradient can be obtained by evaluating the overlaps $\langle \sigma_{k+1} | \partial\hat{U}_k/\partial c_k^{(n)} | \psi_{k-1} \rangle$.

Much like the control amplitudes, the gradient of the fidelity with respect to each control amplitude $\partial\mathcal{F}/\partial c_k^{(n)}$ is flattened into a single vector gradient of length $M \times N$ to make the subsequent optimisation simpler:

$$\nabla\mathcal{F} \equiv \left(\frac{\partial\mathcal{F}}{\partial c_1^{(1)}} \dots \frac{\partial\mathcal{F}}{\partial c_N^{(1)}} \mid \frac{\partial\mathcal{F}}{\partial c_1^{(2)}} \dots \frac{\partial\mathcal{F}}{\partial c_N^{(2)}} \mid \dots \mid \frac{\partial\mathcal{F}}{\partial c_1^{(M)}} \dots \frac{\partial\mathcal{F}}{\partial c_N^{(M)}} \right)^T. \quad (5.13)$$

The power of GRAPE lies in the efficiency with it computes the gradient. What if we were to use a conventional finite-difference method to obtain the derivatives $\partial\mathcal{F}/\partial c_k^{(n)}$? Such an approach would require at least $M \times N + 1$ complete evolutions of the system from time 0 to t_f [9]. GRAPE by contrast needs just 2 such simulations: one forwards, and one backwards. This enables the fast optimisation of pulses with many time-steps.

¹We have dropped the arbitrary function f from our fidelity here. As long as f is a differentiable function of the state overlap, then one can use the chain rule to find the required derivatives.

5.2.3 Propagator derivatives and matrix exponentials

To complete the gradient calculation we need derivatives of the slice propagators. Specifically, we need to compute the derivatives

$$\frac{\partial \hat{U}_k}{\partial c_k^{(n)}} = \frac{\partial}{\partial c_k^{(n)}} \exp \left(-\frac{i}{\hbar} (\hat{H}^{(0)} + \sum_{n=1}^M c_k^{(n)} \hat{H}^{(n)}) \Delta t \right), \quad (5.14)$$

i.e. derivatives of matrix exponentials. This is a tricky operation as the control operators and the drift operators will not necessarily all commute with each other, and the approximation to these derivatives given in the original presentation of GRAPE [9] limited the convergence and accuracy of the algorithm [19]. There are various techniques to compute derivatives of matrix exponentials, each incurring a different computational cost. Examples include simply differentiating the Taylor series expansion of the matrix exponential [19], using commutator series [19], and employing finite-difference approximations with a careful choice of finite-difference step size [19]. As explained by DeFouquieres *et al.* [19], using a finite-difference approximation only for the propagator derivatives is a more reasonable approach than applying finite differencing to obtain derivatives of the fidelity $\partial \mathcal{F} / \partial c_k^{(n)}$.

An alternative approach to evaluate propagator derivatives is the auxiliary matrix method, developed and deployed in the context of NMR optimal control by Goodwin and Kuprov [139]. The method makes use of the following relation [126]:

$$\exp \begin{pmatrix} -\frac{i}{\hbar} \hat{H}_k \Delta t & -\frac{i}{\hbar} \hat{H}^{(n)} \Delta t \\ \mathbf{0} & -\frac{i}{\hbar} \hat{H}_k \Delta t \end{pmatrix} = \begin{pmatrix} \exp(-\frac{i}{\hbar} \hat{H}_k \Delta t) & \frac{\partial}{\partial c_k^{(n)}} \hat{U}_k \\ \mathbf{0} & \exp(-\frac{i}{\hbar} \hat{H}_k \Delta t) \end{pmatrix}. \quad (5.15)$$

If we therefore form the composite matrix in Equation 5.15, and compute its exponential, we obtain both the propagator $\hat{U}_k$ and the derivative $\partial \hat{U}_k / \partial c_k^{(n)}$ in a single step. Using the relationship in Equation 5.15, the expression for the entire gradient may be simplified to

$$\frac{\partial \mathcal{F}}{\partial c_k^{(n)}} = \left(|\sigma_{k+1}\rangle \right)^\dagger \left[\exp \begin{pmatrix} -\frac{i}{\hbar} \hat{H}_k \Delta t & -\frac{i}{\hbar} \hat{H}^{(n)} \Delta t \\ \mathbf{0} & -\frac{i}{\hbar} \hat{H}_k \Delta t \end{pmatrix} \right] \begin{pmatrix} \mathbf{0} \\ |\psi_{k-1}\rangle \end{pmatrix}. \quad (5.16)$$

This is how propagator derivatives are computed in *Spinach*.

Therefore, computing both the fidelity and the gradient requires evaluating matrix exponentials. Matrix exponentials can be computed by evaluating their Taylor approximation until convergence. This entails many matrix-matrix multiplications. However, computing each matrix exponential in isolation is not actually necessary. We observe from Equations 5.11, 5.12, and 5.16 that what we actually need to compute is the action of matrix exponentials on vectors and matrix-vector multiplication is significantly cheaper than matrix-matrix multiplication. This allows us to break up the power series

expansion of a matrix exponential into repeated matrix-vector multiplications and make use of Krylov subspace methods [140], implemented in *Spinach* [125, 141].

We can now summarise the steps of the first-order GRAPE algorithm as follows:

1. Identify the drift and M control operators, choose a time-step size and the number of time-slices N , and provide an initial guess for the control amplitudes. Choose a fidelity of the form given by Equation 5.4, specifying an initial state and target state for the system.
2. Propagate the initial state forward and store the intermediate states for all N time-steps.
3. Propagate the target state backwards and store the intermediate states for all N time-steps.
4. Compute the fidelity using Equation 5.4 by computing the inner-product between the final state after the complete forward propagation and the target state.
5. Compute the $M \times N$ elements of the gradient using Equation 5.16.
6. Carry out a line-search to choose an appropriate step-length (see Section 5.3.2).
7. Update the controls by stepping in the direction of the gradient using the gradient ascent update formula (Equation 5.2).
8. Repeat steps 2 to 7 until the norm of the gradient falls below some tolerance level, or a maximum number of iterations is exceeded.

5.3 Second-order gradient ascent

We can reach a local maximum significantly faster if we can compute not only the direction of steepest ascent (the gradient) but also the curvature of the fidelity landscape (the matrix of second derivatives known as the Hessian). To explain this, let us return to our general numerical optimisation problem: finding the particular set of variables that maximises an objective function $\Phi(\mathbf{x})$. Let us again evaluate the objective function at a position $\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{h}$, where $\mathbf{x}_0$ is an arbitrary initial point, and $\mathbf{h}$ is the displacement. This time we will use the multi-variable Taylor series to approximate $\Phi(\mathbf{x}_1)$ by a quadratic surface

$$\Phi(\mathbf{x}_0 + \mathbf{h}) \approx \Phi(\mathbf{x}_0) + \nabla\Phi(\mathbf{x}_0)^T \mathbf{h} + \frac{1}{2} \mathbf{h}^T \mathcal{H}(\mathbf{x}_0) \mathbf{h}, \quad (5.17)$$

where $\mathcal{H}(\mathbf{x}_0)$ is the matrix of second derivatives of Φ known as the Hessian with elements $\mathcal{H}(\mathbf{x}_0)^{ij}$ given by

$$\mathcal{H}(\mathbf{x}_0)^{ij} = \left. \frac{\partial^2 \Phi(\mathbf{x})}{\partial x^i \partial x^j} \right|_{\mathbf{x}_0}. \quad (5.18)$$

We can compute what displacement we need to maximise our quadratic approximation of Φ in a single step starting at $\mathbf{x}_0$. To do this we need to choose a value of $\mathbf{h}$ that maximises or minimises $\Phi(\mathbf{x}_1)$ and therefore set the derivative of 5.17 with respect to $\mathbf{h}$ to zero:

$$0 = \nabla\Phi(\mathbf{x}_0)^T + \mathbf{h}^T \mathcal{H}(\mathbf{x}_0) \quad (5.19)$$

$$0 = \nabla\Phi(\mathbf{x}_0) + \mathcal{H}(\mathbf{x}_0)\mathbf{h}. \quad (5.20)$$

Equation 5.20 can be rearranged to give the step-direction for Newton's method:

$$\mathbf{h} = -(\mathcal{H}(\mathbf{x}_0))^{-1} \nabla\Phi(\mathbf{x}_0). \quad (5.21)$$

Depending on whether the Hessian is positive- or negative-definite, the Newton step direction will either be one which minimises or maximises Φ , or approaches a saddle point when the Hessian is indefinite. Finally, we arrive at the following update formula for Newton's method:

$$\mathbf{x}_{j+1} = \mathbf{x}_j - (\mathcal{H}(\mathbf{x}_j))^{-1} \nabla\Phi(\mathbf{x}_j). \quad (5.22)$$

We also see that Newton's method reduces to first-order gradient ascent with a step size of unity and where the Hessian matrix is the identity.

5.3.1 The Limited-memory Broyden-Fletcher-Goldfarb-Shanno method

Computing the Hessian at each iteration and then inverting it is a far more costly operation than simply computing the gradient of the objective function, and therefore it is often advantageous to use an approximation to the true Hessian. This is the basis of quasi-Newton methods [123], where during each iteration j an approximation to the Hessian $\mathcal{H}_j$ or its inverse $\mathcal{I}_j = \mathcal{H}_j^{-1}$ is either formed or updated using gradients computed at previous iterations. Using a computationally efficient approximation to the Hessian enables one to combine the comparatively low-cost of computing only gradients with the improved convergence rate of Newton's method.

One well-known Hessian-update method is known as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method [123]. In BFGS, the Hessian is updated according to the formula [123]

$$\mathcal{H}_{j+1} = \mathcal{H}_j - \frac{\mathcal{H}_j \Delta \mathbf{x}_j \Delta \mathbf{x}_j^T \mathcal{H}_j}{\Delta \mathbf{x}_j^T \mathcal{H}_j \Delta \mathbf{x}_j} + \frac{\Delta \mathbf{y}_j \Delta \mathbf{y}_j^T}{\Delta \mathbf{y}_j^T \Delta \mathbf{x}_j}, \quad (5.23)$$

where $\Delta \mathbf{y}_j = \nabla\Phi(\mathbf{x}_{j+1}) - \nabla\Phi(\mathbf{x}_j)$ is the difference between successive gradients and $\Delta \mathbf{x}_j = \mathbf{x}_{j+1} - \mathbf{x}_j$ is the difference between successive positions. The update to the Hessian is computed at the end of each iteration, and the update formula is derived such that the approximate Hessian satisfies the condition known as the *secant equation* [123]:

$$\mathcal{H}_{j+1} \Delta \mathbf{x}_j = \Delta \mathbf{y}_j. \quad (5.24)$$

In other words, the approximate Hessian must reflect how much the gradient changes from position to position. The BFGS method provides a symmetric and positive-definite approximation to the Hessian, assuming an initial positive-definite guess is provided and that the *curvature condition* is satisfied by $\Delta \mathbf{x}_j$ and $\Delta \mathbf{y}_j$:

$$\Delta \mathbf{x}_j^T \Delta \mathbf{y}_j > 0. \quad (5.25)$$

Equation 5.23 may also be reformulated as an update rule for an approximation to the inverse Hessian [123]

$$\mathcal{I}_{j+1} = V_j^T \mathcal{I}_j V_j + \frac{\Delta \mathbf{x}_j \Delta \mathbf{x}_j^T}{\Delta \mathbf{y}_j^T \Delta \mathbf{x}_j}, \quad (5.26)$$

where

$$V_j = \mathbb{1} - \frac{\Delta \mathbf{y}_j \Delta \mathbf{x}_j^T}{\Delta \mathbf{y}_j^T \Delta \mathbf{x}_j}. \quad (5.27)$$

The BFGS quasi-Newton method can therefore be summarised by the following update rule for maximisation of $\Phi(\mathbf{x})$:

$$\mathbf{x}_{j+1} = \mathbf{x}_j + \alpha \mathcal{I}_j \nabla \Phi(\mathbf{x}_j). \quad (5.28)$$

Quasi-Newton techniques such as the BFGS update formula (Equation 5.23) can be memory intensive for high-dimensional optimisation problems, especially if the Hessians are not sparse matrices. The BFGS method successively forms an approximation to the Hessian based on the *entire* history of past gradients. However, an alternative and more memory efficient approach, known as the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm [124], involves forming an approximate Hessian using only the m most recent values of $\Delta \mathbf{x}_j$ and $\Delta \mathbf{y}_j$. This prevents the storage and manipulation of large and dense Hessian approximations. The idea is that the most recent gradients and positions are the most important when approximating the Hessian [123]. Equation 5.26 can be used to obtain a recursive procedure that generates an approximation to $\mathcal{I}_j$ using the m most recent values of $\Delta \mathbf{x}_j$ and $\Delta \mathbf{y}_j$. A ‘two-loop’ recursion L-BFGS algorithm [123] is implemented in *Spinach* where $m = 5$ by default.

5.3.2 Completing the line-search

Computing the gradient, and calculating or estimating the Hessian of an objective function Φ provides the direction in which to step in an attempt to maximise or minimise Φ . However, the step size needs to be carefully chosen and updated during each iteration: if the step size is too large then one could overshoot the target and if it is too small it could take forever to get there. More specifically, we require the particular step size for a given iteration α_j that maximises (or minimises) $\Phi(\mathbf{x}_j + \alpha_j \mathbf{h}_j)$, where the direction

$\mathbf{h}_j$ is provided using straightforward gradient ascent or a quasi-Newton method such as L-BFGS. This one-dimensional optimisation problem is known as the line-search [123].

It is not feasible to find an exact solution to the line-search problem. Instead, an inexact solution is found that satisfies certain conditions. The most obvious condition on α for a maximisation problem is to demand that the value of the objective function at the next step is higher than at the current step: $\Phi(\mathbf{x}_j + \alpha \mathbf{h}_j) > \Phi(\mathbf{x}_j)$. We can rewrite this condition for an ascent direction by requiring that

$$\nabla \Phi(\mathbf{x}_j + \alpha \mathbf{h}_j)^T \mathbf{h}_j > 0. \quad (5.29)$$

However, we must also ensure that there is a sufficient increase in the objective function along the search direction because Equation 5.29 alone does not ensure convergence. We therefore introduce the condition for sufficient increase, also known as the Armijo condition [123]:

$$\Phi(\mathbf{x}_j + \alpha \mathbf{h}_j) \geq \Phi(\mathbf{x}_j) + \alpha C_1 \nabla \Phi(\mathbf{x}_j)^T \mathbf{h}_j. \quad (5.30)$$

The constant C_1 is typically a small number (10^{-2} is the default value in *Spinach*).

The condition for sufficient increase will always be satisfied for very small values of α (this follows from the fact that we have already identified an ascent direction before commencing the line-search). We therefore introduce another condition for an acceptable step-length that will exclude very short step-lengths which are not ruled out by the condition of sufficient increase. This new condition is known as the strong curvature condition [123]:

$$|\nabla \Phi(\mathbf{x}_j + \alpha \mathbf{h}_j)^T \mathbf{h}_j| \leq C_2 |\nabla \Phi(\mathbf{x}_j)^T \mathbf{h}_j|. \quad (5.31)$$

The constant C_2 must be between C_1 and 1 and takes the value of 0.9 in *Spinach*. The curvature condition means that the gradient at the new position must be lower than at the starting position. Combining the conditions in Equations 5.30 and 5.31 yields what are known as the strong Wolfe conditions for α [123].

Now that we have some conditions on acceptable values of α for a given iteration of the optimisation, we need to employ an algorithm to find and trial different step sizes. The particular line-search method implemented in *Spinach* uses two phases known as *bracketing* and *sectioning* and the particular steps are outlined in detail in [126, 142]. The essence of this method is to identify an interval in which a reasonable guess for the step-length should lie, and progressively shrink the interval until a step-length is found that satisfies the strong Wolfe conditions.

5.4 Optimised Raman pulse construction

In this section, we explain how we can optimise tailored Raman laser pulses using the formalism and algorithms presented in this chapter.

The Hamiltonian of our system (Equation 2.49) can be expressed as

$$\hat{H}_R = \frac{\hbar\delta}{2}\hat{\sigma}_z + \frac{\hbar}{2}\Omega_R(t)\left[\cos(\phi_L(t))\hat{\sigma}_x + \sin(\phi_L(t))\hat{\sigma}_y\right], \quad (5.32)$$

where the drift Hamiltonian is $\hat{H}^{(0)} = (\hbar\delta/2)\hat{\sigma}_z$ and we identify two control operators $\hat{H}^{(1)}$ and $\hat{H}^{(2)}$ with the Pauli matrices $\hat{\sigma}_x$ and $\hat{\sigma}_y$ respectively. The two control amplitudes $c^{(1)}(t)$, $c^{(2)}(t)$ are given by $\Omega_R(t)\cos(\phi_L(t))$ and $\Omega_R(t)\sin(\phi_L(t))$ respectively. $\phi_L(t)$ is the phase-profile and $\Omega_R(t)$ is the pulse amplitude.

Using GRAPE we can therefore optimise piece-wise constant functions $\Omega_R(t)\cos(\phi_L(t)) = \{\Omega_1\cos(\phi_1), \dots, \Omega_N\cos(\phi_N)\}$ and $\Omega_R(t)\sin(\phi_L(t)) = \{\Omega_1\sin(\phi_1), \dots, \Omega_N\sin(\phi_N)\}$. In *Spinach*, these two vectors, each with length N equal to the number of time-steps in the pulse, are stacked and combined into a single vector $\mathbf{c}$ of length $2N$ which is fed into the optimisation and updated at each iteration:

$$\mathbf{c} = (\Omega_1\cos(\phi_1), \dots, \Omega_N\cos(\phi_N), \Omega_1\sin(\phi_1), \dots, \Omega_N\sin(\phi_N))^T. \quad (5.33)$$

We can convert the elements of $\mathbf{c}$ into piece-wise constant functions of $\Omega_R(t)$ and $\phi_L(t)$ using the relations [126]

$$\Omega_R = \sqrt{(c^{(1)})^2 + (c^{(2)})^2} \quad (5.34)$$

$$\phi_L = \text{atan2}(c^{(2)}/c^{(1)}), \quad (5.35)$$

where atan2 is the four-quadrant inverse tangent which correctly maps the phase ϕ_L into the range $-\pi < \phi_L \leq \pi$ taking into account the signs of $c^{(1)}$ and $c^{(2)}$.

Since GRAPE actually obtains the gradient of the fidelity $\mathcal{F}$ with respect to control amplitudes $c^{(1)}$ and $c^{(2)}$, if we wish to optimise a purely phase-modulated pulse with a constant or user-defined amplitude profile, for example, we must change the variables of the gradient. This can be done using the following transformations [126]:

$$\frac{d}{d\Omega_R} = \cos(\phi_L)\frac{d}{dc^{(1)}} + \sin(\phi_L)\frac{d}{dc^{(2)}} \quad (5.36)$$

$$\frac{d}{d\phi_L} = c^{(1)}\frac{d}{dc^{(2)}} - c^{(2)}\frac{d}{dc^{(1)}}. \quad (5.37)$$

Therefore, we can ignore the derivative with respect to the pulse amplitude and optimise phase-modulated pulses defined by piece-wise constant functions of laser phases ϕ_L .

5.5 Robustness

An important problem in quantum control is the design of control sequences that implement a desired operation in a way that is robust to variations or noise in the experimental parameters [11]. A common approach to achieve robustness involves replacing the single system fidelity measure by an average over the specific variations present in an ensemble of systems. By maximising the ensemble-average fidelity, one maximises the fidelity for each member of the ensemble, thereby finding a robust solution. This method has been used widely in NMR pulse optimisation in recent years [7–10], and it is the approach we adopt in this thesis because of its simplicity and generality.

There are, however, alternative methods used to design robust control pulses. For example, instead of maximising the ensemble-average fidelity we could repeatedly maximise the “worst-case” fidelity within the ensemble [143–146] (also known as “minimax” optimisation). Alternatively, we could deploy geometric arguments, such as those used in early composite pulse design [5], but these are best suited to simple pulse shapes only. There are also many analytical [147–149] and iterative [120, 150] schemes that can nullify the dependence of the fidelity upon, for example, variations in detuning and coupling strength. Additionally, enforcing adiabatic behaviour [6] can lead to robustness; and in some cases the sensitivity to variations can be directly minimised by modifying the fidelity measure [151].

Interestingly, many optimal control pulses exhibit an inherent robustness to fluctuations in the control parameters [152, 153], a feature which may be understood by analysing the curvature of the control landscape around local maxima [154–156]. This naturally leads to a convenient way to characterise the robustness of solutions to control noise [157], thereby allowing one to select those solutions which are most resilient.

5.5.1 Robust Raman pulses

In atom interferometers, the spread in atomic velocities and positions gives rise to a variation in detuning δ and Rabi frequency Ω_R across the ensemble. In other words, each atom possesses a different drift Hamiltonian and different amplitudes of the control operators. This means our single-atom fidelity (Equation 5.4) is actually a function of the detuning and Rabi frequency experienced by the atom during the pulse:

$$\mathcal{F} \rightarrow \mathcal{F}(\delta, \Omega_R). \quad (5.38)$$

Robustness to off-resonance and pulse-length errors may be obtained by replacing our fidelity measure by an average over fidelities computed each using a different drift Hamiltonian (defined by a discrete sample of N_δ resonant offsets δ , running from $\delta_{\min}$ to $\delta_{\max}$)

and different Rabi frequencies (defined by a discrete sample of N_{Ω_R} amplitudes Ω_R , running from $\Omega_{\min}$ to $\Omega_{\max}$):

$$\mathcal{F} = \frac{1}{N_\delta} \frac{1}{N_{\Omega_R}} \sum_{\delta=\delta_{\min}}^{\delta=\delta_{\max}} \sum_{\Omega_R=\Omega_{\min}}^{\Omega_R=\Omega_{\max}} W_\delta(\delta) W_{\Omega_R}(\Omega_R) \mathcal{F}(\delta, \Omega_R). \quad (5.39)$$

$W_\delta(\delta)$ and $W_{\Omega_R}(\Omega_R)$ are optional weighting factors which can be used to represent particular distributions of δ and Ω . The ensemble-averaged fidelity is normalised to unity to enforce consistency across ensembles with different sizes, and the gradient becomes the average of the gradients for each individual member of the ensemble. The specific distributions of detuning and amplitude errors determine the range of errors for which the pulse will be optimised to compensate. GRAPE will attempt to find a pulse that maximises the ensemble-average fidelity and hence the fidelity for each ensemble member. Furthermore, since computing the fidelities and gradients for each individual ensemble member is independent, this process is parallelised in *Spinach*. Additionally, the approach allows different members of the ensemble to be associated with different fidelities, a feature we exploit in Chapter 9.

5.6 Penalties

There are occasions when we wish to apply constraints to our pulses. For example, we may wish to minimise the pulse duration, constrain the maximum amplitude, or enforce smooth solutions. This can be done by subtracting a penalty term from the fidelity such that the numerical optimisation also attempts to minimise the chosen penalty:

$$\mathcal{F}(\mathbf{c}) \rightarrow \mathcal{F}(\mathbf{c}) - w\mathcal{P}(\mathbf{c}). \quad (5.40)$$

w is a weight that may be adjusted to help the algorithm find a good pulse. Much like the fidelity choice, the gradient of each penalty term must be calculated and the Hessian estimated or calculated to permit second-order quasi-Newton or Newton optimisation.

In this section we present two different penalties that we use throughout this thesis: a spill-out norm-square (SNS) penalty that restricts the maximum amplitude during a pulse; and a derivative-norm-square (DNS) penalty that constraints the jaggedness of a pulse and leads to smoother solutions. For a thorough discussion of these penalties and a derivation of their gradients and Hessians, we refer the reader to the thesis of David Goodwin [126].

Longer pulses are more susceptible to decoherence arising from spontaneous emission and effects caused by the motion of atoms during the pulses. Although it is possible to mitigate these effects by minimising the pulse duration, we choose not to apply such a penalty in this work. Instead, our approach is to optimise pulses for an experimentally

feasible range of durations and compare the performance with composite and ARP pulses - the focus of previous studies [42, 46] - of the same durations. As effects such as spontaneous emission will affect pulses with the same pulse area equally, this approach allows us to compare the expected performance in a range of experimental configurations and decay scenarios.

It is important to limit the maximum amplitude or Rabi frequency $\Omega_R(t)$ during a pulse. This is because the maximum power of our Raman lasers represents a physical limitation of the experiment. We could therefore penalise the amplitude during the pulse by subtracting the following penalty from the fidelity

$$\mathcal{P}_{\text{NS}} = \frac{1}{N} \sum_{i=1}^{2N} c_i^2 \quad (5.41)$$

$$= \frac{1}{N} \sum_{k=1}^N \Omega_k^2, \quad (5.42)$$

where c_i are the elements of the flattened vector of $2N$ control amplitudes (Equation 5.33), and Ω_k is the amplitude Ω_R for the time-step k . We refer to this as the norm-square (NS) penalty [126]. One problem with this penalty is that it will constantly drive the amplitude during the pulse to zero and prevent any algorithm converging to a high-fidelity solution: a pulse with no amplitude will have zero fidelity.

A better option, implemented in *Spinach*, is only to apply the NS penalty if the amplitude for a given time-step and iteration passes beyond a user-defined limit - proportional to the maximum Rabi frequency available experimentally. Hence, we define the spill-out norm-square (SNS) penalty by [126]

$$\mathcal{P}_{\text{SNS}} = \begin{cases} \frac{1}{N} \sum_{i=1}^{2N} (|c_i| - A)^2, & |c_i| > A \\ 0, & |c_i| \leq A. \end{cases} \quad (5.43)$$

Application of this penalty will constrain the magnitudes of the control amplitudes $\Omega_R(t) \cos(\phi_L(t))$ and $\Omega_R(t) \sin(\phi_L(t))$ such that the maximum Rabi frequency $\Omega_R(t) \leq A\sqrt{2}$. The derivative of the SNS penalty with respect to the controls is required in gradient-based optimisations and given by

$$\frac{\partial \mathcal{P}_{\text{SNS}}}{\partial c_i} = \begin{cases} \frac{2}{N} \sum_{i=1}^{2N} (|c_i| - A), & |c_i| > A \\ 0, & |c_i| \leq A. \end{cases} \quad (5.44)$$

We may sometimes wish to ensure that our optimised pulses possess smooth waveforms (a specific example relating to phase noise is given in Chapter 8), and we therefore introduce a derivative-norm-square (DNS) penalty that acts to smooth the control pulse throughout the optimisation. The DNS penalty is proportional to the squared-norm of

the second derivatives of the control amplitudes $c^{(1)}(t)$ and $c^{(2)}(t)$,

$$\mathcal{P}_{\text{DNS}}^{(n)} = \frac{1}{N} \sum_{k=1}^N \left(\sum_{l=1}^N D_{kl} c_l^{(n)} \right)^2, \quad (5.45)$$

where D_{kl} is an operator that obtains an approximation to the second-derivative of $c^{(n)}(t)$, for example by using finite differencing. We can find the derivative of this penalty with respect to the control amplitudes as follows:

$$\frac{\partial \mathcal{P}_{\text{DNS}}^{(n)}}{\partial c_j^{(n)}} = \frac{2}{N} \sum_k \left(\sum_l D_{kl} c_l^{(n)} \right) \left(\frac{\partial}{\partial c_j^{(n)}} \sum_l D_{kl} c_l^{(n)} \right) \quad (5.46)$$

$$= \frac{2}{N} \sum_k \left(\sum_l D_{kl} c_l^{(n)} \right) \left(\sum_l D_{kl} \delta_{jl} \right) \quad (5.47)$$

$$= \frac{2}{N} \sum_k \sum_l D_{kl} c_l^{(n)} D_{kj}. \quad (5.48)$$

The derivatives of the DNS penalty with respect to each of the control amplitudes $c^{(1)}(t)$ and $c^{(2)}(t)$ may then be flattened into a single vector of length $M \times N$, and fed back into the optimisation. In *Spinach*, D represents a second-order finite-difference approximation the second-derivative using five stencil points with left- and right-sided finite-difference schemes for the edges.

Part III

Experimental setup and procedures

Chapter 6

Experimental procedures

To implement and test the optimised pulses which are the main focus of this thesis we require a cold sample of ^{85}Rb that we can localise in space and in which we can induce Raman transitions between the two hyperfine ground states $|g\rangle = |5^2S_{1/2}, F=2\rangle$ and $|e\rangle = |5^2S_{1/2}, F=3\rangle$. We also need ways to prepare all the atoms in the same state, measure the temperature of the atom cloud, apply the optimised Raman pulses to the atoms, and detect the resulting fraction of the atoms in a given state after a pulse sequence. We employ a combination of optical and magnetic fields to trap and cool ^{85}Rb in a magneto-optical trap (MOT) [70]. In this chapter, we first provide an overview of how a MOT works, before providing a description of the experimental procedures and systems used to implement Raman pulses in the atom cloud.

I developed the pulse designs, optimisations, and simulations which form the key results of this thesis. The experiments were performed by Max Carey, and I advised and assisted with the experimental implementation of the pulses; the results of this collaboration are presented in [3].

6.1 Magneto-optical trapping of atomic rubidium

^{85}Rb atoms may be captured and cooled within a magneto-optical trap (MOT) [70]. A MOT works by using a combination of optical and magnetic fields which lead to velocity- and position-dependent forces on individual atoms. In this section we provide an overview of the theory underlying a MOT and explain how such a trap is implemented experimentally.

Photons have momentum and therefore a laser beam can impart a scattering force $\mathbf{F}$ to a two level atom, given by the product of the imparted photon momentum $\hbar\mathbf{k}$ and

the rate of absorption of photons by the atom¹. If the atom is moving towards the laser source, this force will be in the direction opposite to the atomic motion since the atoms are given a kick by the momentum of absorbed photons. Upon absorption of the incident photons, the atoms will spontaneously emit the photons randomly in all directions, thus resulting in a net force on the atoms. This force on our two-level atom can be written as

$$\mathbf{F} = (\hbar\mathbf{k}) \times \frac{\Gamma}{2} \frac{\Omega^2/2}{\delta^2 + \Omega^2/2 + \Gamma^2/4}, \quad (6.1)$$

where $\hbar\mathbf{k}$ is the photon momentum, Γ is the spontaneous decay rate of the transition, Ω is the Rabi frequency, and δ is the detuning of the laser beam from resonance. The detuning is given by

$$\delta = (\omega_L - \omega_0) - \delta_{\text{Doppler}} - \delta_{\text{Zeeman}}, \quad (6.2)$$

where ω_0 is the transition frequency, ω_L is the laser frequency, $\delta_{\text{Doppler}} = \mathbf{k} \cdot \mathbf{v}$ is the Doppler shift dependent on the atomic velocity $\mathbf{v}$, and $\delta_{\text{Zeeman}} = m_F \mu_B g_F B / \hbar$ is the Zeeman shift (Equation 2.42), which depends on the magnetic field B .

If atoms are illuminated by counter-propagating beams which are red-detuned from resonance ($\delta_L = \omega_0 - \omega_L > 0$), then δ^2 decreases when the wavevectors $\mathbf{k}$ and the atomic velocities are in opposite directions. This means that although for an atom at rest the forces from each beam are balanced, if the atom is travelling towards a given beam it becomes resonant with the beam and the forces become imbalanced. This means there is a velocity-dependent damping force which will cool atoms along the beam axis. This technique is known as optical molasses, and if counter-propagating beams are aligned with all three Cartesian axes, the result is cooling in three dimensions.

Optical molasses cannot cool atoms without limit; the cooling process is opposed by heating due to fluctuations in the random momentum kicks associated with spontaneous emission. This leads to a theoretical minimum temperature known as the Doppler limit, determined by the linewidth of the atomic transition. For ^{85}Rb , it is approximately 146 μK . However, in practice we can reach considerably lower temperatures by exploiting sub-Doppler cooling mechanisms [71, 158], which rely on factors such as regular variations in the polarisation of the electric field vector along the cooling axis. A detailed treatment of the sub-Doppler cooling mechanisms at play in our experiment is beyond the scope of this work.

No real atoms behave as perfect two-level systems and to implement optical molasses one must identify a ‘closed transition’ within a given atom. The cooling process requires many cycles of absorption and spontaneous emission. If the atoms spontaneously emit

¹In the steady state the rate of absorption is equal to the rate of scattering, and is given by the product of the spontaneous decay rate and the steady-state population of the excited state, which may be found by solving the optical Bloch equations in the presence of a damping term to account for spontaneous decay [53].




FIGURE 6.1: Diagram showing the cooling and repump optical transitions used in a ^{85}Rb MOT. The solid lines show the absorption routes and the dotted lines show the allowed routes of spontaneous emission. The cooling laser is red-detuned, δ_L , from the $|5^2S_{1/2}, F = 3\rangle \rightarrow |5^2P_{3/2}, F' = 4\rangle$ transition. The cooling beam will also excite atoms into the $|5^2P_{3/2}, F' = 3\rangle$ state with a small non-zero probability and these atoms can spontaneously decay into the lower level $|5^2S_{1/2}, F = 2\rangle$ [56]. The repump beam is therefore tuned to the $|5^2S_{1/2}, F = 2\rangle \rightarrow |5^2P_{3/2}, F' = 3\rangle$ transition to return atoms to the cooling cycle.

to ‘dark’ states from which they cannot be excited back by the same lasers then they can no longer be cooled.

^{85}Rb is a reasonably good choice of atomic species to trap and cool in a MOT because of the existence of a closed transition route to implement optical molasses. In order to cool ^{85}Rb , the MOT cooling laser is red-detuned from the $|5^2S_{1/2}, F = 3\rangle \rightarrow |5^2P_{3/2}, F' = 4\rangle$ transition on the D_2 line, as shown in Figure 6.1. The solid line shows the route of absorption of a photon from the beam, and the dotted line shows the dipole-allowed route of spontaneous emission. However, atoms are lost from this cycle because the cooling laser can, with a small probability, also excite atoms to the $|5^2P_{3/2}, F' = 3\rangle$ level. Atoms excited to this state can spontaneously emit to the ‘dark’ lower ground hyperfine state, $|5^2S_{1/2}, F = 2\rangle$, which is no longer coupled by the cooling laser, and therefore no longer participate in the cooling. At least one repump beam is therefore required to pump atoms out of this state and back onto the cooling cycle, also shown in Figure 6.1. This repump beam is tuned to the $|5^2S_{1/2}, F = 2\rangle \rightarrow |5^2P_{3/2}, F' = 3\rangle$ transition.

Although optical molasses slows atoms down through a velocity-dependent force, it

does not trap them. A position-dependent force is needed to confine atoms to a region of space. In a MOT, this is done by engineering a position-dependent magnetic field and exploiting the Zeeman splitting of the atomic energy levels [70]. A pair of current-carrying coils in which the current flows in alternating directions (known as an anti-Helmholtz configuration) generate a quadrupole magnetic field in the region between the coils. In a quadrupole magnetic field, the field strength increases linearly from zero as an atom moves away from the centre in any direction. Thus, the Zeeman splitting also increases linearly in all directions from the centre. If the 3 pairs of counter-propagating red-detuned cooling beams carry opposite circularly-polarised light, the result is a position-dependent absorption of light in all directions.

To illustrate how this provides a trapping force, we can consider what happens in one direction, say the z direction. As an atom moves away from the centre of the quadrupole field in the positive z direction, the energy splitting induced between the m_F levels increases, and the $\Delta m_F = -1$ transition moves closer to resonance with red-detuned σ^- light propagating in the $-z$ direction, leading to a restoring force directed towards $z = 0$. Conversely, for an atom at a position $z < 0$, the opposite direction of the magnetic field leads to the opposite shift in the m_F levels. As a result, the $\Delta m_F = +1$ transition becomes resonant with red-detuned σ^+ polarised light travelling in the $+z$ direction. Thus, we have a position-dependent restoring force directed towards the centre of the magnetic field.

Figure 6.2 depicts the experimental arrangement of the MOT, showing how the cooling and repump beams are delivered to the vacuum chamber. The cooling and repump beams are generated using two *Vescent D2-100-DBR* distributed Bragg reflector (DBR) lasers. Both lasers are locked to the correct transition frequencies as described in [103] and each beam is coupled into fibres and passed twice through separate acousto-optic modulators (AOMs) to control the amplitude of the light sent to the vacuum chamber without applying an additional frequency shift to each beam: by double-passing through an AOM, the shift in frequency is added and then subtracted.

The cooling light is divided equally between 3 perpendicular axes by a combination of half-waveplates (HWPs), polarising beamsplitting cubes (PBSCs), and mirrors. Rotating the HWPs allows us to finely balance the intensity of the light along each axis. Quarter-wave plates (QWPs) are used to polarise the light before it enters the vacuum chamber. The light along each axis is retro-reflected back into the chamber and an additional QWP ensures that the counter-propagating beams have the correct opposite circular polarisations needed for the MOT. The repump beam is divided using a PBSC between the two horizontal axes. Finally, current carrying coils in an anti-Helmholtz configuration above and below the vacuum chamber are used to generate the required quadrupole field for the MOT's trapping force.




FIGURE 6.2: Experimental arrangement of MOT laser beams and vacuum chamber. (a) is a side-view of the vacuum chamber; while (b) depicts a top-down view. The cooling light is divided by a polarising beamsplitting cube (PBSC) below the chamber and one-third of the light is sent upwards through the chamber, and subsequently reflected down again. The quarter waveplates (QWPs) ensure that the beams have opposite circular polarisations. The remaining two-thirds of the cooling light is directed to the upper level and split by another PBSC which divides the light evenly between the two horizontal axes. The half waveplates (HWPs) are used to control the balance of the cooling light between the axes. The repump beam enters on the upper level and, after passing through the same PBSC as the cooling light, is split between the two horizontal axes. The MOT coils in anti-Helmholtz configuration are also shown in (a) along with a representation of the quadrupole magnetic field. The fluorescence from atomic decays is captured by a photo-multiplier tube (PMT), shown in (b).

6.2 Raman laser system

The coherent control schemes in this thesis rely on the ability to induce stimulated Raman transitions in atomic rubidium. To implement such Raman transitions between the ground hyperfine levels $|5^2S_{1/2}, F = 2\rangle$ and $|5^2S_{1/2}, F = 3\rangle$ in ^{85}Rb we require two phase-coherent beams that we can separate spatially with a frequency difference matching the hyperfine splitting $\omega_{\text{eg}} = 2\pi \times 3.03573244 \text{ GHz}$. We also need the ability to apply a tunable additional detuning δ_L to the Raman beams. Furthermore, we require precise control of the relative phase ϕ_L between the two frequency components, as this is the main control parameter used to construct our optimised Raman pulses. In this section we describe how we generate two Raman beams with the correct frequencies, and how we can control their relative phase.

Figure 6.3 provides a simplified overview of how the Raman beams are generated and controlled in the experiment. Both Raman beams are generated from a single 780 nm *Ea-gleyard* distributed feedback (DFB) diode laser. The light at frequency ω_0 is first passed through *one of two* acousto-optic modulators (AOMs) (selected by a radio-frequency

switch) driven at a fixed frequency of $\omega_A = 310$ MHz. One AOM will diffract the incoming light into the first positive order with frequency $\omega_0 + \omega_A$ and the other will diffract the light into the first negative order with frequency $\omega_0 - \omega_A$. The effective wavevector of the Raman transition may be reversed by alternating which AOM is used to diffract the light from the DFB diode laser. This is explained in subsection 6.2.1.

Suppose the positive first-order deflected beam is selected to form one of the Raman beams. The undeflected light is sent through an electro-optic modulator (EOM) driven at a frequency of $\omega_E = \omega_{\text{eg}} - \omega_A + \delta_L \approx 2.7$ GHz. The EOM produces spectral sidebands at frequencies $\omega_0 \pm n\omega_E$ where n is an integer.

Unlike the AOM, the different frequency components at the output of the EOM are not spatially separated, and the central or carrier frequency component at ω_0 must be removed. This is done using an optical fibre Mach-Zehnder Interferometer (MZI), described in [103]. After the carrier has been removed from the EOM output, the frequency difference between the deflected AOM output light and the $n = -1$ sideband from the EOM output, is given by

$$\omega_A + \omega_0 - [\omega_0 - (\omega_{\text{eg}} - \omega_A + \delta_L)] = \omega_{\text{eg}} + \delta_L, \quad (6.3)$$

hence forming a Raman pair with the correct frequency difference². The light from the AOM and EOM is then sent through two separate tapered amplifiers (TA) to increase the optical power, and combined using a polarising beamsplitter cube (PBSC).

The output intensity of the Raman beams is controlled by passing both beams through a final AOM. The undeflected beam is dumped, and the first diffracted order is only present when the radio-frequency (RF) signal generator driving the AOM is switched on, thus providing an effective switch for the Raman beams.³ This AOM also shifts the frequency of both Raman beams but since the shift is common to both components the resonance condition of the Raman transition is unaffected.

Following the final AOM, both beams are fibre-coupled and sent through refractive beam-shaping optics [56]. These beam-shapers produce an output beam with a top-hat beam intensity profile with a cross-sectional area of approximately 1.2 mm^2 . Although the resulting intensity profile of the beam-front is approximately flat, there are variations in intensity of the order of $\pm 5\%$, and these variations cause the Rabi frequency of the Raman transition to vary across the atom cloud [56].

²We note that, in the experimental arrangement used to obtain the data in this thesis, the off-resonant $n = 1$ sideband from the EOM was not removed, leading to a small additional light shift in the Raman detuning [103].

³In principle this AOM can provide more precise continuous control of the amplitude of the deflected beams but, at present, it only functions as a switch. This switch is not instantaneous, however, and in reality there is a rise-time of approximately 100 ns before the output light reaches maximum intensity.




FIGURE 6.3: Simplified diagram of the Raman laser system used to generate two spatially separated Raman beams with a frequency difference matching the ground hyperfine splitting in ^{85}Rb . The light from a single 780 nm *Eagleyard* distributed feedback (DFB) diode laser is passed through *one of two* acousto-optic modulators (AOMs). The selected first-diffracted order (dashed red line) is passed through a non-polarising beamsplitter cube (BSC) and forms one component of the two Raman beams. The remaining component is formed by passing the undeflected light from the AOM through an electro-optic modulator (EOM). After the carrier frequency is removed from the EOM output using an optical Mach-Zehnder interferometer (MZI), both beams are amplified using tapered amplifiers (TA) and sent through single-mode (SM) fibres, recombined with a polarising beamsplitter cube (PBSC), and passed through another AOM which controls whether the Raman light enters the vacuum chamber or not. Both beams are then separated again with another PBSC and sent through polarisation maintaining (PM) fibres. They are then passed through refractive beam-shaping (BS) optics and directed into the MOT chamber.

6.2.1 Switching the effective wavevector direction

We can reverse the effective wavevector direction $\mathbf{k}_{\text{eff}}$ by using a radio-frequency switch to alternate between two AOMs both driven at the same frequency ω_A . The first AOM diffracts the incoming light at frequency ω_0 into the first positive order ($\omega_0 + \omega_A$) and the second AOM diffracts incoming light into the first negative order ($\omega_0 - \omega_A$). If the first positive order AOM is selected, a Raman pair is formed between $\omega_0 + \omega_A$ and the $n = -1$ sideband from the EOM output (as shown in Equation 6.3). This means the EOM-output beam becomes the lower frequency Raman beam (ω_2) and the AOM-shifted beam becomes the higher frequency beam (ω_1). In this situation the $n = +1$ EOM sideband is far off-resonance. However, if the first negative order AOM is selected, a Raman pair is formed between $\omega_0 - \omega_A$ and the $n = +1$ EOM sideband. This means the EOM-output beam becomes the higher frequency component (ω_1) and the AOM-output beam becomes the lower frequency beam (ω_2). The $n = -1$ EOM sideband is now far

off-resonance and the two Raman frequencies each enter the chamber from the opposite direction. Consequently, the atoms experience a reversal of effective wavevector.

6.2.2 Implementing optimal pulses

As explained in Chapter 5, our optimised pulses are described by piece-wise constant waveforms $\phi_L(t) = \phi_1, \phi_2, \dots, \phi_N$ where the relative laser phase between the two Raman beams takes a different value for each time-step in the pulse.

We achieve control of the relative phase between the two Raman beams by modulating the phase of the radio-frequency (RF) signal driving the EOM. The phase of this RF signal is controlled using a *Miteq SDM0104LC1CDQ* I&Q modulator, where the in-phase (I) and quadrature (Q) inputs are provided by the two outputs of a *Keysight 33612A* arbitrary waveform generator (AWG). Given an optimised pulse $\phi_L(t)$, the AWG is programmed to output waveforms $I = A_0 \sin \phi_L(t)$ and $Q = A_0 \cos \phi_L(t)$, where A_0 is the modulation amplitude. This leads to a phase shift in the signal driving the EOM⁴ of $\phi(t) = \arctan(I/Q)$. The phase profile is therefore applied to the phase of the Raman frequency component at the EOM output.

6.3 Experimental procedure

A typical experimental sequence, illustrated in Figure 6.4, may be summarised as follows:

1. **MOT:** The MOT B-field, cooling lasers, and repump beams are all switched on for approximately 1 second. Atoms are consequently trapped and cooled to a temperature between 80 and 230 μK [56]. A photo-multiplier tube (PMT) with a fixed lens system [56] focused on the trap centre captures a proportion of the photons associated with the decays from the excited states $|5^2P_{3/2}, F' = 4\rangle$ and $|5^2P_{3/2}, F' = 3\rangle$ that occur during the MOT cooling cycle.
2. **Molasses cooling:** In this stage, the magnetic quadrupole field is extinguished, and the intensity of the cooling light is lowered gradually. Atoms continue to cool in an optical molasses for approximately 10 milliseconds. Varying the intensity of the cooling light in this stage allows us to control the temperature of the cloud in the range 10-100 μK . Lower levels of cooling intensity during this stage result in more efficient sub-Doppler cooling and hence lower cloud temperatures. A full characterisation of this temperature control can be found in [103].

⁴The phase mapping provided by the I&Q modulator does not follow a perfect linear relationship [56], and in fact becomes less linear as the amplitude voltage A_0 is increased. In this work, we did not apply a correction to our optimised pulse sequences to account for this, but instead decreased the modulation amplitude until the mapping became approximately linear.




FIGURE 6.4: Diagram showing the sequence of experimental operations outlined in the text. The timings are not to scale, and the particular sequence of Raman pulses depends on the experiment performed.

3. **State preparation:** Atoms are prepared in the lower ground hyperfine state $|g\rangle = |5^2S_{1/2}, F = 2\rangle$ by applying the cooling beams for approximately 4 milliseconds. In the absence of the repump beam, atoms are eventually lost from the cooling cycle and decay to the state $|5^2S_{1/2}, F = 2\rangle$. The signal on the PMT decays exponentially as atoms are lost from the cooling cycle.
4. **Raman pulses:** Now we have a cold localised sample of ^{85}Rb all in the state $|g\rangle = |5^2S_{1/2}, F = 2\rangle$ and our experiments with Raman pulses can begin. During this stage of the experiment, the cooling and repump beams are switched off and the atom cloud expands and falls under gravity. Although typically this stage lasts 1 millisecond, we have a maximum interaction window of approximately 5 milliseconds before the atoms fall out of the 1.2 mm^2 region illuminated by the horizontal Raman beams.
5. **State read-out:** Following a Raman pulse sequence, we must determine what fraction of atoms has been transferred to the excited state $|e\rangle = |5^2S_{1/2}, F = 3\rangle$. First, the cooling beams are turned on for $\tau_{SR} \approx 500 \mu\text{s}$. This excites all atoms in the state $|5^2S_{1/2}, F = 3\rangle$ to the upper state $|5^2P_{3/2}, F' = 4\rangle$. The resulting fluorescence is captured using a photo-multiplier tube (PMT). We fit an exponential curve to the resulting decay in PMT signal, and determine the amplitude A_1 of this decay. In order to normalise this signal by the total number of atoms and thereby estimate the fraction of atoms transferred to the state $|5^2S_{1/2}, F = 3\rangle$ by the Raman pulse sequence, we apply the repump beam for $\tau_R \approx 100 \mu\text{s}$. This pumps all atoms into the state $|5^2S_{1/2}, F = 3\rangle$, and we perform another exponential fit following a second application of the cooling light for τ_{SR} . The amplitude




FIGURE 6.5: Illustration of the principles behind Raman velocimetry. (a) The simulated transfer efficiency of a long and low-power π pulse with Rabi frequency $\Omega_R = 20$ kHz and the velocity distribution for a 32 μK cloud of ^{85}Rb . (b) Data showing the measured fraction of atoms transferred to the excited state by a selective π pulse as the laser detuning is scanned. From this measured distribution we estimate a cloud temperature of approximately 32 μK .

of this decay, A_2 , is proportional to the total number of atoms. Therefore, the fraction of atoms transferred to the excited state by the Raman pulse sequence is given by A_1/A_2 .

6.4 Raman velocimetry

Counter-propagating Raman beams induce velocity-sensitive two-photon transitions between states $|g\rangle = |5^2S_{1/2}, F=2\rangle$ and $|e\rangle = |5^2S_{1/2}, F=3\rangle$. The two-photon detuning δ (Equation 2.26) depends on the component of the atomic velocity in the direction of the effective wavevector. Thus, as discussed in Chapter 2, the probability for a π pulse to transfer an atom from $|g\rangle$ to $|e\rangle$ depends on the atomic velocity. In Raman velocimetry [92], we exploit this sensitivity to obtain a measurement of the velocity distribution and hence temperature of the atomic cloud.

If the two-photon Rabi frequency Ω_R is lowered, the π pulse becomes more velocity-selective: a very long π pulse will transfer atoms to the state $|e\rangle$ only within a narrow range of velocities. By scanning the laser detuning δ_L and measuring the fraction of atoms transferred to the excited state, we can therefore reveal the velocity distribution of the atom cloud. The width of the measured velocity distribution is determined

by the temperature of the atom cloud along the Raman beam axis, and assuming the distribution is well-modelled by a Maxwell-Boltzmann distribution, we calculate the temperature T using the relation $T = m\sigma_v^2/k_B$, where σ_v is the width of a Gaussian fit to the measured distribution. Figure 6.5 illustrates Raman velocimetry in action. Figure 6.5a shows the simulated transfer efficiency for a long, low-power π pulse, highlighting its velocity selectivity, and Figure 6.5b shows velocimetry data taken in our experiment for a cloud with a calculated temperature of $\sim 32 \mu\text{K}$.

When using Raman velocimetry one isn't actually measuring the atomic velocity distribution directly, but rather the convolution of the π pulse excitation spectrum and the atomic velocity distribution. This convolution exhibits artefacts that come from the sinc-like lobes of the π pulse excitation, and although in principle the resolution of the measurement can be increased by using longer and hence more selective π pulses with lower beam intensities, the fraction of atoms transferred at each laser detuning in the scan is reduced, incurring a loss in the signal-to-noise ratio of the measurement.

We note that more advanced techniques such as interferometric velocimetry [2, 102, 103] can provide much cleaner and more accurate measurements of atomic velocity distributions and overcome the limitations present in Raman velocimetry.

Part IV

Results

Chapter 7

Robust mirror pulses for atom interferometry

Mirror, or inversion, pulses are crucial for atom interferometry. Every closed interferometer has at least one mirror to send the diverging paths back towards each other; LMT interferometers may have many mirror pulses; and we also use inversion pulses in the readout after the interferometer sequence has finished.

In each case, the mirror pulse fidelity affects the interferometer performance. We always care about the fraction of the atomic population transferred to the target state; but we may also care about the effect of the mirror pulse upon the phase of the superposition, and its dependence upon external contributions to the Lagrangian. We may know the initial state, or we may need pulses that work for a variety of initial states. Whereas mirror pulse always refers to inversion, its meaning can therefore be different according to the application. This means that, for optimisation, the choice of fidelity measure is important.

Most optimisations are trade-offs: a better compensation of one type of error might be possible if we do not also need to accommodate another error; and we might expect that having more control parameters at our disposal (more time-steps, longer pulse duration, amplitude as well as phase) would allow for higher terminal fidelities to be reached. The way that constraints are implemented can also change the optimisation landscape, possibly making it harder to find a good solution. We therefore expect the optimum solution to vary as the design parameters and constraints are varied. The dependence of the sensitivity and performance of the solution upon these parameters and constraints, and the nature of this dependence (smooth or discontinuous) are important in determining the generality of the solution.

In this chapter, we therefore review previous implementations of mirror pulses, and a range of applications in which they are used, mentioning the effects of past shortcomings

upon performance in these applications. We consider different fidelity measures for the mirror operation, and identify the constraints and optimisation parameters for typical applications. We then present the results of GRAPE optimisations, and identify some interesting features in the computational solutions. We compare our pulses with previous approaches; and we establish the dependence of the solutions and their predicted performance upon these parameters and constraints, and their sensitivity to the initial guess.

Finally, we demonstrate experimentally a pulse design that achieves a ground hyperfine state-transfer efficiency of 99.8(3)% in a 35 μ K thermal cloud of atomic rubidium, compared with a conventional π pulse efficiency of 75(3)%. This inversion pulse is highly robust to variations in laser intensity and detuning, with a fidelity that remains above 90% at detunings for which the π pulse fidelity is below 20%, and is thus suitable for large-momentum-transfer interferometers using thermal atoms or operating in non-ideal environments. Sections of this chapter have been published [1, 3] and the experimental data were taken in collaboration with Max Carey.

7.1 Mirror pulse fidelities and design objectives

The choice of a mirror pulse fidelity measure depends on the application. We consider below two examples in which mirror pulses play a different role in the interferometer sequence, and consider what fidelity measure is appropriate in each case.

First we examine the LMT interferometer, in which the beamsplitter and mirror operations are extended by multiple ‘augmentation’ pulses with alternating effective wavevectors that are designed to swap the population between the internal states in each interferometer arm whilst imparting additional momentum to the atoms [40, 42, 45]. Each successive augmentation pulse must efficiently transfer atomic population within interferometer arms that become increasingly separated in resonance frequency due to the momentum imparted by the successive (two-photon) recoils. To optimise an augmentation pulse for LMT interferometry, which we represent by the propagator $\hat{U}$, it may be sufficient to consider the point-to-point (PP) fidelity

$$\mathcal{F}_{\text{square}} \equiv |\langle e|\hat{U}|g\rangle|^2 \quad (7.1)$$

without concern for the relative phase introduced between the two states. This is because, in LMT interferometers, the augmentation pulses appear in pairs within the extended pulse sequence so that the interferometer phase introduced by each pulse is, to first order, cancelled out by that introduced by a subsequent one [42, 46, 50]. Choosing to optimise a PP operation as opposed to a universal rotation (UR) pulse effectively gives the pulse optimisation algorithm a larger target to shoot at, allowing impressive fidelity to be achieved with a modest pulse area.

However, in our second example, for the mirror pulse at the centre of the three-pulse Mach-Zehnder interferometer sequence, a PP pulse is expected to lead to a loss in interferometer contrast and sensitivity. This is because the central mirror pulse must be applied to atoms that lie ideally at a range of points on the equator of the Bloch sphere. The mirror pulse in this case must provide a π rotation about a fixed axis in the equatorial plane for atoms in the ensemble that naturally have different initial states. If the mirror pulse provides a rotation about different axes for different atoms, the atoms will exit the interferometer with a spread in phase. This will scramble the interferometric measurement and lead to fringes washing out when averaged over the entire atom cloud. One option, therefore, is to optimise the mirror pulse to be a universal rotation (UR) 180° pulse with a fixed rotation axis in the xy plane for all atoms with the range of Raman detunings and coupling strengths found in the ensemble. We therefore consider the following UR fidelity measure [10]

$$\mathcal{F}_{\text{UR}} = \frac{1}{2} \text{Tr}(\hat{U}_\pi^\dagger \hat{U}), \quad (7.2)$$

which is maximised when $\hat{U} = \hat{U}_\pi$, where $\hat{U}_\pi = \exp(-i\frac{\pi}{2}\hat{\sigma}_y)$ is the propagator for an ‘ideal’ π rotation about the y -axis of the Bloch sphere.

We can also simplify this UR fidelity in a useful way as follows:

$$\mathcal{F}_{\text{UR}} = \frac{1}{2} \text{Tr}(\hat{U}_\pi^\dagger \hat{U}) \quad (7.3)$$

$$= \frac{1}{2} \text{Tr} \left[\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}^\dagger \begin{pmatrix} \langle g|\hat{U}|g\rangle & \langle g|\hat{U}|e\rangle \\ \langle e|\hat{U}|g\rangle & \langle e|\hat{U}|e\rangle \end{pmatrix} \right] \quad (7.4)$$

$$= \frac{1}{2} (\langle e|\hat{U}|g\rangle - \langle g|\hat{U}|e\rangle) \quad (7.5)$$

$$= \text{Re}(\langle e|\hat{U}|g\rangle) \quad (7.6)$$

$$\equiv \mathcal{F}_{\text{real}}. \quad (7.7)$$

We therefore see that maximising the UR mirror pulse fidelity is equivalent to maximising the real part of the overlap $\langle e|\hat{U}|g\rangle$, which we denote $\mathcal{F}_{\text{real}}$. Maximising $\mathcal{F}_{\text{real}}$ will therefore maximise the state-transfer efficiency $\mathcal{F}_{\text{square}}$ and fix the phase of the overlap $\langle e|\hat{U}|g\rangle$. These two conditions must be met by any mirror pulse because the contrast of the fringes depends on the state-transfer efficiency (Equation 3.13) and the contribution of the mirror pulse to the interferometer phase is given by $2\phi(\langle e|\hat{U}|g\rangle)$ (Equation 3.14). Therefore, maximising $\mathcal{F}_{\text{real}}$ over an ensemble of off-resonance and pulse-length errors both minimises variations in the interferometer phase over that ensemble and maximises the visibility of the interference fringes.

We note here that $\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$ are only two choices of mirror pulse fidelities out of many possibilities. For example, in the Mach-Zehnder interferometer the mirror pulse is expected to operate on atoms whose states are distributed close to the equatorial plane

of the Bloch sphere (assuming the initial beamsplitter pulse is reasonably efficient), and therefore only certain initial states need be rotated about a fixed axis. The mirror pulse therefore need not strictly be a universal π rotation for all initial states - only for a specified initial distribution of states. Such a ‘not-quite’ UR pulse may be easier to optimise than a full UR designed to work for any and all initial states of a two-level atom. Also, the fidelity $\mathcal{F}_{\text{real}}$ treats phase uniformity and state-transfer efficiency on an equal footing, and it may be advantageous to uncouple these two properties such that they may be weighted differently in the optimisation: it may be more advantageous in an interferometric measurement application to prioritise phase uniformity over state-transfer efficiency, for example.

As discussed in Chapter 5, we construct our control pulses as piece-wise constant waveforms where the combined Raman laser phase ϕ_L (and potentially amplitude or Rabi frequency Ω_R) is varied for each step in the pulse. The goal of our optimisation procedure is to find the sequence of laser phases that maximises our chosen inversion fidelity ($\mathcal{F}_{\text{square}}$ or $\mathcal{F}_{\text{real}}$) over the entire atomic cloud.

We build robustness into our optimal control designs by replacing our single-atom inversion fidelity by an average over the detuning and amplitude errors present in the atomic sample,

$$\mathcal{F} \rightarrow \sum_{\text{ensemble}} \mathcal{F}. \quad (7.8)$$

The ensemble of off-resonance and pulse-length errors supplied to the optimisation must reflect those present in the experimental atom sample. The dominant contribution to off-resonance errors in our experiment comes from the temperature of our atom cloud, which is typically in the range of (10 - 200) μK [56]. For these temperatures, the spread in atomic velocity leads to a spread in Raman detuning that can be many hundreds of kHz and comparable in size to the Rabi frequency - this can thus lead to significant reductions in the fidelity of π and $\pi/2$ pulses in our interferometer. In LMT applications, however, the dominant contribution to off-resonance errors comes from the increased momentum separation of the diffracting wave-packets.

Similarly, the ensemble of pulse-length errors fed into the optimisation should ideally reflect the distribution of such errors found in the atom sample. Sources of pulse-length errors for our system include spatial inhomogeneity in the Raman beam intensity profiles, and the occupation of different m_F sub-levels which each have a different two-photon Rabi frequency. Rather than optimising pulses for a system-specific distribution of coupling strengths, in this chapter we focus on designing pulses that can withstand variation of the two-photon Rabi frequency within an experimentally feasible range of the nominal field strength of $\pm 10\% \Omega_{\text{eff}}$, where Ω_{eff} is the nominal or effective Rabi frequency.

In this chapter we focus on designing mirror pulses using optimal control that mitigate off-resonance errors that arise from the temperature of a cloud of atomic ^{85}Rb . In this case the distribution of detuning, $f(\delta)$, comes from the distribution of velocities along the Raman beam axis $f(v)$ of an atomic cloud at a given temperature T - which we take to be a Maxwell-Boltzmann or Gaussian distribution:

$$f(\delta) = k_{\text{eff}} f(v) \quad (7.9)$$

$$f(v) = \frac{1}{\sigma_v \sqrt{2\pi}} e^{-\frac{v^2}{2\sigma_v^2}}. \quad (7.10)$$

Here, $\sigma_v = \sqrt{k_B T / m}$ is the width or standard deviation of the velocity distribution where m is the atomic mass, v is the atomic velocity component parallel to the Raman beams, and k_B is the Boltzmann constant. We define our ensemble of off-resonance errors to be a sample of detunings in the range $(\pm 3k_{\text{eff}}\sigma_v)$, where each ensemble member's contribution to the final fidelity is weighted in the optimisation according to Equation 7.10. This ensures that the optimisation prioritises maximising the fidelity for the region in frequency space where most of the atoms in the ensemble are expected to be.

The number of detunings used in the ensemble of off-resonance errors and the number of amplitude errors used in the ensemble of pulse-length errors is limited by the available computational resources because the single-atom fidelity must be averaged over both distributions during each iteration of the optimisation. In practice, therefore, we only optimise pulses to compensate error distributions that are approximations to the true distribution.

7.2 Mirror pulse performance

Our approach is to optimise inversion pulses using our two mirror fidelities ($\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$) in the presence of typical off-resonance and pulse-length errors found in a thermal sample of ^{85}Rb , before comparing the resulting performance with that achieved using composite pulses and adiabatic rapid passage schemes. In general, given a fixed available Raman beam power, we seek to find the shortest mirror pulse that reaches the highest inversion fidelity in our atom cloud. This is important to both reduce the effect of spontaneous emission and minimise effects caused by atomic motion through non-uniform fields during the pulse. We are therefore interested in how the maximum achievable fidelity varies with the duration of the pulse and also how the performance compares with existing composite and ARP pulses.

We have optimised phase-modulated inversion pulses using both the $\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$ fidelities for a range of pulse durations from $1t_\pi$ to $8t_\pi$ varied in increments of t_π , where t_π is the duration of a rectangular π pulse. We have repeated each individual optimisation 50 times to examine the dependence of the final pulse on the initial guess which is taken

to be a random series of laser phases¹: our gradient-based optimisation algorithm has no guarantee of finding a global maximum of the fidelity and we therefore expect that multiple attempts will be required to obtain a good pulse. Each pulse was optimised using the L-BFGS GRAPE method (described in Chapter 5) to mitigate off-resonance errors arising from a thermal sample of ^{85}Rb at a typical experimental temperature of $40\text{ }\mu\text{K}$ and a range of pulse-length errors of $\pm 10\% \Omega_{\text{eff}}$. The pulse time-step was fixed to be $t_\pi/50$, and the nominal Rabi frequency, Ω_{eff} , was set to 200 kHz in line with previous experimental values [45]. Each optimisation was set to terminate after 1000 iterations, or when the norm of the gradient fell below 10^{-6} , or the norm of the step size fell below 10^{-3} .

Figure 7.1 shows histograms corresponding to the distributions of terminal fidelities found for both the UR and PP fidelity optimisation. We see that, starting from random phase profiles, the UR pulses converge in general to a smaller number of solutions than in the PP case and that these solutions have lower terminal fidelities. With the exception of the shorter pulses with durations less than $2t_\pi$ where each fidelity measure always converged to a given solution, the PP fidelity distributions are broader implying that it is worth trying multiple guesses to find the best pulse for a given set of parameters. The maximum and average terminal fidelity increases as the pulse duration and available pulse area is increased.

A selection of the best pulses found for each fidelity choice is depicted in Figure 7.2. Even in the case of random initial guesses for the phase-profiles and no constraint on the smoothness of the resulting waveforms, we typically find smooth pulses appearing for each fidelity choice. Furthermore, different classes of $\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$ mirror pulses emerge from the optimisation depending on the overall sequence length: some of these pulses are symmetrical about the temporal midpoint while some are highly asymmetric.

We find that a specific class of high fidelity $\mathcal{F}_{\text{square}}$ pulses may be found consistently by GRAPE for each duration when starting from a parabolic guess for the phase profile². Although the initial guess was a parabolic phase profile for these pulses, the resulting shape is not well-fitted by a parabolic approximation. These pulses belong to a particular class of non-linear frequency sweeps redolent of ARP pulses, a feature also found in the NMR optimal control literature [8, 159]. Despite their appearance, these non-linear frequency sweeps are not of high adiabaticity. Figure 7.3 shows the action of the non-linear frequency swept pulse from Figure 7.2g on the Bloch sphere. Even when detuned by a considerable amount ($\delta = 1.5\Omega_{\text{eff}}$), the final state ends up much closer to the intended target state than a simple π pulse, thereby demonstrating a considerable resilience to off-resonance errors.

¹Initial random phase profiles are obtained using MATLAB's *randn* function, which is used to generate normally distributed random numbers for each time-slice in a pulse.

²We define this parabolic guess by $\phi_{\text{guess}}(t) = -2\pi \times \left(\frac{(t-\tau/2)}{\tau/2}\right)^2$, where τ is the pulse duration.




FIGURE 7.1: Histograms showing the distribution of terminal fidelities for a selection of pulse durations ($2t_\pi$, $4t_\pi$, $6t_\pi$, and $8t_\pi$) where the initial guess for the waveforms was a random series of phases. With the exception of shorter pulses where the duration is less than $2t_\pi$, the PP fidelity distribution is broader than the UR one, meaning the UR optimisations in general converge to a smaller number of solutions. The red vertical lines indicate the terminal $\mathcal{F}_{\text{square}}$ fidelity that may be obtained when starting from a parabolic guess for the phase profiles. Each pulse was optimised for a $40 \mu\text{K}$ ^{85}Rb cloud with pulse-length errors in the range $\pm 10\%$ Ω_{eff} . The nominal Rabi frequency was 200 kHz .

7.2.1 Robust state-transfer

To compare the state-transfer efficiency of the GRAPE pulses with ARP and composite pulse alternatives, we took the GRAPE-optimised pulses optimised for each fidelity measure and duration choice and calculated the expected fraction of atoms transferred to the excited state, $\langle P_e \rangle$, in a $40 \mu\text{K}$ cloud with pulse-length errors within the design range of $\pm 10\%$. This involved drawing a sample of atomic velocities (and hence detunings) from a Maxwell-Boltzmann distribution corresponding to a temperature of $40 \mu\text{K}$ and a sample of Rabi rates from a uniform distribution with a range of 20% the nominal Rabi rate of 200 kHz . The final fraction of atoms in the excited state after application of a




FIGURE 7.2: Selection of the highest fidelity GRAPE-optimised phase-modulated inversion pulse waveforms found for a range of pulse durations optimising fidelities $\mathcal{F}_{\text{square}}$ ((a) to (h)) and $\mathcal{F}_{\text{real}}$ ((i) to (l)) for a 40 μK ^{85}Rb cloud with pulse-length errors in the range $\pm 10\%$ Ω_{eff} . Pulses (e) to (h) represent the terminal $\mathcal{F}_{\text{square}}$ waveforms starting from a parabolic guess for the phase profiles. The remaining pulses represent the best pulse found from 50 individual optimisations at each duration starting from a random initial phase profile. The terminal fidelity is shown above each panel.

given pulse to an atom in the ground state was evaluated numerically for each detuning and amplitude in the sample, and the average transfer efficiency (final fraction of atoms in the excited state) computed³. The results for the best GRAPE pulses for each fidelity measure are shown in Figure 7.4 along with those for a selection of the best-performing ARP and composite pulses taken from [42, 44–46].

The particular selection of composite pulses we have considered includes: WALTZ [43], Knill [160], CORPSE [61], Levitt [5], 90-225-315 [161], and 90-240-90 [162]. In agreement with previous studies [42, 45], we find that the WALTZ pulse obtains the highest fidelity out of the composite pulses tested. Alongside the selection of composite pulses, Figure 7.4 shows the simulated transfer efficiency for the well-known ARP tanh/tanh pulse [51] with the same sweep parameters as those used by Kotru *et al.* [46]. The tanh/tanh ARP pulse is a particularly efficient frequency-swept pulse that is able to maintain

³More details on the simulations can be found in Appendix B.




FIGURE 7.3: Bloch sphere trajectories for the $\mathcal{F}_{\text{square}}$ (PP) pulse of duration $6t_\pi$ taken from Figure 7.2g acting on (a) a resonant atom initially in the state $|g\rangle$ and (b) a detuned atom starting in the same initial state. (b) also depicts the state trajectory for the standard π pulse (red dotted curve) under the same detuning of $\delta = 1.5\Omega_{\text{eff}}$.




FIGURE 7.4: Simulated fraction of excited atoms following different inversion pulses with durations ranging from t_π to $8t_\pi$ in a $40\ \mu\text{K}$ sample with a range of pulse-length errors $\pm 10\%$ Ω_{eff} . The nominal Rabi frequency, Ω_{eff} , was 200 kHz. The composite pulses were taken from [45] and the tanh/tan ARP pulse from [44]. Phase-modulated GRAPE inversion pulses found using fidelity $\mathcal{F}_{\text{square}}$ obtain the highest ensemble fidelity at all pulse durations.

100% inversion efficiency on resonance at considerably shorter durations than alternative ARP pulses [44].

We find that at all pulse durations explored GRAPE $\mathcal{F}_{\text{square}}$ pulses reach a higher average inversion fidelity compared with the selected ARP and composite pulses. The largest improvement shown by GRAPE is at durations less than $4t_\pi$ where t_π is the duration of a π pulse; the ARP tanh/tan pulse becomes more robust to both off-resonance and pulse-length errors as it becomes longer and thereby more adiabatic [44]. The $\mathcal{F}_{\text{real}}$ fidelity is

more restrictive than the PP $\mathcal{F}_{\text{square}}$ fidelity (restricting variations in the phase of the final state across the ensemble in addition to purely maximising state-transfer) and the UR pulses must therefore be longer to reach similar average state-transfer efficiency to the PP pulses.

The results in Figure 7.4 conflate resilience to off-resonance and pulse-length errors by averaging over an ensemble of both error sources. We can instead simulate how the state-transfer efficiency of a given pulse varies as a function of off-resonance and pulse-length errors to examine the robustness to both error classes. The results are depicted using contour plots in Figure 7.5 where we have shown the results for a non-linear frequency swept GRAPE pulse of duration $6t_\pi$ (derived from a parabolic guess for the phase profile with terminal waveform depicted in Figure 7.2g) and also depict results for the WALTZ composite pulse and the ARP tanh/tan pulse (the duration of ARP was fixed to be $6t_\pi$ to ensure a fair comparison with GRAPE). The red hatched regions show the region of optimisation of both error sources ($\pm 3k_{\text{eff}}\sigma_v$ for δ at $40\ \mu\text{K}$ and a 20% offset range for amplitude). The GRAPE solution proves to be significantly more robust to off-resonance errors than the WALTZ and ARP pulses within the ensemble range and, interestingly, we find that the GRAPE solution can withstand a significantly larger range of pulse-length errors than it is designed to compensate for.

We conclude that phase-modulated GRAPE pulses can yield significantly more robust state-transfer in a two-level atom than may be achieved with the best-performing ARP and composite pulses that have been used in atom interferometry to-date. Furthermore, the increased robustness to off-resonance errors afforded by GRAPE $\mathcal{F}_{\text{square}}$ pulses makes them an attractive option for use as augmentation pulses within LMT interferometers. Using such GRAPE augmentation pulses should enable larger interferometer areas and hence greater intrinsic sensitivity gains than can be achieved with conventional π pulses, composite pulses, or ARP pulses.

7.2.2 Interferometric phase variation and contrast

What is the effect on the interferometer phase when replacing the π pulse in a three-pulse interferometer with an optimised alternative? Previous experimental studies of frequency-swept Raman ARP and composite inversion pulses in atom interferometry have focused on the problem of providing robust state-transfer [41, 42, 45, 46], and not on the effect of a pulse on the phase of the interferometer when employed as the central mirror operation within a sequence. We expect that mirror pulses found by optimising $\mathcal{F}_{\text{real}}$ will minimise variation in the interferometer phase for an ensemble of atoms with a distribution of off-resonance and pulse-length errors. This will be crucial if these pulses are to find application in interferometric sensors because any variation in the interferometer phase across an atomic sample will result in the fringes washing out when the output is averaged over the atomic cloud at the end of the interferometer.




FIGURE 7.5: State-transfer efficiency ($\mathcal{F}_{\text{square}}$) of different pulses as a function of off-resonance and pulse-length errors. The hatched region represents the region of optimisation and the contours correspond to fidelities of 0.8, 0.9, and 0.95. The GRAPE pulse is taken from Figure 7.2g.

Using Equation 3.14 from Chapter 3, we can simulate the dependence of the interferometer phase on the Raman detuning for the $\pi/2 - \pi - \pi/2$ sequence where the central π pulse is replaced by a robust alternative such as an ARP or composite pulse but the initial and final $\pi/2$ pulses remain as fractional Rabi oscillations. The results are shown in Figure 7.6. We see that the GRAPE $\mathcal{F}_{\text{real}}$ pulse of duration $8t_\pi$ (taken from Figure 7.2) minimises variation in the interferometer phase with the Raman detuning δ but the $\mathcal{F}_{\text{square}}$ pulse of identical duration introduces a large unwanted variation in phase with detuning much like the ARP pulse. This means that although the PP $\mathcal{F}_{\text{square}}$ pulses can provide robust state-transfer between states $|g\rangle$ and $|e\rangle$, when applied within a three-pulse sequence different atoms (with different velocities) will exit the interferometer with different phases. The same is true of the PP WALTZ pulse and the ARP schemes. The phase shift following the conventional $\pi/2 - \pi - \pi/2$ sequence is Doppler-insensitive providing the Rabi rate does not vary from pulse to pulse [109].




FIGURE 7.6: Variation of single-atom interferometer phase $\Delta\phi$ with Raman detuning simulated for different sequences where the central mirror pulse is replaced by a robust alternative but the initial and final $\pi/2$ pulses remain as rectangular fractional Rabi oscillations. We have applied a uniform phase shift to the $\mathcal{F}_{\text{real}}$, $\mathcal{F}_{\text{square}}$, and ARP curves to aid visual clarity.

If the interferometer phase varies with the Raman detuning, this implies that in an ensemble of atoms with a velocity spread, different atoms will exit the interferometer with different phases [108]. This can lead to a reduction in fringe visibility when the output signal is averaged over the entire atom cloud. We can simulate the thermally-averaged interferometer signal for a given pulse sequence and atomic cloud temperature by integrating the final excited state probability P_e (Equation 3.11) over the atomic velocity distribution corresponding to a specific temperature [108]. We can then generate a simulated fringe by varying the phase of the final beamsplitter, which we take to be a conventional $\pi/2$ pulse. We then find the average contrast by fitting a sinusoidal function to the thermally-averaged output fringe.

The results of such a simulation are depicted in Figure 7.7 for atomic cloud temperatures for ^{85}Rb in the range 1 to 150 μK . As expected, the $\mathcal{F}_{\text{real}}$ GRAPE pulse leads to a significant improvement in the ensemble-averaged contrast as the temperature and hence velocity spread of the cloud is increased. This is because $\mathcal{F}_{\text{real}}$ pulses minimise the dependence of the interferometer phase on the detuning, allowing for a robust central mirror pulse. The black dotted curve in Figure 7.7 corresponds to a sequence with a ‘perfect’ π pulse that is resonant for all atoms in the ensemble - thus illustrating the decay in contrast due to the Doppler sensitivity of the rectangular beamsplitters alone and providing a target performance when optimising the mirror operation in isolation. GRAPE $\mathcal{F}_{\text{real}}$ pulses approach the ‘perfect’ mirror pulse as their duration and fidelity increase.




FIGURE 7.7: Simulated interferometer contrast $\mathcal{B}$ for different pulse sequences as a function of atomic cloud temperature. In each case the mirror pulse is substituted by either a GRAPE pulse, a composite inversion pulse, or the ARP tanh/tan pulse. The GRAPE $\mathcal{F}_{\text{square}}$ pulse was the $8t_\pi$ pulse depicted in Figure 7.2c, and the $8t_\pi$ GRAPE $\mathcal{F}_{\text{real}}$ pulse was also taken from Figure 7.2. The ARP pulse duration was $8t_\pi$ to ensure a fair comparison with GRAPE alternatives. The black dotted curve corresponds to a sequence with a “perfect” π pulse that is resonant for all atoms and the associated decay highlighted by the shaded region is due entirely to the rectangular $\pi/2$ pulses.

The nominal Rabi frequency was 200 kHz

Table 7.1 summarises the results for the thermally-integrated interferometer contrast at three temperatures of 30, 60, and 80 μK for a selection of composite and ARP pulses. The only composite pulses that improve the contrast relative to the conventional sequence of $\pi/2$ and π pulses are the Knill and ‘90-240-90’ sequences. This can be explained by recalling that these sequences were designed as UR pulses [56], and as a result the phase of the overlap $\langle e|\hat{U}|g\rangle$ is uniform near resonance.

Although substituting the π pulse in the $\pi/2 - \pi - \pi/2$ sequence with a PP composite, optimised $\mathcal{F}_{\text{square}}$, or adiabatic pulse leads to different atoms exiting the interferometer with different phases, it may be possible to cancel the detuning-dependent phase errors by using an alternative pulse sequence with a pair of inversion pulses, much like in LMT interferometer sequences that employ additional mirror pulses in pairs to increase the interferometer area. For example, using a $\pi/2 - \pi - \pi - \pi/2$ sequence [159] where the first mirror pulse occurs immediately after the first beamsplitter and is used only to cancel the phase shift of the second mirror pulse. This would only work if the Rabi frequency for a given atom did not vary much during the first dwell-time, otherwise the phase compensation would be imperfect. Indeed this is also a problem when adiabatic pulses are used in LMT interferometers [46], and we explore this aspect of LMT interferometers in greater detail in Chapter 9.

TABLE 7.1: Simulated contrast values for different interferometer sequences at temperatures of 20, 40, and 80 μK . In each sequence the beamsplitters were $\pi/2$ pulses and the mirror pulse was replaced by either: a GRAPE inversion pulse, a composite pulse, or an ARP pulse. The “perfect” π pulse sequence contains a π pulse that is fixed to be resonant with every atom in the velocity distribution thus giving an indication of the decay in contrast with temperature that can be attributed to the imperfections in the beamsplitters only. Bold values indicate the best-performing pulse at each temperature.

Inversion pulse	Type	Length (t/t_π)	Interferometer contrast		
			20 μK	40 μK	80 μK
π	UR	1	0.76	0.64	0.50
“Perfect” π	UR	1	0.99	0.97	0.91
WALTZ [43]	PP	3	0.27	0.22	0.20
Knill [160]	UR	5	0.88	0.75	0.58
Levitt [5]	PP	2	0.75	0.64	0.52
90-225-315 [161]	PP	3.5	0.67	0.59	0.48
90-240-90 [162]	UR	2.3	0.80	0.68	0.56
CORPSE [61]	UR	4.3	0.73	0.56	0.42
ARP tanh/tan [51]	PP	6	0.60	0.44	0.31
		8	0.53	0.38	0.27
GRAPE $\mathcal{F}_{\text{square}}$	PP	6	0.46	0.33	0.22
		8	0.33	0.22	0.15
GRAPE $\mathcal{F}_{\text{real}}$	UR	6	0.95	0.91	0.82
		8	0.97	0.93	0.85

7.3 Dependence on optimisation parameters

We find that optimised Raman pulses yield significant improvements in simulated fidelity, phase uniformity, and contrast when compared with composite and adiabatic pulse alternatives. However, it is also useful for practical users and for fundamental interest to know how sensitive the optimisation is to the choices made in performing it. In this section we therefore present the results of a preliminary investigation into the dependence of GRAPE-optimised solutions on the parameters used in the optimisation and highlight considerations that are important when designing robust mirror pulses using optimal control. We explore the dependence of the solutions on the number of time-steps used to form the piece-wise constant waveforms and how the solutions depend on the ensemble of off-resonance and pulse-length errors used.

7.3.1 Number of time-steps

How does the terminal fidelity reached by GRAPE depend on the number of time-steps N used to define the piece-wise constant waveforms? Do different solutions emerge as the number of time-steps is increased? To answer these questions we have optimised mirror pulses for a 40 μK cloud with pulse-length errors in the range $\pm 10\%$ Ω_{eff} using fidelities $\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$, varying the number of time-steps used to form the waveforms from 10 to 300 in increments of 10. The total duration of the pulses was fixed to be $4t_\pi$ and the optimisation was repeated for each time-step choice 12 times where in each case




FIGURE 7.8: Terminal fidelity for GRAPE inversion pulses of duration $4t_\pi$ as the number of time-steps N is varied from 10 to 300 in increments of 10. Each point corresponds to the best pulse found out of 12 optimisations of each mirror fidelity measure starting from random guesses. No significant increase is observed beyond 30 time-steps for each fidelity choice.

the initial guess was a different random series of laser phases. The maximum terminal fidelities reached by GRAPE for each fidelity choice and time-step are displayed in Figure 7.8.

Figure 7.8 shows that beyond approximately 30 time-steps, there was no significant increase in terminal fidelity obtained from further increasing the number of time-steps. $\mathcal{F}_{\text{square}}$ pulses exhibit what look like oscillations in the maximum fidelities indicating that some time-step periodicities may be better than others although since each optimisation was repeated only 12 times this could be due to choosing a poorer sample of initial guesses for certain values of N . The $\mathcal{F}_{\text{real}}$ pulses all converge to the same shape beyond 30 time-steps, and the terminal fidelity reached with 300 steps is only 0.003 greater than that found using 30.

We observe significant stability in the shape of the optimal solutions as the number of time-steps is varied. For example, Figure 7.9 depicts terminal $\mathcal{F}_{\text{square}}$ waveforms optimised from the same parabolic initial guess as the number of time-steps, N , is varied. We observe a smooth convergence to the non-linear frequency swept pulses discussed in Section 7.2 and there are no sudden changes in fidelity or pulse shape as N is increased. Sometimes, however, different shapes can be found for different numbers of time-steps starting from the same guess. This is illustrated in Figures 7.10 and 7.11, where we depict colourmap plots of terminal waveforms found optimising our two mirror pulse fidelities starting this time from random initial guesses. For the $\mathcal{F}_{\text{real}}$ pulses shown in Figure 7.10, we find that the majority of time-step choices lead to a symmetrical pulse shape, but for a few specific values of N , a slightly asymmetrical solution is obtained with similar fidelity to the symmetrical case. In the case of the $\mathcal{F}_{\text{square}}$ pulses from Figure 7.11 some solutions that appear to be different shapes are actually related by a symmetry transformation (e.g. Figure 7.11d), while others yield almost identical fidelities yet appear to have unrelated shapes (e.g. Figure 7.11c).




FIGURE 7.9: (a) Colour-map depicting how the terminal phase profile for $\mathcal{F}_{\text{square}}$ pulse optimisation varies with the number of time-steps, N . Each pulse was optimised starting from a parabolic guess for a cloud of temperature $40 \mu\text{K}$ with pulse-length errors in the range $\pm 10\% \Omega_{\text{eff}}$. The terminal fidelities for each time-step are shown in panel (b). Panels (c) and (d) show the shape of the solutions at $N = 8, 12, 44$, and 64 . The waveforms converge to a smooth shape and no new features are found in the solutions as N is increased.

7.3.2 Ensemble size and specification

7.3.2.1 Ensemble density

The ensemble of off-resonance errors used to optimise a robust solution must be carefully constructed. If the ensemble consists of too few individual detunings within the desired range of error compensation, the optimal control algorithm may obtain a superficially high-fidelity solution that, upon closer inspection, only obtains high-fidelity for those specific detunings included in the ensemble and not for the region in frequency between them. An example of this is depicted in Figure 7.12. Using an ensemble of only 41 detunings to represent the distribution in detuning arising from a $160 \mu\text{K}$ atom cloud,




FIGURE 7.10: (a) Colour-map depicting how the terminal phase profile for $\mathcal{F}_{\text{real}}$ pulse optimisation varies with the number of time-steps, N . Each pulse was optimised starting from the same random guess for a cloud of temperature $40 \mu\text{K}$ with pulse-length errors in the range $\pm 10\% \Omega_{\text{eff}}$. The terminal fidelities for each time-step are shown in panel (b). Panel (c) shows the solutions for $N = 16$ and $N = 20$ and panel (d) shows the solutions at $N = 24$ and $N = 28$. The solutions for $N = 16$ and $N = 20$ have an asymmetry which is not apparent in the other solutions, which all resemble smoother versions of the symmetrical $N = 24$ and $N = 28$ waveforms. Despite the different shapes of the solutions in (c) they have a similar fidelity to the symmetrical shapes.

we obtained a mirror pulse that introduced high-frequency variations in the transfer efficiency. The peaks in fidelity occur at exactly the specific detunings of the ensemble and are clearly sub-optimal elsewhere. Re-running this particular optimisation with twice as many detunings in the sample resulted in a smoother pulse waveform that maintains high fidelity for the entire range of detuning in the ensemble. It is also possible that this problem could be mitigated by applying a smoothing penalty in the optimisation.




FIGURE 7.11: (a) Colour-map depicting how the terminal phase profile for $\mathcal{F}_{\text{square}}$ pulse optimisation varies with the number of time-steps, N , when the initial guess is a random phase profile. Each pulse was optimised for a cloud of temperature $40 \mu\text{K}$ with pulse-length errors in the range $\pm 10\% \Omega_{\text{eff}}$. The terminal fidelities for each time-step are shown in panel (b). Panel (c) shows the solutions for $N = 52$ and $N = 56$ and panel (d) shows the solutions at $N = 60$ and $N = 64$. While the pair of solutions in (c) have near-identical fidelities they appear to have fundamentally different shapes. The pair of solutions depicted in panel (d) appear to be different solutions but are in fact almost identical and related by a symmetry transformation: the solution for $N = 64$ is approximately the time-reversed and inverted form of the $N = 60$ solution with an additional phase offset. Such a symmetry transformation leaves the ensemble fidelity $\mathcal{F}_{\text{square}}$ of the pulse unchanged.

7.3.2.2 Ensemble distribution

Throughout this chapter we have optimised pulses using a Gaussian distribution of detunings whereby each offset in the sample is weighted according to the distribution arising from a specific atomic temperature (described by Equation 7.10). However, we could alternatively choose a uniform distribution that encompasses the majority of the atomic velocity distribution. Figure 7.13 shows the waveforms and transfer efficiencies




FIGURE 7.12: (a) Waveforms for $\mathcal{F}_{\text{square}}$ pulses optimised for an ensemble of atoms at a temperature of 160 μK . The blue solid line represents a pulse optimised for an ensemble consisting of 41 detunings and the black dotted curve represents an optimisation of the same fidelity but where the ensemble consisted of 81 detunings. (b) and (c) show the simulated transfer efficiency as a function of detuning for each pulse shown in (a) along with the specific detunings included in each ensemble. The shaded blue regions correspond to the distribution of detuning arising from a cloud of atoms at a temperature of 160 μK .

as a function of detuning for two pulses found optimising $\mathcal{F}_{\text{square}}$ for a 40 μK cloud. For one optimisation we used a linear ensemble with detunings sampled within the




FIGURE 7.13: Phase profiles (a) and transfer efficiencies (b) of two PP state-transfer pulses optimised using GRAPE for a $40 \mu\text{K}$ cloud of ^{85}Rb . The dashed orange curve represents a pulse found optimising $\mathcal{F}_{\text{square}}$ for a uniform distribution of detunings spanning $6 k_{\text{eff}}\sigma_v$ of the true detuning distribution (represented by the blue shaded region). The solid blue curve represents an optimisation from the same initial guess but for a Gaussian-weighted ensemble. This Gaussian-weighted pulse prioritises the transfer efficiency on resonance, and will therefore maintain efficiency even if the temperature of the distribution is lower than expected.

$6 k_{\text{eff}}\sigma_v$ range of the true distribution and for the second pulse we used a Gaussian-weighted sample spanning the same range. The initial guess was the same in both cases. We see from Figure 7.13 that the Gaussian-weighted distribution is clearly the better ensemble choice for this pulse duration. This is because the linear-ensemble choice leads to a pulse that sacrifices the transfer efficiency for resonant atoms to increase the fidelity of atoms further from resonance. The Gaussian-weighted pulse will clearly also maintain high fidelity if the temperature of the distribution is lower than expected. We therefore expect that pulses designed using Gaussian-weighted ensembles of detuning will provide an improvement in inversion performance in a thermal atom cloud, especially for short pulse durations where it is harder to provide near-perfect transfer for the entire range of resonant offsets within the atomic ensemble. We note in passing that linear ensembles of resonant offsets are used in many NMR optimal control studies of robust pulse design [8, 10] and, while we sometimes obtain similar waveforms, the ‘best’ pulse in the atomic context will therefore likely differ.

7.4 Amplitude modulation

In our pulse design we have the freedom to allow both the amplitude $\Omega_R(t)$ and laser phase $\phi_L(t)$ to vary with each time-step. Allowing amplitude modulation forces us to impose a constraint on the maximum amplitude during a pulse. If no such constraint




FIGURE 7.14: Panels (a) and (b) represent the phase and amplitude profiles of two pulses optimised using GRAPE. The dotted blue curve represents a phase-only optimisation and the solid red curve represents the results of an optimisation where both phase and amplitude were allowed to vary in each time-step with no limit on the maximum amplitude. Panel (c) shows the simulated transfer efficiency as a function of the Raman detuning for each pulse including a re-scaled version of the amplitude-modulated pulse (purple dot-dashed curve). The blue shaded region in (c) represents the velocity distribution due to a 40 μK atom cloud.

is applied, the algorithm will increase the pulse amplitude to become robust to off-resonance errors via power broadening [56]: the higher the Rabi frequency, the less the atomic state trajectory is deflected by a given detuning. To illustrate this problem, we repeated the 40 μK inversion pulse optimisation using fidelity $\mathcal{F}_{\text{square}}$ for a fixed duration of $4t_\pi$ from Section 7.2 using an initial guess of a random phase profile and a rectangular amplitude envelope but this time allowed $\Omega_R(t)$ to vary. As expected, and after only 100 iterations, GRAPE was able to reach a fidelity of .999 (the peak fidelity obtained in Section 7.2 for the same temperature and duration with no amplitude modulation was 0.981) by increasing the amplitude and hence available pulse area. Figure 7.14 shows the resulting phase and amplitude profiles in comparison to the best phase-only pulse of the same duration from Section 7.2. The robustness to off-resonance errors is significantly greater than the constant-amplitude pulse. However, the maximum amplitude is approximately three times that of the original pulse! If we re-scale the maximum amplitude back to 1 (corresponding to the nominal Rabi frequency Ω_{eff}) and scale the duration of each time-step accordingly to ensure the same rotation on resonance, we see a significant reduction in bandwidth. This simple example demonstrates the necessity of limiting the maximum amplitude when using optimal control.

In practice there is a maximum attainable laser power and Rabi frequency for a given




FIGURE 7.15: Initial and final waveform for a $\mathcal{F}_{\text{square}}$ pulse optimised for a temperature of $40 \mu\text{K}$ and a range of pulse-length errors of $\pm 10\% \Omega_{\text{eff}}$. Both the laser phase (a) and amplitude (b) were allowed to vary and the SNS penalty was imposed to keep the Rabi frequency from exceeding Ω_{eff} .

experimental arrangement. Furthermore, reducing peak laser power requirements is a necessary step in the development of mobile quantum sensors. We therefore seek to make a fair comparison of pulse performance by fixing the maximum amplitude of each pulse to be identical. There are different ways of imposing penalties on the pulse amplitude $\Omega_R(t)$. As explained in Chapter 5, we could minimise the amplitude during the pulse by adding to the fidelity a penalty proportional to the sum of the amplitudes for each time-step. This approach is often used in NMR radio-frequency pulse optimisation where it is advantageous to minimise the power-requirements of a pulse [8–10, 21, 23, 126, 163]. However, applying this type of penalty can be problematic as it will simultaneously act to reduce the amplitude (satisfying the penalty reduction) and lower the inversion fidelity.

A preferable option, therefore, is to limit the amplitude to within some user-defined experimental limits, and only constrain the pulse if the amplitude passes beyond the limit in a single step of the optimisation. This could be achieved by “clipping” the waveform if it passes beyond the amplitude limit [164], or by imposing the spill-out norm-square (SNS) penalty [23, 126] introduced in Chapter 5, applied only if the amplitude limit is exceeded and which is proportional to the norm square of the amplitude beyond the limit. The SNS penalty will immediately steer the optimisation back to a solution within the limits if it passes beyond them.

A further complication that could arise when employing pulses with variable amplitude is the effect of an uncompensated AC Stark shift. Varying the intensity of the light as a function of time would lead to a time-dependent Stark shift, thus resulting in a time-dependent off-resonance error.

We have repeated the optimisations carried out in Section 7.2, but this time allowed the pulse amplitude $\Omega_R(t)$ to vary whilst implementing a SNS penalty. We find the following: solutions appear to converge to the constant-amplitude phase-modulated pulses obtained in Section 7.2 and no improvement is found in terms of fidelity. Indeed, we find very

similar pulse shapes as depicted in Figure 7.15. Figure 7.15 shows the results of an individual pulse optimisation using fidelity $\mathcal{F}_{\text{square}}$, where the amplitude has converged to a rectangular envelope following 1000 iterations of GRAPE. This shows that there is no clear advantage to introducing amplitude modulation in terms of terminal fidelity because our optimisations converge to constant-amplitude solutions. It may be that using a different penalty on the amplitude such as a sine/cosine mapping allows for better solutions but this is a question left for future work.

We note here that optimising amplitude-modulated pulses does prove useful in some situations such as designing interferometer sequences to suppress high-frequency laser phase noise (explored in Chapter 8) and to design frequency-selective LMT pulses that suppress excitation of atoms within certain detuning bands (explored in Chapter 9).

7.5 Experimental results

In this section we demonstrate experimentally a robust GRAPE-optimised inversion pulse design in our Raman system and compare its performance with the rectangular π pulse and the WALTZ composite pulse, which was the best composite pulse tested in this system previously. In both cases we find large improvements in peak fidelity and robustness to Raman detuning with GRAPE. Our hardware was limited to realising pulses of constant amplitude when these experiments were performed and therefore ARP pulses were not tested.

We have designed a population inversion pulse using GRAPE that maximises the transfer of atoms initially in the state $|g\rangle$ to the state $|e\rangle$ for a cloud with a temperature of $120\ \mu\text{K}$ and a large variation in Rabi frequency of $\pm 10\% \Omega_{\text{eff}}$. The pulse duration was chosen to be $12\ \mu\text{s}$ for a Rabi frequency of $310\ \text{kHz}$, making it 7.4 times longer than a rectangular π -pulse, and allowing for a high terminal optimisation fidelity. The Rabi frequency used in the optimisation was determined empirically from the duration of a rectangular π pulse.

This pulse had 100 time-steps and the algorithm converged to the symmetric waveform shown in Figure 7.16 when optimising the point-to-point fidelity measure (Equation 7.1) with the derivative-norm-square (DNS) penalty term added (see Chapter 5), proportional to the difference between adjacent pulse steps, to enforce waveform smoothness [23]. We found that increasing the number of time-steps in this pulse led only to a negligible increase in fidelity.

As explained in Chapter 6, and recalled here for convenience, we realise our pulses on a thermal cloud of $\sim 10^7$ ^{85}Rb atoms, released from a 3-D magneto-optical trap (MOT), cooled in an optical molasses for $\sim 10\ \text{ms}$ and optically pumped into the




FIGURE 7.16: Optimal state-transfer pulse designed using GRAPE to transfer atoms between levels $|g\rangle$ and $|e\rangle$. The fidelity choice was $\mathcal{F}_{\text{square}}$ (Equation 7.1) and the pulse was optimised to provide efficient state-transfer for a $120\ \mu\text{K}$ cloud of ^{85}Rb with pulse-length errors in the range $\pm 10\%$ Ω_{eff} . The effective Rabi frequency, Ω_{eff} , was 310 kHz.

$|5S_{1/2}, F=2\rangle \equiv |g\rangle$ state with a distribution over the five m_F sub-levels which, as the pumping is performed with the MOT light along all axes, is assumed to be uniform [45].

The cloud temperature is tuned by adjusting the intensity of the cooling light during the optical molasses and is measured by performing Raman Doppler velocimetry with the Raman beams at low power [92]. The temperature can be adjusted in the range of $20 - 200\ \mu\text{K}$ and, combined with the multiple Rabi frequencies due to the different coupling strengths of the five m_F levels present in the sample, provides an inhomogeneous system with which to explore the performance of pulses designed to provide robustness against the resulting off-resonance and pulse-length errors.

As explained in Chapter 6, control of the relative laser phase between the two Raman beams is achieved by modulating the phase of the radio-frequency signal driving the EOM (electro-optic modulator) using an I&Q modulator. The I and Q inputs of the I&Q modulator are controlled by the dual outputs of an arbitrary waveform generator (AWG). To realise a phase sequence $\phi_n, n = 1, 2, \dots, N$, the outputs are programmed with the waveforms $I_n = \sin(\phi_n)$ and $Q_n = \cos(\phi_n)$ and are configured to hold the final phase value ϕ_N until a hardware trigger is received and the waveforms are played at a sample rate that is adjusted to set the total duration of the modulation τ_{mod} ⁴.

⁴In practice, the time required for the AWG to respond to a hardware trigger is determined by the sample rate. We therefore oversample waveforms, so that each ϕ_n comprises many “subsamples” of equal phase, in order to keep the AWG sample period much less than the $\sim 100\ \text{ns}$ rise-time of the AOM that shutters the Raman beams.




FIGURE 7.17: Illustration, not to scale, of the timing parameters optimised to realise optimal inversion 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 ~ 100 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. When the phase waveform has finished, the phase is maintained at its final value ϕ_N .

By concurrently triggering the phase modulation AWG and the AOM that shutters the Raman beams to initiate a pulse, then measuring $\langle P_e \rangle$ once the AOM shutter has been turned off again after a variable time τ , we can track the temporal evolution of the atomic state during a phase sequence. The hardware trigger delay and sample rate of the AWG are adjusted in order to vary the start time τ_{offset} and duration τ_{mod} of the phase waveform respectively (Figure 7.17) and maximise the peak excited fraction in these temporal scans.

Having chosen the experimental pulse timings such that the fidelity is maximised at a single value of the Raman detuning δ , set to the centre of the light-shifted resonance determined from the spectral profile of a rectangular π -pulse, a spectral profile can be measured by measuring $\langle P_e \rangle$ upon completion of the pulse at a range of detuning values.

7.5.1 Temporal scans

The temporal profile, after optimising τ_{offset} , is shown in Figure 7.18, showing a peak in the excited fraction at the end of the phase sequence, represented by the vertical dashed line, after which damped Rabi oscillations are evident as the phase is kept constant.




FIGURE 7.18: (a) Optimal state-transfer pulse designed using GRAPE to transfer atoms between levels $|g\rangle$ and $|e\rangle$. (b) Measured fraction (circles) of atoms in the excited state $|e\rangle$ after the Raman light is extinguished at various times during a pulse. The solid line is a theoretical curve produced by the model from [45], 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, assuming 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 $35 \mu\text{K}$.

We find that optimised pulses demonstrate a considerable resilience to variations in the trigger delay τ_{offset} , and can be started as much as a quarter of a π -pulse duration after the light is turned on with little change to the peak transfer.

The highest excited fraction for all pulses is achieved when the light is kept on for slightly longer than the phase sequence, with $\tau > \tau_{\text{mod}}$ by $\sim 200 \text{ ns}$. The AOM rise-time is not factored into the optimisation process and this is the only effect of it that we observe experimentally, with the temporal data being well fitted by a numerical model that assumes the light reaches full intensity instantaneously at the start of the phase sequence as shown in Figure 7.18.

The temporal scan shows good agreement between our model (outlined in Appendix B) and experiment, which gives us reason to believe that we have a) correctly compensated for the pulse errors in our atom cloud and b) that the atoms are ‘seeing’ the intended and undistorted waveform.

7.5.2 Robustness

We now turn our attention to exploring the robustness of our pulses to variations in the Raman detuning. In order to do this, we prepare all atoms in the lower ground hyperfine state $|g\rangle$, apply our chosen inversion pulse, and measure the final fraction of atoms in




FIGURE 7.19: Fraction of atoms transferred to the excited state as a function of the laser detuning δ_L in (a) a $35 \mu\text{K}$ atom cloud and (b) a $150 \mu\text{K}$ cloud. Data are shown for the GRAPE pulse from Figure 7.16 (diamonds), the WALTZ pulse (empty circles) and a rectangular π -pulse (filled circles). The effective Rabi frequency was 270 kHz . Solid lines show theory curves produced from the model used by Dunning *et al.* [45], which assumes a Maxwell-Boltzmann atomic velocity distribution. The offset of the peaks from $\delta_L = 0$ is due to the light shift.

the excited state. This procedure is repeated while varying a fixed laser detuning δ_L between the Raman beams and the atoms. We obtain these ‘spectral scans’ for two different atomic cloud temperatures, one at $\sim 35 \mu\text{K}$, and the other one at $\sim 150 \mu\text{K}$. We compare our GRAPE pulse with the WALTZ composite pulse and the familiar π pulse.

In a sub-Doppler cooled cloud at $35 \mu\text{K}$, GRAPE and WALTZ pulses achieve close to $99.8(3)\%$ and $96(2)\%$ transfer respectively at the light-shifted resonance, while a rectangular π -pulse achieves just $75(3)\%$. This is shown in Figure 7.19, where the

broadband nature of the GRAPE pulse is evident: while the fidelity of the WALTZ pulse drops below 90% when detuned 100 kHz from resonance, the GRAPE pulse can be detuned by 380 kHz for the same fidelity. This broad spectral profile can also be a signature of a good LMT augmentation pulse, which must work equally well for atoms that have received a large number of recoil kicks [42].

In a cloud with a temperature much closer to the Doppler cooling limit $\sim 150 \mu\text{K}$, for which the peak transfer of a rectangular π -pulse is just 54(2)%, the broader spectral width of the WALTZ and GRAPE pulses results in more efficient state-transfer on resonance. This is also shown in Figure 7.19, where the GRAPE pulse achieves 89(4)% transfer on resonance compared with 81(4)% for WALTZ.

7.6 Conclusions

Tolerance of pulse-length and off-resonance errors is essential for the pulse operations in atom interferometers, where a range of velocities, beam intensities and Zeeman sub-states may be encountered. We have used the optimal control technique of gradient ascent pulse engineering (GRAPE) to obtain robust mirror pulses, tailored to accommodate the inhomogeneities found in cold-atom matter-wave interferometers, and find such pulses to outperform all the other composite pulses and adiabatic rapid passage pulses that we have tested. The improved fidelity should allow improvements in the sensitivity of interferometric measurements by increasing the fringe contrast and permitting greater use of augmentation pulses for large-momentum-transfer interferometers. The tolerance of atom velocity variations will lower measurement noise by allowing the use of warmer atom clouds, and hence higher atom numbers, without incurring the losses of further cooling or filtering.

Our optimal control approach depends upon an appropriate choice of the measure of performance. We have presented and analysed two different choices of mirror pulse fidelity: a point-to-point (PP) state-transfer fidelity (Equation 7.1) and a phase-preserving universal-rotation (UR) fidelity (Equation 7.3). Pulses designed using the PP fidelity maximise the transfer of atoms between the two states of the interferometer in the presence of large off-resonance and pulse-length errors, and should therefore improve the sensitivity of large-momentum-transfer (LMT) interferometers. Alternatively, our UR pulse designs are tailored to replace the central pulse in the three-pulse Mach-Zehnder atom interferometer sequence used for inertial sensing applications. While many composite and adiabatic pulses can lead to improvements in state-transfer efficiencies when compared with the π pulse, their effect on the interferometer phase can be detrimental, leading different atoms to exit the interferometer with different phases. Our optimised UR pulses are designed specifically to maximise the interferometer contrast by reducing

unwanted variation in the interferometer phase with detuning and coupling strength errors.

We have also experimentally demonstrated a PP inversion pulse that achieves 99.8(3)% transfer between the hyperfine ground states in a thermal cloud of ^{85}Rb atoms, compared with a conventional π pulse efficiency of 75(3)%. Our optimised inversion pulse is robust to variations in laser intensity and detuning, maintaining a transfer efficiency of 90% at detunings for which the π pulse fidelity is below 20%. Furthermore, we find excellent agreement between the predicted state-evolution during the pulse and that observed experimentally, thus confirming the validity of the optimisation approach and our model of the specific inhomogeneities present. The best results of Raman LMT interferometers to date have been achieved with adiabatic augmentation pulses [46, 48], but optimal control pulses such as this have the potential to realise similarly robust transfer with significant reductions in pulse duration [8].

The optimal control solutions depend on the parameters used in the optimisation, and we have conducted a preliminary investigation into the nature of this dependence. We have found that GRAPE-optimised mirror pulses do not always converge to the same solution when the initial guess for the phase profile is varied. This implies the existence of multiple peaks on the fidelity landscape and that, while many solutions differ only slightly in terms of terminal fidelity, they may differ in other respects. For example, some solutions may be more robust to off-resonance and pulse-length errors outside the desired compensation range; and others may be robust to control noise and therefore the best choice to implement experimentally.

Chapter 8

Optimised interferometer sequences

In this chapter we extend optimal control from the design of mirror pulses addressed in the previous chapter to the optimisation of the beamsplitter operation and the entire interferometer sequence. We present tailored pulse sequences that improve the contrast and phase-sensitivity of atom interferometers subject to systematic inhomogeneities in detuning and Rabi frequency. These pulse sequences are investigated through simulations, and preliminary experimental data are presented which show a threefold improvement in fringe visibility in a thermal cloud of ^{85}Rb atoms with our optimised pulse sequence over that obtained using conventional π and $\pi/2$ pulses. We also investigate how the symmetry of the pulses is related to their performance, proposing a symmetry-based “flip-reverse” pulse design, and examine the susceptibility of optimised pulse sequences to laser phase noise. Parts of this chapter have been published [1, 3], and the results represent the first reported application of optimal control theory to design tailored Raman pulse sequences for atom interferometry. The experimental data presented in this chapter were taken in collaboration with Max Carey.

8.1 Beamsplitter pulses

Following application of the Mach-Zehnder three-pulse $\pi/2 - \pi - \pi/2$ sequence to an atom initially in the state $|g\rangle$, the probability to exit the interferometer in the state $|e\rangle$ varies according to Equation 3.11,

$$P_e = \frac{1}{2}(\mathcal{A} - \mathcal{B} \cos(\Phi + \Delta\phi)), \quad (8.1)$$

where $\mathcal{A}$ is the fringe offset, $\mathcal{B}$ is the fringe visibility or contrast, and Φ is the interferometer phase which depends on the acceleration or rotation rate of the atom relative to the fixed laser beams. $\Delta\phi$ is a shift to the interferometer phase.

An optimal three-pulse interferometer sequence, represented by propagators $\hat{U}_1, \hat{U}_2, \hat{U}_3$ should maximise the fringe contrast $\mathcal{B}$, minimise any unwanted variation of/in the inertial phase, $\Delta\phi$, and fix the offset $\mathcal{A}$ for all atoms with the range of detunings and coupling strengths found in the atomic cloud. Ideally, sequences should be robust to variations in these quantities that occur both spatially (i.e. across the atom sample) and temporally (i.e. from pulse to pulse).

In Chapter 3 we showed that in order to obtain the optimum offset and interferometer contrast the following condition on the beamsplitter pulses (pulses 1 and 3) must be met:

$$|\langle e|\hat{U}_{1,3}|g\rangle|^2 = \frac{1}{2}. \quad (8.2)$$

This condition is equivalent to requiring that the first beamsplitter pulse produce an equal superposition of the two states of the interferometer - effecting a 90° rotation of the quantum state. Also, in addition to maximising the visibility of the interference fringes, it is crucial that the contribution of each pulse to the interferometer phase shift $\Delta\phi$ is fixed for all atoms in the ensemble:

$$\Delta\phi = \phi_1 - 2\phi_2 + \phi_3 \quad (8.3)$$

$$\phi_1 \equiv \phi(\langle e|\hat{U}_1|g\rangle) - \phi(\langle g|\hat{U}_1|g\rangle) = \text{constant} \quad (8.4)$$

$$\phi_2 \equiv \phi(\langle e|\hat{U}_2|g\rangle) = \text{constant} \quad (8.5)$$

$$\phi_3 \equiv \phi(\langle g|\hat{U}_3|g\rangle) + \phi(\langle e|\hat{U}_3|g\rangle) = \text{constant}. \quad (8.6)$$

ϕ_1 is equivalent to the phase of the resultant superposition state, or ‘superposition phase’, formed by application of the first beamsplitter to the initial state $|g\rangle$.

Our first interferometer pulse $\hat{U}_1$ should take atoms from the initial state $|g\rangle$, to an equal superposition with well-defined phase on the equator of the Bloch sphere, e.g. $|\psi_T\rangle = (|g\rangle + |e\rangle)/\sqrt{2}$. This means our fidelity measure, $\mathcal{F}_1$, may be written as

$$\mathcal{F}_1 = |\langle \psi_T|\hat{U}_1|g\rangle|^2, \quad (8.7)$$

yielding a PP 90° state-transfer pulse. This choice of fidelity measure, if maximised, results in a beamsplitter which satisfies the conditions $|\langle e|\hat{U}_1|g\rangle|^2 = 1/2$ and $\phi_1 = \phi(\langle e|\hat{U}_1|g\rangle) - \phi(\langle g|\hat{U}_1|g\rangle) = 0$. We note that it is not necessary for the first pulse to perform a universal 90° rotation as all atoms start in the same state, namely $|g\rangle$. This makes the optimisation less restrictive, as PP pulses are easier to find and shorter than UR pulses [10].

Our chosen PP beamsplitter fidelity measure treats population split and phase uniformity in equal measure, but it is possible to modify the fidelity measure to prioritise, for example, phase uniformity in the resulting superposition over producing an equal population split. This approach was adopted by Skinner *et al.* [165] in the design of 90° rotation pulses for NMR applications.

To investigate how the pulse design varies with the desired compensation range of off-resonance and pulse-length errors, we have optimised two PP beamsplitter pulses using the fidelity measure given by Equation 8.7 to compensate different combinations and magnitudes of control errors. The first pulse, which we denote beamsplitter 1 (BS 1) was designed to compensate off-resonance errors arising from a cloud at $10\ \mu\text{K}$ and amplitude errors in the range $\pm 5\% \Omega_{\text{eff}}$, while the second pulse, beamsplitter 2 (BS 2), was designed to compensate off-resonance errors arising from a much colder cloud at $1\ \mu\text{K}$ and a larger range of amplitude errors in the range $\pm 60\% \Omega_{\text{eff}}$. The off-resonance errors used in the optimisation corresponded to a uniform distribution spanning $6k_{\text{eff}}\sigma_v$ rounded to the nearest 10 kHz, where σ_v is the width of the ‘true’ Maxwell-Boltzmann velocity distribution for ^{85}Rb at each temperature. The pulse durations were $12t_\pi$ each with 400 time-slices, leading to terminal fidelities of 0.999 for BS 1 and 0.998 for BS 2. The nominal Rabi frequency Ω_{eff} was 200 kHz. The L-BFGS GRAPE algorithm was used to optimise each pulse. The algorithm was set to terminate if the number of iterations passed 1000 or if the norm of the gradient fell below 10^{-6} . The initial guess for each pulse was a random series of laser phases and the algorithm converged to the asymmetric waveforms shown in Figure 8.1.

Figures 8.1c and 8.1d show the simulated probability for an atom to be transferred to the excited state by each pulse as a function of the Raman detuning in the case of no amplitude error and Figures 8.1g and 8.1h show the probability as a function of amplitude error for a resonant atom. The resulting superposition phase, ϕ_1 , defined by Equation 8.4, is depicted as a function of detuning and amplitude offset in Figures 8.1e, 8.1f, 8.1i, and 8.1j. The same quantities are depicted for the rectangular $\pi/2$ beamsplitter pulse for comparison and the shaded blue regions correspond to the ranges of off-resonance and pulse-length errors included in the optimisation ensemble.

We can see that the GRAPE optimised pulses satisfy the design goals by minimising variation in the superposition phase and producing an equal population split over the design ranges of detunings and pulse-length errors. We can also see from Figures 8.1e and 8.1f that the superposition phase, ϕ_1 , following the rectangular $\pi/2$ exhibits an (almost) linear dependence on the detuning. This is easy to understand when one recalls that an off-resonance error corresponds to a tilt in the rotation axis in the Bloch sphere picture. This linear dependence of the superposition phase is cancelled by the final beamsplitter in the conventional $\pi/2 - \pi - \pi/2$ sequence, providing the coupling strength does not vary between the pulses [109].




FIGURE 8.1: Optimisation results for two PP 90° beamsplitter pulses: (a) a pulse optimised for a $10\ \mu\text{K}$ cloud with pulse-length errors in the range $\pm 5\% \Omega_{\text{eff}}$ (BS 1), and (b) a pulse optimised for a $1\ \mu\text{K}$ cloud with pulse-length errors in the range $\pm 50\% \Omega_{\text{eff}}$ (BS 2). (c) and (d) show the state-transfer efficiency of BS 1 and BS 2, respectively, as a function of the Raman detuning. (e) and (f) show the resulting superposition phase ϕ_1 for each pulse as a function of the Raman detuning. The state-transfer efficiency on resonance as a function of amplitude error is shown in (g) and (h) and, finally, (i) and (j) show the superposition phase ϕ_1 for each pulse as a function of amplitude error on resonance. Phases have been unwrapped and shifted by a constant amount to aid visual clarity. The shaded regions correspond to the compensation range of off-resonance and pulse-length errors.

There exists a range of composite 90° rotation pulses designed to compensate various degrees of off-resonance and pulse-length errors in NMR [63, 166, 167]. Like composite 180° pulses, 90° rotation pulses may be categorised into belonging to the class of point-to-point (PP) or universal rotation (UR) transformations. A UR 90° pulse is designed to effect a 90° rotation about a fixed axis for any initial spin state, and a PP 90° pulse is designed to transform a specific initial state to a specific target e.g. rotate a Bloch vector initially aligned along the z -axis to a specific point on the equator. However, 90° pulses may also be designed to transform an initial Bloch vector aligned with the z -axis to a range of states on the equator where the superposition phase has a certain dependence on the detuning [63, 168]. The family of ICEBERG (Inherent Coherence Evolution optimised Broadband Excitation Resulting in constant phase Gradients) pulses [169] is an example of this class of pulse, designed in this case to produce a superposition phase with a linear dependence on offset frequency. While strictly not a PP operation in isolation, such pulses may be easier to optimise than PP or UR pulses and may be combined with a delay period with the control fields extinguished to reproduce an overall PP transformation.

Frequency swept ARP pulses can also be used as beamsplitters [49, 101], known as adiabatic half-passages or AHP (in contrast to the adiabatic full passage AFP for inversion applications). By terminating an ARP inversion pulse mid frequency sweep, a resonant atom will end up in an equal superposition state with well-defined phase. Unlike ARP inversion pulses that impart an intensity dependent dynamic phase when applied to superposition states, the ARP beamsplitter leads to a superposition phase that is independent of intensity and detuning in the adiabatic limit. However, ARP beamsplitters cannot take atoms with different detunings to the equator: if an atomic ensemble contains atoms with a spread in detunings, different atoms will under- and over-shoot the equator, thereby introducing a Doppler sensitivity to the pulse. ARP beamsplitters therefore trade robustness against off-resonance errors for robustness against pulse-length errors.

Figure 8.2 shows fidelity contour plots for a range of beamsplitter pulses, including two composite 90° pulses and the tanh/tan AHP pulse. The 90° composite pulses correspond to those with the highest degree of compensation of off-resonance and pulse-length errors taken from [150, 166]. Fidelity contours are also shown for the two GRAPE pulses (BS 1 and BS 2) taken from Figure 8.1. GRAPE-optimised beamsplitters clearly demonstrate an increased robustness to both error sources compared with the ARP and composite alternatives.

8.1.1 “Flip-reverse” operation

It is clear that the final beamsplitter cannot be produced using the same fidelity choice as the first because it must be applied to atoms that have different initial states (ideally




FIGURE 8.2: State-transfer beamsplitter efficiency (Equation 8.7) of different pulses as a function of off-resonance and pulse-length errors. The hatched regions represent the regions of optimisation (included only for the GRAPE pulses) and the contours correspond to fidelities of 0.8, 0.9, 0.95, and 0.99.

at a range of points on the Bloch sphere equator). How then do we optimise the final pulse in the Mach-Zehnder sequence?

The third pulse must accurately map the relative phase acquired between the two internal states at the end of the second dwell-time, onto a difference in atomic population. This pulse does not need to be a universal 90° rotation, as only the z-component of the final Bloch vector matters when measuring the excited state population at the end of the interferometer. The action of this pulse can be thought of as a phase sensitive $\pi/2$ rotation, and may be obtained by taking the pulse profile for the first pulse $\phi(t)$ with duration τ , and reversing it in time about the temporal midpoint and inverting the pulse profile to obtain $-\phi(\tau - t)$. We refer to this as the “flip-reverse” operation.

Our pulses are represented by waveforms $\phi(t)$ in the combined Raman laser phase which are piece-wise constant functions $\phi(t) = \phi_1, \phi_2, \dots, \phi_n$. They can therefore be represented by a series of propagators, $\hat{U}^i$, or series of rotations $\mathbf{R}_i$, where $i = 1, 2, \dots, n$ for a pulse with n steps. The overall effect is a combined pulse propagator and effective rotation, $\hat{U}_{\text{pulse}}$ and $\mathbf{R}_{\text{pulse}}$, which are functions of the detuning and Rabi frequency. If the laser phase $\phi_i \rightarrow -\phi_i$ is inverted for an individual rotation element in the pulse, $\mathbf{R}_i$, then this is equivalent to inverting the y-component of that step’s rotation axis (this follows from the fact that during a time-step, the quantum state rotates about the field vector $\mathbf{\Omega} = \Omega_R \cos(\phi)\hat{x} + \Omega_R \sin(\phi)\hat{y} + \delta\hat{z}$). We denote this phase inversion operation by $\mathbf{R}_i \rightarrow \mathbf{R}_i^Y$.

Let $\overset{1}{\mathbf{R}}$ represent the effective rotation of the first beamsplitter, and $\overset{3}{\mathbf{R}}$ the rotation of the last pulse. If $\phi_3(t) = -\phi_1(\tau - t)$, then it follows that

$$\overset{3}{\mathbf{R}} = \overset{1}{\mathbf{R}}_1^Y \overset{1}{\mathbf{R}}_2^Y \dots \overset{1}{\mathbf{R}}_n^Y \quad (8.8)$$

$$= \mathbf{R}_y^{-1}(\pi) \mathbf{R}_y(\pi) (\overset{1}{\mathbf{R}}_1^Y \overset{1}{\mathbf{R}}_2^Y \dots \overset{1}{\mathbf{R}}_n^Y) \mathbf{R}_y^{-1}(\pi) \mathbf{R}_y(\pi) \quad (8.9)$$

$$= \mathbf{R}_y^{-1}(\pi) \mathbf{R}_y(\pi) \overset{1}{\mathbf{R}}_1^Y \mathbf{R}_y^{-1}(\pi) \mathbf{R}_y(\pi) \overset{1}{\mathbf{R}}_2^Y \mathbf{R}_y^{-1}(\pi) \dots \mathbf{R}_y(\pi) \overset{1}{\mathbf{R}}_n^Y \mathbf{R}_y^{-1}(\pi) \mathbf{R}_y(\pi) \quad (8.10)$$

$$= \mathbf{R}_y^{-1}(\pi) (\overset{1}{\mathbf{R}}_1^{-1} \overset{1}{\mathbf{R}}_2^{-1} \dots \overset{1}{\mathbf{R}}_n^{-1}) \mathbf{R}_y(\pi) \quad (8.11)$$

$$= \mathbf{R}_y^{-1}(\pi) (\overset{1}{\mathbf{R}}_n \dots \overset{1}{\mathbf{R}}_2 \overset{1}{\mathbf{R}}_1)^{-1} \mathbf{R}_y(\pi) \quad (8.12)$$

$$= \mathbf{R}_y^{-1}(\pi) \overset{1}{\mathbf{R}}^{-1} \mathbf{R}_y(\pi) \quad (8.13)$$

$$= \overset{1}{\mathbf{R}}^Y, \quad (8.14)$$

where in lines 8.10 and 8.13 we have applied the geometric identity [113]

$$\mathbf{R}^{-1} \equiv \mathbf{R}_y(\pi) \mathbf{R}^Y \mathbf{R}_y^{-1}(\pi). \quad (8.15)$$

$\mathbf{R}_y(\pi)$ represents a rotation by π about the y -axis. Therefore, the rotational axes for pulses 1 and 3 are related by inverting the y component of the effective rotation axis. Suppose pulse 1 performs an effective rotation about the axis $\hat{\mathbf{n}} = (n_x, n_y, n_z)$ by an angle θ . Its propagator may then be written as¹

$$\hat{U}_1 = \begin{pmatrix} \cos(\frac{\theta}{2}) - i \sin(\frac{\theta}{2})n_z & -\sin(\frac{\theta}{2})n_y - i \sin(\frac{\theta}{2})n_x \\ \sin(\frac{\theta}{2})n_y - i \sin(\frac{\theta}{2})n_x & \cos(\frac{\theta}{2}) + i \sin(\frac{\theta}{2})n_z \end{pmatrix}. \quad (8.16)$$

Applying the flip-reverse procedure gives us

$$\hat{U}_3 = \begin{pmatrix} \cos(\frac{\theta}{2}) - i \sin(\frac{\theta}{2})n_z & +\sin(\frac{\theta}{2})n_y - i \sin(\frac{\theta}{2})n_x \\ -\sin(\frac{\theta}{2})n_y - i \sin(\frac{\theta}{2})n_x & \cos(\frac{\theta}{2}) + i \sin(\frac{\theta}{2})n_z \end{pmatrix}. \quad (8.17)$$

From this we can see that $|\langle e|\hat{U}_3|g\rangle|^2 = |\langle e|\hat{U}_1|g\rangle|^2$, $\phi(\langle g|\hat{U}_3|g\rangle) = \phi(\langle g|\hat{U}_1|g\rangle)$, and $\phi(\langle e|\hat{U}_3|g\rangle) = -\phi(\langle e|\hat{U}_1|g\rangle) + \pi$, meaning $\phi_1 = -\phi_3 + \pi$. Therefore, if pulse 1 is optimised to be a 90° point-to-point pulse using the fidelity measure given by Equation 8.7, ϕ_1 will be constant and pulse 3 will fulfil the optimal conditions on $\Delta\Phi$, the contrast, and the offset.

Optimising the three-pulse interferometer with these symmetry constraints (the final pulse has the time-reversed and inverted profile of the first beamsplitter) minimises the unwanted modification to the inertial phase term Φ and maximises the contrast of resulting fringes. Furthermore, the optimisation procedure is simplified because instead of searching for 3 pulses, only 2 are required to form a complete Mach-Zehnder sequence.

Using the flip-reverse construction introduces a cancellation in the contributions of the first and final beamsplitter pulses to the interferometer phase: any variation in the superposition phase introduced by the first pulse can be cancelled by the final one. This means that the beamsplitter does not strictly need to be a PP pulse, but could instead simply be designed to bring atoms down to the Bloch sphere equator with some phase spread. However, in reality, the cancellation between the phase terms relies on the Rabi frequency and detuning for a given atom remaining constant throughout the sequence, which may not be the case especially in situations where the atom cloud expands over significant variations in optical intensity. We therefore build in robustness to this imperfect cancellation by optimising PP pulses and thereby minimising variation in ϕ_1 and ϕ_2 .

Figure 8.3 depicts the simulated interferometer phase following sequences formed of GRAPE beamsplitters (BS 1 and BS 2) combined using the flip-reverse procedure with π mirror pulses. Figure 8.3a shows the case where the Rabi frequency for a given atom is the same for the first and final pulses, i.e. $\Omega_1 = \Omega_3$. In this situation the cancellation

¹This follows from the fact that the propagator for a rotation by angle α about an arbitrary axis described by the unit vector $\hat{\mathbf{n}}$, is given by $\exp(-i\alpha\hat{\mathbf{n}} \cdot \hat{\boldsymbol{\sigma}}/2)$, where $\hat{\boldsymbol{\sigma}} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)$ is the vector of Pauli spin matrices.

of phases ϕ_1 and ϕ_2 is perfect. If the flip-reverse transformation is not applied and the final beamsplitter is simply the same as the first pulse, then the result is a large and unwanted variation in the interferometer phase as a function of the Raman detuning. This variation with detuning will wash out the fringes in a cloud of atoms with non-zero temperature, and different atoms will exit the interferometer with different phases.

Figure 8.3b and Figure 8.3c illustrate what happens when the Rabi frequency for the final pulse Ω_3 is decreased relative to the Rabi frequency for the first pulse Ω_1 . In this situation the cancellation of beamsplitter phases is imperfect, and ripples reappear in the phase response. Although the magnitude of unwanted phase variation increases as the discrepancy in Rabi frequencies between the pulses is increased, the beamsplitter pulse optimised to withstand large variations in amplitude for a cold atom cloud, BS 2, exhibits a far more uniform phase response near resonance than BS 1.

Figure 8.4 shows how the simulated interferometer contrast $\mathcal{B}$ following different pulse sequences varies with the detuning of an atom throughout the sequence and when the coupling strength varies from pulse to pulse. The conventional π mirror pulse, with its increased Doppler sensitivity compared with the $\pi/2$ pulse, is the primary source of contrast loss due to off-resonance errors within the $\pi/2 - \pi - \pi/2$ sequence. This explains why - as demonstrated in the previous chapter - significant increases in interferometer contrast may be obtained by optimising the mirror operation alone. As a result, the effect of the beamsplitters on the contrast was isolated in the simulation by replacing the central π pulse by a ‘perfect’ π pulse that effects the same rotation regardless of detuning. The ARP beamsplitter sequence shows an increased sensitivity to detuning with a narrow corridor of high contrast extending over a large range of pulse-length errors. The two GRAPE sequences display a large increase in the robustness of the interferometer contrast, with the BS 2 sequence capable of maintaining contrast above 99% even when the Rabi frequency during the final pulse is $0.4\times$ that of the first.

8.2 Antisymmetric mirror pulses

In this section we revisit the mirror pulse and use symmetry arguments to show that mirror pulses with antisymmetric phase profiles automatically fulfil the condition that the interferometer phase should be unaffected by detuning and/or variations in coupling strength. This property is useful because even the highest-fidelity UR pulses found using the mirror pulse fidelity measure $\mathcal{F}_{\text{real}}$ from Chapter 7 lead to small ripples in the contribution of the pulse to the interferometer phase. These ripples could lead to a bias in the interferometer phase when the signal is averaged over a distribution of atoms, and introduce a sensitivity to any asymmetry in the velocity distribution.

The symmetry of a pulse determines how the effective rotation axis behaves as a function of the offset from resonance and variation in the field strength [10]. Consider an




FIGURE 8.3: The effect of a variation in Rabi frequency between pulses on the interferometer phase $\Delta\phi$ following different interferometer sequences. The interferometer phase is simulated as a function of a constant detuning. (a) shows the interferometer phase following the standard $\pi/2 - \pi - \pi/2$ sequence and the phase following two sequences formed of the GRAPE beamsplitters from Figure 8.1 combined in flip-reverse configurations with conventional π mirror pulses. The red dotted line shows the interferometer phase that results from a sequence where the flip-reverse transformation is not applied and the final beamsplitter is identical to the first one. (b) and (c) show the interferometer phase in the situation where the Rabi frequency during the final pulse Ω_3 is $0.75\times$ and $0.5\times$ the Rabi frequency during the first pulse Ω_1 , respectively. To aid visibility, the curves have been separated by a constant phase shift.

antisymmetric pulse profile $\phi(t)$, where $\phi(t) = -\phi(\tau - t)$. Recalling our definition of the phase inversion operation for a particular rotation segment, the total effective rotation




FIGURE 8.4: Simulated contrast $\mathcal{B}$ following different interferometer sequences as a function of constant Raman detuning throughout the sequence (y -axis) and Rabi frequency mismatch between pulse 3 and 1 (x -axis). The mirror in each case is a ‘perfect’ π pulse which is completely Doppler-insensitive and therefore the contribution to the contrast is purely due to the effect of the beamsplitter pulses. Both the GRAPE-optimised beamsplitters and the ARP beamsplitter have the same duration of $12t_\pi$. The contours correspond to contrast values of 0.8, 0.9, 0.95, and 0.99. The region of error compensation is shown by the hatched box for each GRAPE pulse.

represented by this pulse is [10]

$$\mathbf{R}_{\text{pulse}} = \mathbf{R}_n \dots \mathbf{R}_2 \mathbf{R}_1 \quad (8.18)$$

$$= \mathbf{R}_1^Y \mathbf{R}_2^Y \dots \mathbf{R}_n^Y \quad (8.19)$$

$$= \mathbf{R}_{\text{pulse}}^Y. \quad (8.20)$$

This result implies that the y -component of the pulse’s effective rotation axis must be 0. This leads to the well-known result in NMR that any pulse with symmetric amplitude profile and an antisymmetric phase profile must have an effective rotation axis confined to the xz plane [10, 93, 113].

Given that the effective rotation axis of a pulse with this symmetry is confined to the xz plane of the Bloch sphere, we can show that an antisymmetric mirror pulse used in a three-pulse interferometer will ensure the inertial phase is not detrimentally modified for different atoms.

Consider again the form of a general pulse propagator for the spin 1/2 system, representing a rotation by angle θ about an axis $\hat{\mathbf{n}} = (n_x, n_y, n_z)$,

$$\hat{U}_{\text{pulse}} = \exp\left(-\frac{i\theta\hat{\mathbf{n}} \cdot \hat{\boldsymbol{\sigma}}}{2}\right). \quad (8.21)$$

The effective angle and axis depend on both the detuning δ and Rabi frequency Ω_R during the pulse.

If the pulse profile for this pulse is antisymmetric, then $\hat{\mathbf{n}} = (n_x, 0, n_z)$, which implies

$$\hat{U}_{\text{pulse}} = \begin{pmatrix} \cos(\frac{\theta}{2}) - i\sin(\frac{\theta}{2})n_z & -i\sin(\frac{\theta}{2})n_x \\ -i\sin(\frac{\theta}{2})n_x & \cos(\frac{\theta}{2}) + i\sin(\frac{\theta}{2})n_z \end{pmatrix}. \quad (8.22)$$

Finally, we see that $\phi(\langle e|\hat{U}_2|g\rangle) = \text{constant}$, and therefore for a mirror pulse with this symmetry $\Delta\phi$ has no contribution from the mirror pulse for any detuning or Rabi frequency error that remains constant during the pulse. Mirror pulses with antisymmetry automatically fulfil the condition that the interferometer phase should be unaffected by detuning and/or variations in coupling strength. Antisymmetric UR 180° pulses may be constructed by first optimising a single PP 90° pulse and following the steps outlined by Luy *et al.* [170]².

To illustrate this useful property of antisymmetric mirror pulses and compare them with those designed following the procedure outlined in Chapter 7, we have optimised two distinct UR mirror pulses: one using $\mathcal{F}_{\text{real}}$ (unconstrained UR), and the other using the construction procedure outlined by Luy *et al.* [170] which results in an antisymmetric UR 180° pulse (antisymmetric UR). The temperature was 20 μK and the pulse-length error compensation range was $\pm 10\% \Omega_{\text{eff}}$, where the effective Rabi frequency was 200 kHz. The phase profiles obtained using GRAPE for each pulse type are shown in Figure 8.5. Figure 8.5 also shows the simulated transfer efficiency of each pulse as a function of detuning and pulse-length errors, and the variation in the contribution of each pulse to the interferometer phase. While it is clear that both pulses optimise the transfer efficiency over the design range, and that the $\mathcal{F}_{\text{real}}$ pulse has minimised variation in the phase, the antisymmetric pulse results in no phase variation by symmetry.

²Alternatively, we could impose antisymmetry as a constraint and optimise the state-transfer mirror pulse fidelity measure given by Equation 7.1.




FIGURE 8.5: Phase profiles (a) and (b) for UR mirror pulses designed without and with antisymmetric symmetry constraints for a cloud of ^{85}Rb at $20\ \mu\text{K}$ and a range of pulse-length errors of $\pm 10\% \Omega_{\text{eff}}$. (c) and (d) depict the transfer efficiency of each pulse as functions of detuning and pulse-length errors respectively. (e) and (f) show the simulated contribution of each mirror pulse to the interferometer phase, $\Delta\phi$, as functions of detuning and pulse-length errors respectively. There is no phase variation caused by the antisymmetric pulse. The optimisation ranges for off-resonance errors and pulse-length errors are shown by the shaded blue regions.

8.3 Interferometric phase noise

We have so far focused our attention on the design of pulse sequences that mitigate errors in the control field such as detuning and amplitude offsets that do not vary on the time-scale of individual pulses or interferometer sequences. This is often a good assumption if the atoms do not move much during individual pulses or interferometer sequences and therefore see an approximately constant laser intensity and phase. In this section we investigate the effect of time-dependent noise in the control parameters themselves, and explore the effect of phase noise in the Raman beams on the interferometer sensitivity when employing GRAPE-optimised pulses.

The sensitivity of a conventional atom interferometer is degraded in the presence of laser phase noise in the Raman beams. Fluctuations in the laser phase throughout the interferometer manifest in fluctuations in the phase of resulting fringes, thus limiting the sensitivity of the interferometer [94]. Cheinet *et al.* [94] demonstrated that the $\pi/2 - \pi - \pi/2$ interferometer sequence of rectangular pulses acts like a low-pass filter for phase noise, suppressing the influence of noise on the interferometer phase at frequencies above the Rabi frequency.

Replacing the conventional π and $\pi/2$ pulses with alternatives in which the amplitude and/or phase vary in time changes the response of the interferometer to phase noise at high frequencies. For example, Fang *et al.* [96] studied the effect of shaped pulses on the phase sensitivity of a Mach-Zehnder atom interferometer. They found that using pulses with smooth amplitude profiles led to a rejection of high-frequency phase noise by the interferometer. Among the pulse shapes tested, the Gaussian pulse [118], where the temporal amplitude of the pulses is modulated according to a truncated Gaussian function, demonstrated the best noise rejection at high frequencies. This behaviour was attributed to the gradual rise and fall in the pulse amplitude when compared with rectangular pulses. Although the Gaussian pulses yielded the best noise rejection, their robustness to amplitude and detuning offsets is similar to the standard rectangular pulses [96, 118]. We are therefore interested in how robust phase-modulated sequences obtained through optimal control perform in the presence of phase noise. Phase noise will manifest as a distortion of the pulse waveforms, lowering the individual pulse fidelities, and therefore it is important to investigate the sensitivity or otherwise of GRAPE-optimised pulses to this effect.

Our aim is to calculate the variance of the interferometer phase σ_Φ^2 following a specific pulse sequence in the presence of phase noise at different frequencies. To do this we draw a sample of noise realisations at a given frequency ω , simulate the temporal evolution of the quantum state throughout the Mach-Zehnder interferometer for each noise realisation in the sample and, finally, compute the mean-squared variation in the interferometer phase compared with its noise-free value. Variations in the interferometer phase can be determined by evaluating variations in the final transition probability at the midpoint of a fringe [94, 171].

We model Raman phase noise at fixed frequency ω by the sinusoidal modulation $\phi_{\text{noise}}(t) = A \sin(\omega t + B)$ [94, 171, 172]. Furthermore, we assume that the amplitude of this noise A is constant, but that the offset phase B is uniformly distributed in the interval $[-\pi, \pi]$. To accurately model the effect of high-frequency noise during interferometer pulses, the pulse propagators used to evolve the atomic state must be subdivided e.g. using Python's *numpy.repeat* function [173] into a large number of time-steps, making the calculation




FIGURE 8.6: Interferometer sensitivity to phase noise with amplitude $A = \pi/10$ for a range of noise frequencies computed for an interferometer composed of rectangular $\pi/2$ and π pulses (solid red curve) and Gaussian pulses (dashed purple curve). (a) shows the amplitude profiles of all pulses in the interferometer (pulse separations are not to scale). (b) shows the phase response to noise. The dwell-time of the interferometer was $250 \mu\text{s}$ and the effective Rabi frequency was 200 kHz . The noise sample size was 500.

highly resource intensive³. Since we are only interested in the effect of phase noise on the interferometer, we consider only resonant interactions ($\delta = 0$).

The results for the $\pi/2 - \pi - \pi/2$ sequence and for a sequence composed entirely of Gaussian pulses are shown in Figure 8.6. Our approach reproduces the analytical results obtained in [94] for rectangular pulses and the high-frequency suppression obtained with Gaussian pulses is also evident. Our Gaussian pulses are truncated such that the total length of each pulse is 12 standard deviations of the width. Smoothing the mirror pulse alone results in little improvement to the response compared with the $\pi/2 - \pi - \pi/2$ sequence and smoothing the beamsplitters is required to achieve maximum suppression [96]. The low frequency behaviour remains identical for rectangular and Gaussian pulses and is due to the overall three-pulse structure of the Mach-Zehnder sequence. Gaussian pulses were simulated using piece-wise constant approximations to the ideal pulse shapes with 500 time-slices per pulse. As expected, the coarseness of these

³To model the effect of phase noise up to a frequency of $10^2 \Omega_{\text{eff}}$, we ensure each pulse waveform is subdivided such that the duration of each time-step is $< 1/(10^3 \Omega_{\text{eff}})$.

waveforms affects the results: if the number of time-steps used to realise the Gaussian pulse is small, then peaks in the response appear at high-frequencies. Similarly, if the truncation of the Gaussian is changed such that the length is reduced to 6 standard deviations instead of 12, for example, the response becomes similar to that for the rectangular sequence at frequencies greater than $10 \Omega_{\text{eff}}$. This suggests that to see the full benefit of using Gaussian pulses, the individual waveforms must be truncated such that the pulse lengths are as large as possible.

8.3.1 Phase noise and GRAPE sequences

We now examine the effect of phase noise on optimised phase-modulated GRAPE sequences. We also wish to compare the performance of constant-amplitude GRAPE pulses with those in which the amplitude profile has been smoothed. There are several approaches we could take to obtain pulses with smooth amplitude profiles. One option is to optimise only the phase of our pulses within a user-defined smooth amplitude envelope. Another is to optimise both the pulse phase and amplitude simultaneously but impose penalties which enforce waveform smoothness, limit the maximum amplitude during the pulse, and fix the initial and final points in the waveforms to be 0. We adopt the first strategy because it makes the optimisation procedure simpler and allows us to investigate the effect of pulses with a well-defined and reproducible amplitude profile. We evaluate the effect of a smooth “Gaussian-flat” envelope, defined by the convolution of a Gaussian and rectangular profile. The resulting pulse has an approximately rectangular profile but with a gradual rise and fall.

Figure 8.7 shows the simulated interferometer response, σ_{Φ}^2 , to phase noise for two different GRAPE flip-reverse sequences: a constant rectangular amplitude pulse and a pulse with a smooth amplitude profile. The duration of each pulse was $12t_{\pi}$ and the number of time-steps in each pulse was 1000. The pulses in each sequence were optimised using fidelities from Equations 8.7 and 7.3 for the beamsplitter and mirror pulses respectively. Each pulse was optimised for pulse-length errors of $\pm 10\% \Omega_{\text{eff}}$, where $\Omega_{\text{eff}} = 200$ kHz, and off-resonance errors of $\delta = \pm 0.5 \Omega_{\text{eff}}$ (approximately corresponding to the 4σ width of a $4 \mu\text{K}$ sample of ^{85}Rb).

The behaviour of both the constant-amplitude and Gaussian-flat amplitude sequences is similar to that of the rectangular sequence - demonstrating that while GRAPE pulses are not necessarily more susceptible to phase noise than conventional rectangular pulses, smoothing the amplitude profile alone is no guarantee of noise suppression. We therefore repeated the optimisation and imposed a smoothing penalty on the phase profiles proportional to the squared difference between adjacent phase values. When the derivative-norm-square (DNS) smoothing penalty is applied in combination with a smooth Gaussian-flat amplitude profile, the resulting GRAPE sequence demonstrates a




FIGURE 8.7: Interferometer sensitivity to phase noise with amplitude $A = \pi/10$ for a range of noise frequencies computed for interferometers composed of flip-reverse GRAPE sequences. The dot-dashed blue curves represent a sequence where the amplitude remains constant, and the solid orange curve represents a phase-modulated pulse sequence optimised using the Gaussian-flat amplitude profile. There was no smoothing penalty applied to the phase profiles. The dwell-time of the interferometer was $250 \mu\text{s}$ and the effective Rabi frequency was 200 kHz. The noise sample size was 500.

suppression of phase noise at frequencies above the Rabi frequency, though not to the same extent as the Gaussian pulses. This is depicted in Figure 8.8.

We observe a peak in the response near a noise frequency of $80 \Omega_{\text{eff}}$, which we believe is due to the number of time-slices in the GRAPE pulses. For example, if we replace each of the 1000-step pulse waveforms with a smooth 10,000-step spline interpolation, this peak in the response vanishes. Conversely, if we repeat the optimisation with pulses of the same duration but only 200 time-steps, we observe multiple peaks in the response at frequencies beyond $10 \Omega_{\text{eff}}$, suggesting that the number of time-steps used to build the pulses plays a significant role in the suppression of high-frequency phase noise.

These results represent a preliminary investigation into the potential of phase noise suppression with optimised pulses. In conclusion, constant-amplitude phase-modulated GRAPE pulses do not necessarily result in a worse phase noise susceptibility when compared with the conventional rectangular sequence, and smoothing the amplitude and phase profiles can suppress phase noise at frequencies above the Rabi frequency which may be useful if the experimental noise spectrum includes high-frequency phase noise. Future work will focus on building noise robustness into the optimisation directly by modifying the fidelities to suppress noise given knowledge of the experimental noise spectrum.

8.4 Preliminary experimental results

Using fidelity measures for optimal beamsplitter and mirror pulses (Equations 8.7 and 7.3, respectively) we optimised all three pulses of the Mach-Zehnder interferometer sequence for an atomic sample with a temperature of $120 \mu\text{K}$ and a coupling strength variation of $\pm 10\% \Omega_{\text{eff}}$. The resulting pulse profiles are shown in Figure 8.9. The phase sequence of the final pulse was taken to be the inverted and time-reversed profile of the first according to the flip-reverse design procedure outlined in Section 8.1.1. We expect our optimal Mach-Zehnder sequence of pulses to be capable of maintaining a higher contrast than conventional rectangular pulses despite significant variations in detuning and Rabi frequency in the atomic cloud.

We have started to test these three-pulse interferometer sequences in our experimental apparatus, comparing the performance with that obtained using a sequence of conventional rectangular $\pi/2$ and π pulses. Interferometer fringes are obtained by scanning the phase offset ϕ_{bs} applied to the final pulse in the sequence (shown in Figure 8.10). The dwell-time between pulses is limited, at present, to $100 \mu\text{s}$ by suspected phase noise between the counter-propagating beams that also limits the overall contrast. The initial results are promising, and the relative improvement in contrast provided by the GRAPE-optimised sequence as the cloud temperature is varied is shown in Figure 8.11. GRAPE improved the contrast of the fringes at all temperatures investigated and, in a




FIGURE 8.8: Interferometer sensitivity to phase noise for different noise frequencies computed for interferometers composed of GRAPE sequences with the DNS smoothing constraint applied to the phase profiles. The dot-dashed blue curves represent a sequence where the amplitude remains constant, and the solid orange curve represents a phase-modulated pulse sequence optimised using the Gaussian-flat amplitude profile. The interferometer dwell-time was $250 \mu\text{s}$ and the effective Rabi frequency was 200 kHz . The noise sample size was 500 and the noise amplitude was $A = \pi/10$.




FIGURE 8.9: Optimisation results for a PP 90° beamsplitter pulse (a), and an anti-symmetric UR 180° mirror pulse (b). The pulse profiles are plotted against time as a fraction of the duration of a rectangular π pulse, t_π . 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 $120\ \mu\text{K}$ and a range of coupling strengths of $\pm 10\%$ Ω_{eff} . Ω_{eff} , or the effective Rabi frequency, was $420\ \text{kHz}$.

hot $94(4)\ \mu\text{K}$ 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 8.12. This enhancement is also apparent in the uncertainties in the phases of the fitted sinusoids. At $94(4)\ \mu\text{K}$ the average uncertainty in the fitted phase due to the GRAPE sequence was a factor of 2.9 smaller than the uncertainty in the fitted phase from the rectangular pulse sequence.

Applying the flip-reverse operation to obtain the final beamsplitter is necessary to observe an increase in contrast with GRAPE. For example, at a temperature of $\sim 25\ \mu\text{K}$, the contrast following the full antisymmetric flip-reverse sequence depicted in Figure 8.10 was higher by a factor of 2 than the contrast following a sequence which was in all respects identical but where the flip-reversal procedure was not applied to the final beamsplitter. In this case, the sequence will lead to a large variation of the interferometer phase with detuning thereby causing atoms with different velocities to exit the interferometer with different phases and the interference to wash out when the signal is averaged over the ensemble.

Some experimental evidence suggests that GRAPE sequences are less susceptible to drifts in offset and phase of the interferometer fringes. The variation of the fringe offsets and phases from their respective means for a range of temperatures for GRAPE and rectangular interferometers is shown in Figures 8.13b and 8.13c, 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 which could




FIGURE 8.10: Illustration, not to scale, of an optimised interferometer sequence. Three pulses, as depicted in figure 8.9, are combined, with a time τ_{dwell} between the light of one pulse turning off and the subsequent pulse turning on. Fringes are generated 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.

be caused by an offset of the laser detuning from the centre of the atomic momentum distribution [109]. It is notable that the GRAPE interferometer appears less susceptible to this shift.

Another route of inquiry will be to explain why the increase in contrast is quite so large 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 [1], 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. When this constraint was met, the contrast improvement far exceeded that of replacing just the mirror, or indeed the beamsplitters, in isolation.

8.5 Conclusions

We have introduced and experimentally demonstrated a design for optimal pulses that improves the contrast of the three-pulse Mach-Zehnder sequence. We expect such optimal pulse sequences to have applications in improving the sensitivity and robustness of atom interferometric sensors 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.




FIGURE 8.11: 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 8.12: 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.




FIGURE 8.13: 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 8.11. The vertical error bars represent the fitting uncertainty of the sinusoidal curves to each experimental run; for the contrast and fringe offset measurements these have magnitudes $\sim \pm 0.005$, while the phase becomes more uncertain with reduced contrast. Consequently, at $25 \mu\text{K}$ and $94 \mu\text{K}$ the phase variation errors for GRAPE pulses are $\sim \pm 0.04$ rad and $\sim \pm 0.1$ rad respectively, while those for rectangular pulses are $\sim \pm 0.07$ rad and $\sim \pm 0.3$ rad. Data points for several experimental runs at each temperature are shown, and their vertical spread gives an indication of the variation of each parameter over longer time-scales measured in tens of minutes.

Specifically, we have detailed a 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. We have discussed how a careful choice of pulse fidelities that preserve the interferometer phase across an ensemble is necessary to prevent the interference washing out when measuring the ensemble-average signal. Also, we have highlighted how restricting the symmetry of the mirror pulse leads to a tolerance of inhomogeneities in detuning and Rabi frequency.

We have shown up to a threefold contrast improvement in a proof-of-principle interferometer with hot $94(4) \mu\text{K}$ atoms, although our current investigations have been limited by experimental phase noise. Although the interferometer dwell-time is limited to $100 \mu\text{s}$ at present, our technique can mitigate errors in coupling strength and detuning that are often present in more sensitive interferometers, for example those which increase the scale-factor by increasing the momentum splitting of the diffracting wave-packets [42, 46].

Future work will extend our experimental study of optimised interferometer sequences to test their efficacy and robustness when experimental noise is no longer such a limiting constraint.

Chapter 9

Bi-selective pulses for large-area atom interferometry

Most atom interferometers for inertial sensing use Bragg [89] or Raman [36, 92] transitions driven by counter-propagating laser pulses as the beamsplitters and mirrors that split and direct the atomic wave-packets. The two-photon recoil accompanying these stimulated scattering processes imparts a momentum difference of $2 \hbar k$ between the two interferometer paths, where k is the single-photon wave number. As with an optical interferometer, the measurement sensitivity depends upon the spatial area enclosed. This is proportional to the momentum separation, and can hence be increased by using additional mirror pulses - known as augmentation pulses - to impart further impulses to the atomic wave-packets forming a large-momentum-transfer (LMT) interferometer [40].

In practice, inhomogeneities such as variations in beam intensity and atomic velocity reduce the fidelity of the augmentation pulses. The accrued effect of such imperfections limits the fringe visibility and can reverse the LMT sensitivity gains [40, 42, 46]. Cooling the atoms or filtering their velocity distribution can improve the initial fidelity at the expense of a weaker signal, but a velocity spread is inevitable in LMT interferometry because of the increasing momentum separation introduced between the primary interfering arms. Techniques such as increasing the Rabi frequency [174] and employing beam shaping optics [110] increase the experimental complexity and power requirements of the interferometer.

High-order Bragg diffraction pulses [89, 175, 176] and Bloch oscillations within optical lattices [91, 177, 178] can impart multi-photon recoils with a single pulse, and have allowed interferometry with wave-packets separated by as much as $408 \hbar k$ [179, 180]. However, both techniques require sub-recoil cooling or velocity selection [90, 181], limiting the signal-to-noise ratio (SNR). Furthermore, Bloch oscillations require long pulses that can constitute a considerable fraction of the total interferometer duration and are therefore sensitive to laser intensity and wavefront inhomogeneities [179].

An alternative approach is to seek robust Raman augmentation pulses that retain high fringe visibility for large momentum separations while enabling the use of large thermal atomic samples and electronic state read-out adopted by many inertial sensor prototypes [39, 182]. Composite [45, 111] and adiabatic rapid passage (ARP) [49] Raman pulses have been used to achieve momentum splittings of $18 \hbar k$ [42] and $30 \hbar k$ [46] respectively. Both these techniques can increase the velocity acceptance of the augmentation pulses at the expense of an increase in pulse duration, but in practice fringe visibility was lost after 4-7 augmentation pulses.

In this chapter, we apply optimal control to the specific case of LMT interferometry. We present designs for the augmentation pulses of LMT atom interferometers that maintain their fidelity as the wave-packet momentum difference is increased. We do so by tailoring these “bi-selective” pulses using optimal control to the evolving bi-modal atomic momentum distribution, which should allow greater interferometer areas and hence increased inertial measurement sensitivity, without requiring elevated Rabi frequencies or extended frequency chirps. This technique is a departure from previous approaches in which the same robust pulse has been used for every augmentation pulse in the interferometer sequence [40, 42, 46], ultimately limiting the achievable momentum splitting to the velocity acceptance of the individual pulse.

Using a previously validated model [46], we have simulated the application of our pulse designs to large-momentum-transfer atom interferometry using stimulated Raman transitions in a laser-cooled atomic sample of ^{85}Rb and compared their performance with the best composite and ARP schemes. After the wave-packets have separated by 42 photon recoil momenta at $1 \mu\text{K}$, our pulses maintain a fringe contrast of 90% whereas, for adiabatic rapid passage and conventional π pulses, the contrast is less than 10%. Furthermore, we show how bi-selective pulses can tolerate variations in Rabi frequency that occur throughout an interferometer and how the technique may be adapted to suppress the detrimental off-resonant excitation that limits other broadband pulse schemes. The bi-selective technique has broad applicability because many interferometer arrangements are already set up to implement similar modulation sequences and, as the algorithm optimises tolerance of variations, the designs do not depend critically upon experimental parameters.

The material in this chapter is adapted from [4].

9.1 LMT using robust pulses

Figure 9.1 shows the space-time diagram of a typical Raman LMT interferometer. While the standard $\pi/2 - \pi - \pi/2$ interferometer sequence produces a momentum splitting between the interferometer arms of $\hbar k_{\text{eff}}$, where $k_{\text{eff}} = 2k$ is the effective wavevector of the atom-optics, a greater momentum splitting can be achieved by augmenting the




FIGURE 9.1: Positions of atomic trajectories (left) as a function of time within a LMT interferometer sequence showing the case where the beamsplitter and mirror operations are extended by a sequence of 4 augmentation pulses (18 $\hbar k$ trajectory). The initial momentum distribution (right), represented by the red shaded region, is separated into two arms. The momentum distribution seen by each pulse throughout the LMT sequence for ^{85}Rb atoms with a Maxwell-Boltzmann temperature of $\sim 1 \mu\text{K}$ is shown by the blue shaded regions on the right. As more pulses are added to increase the interferometer area, the separation in the resonance conditions for the arms begins to exceed the velocity acceptance of a π -pulse, shown bottom right for a Rabi frequency of 200 kHz.

mirror and beamsplitter pulses with additional pulses with alternating wavevectors [40]. These augmentation pulses are designed to transfer atomic population between two internal states ($|g\rangle$, $|e\rangle$) and increase the momentum splitting between the wave-packets in the interferometer. By extending the beamsplitter operation by N augmentation pulses, and the mirror operation by $2N$ augmentation pulses, the momentum splitting may be increased to $(2N + 1)\hbar k_{\text{eff}}$, thereby increasing the intrinsic phase sensitivity of the interferometer.

The increased momentum splitting between the interferometer arms introduces differential Doppler shifts between them, quantised in multiples of the two-photon recoil shift $\delta_{\text{recoil}} = \hbar k_{\text{eff}}^2 / 2m$, that depend on the atomic mass m . Indeed, this is automatic, since the Doppler shift and inertial sensitivity have the same origin. For the n th augmentation pulse ($n = 1, 2, 3, \dots, N$), the arms of the interferometer have Raman resonance conditions separated by $4n\delta_{\text{recoil}}$ [42, 44, 50] (for Raman transitions on the ^{85}Rb D2 line, $\delta_{\text{recoil}} \approx 2\pi \times 15.4 \text{ kHz}$). The distribution of detunings for the n th augmentation pulse may thus be visualised as two Gaussian distributions, corresponding to the Maxwell-Boltzmann temperature of the atomic source, separated in frequency space by $4n\delta_{\text{recoil}}$.

As illustrated in Figure 9.1, the augmentation pulses are only efficient as long as this split distribution fits within the velocity acceptance of the pulses. For conventional π -pulses this can span just a few recoil momenta¹, fundamentally limiting the momentum transfer achievable before fringe visibility is lost [40, 42].

Atom-optics that are robust to off-resonance errors and hence possess large velocity acceptance are therefore of particular interest in LMT applications. ARP and composite pulses can achieve efficient state transfer in the presence of large detuning offsets, making them potentially attractive choices for LMT augmentation pulses. Butts *et al.* [42] employed the WALTZ composite pulse [43] to improve the contrast within a LMT interferometer using a $0.5\ \mu\text{K}$ sample of ^{133}Cs atoms obtained using laser-cooling and velocity selection. WALTZ enabled a momentum splitting of $18\ \hbar k$, increasing the intrinsic sensitivity over that obtained using simple π pulses. However, visibility was lost after 4 augmentation pulses and when WALTZ was used in a similar experiment with a hotter atomic sample, there was no improvement in contrast compared with that achieved using rectangular π pulses [44]. The ARP tanh/tan pulse was implemented in a large-area atom interferometer by Kotru *et al.* [46] increasing the interferometer contrast over that obtained using simple π pulses and achieving a momentum splitting of $30\ \hbar k$ in a $9\ \mu\text{K}$ ^{133}Cs atomic sample.

Although ARP can obtain impressive population transfer efficiency, with a detuning bandwidth that increases with the pulse duration, its potential utility in interferometry is limited by the dynamic phase imprinted on the diffracting wave-packets and the variation in effective time origin [44, 48–50]. The dynamic phase depends on the optical intensity, and rapid dephasing is therefore inevitable when the atom cloud expands through variations in laser intensity. In practice this effect is limited because ARP pulses applied in quick succession approximately cancel the dynamic phase, but it leads to a trade-off between longer and theoretically more efficient pulses, and dephasing caused by imperfect dynamic phase cancellation when the beam quality is non-ideal [46]. There is therefore a need for pulses which match or improve upon ARP in terms of state transfer efficiency, but which are robust to temporal variations in the Rabi frequency during the interferometer sequence.

LMT interferometry requires augmentation pulses that provide efficient population transfer across the atom cloud and throughout the sequence. To maintain sensitivity and prevent loss of fringe visibility, the interferometer sequence itself should impart the same phase to all atoms, and all atoms should have the same sensitivity to the external influence being measured. Pulses need not satisfy these conditions individually, provided that subsequent cancellation achieves them for the sequence as a whole.

¹The velocity acceptance of the π pulse is determined by the Rabi frequency Ω_R . For example, when $\delta = \Omega_R/3$, the state-transfer efficiency of the π pulse is approximately 90%. For a feasible experimental Rabi frequency of 200 kHz, this corresponds to $\delta \approx 4\delta_{\text{recoil}}$ for ^{85}Rb .

9.2 Bi-selective pulses

In this section we introduce the concept of “bi-selective” augmentation pulses for LMT atom interferometry. Specifically, we discuss how these pulses may be obtained using optimal control and introduce figures of merit that characterise their performance.

Consider the following two possible fidelity measures for augmentation pulses,

$$\mathcal{F}_{\text{square}} = |\langle e|\hat{U}|g\rangle|^2 \quad (9.1)$$

$$\mathcal{F}_{\text{real}} = \text{Re}(\langle e|\hat{U}|g\rangle). \quad (9.2)$$

Both of these measures of fidelity, if maximised, will yield pulses that efficiently transfer population from one basis state to the other, the essential requirement of an augmentation pulse. If the fidelities are averaged over an ensemble of detunings and a range of Rabi frequencies, the resulting pulse will be made robust to these specific errors. Maximising $\mathcal{F}_{\text{square}}$ leads to pulses where the phase of the overlap $\langle e|\hat{U}|g\rangle$ is unimportant but the quantum state is rotated by 180° from pole to pole on the Bloch sphere. Conversely, maximising $\mathcal{F}_{\text{real}}$ leads to pulses where the phase of the overlap $\langle e|\hat{U}|g\rangle$ is well-defined.

If the phase of the overlap $\langle e|\hat{U}|g\rangle$ varies from pulse to pulse across an atomic sample in an LMT sequence, the resulting interference fringes can be washed out when the contributions from each atom in the cloud are averaged at the end of the interferometer. Providing the Rabi frequency does not change much between pulses, these phases approximately cancel in LMT interferometers because the augmentation pulses occur in pairs. However, if the Rabi frequency varies temporally throughout the interferometer, rapid dephasing can occur. This is readily observed with ARP, which requires a high degree of cancellation in the phase factors introduced by each pulse in the sequence [46, 48, 50]. Optimising $\mathcal{F}_{\text{real}}$ for a range of Rabi frequencies will mean the phase of the overlap $\langle e|\hat{U}|g\rangle$ remains fixed even if the coupling strength varies within that range². Therefore, we expect $\mathcal{F}_{\text{real}}$ to yield augmentation pulses that are insensitive to dephasing caused by temporal changes in the Rabi frequency due, for example, to noisy beam intensity profiles and/or the ballistic expansion of the atom cloud.

For the nominal three-pulse $\pi/2 - \pi - \pi/2$ interferometer, the Raman detunings arise from the velocity components parallel to the Raman beam axis and hence the temperature of the atomic cloud. As the sequence is extended by augmentation pulses with alternating effective wavevector directions, the resonant frequencies of each interferometer arm separate due to the imparted photon momentum (shown on the right-hand side of Figure 9.1). This means the detuning distribution splits in two, and continues to separate as more momentum is imparted. Optimal control can optimise a pulse for any assumed velocity distribution [183]. We use this feature to design bi-selective pulses,

²Fixing the phase of the overlap $\langle e|\hat{U}|g\rangle$ also fixes the phase of the overlap $\langle g|\hat{U}|e\rangle$ because $\langle g|\hat{U}|e\rangle = e^{i\pi} \langle e|\hat{U}|g\rangle^*$ (Equation 2.52).

tailored to yield efficient population transfer only in the frequency ranges occupied by the primary interferometer arms. By optimising each pulse in the sequence individually, as opposed to employing a single augmentation pulse design throughout the interferometer [40, 42, 46], we can therefore take advantage of the fact that the transfer efficiency in the frequency space in between the two arms is unimportant. This allows bi-selective pulses to be shorter than ARP, composite, or optimised shaped pulses of the equivalent efficiency.

A LMT interferometer of ‘order’ N includes a total of $4N$ augmentation pulses in order to separate and recombine interferometer arms whose momenta differ by $(2N + 1)\hbar k_{\text{eff}}$ resulting in a corresponding increase in phase sensitivity, although the separation is often also quoted for Raman interferometers in units of the single-photon recoil momentum $\hbar k \approx \hbar k_{\text{eff}}/2$. In a LMT interferometer with optimised bi-selective pulses, each of the $4N$ augmentation pulses will be one of N individually optimised pulses indexed $n = 1, 2, \dots, N$, concatenated in the manner shown in Figure 9.1 and optimised to be resonant with two Doppler momentum distributions centred about $\delta = \pm 2n\delta_{\text{recoil}}$.

We express the bi-selective fidelity measure for the n th augmentation pulse as

$$\mathcal{F}_n^A = \sum_{\delta, \Omega_R \in L_n} \mathcal{F}(\delta, \Omega_R) + \sum_{\delta, \Omega_R \in U_n} \mathcal{F}(\delta, \Omega_R), \quad (9.3)$$

where $\mathcal{F} = \mathcal{F}_{\text{real, square}}$ is the single-atom augmentation pulse fidelity measure (Equation 9.1 or 9.2) and L_n and U_n are the ensembles representing the atomic frequency distributions of the lower and upper arms during the n th augmentation pulse respectively. In order to normalise the fidelity measure such that the maximum value is unity, we divide Equation 9.3 by the number of detunings and amplitude errors included in the entire ensemble.

We compose the detuning ensembles for the lower and upper interferometer arms for the n th augmentation pulse from two uniform discrete distributions centred at $\pm 2n\delta_{\text{recoil}}$. The ensembles used in the optimisation should represent the velocity distribution of the atoms. However, in order to reduce computation time we approximate the true distribution using two uniform distributions each with a sample size of 20. The range spanned by each distribution is given by 4 standard deviations of detuning arising from the velocity distribution along the Raman beam axis, which we assume follows a Maxwell-Boltzmann distribution.

Our LMT pulses are individually optimised, which assumes there are no correlations in residual errors between pulses. However, it may be possible to develop alternative measures of performance which reflect the fidelity of the entire interferometer and allow pulses within a sequence to be optimised cooperatively [184], for example by compensating each other’s imperfections.

An alternative approach to achieve bi-selectivity involves concatenating two inversion pulses with the appropriate frequency shifts to address each interferometer arm separately. While this makes the individual optimisations simpler, the resulting pulses are longer, give different velocity classes different pulse origins, and lose fidelity as the classes begin to overlap. A simple superposition of the two components would cure the problems of length and origin, though not overlap, and requires amplitude modulation in addition [185].

The concept of bi-selective pulses for LMT interferometry bears resemblance to the technique of band-selective pulses in NMR spectroscopy [137, 185–187], where pulses have been designed to excite or invert nuclear spins within single or multiple frequency bands, but suppress the response of spins outside the desired frequency range. Typically, these pulses require smooth waveforms because the response at large resonance offsets is determined by the Fourier transform of the pulse shape [188]. As a result, composite pulses, which are composed of concatenated sequences of constant-amplitude pulses with discrete phase shifts, lead to non-negligible excitation far off-resonance.

The suppression of the action of a pulse outside a specific frequency range may be useful in interferometer geometries where, for example, the counter-propagating Raman beams are obtained by retroreflection of a single beam with both Raman frequencies and there are necessarily four frequency components that may interact with the atom cloud [108]. When there is no acceleration of the atoms along the beam axis then double-diffraction interferometry schemes can be employed that make use of all of the frequency components simultaneously [189, 190]. However, in vertically-orientated, ground-based atom interferometers, the Doppler shift caused by gravitational acceleration is commonly used to isolate a single frequency pair by shifting the other pair off resonance [42]. However, broadband pulses such as composite or ARP pulses often have non-zero transfer efficiency at large detunings [44, 188], meaning one must wait longer to ensure negligible excitations from the off-resonant pair, costing time which may otherwise be used to enhance the sensitivity.

By directly suppressing the transfer efficiency outside a specific range of frequencies, such off-resonant excitation with broadband pulses may be avoided. We achieve this suppression of unwanted excitation by modifying our bi-selective pulse fidelity measure (Equation 9.3), adding the following penalty term which represents the suppression band of detunings where the transfer efficiency should be minimised,

$$\mathcal{F}_{\text{suppress}} = \sum_{\Omega_R, \delta=\delta_{\min}}^{\delta_{\max}} |\langle g|\hat{U}|g\rangle|^2 + \sum_{\Omega_R, \delta=-\delta_{\min}}^{-\delta_{\max}} |\langle g|\hat{U}|g\rangle|^2. \quad (9.4)$$

Here, $\delta_{\min, \max}$ represent the initial and final Raman detuning values for the two suppression bands.




FIGURE 9.2: Bi-selective phase-modulated augmentation pulses optimised with the phase-sensitive fidelity measure $\mathcal{F}_{\text{real}}$ for a temperature of $1 \mu\text{K}$. The resulting phase profiles for the 1st, 4th, 7th, and 10th augmentation pulses are shown in the left panel. The right panel shows the simulated transfer efficiency of each bi-selective pulse (red solid curve), the ARP tanh/tan pulse of equivalent length (orange dashed curve), and the WALTZ composite pulse (purple dotted curve) as a function of the Raman detuning. The corresponding detuning distributions seen by each augmentation pulse for the primary interferometer arms are shown by the shaded regions. The range of detunings included in each optimisation is shown by the hatched regions.

9.3 Characterisation of bi-selective pulses

We have optimised phase-modulated, constant-amplitude, bi-selective augmentation pulses with GRAPE for LMT interferometers of orders up to and including $N = 10$, corresponding to a momentum separation of $42 \hbar k$. We have performed the optimisations using both the phase insensitive fidelity measure $\mathcal{F}_{\text{square}}$ and the phase sensitive fidelity measure $\mathcal{F}_{\text{real}}$, and have investigated their performance through simulation of the resulting interferometer contrast in a laser-cooled sample of ^{85}Rb at $1 \mu\text{K}$. In all optimisations, the timestep of the pulses was 25 ns, and the effective Rabi frequency was 200 kHz, meaning




FIGURE 9.3: Bi-selective phase-modulated augmentation pulses optimised with the phase-insensitive fidelity measure $\mathcal{F}_{\text{square}}$ for a temperature of $1 \mu\text{K}$. The resulting phase profiles for the 1st, 4th, 7th, and 10th augmentation pulses are shown in the left panel. The right panel shows the simulated transfer efficiency of each bi-selective pulse (blue solid curve), the ARP tanh/tan pulse of equivalent length (orange dashed curve), and the WALTZ composite pulse (purple dotted curve) as a function of the Raman detuning. The corresponding detuning distributions seen by each augmentation pulse for the primary interferometer arms are shown by the shaded regions. The range of detunings included in each optimisation is shown by the hatched regions.

the duration of a rectangular π pulse (t_π) should be $2.5 \mu\text{s}$ in the absence of any off-resonance or pulse-length errors. The pulses were optimised for an atomic temperature of $1 \mu\text{K}$, and a range of amplitude errors of $\pm 10\%$ the effective Rabi frequency. The length of the augmentation pulses was fixed to be $4t_\pi$. This duration allowed for a sufficiently good pulse to be obtained when using both fidelities, although optimising $\mathcal{F}_{\text{real}}$ typically requires a longer pulse to reach an equivalent fidelity to the phase-insensitive case because of the extra constraint on phase. The initial guess for the phase profile of the pulses was in each case a sequence of random phases.

The NMR spin simulation software suite for MATLAB, *Spinach* [125], and its optimal

control module, were modified to optimise bi-selective pulses using the L-BFGS GRAPE method [19]. Each pulse optimisation was set to terminate following 300 iterations or when either the norm of the fidelity gradient became smaller than 10^{-7} or the norm of the step size dropped below 10^{-3} . The resulting waveforms, optimising the fidelities $\mathcal{F}_{\text{real}}$ and $\mathcal{F}_{\text{square}}$, are shown in Figures 9.2 and 9.3 respectively. The waveforms are shown for pulses tailored to the velocity distributions expected for the 1st, 4th, 7th, and 10th augmentation pulses of the LMT sequence. Efficient population transfer is achieved in each case. Interestingly, smooth and symmetrical waveforms are found despite there being no constraint on symmetry or waveform smoothness included in the optimisation.

As n is increased, the magnitudes of the individual phase excursions within each pulse grows, and phase variations appear symmetrically about the temporal midpoint for the $\mathcal{F}_{\text{real}}$ pulses. Interestingly, we observe that for the $n = 1$ $\mathcal{F}_{\text{square}}$ optimisation, the solution strongly resembles the class of non-linear frequency swept pulses discussed in Chapter 7.

Figure 9.4 shows the phase profile and simulated transfer efficiency of the 10th augmentation pulse when the fidelity measure $\mathcal{F}_{\text{square}}$ is optimised. Results are also shown for the WALTZ composite pulse and the tanh/tan ARP pulse with durations of $4t_\pi$ and $8t_\pi$. The corresponding detuning distributions for the upper and lower interferometer arms are also shown as shaded regions. Whereas the efficiencies of the WALTZ and ARP augmentation pulses are limited by the finite velocity acceptances, the bi-selective pulse is tailored for the split momentum distribution of the two primary interferometer arms, allowing it to achieve a higher efficiency than a tanh/tan ARP pulse of twice its duration.

$\mathcal{F}_{\text{real}}$ pulses are designed to minimise variation in the phase of the overlap $\langle e|\hat{U}|g\rangle$ over the range of Rabi frequencies and detunings included in the optimisation ensemble. This property is illustrated in Figure 9.5, where we have simulated the phase of the overlap $\langle e|\hat{U}|g\rangle$ for the 10th $\mathcal{F}_{\text{real}}$ augmentation pulse as a function of amplitude error (pulse-length error) $\Omega_R/\Omega_{\text{eff}}$ and the Raman detuning (off-resonance error) δ . The results are also shown for the 10th $\mathcal{F}_{\text{square}}$ pulse, the tanh/tan ARP pulse, and WALTZ. The detuning was set to the centre of the Doppler-shifted resonance of the upper interferometer arm during the 10th pulse. Both the tanh/tan and $\mathcal{F}_{\text{square}}$ pulses exhibit a large variation of this phase with amplitude. We find that the $\mathcal{F}_{\text{real}}$ pulse, however, minimises variation in the phase over a range of amplitude errors larger than the design range ($\pm 10\% \Omega_{\text{eff}}$). This should make $\mathcal{F}_{\text{real}}$ pulses less susceptible to the dephasing that can occur with ARP and $\mathcal{F}_{\text{square}}$ pulses if the Rabi frequency varies from pulse to pulse across the atom cloud.

We have kept the amplitude constant when optimising the bi-selective fidelity measure $\mathcal{F}_n^A$ (Equation 9.3) because allowing the amplitude to vary does not lead to a higher terminal fidelity. When the amplitude is allowed to vary and the maximum number of iterations is extended, both the phase-sensitive and phase-insensitive fidelities lead to




FIGURE 9.4: Phase profile (a) and simulated transfer efficiency (b) for the 10th bi-selective pulse found optimising the fidelity measure $\mathcal{F}_{\text{square}}$ (blue solid line). The transfer efficiency for the ARP tanh/tan pulse (orange dashed line for duration $4t_\pi$ and green dot-dashed line for duration $8t_\pi$) and the WALTZ pulse (purple solid line) are shown in (c). The detuning distributions of the two primary interferometer arms, corresponding to a temperature of $1 \mu\text{K}$ during the 10th augmentation pulse are shown by shaded regions. The hatched regions represent the bands of detuning where the transfer efficiency is maximised using GRAPE.

solutions that approach constant-amplitude, purely phase-modulated pulses. However, when the penalty to suppress off-resonant excitation is included in the optimisation (Equation 9.4) we find this is not the case and that allowing amplitude modulation can lead to better solutions.




FIGURE 9.5: Phase of overlap $\langle e|\hat{U}|g\rangle$ for different pulses as a function of the Raman detuning (a) and amplitude error (b). The solid red and dot-dashed blue curves correspond to the 10th $\mathcal{F}_{\text{real}}$ and $\mathcal{F}_{\text{square}}$ bi-selective pulses, respectively. The orange dashed curve represents the tanh/tan pulse of the same duration ($4t_\pi$) and the purple dotted curve shows the results for WALTZ. The hatched regions represent the design ranges in detuning and amplitude errors over which the phase is minimised when using $\mathcal{F}_{\text{real}}$. In (a) $\Omega_R = \Omega_{\text{eff}}$ and in (b) the detuning corresponds to the centre of the Doppler-shifted resonance of the upper interferometer arm during the 10th augmentation pulse. The curves have been unwrapped and a uniform phase shift has been applied to enhance visual clarity.

9.3.1 Parabolic approximation of $\mathcal{F}_{\text{square}}$ bi-selective pulses

Optimal control can sometimes find particularly simple solutions which prompt us to wonder about their deeper significance. For example, many of the profiles found optimising $\mathcal{F}_{\text{square}}$ have a strong parabolic component corresponding to a frequency chirp - although we are far from the adiabatic regime, and some include a π phase step near the temporal midpoint. Consider the phase profile shown in Figure 9.4 for the $n = 10$ $\mathcal{F}_{\text{square}}$ pulse. How much of the bi-selective behaviour of this pulse is captured by these two features? To investigate this question we have fitted a parabola ($\phi = -a(t/t_\pi - b)^2 + c$) to this phase profile with a π jump midway and simulated the resulting transfer efficiency as a function of detuning. We have ignored the initial phase jump in the parabolic fit. The results are shown in Figure 9.6. The fitting parameters were $a = 0.88\pi$, $b = 2.10$, and $c = 0.31$.

We find that the parabolic approximation captures most of the bi-selective behaviour: the π -flip leads to zero transfer efficiency on resonance, and the parabolic curve leads to good - although imperfect - state transfer at two symmetric off-resonant frequencies. Figure 9.6 also shows the resulting waveform when this parabolic fit is used as an initial guess in the optimisation. Following 300 iterations, the basic parabolic shape and phase jump remain, but symmetric variations are superimposed on the profile that flatten out and improve the efficiency off-resonance. The initial rapid phase sweep present in




FIGURE 9.6: Phase profiles ((a) and (c)) and simulated transfer efficiency ((b) and (d)) corresponding to $n=10$ $\mathcal{F}_{\text{square}}$ bi-selective pulses. The orange dashed curve represents a simple parabolic approximation to the $\mathcal{F}_{\text{square}}$ pulse from Figure 9.5 and captures most of the spectral behaviour. The dot-dashed black curve represents the results of an optimisation using the parabolic approximation as a starting point.

the original waveform seems not to affect the pulse performance: as the phase is quickly swept from 0 to $\approx -3\pi$, the field vector rapidly rotates around the z -axis thereby keeping atomic states near the north pole of the Bloch sphere. When the pulse is simulated without this initial phase sweep, the transfer efficiency as a function of detuning remains almost unchanged.

The simulated trajectory of the Bloch vector under the evolution of the $n = 10$ $\mathcal{F}_{\text{square}}$ pulse from Figure 9.4 is shown in Figure 9.7. The trajectory is shown for a resonant atom, and two off-resonant detunings (± 300 kHz). The phase jump midway through the pulse leads the resonant atom to approximately retrace its steps during the second half of the pulse. The positively detuned atomic trajectory is inverted during the first half of the pulse, then kept near the pole for the second half. This happens in reverse for the negatively detuned atom, and the pulse does most of its work inverting the atomic state in the second half of the pulse.

Although we are able to gain some insight into how this particular class of bi-selective pulse works, there are remaining questions. For example, although we can identify that specific sections and features of the pulse waveforms act on atoms with different resonant frequencies leading to the desired bi-selectivity, it is unclear how features in the waveforms give robustness to fluctuations in amplitude and detuning. This will be the subject of future investigations.




FIGURE 9.7: Trajectories of the state vector under the $n=10$ $\mathcal{F}_{\text{square}}$ pulse in Figures 9.4 and 9.6 for an atom with $\delta = +300$ kHz (a), $\delta = -300$ kHz (b), and $\delta = 0$ (c).

9.4 LMT contrast

We have investigated the performance of bi-selective pulses through simulation of the resulting interferometer contrast in a laser-cooled sample of ^{85}Rb at cloud temperatures of 1 and 5 μK . We have simulated LMT interferometers of different LMT orders, and compared the performance using our bi-selective augmentation pulses with that using rectangular π and tanh/tan ARP augmentation pulses. The contrast in each case is averaged over a thermal distribution of atoms and a uniform distribution of Rabi frequencies that is either kept temporally constant – to represent a non-uniform laser intensity distribution – or else varied from pulse to pulse – to represent the motion of atoms across such a distribution. Following the approach taken by Kotru [44], only the primary interfering paths are included in the calculation to reduce computation time. The numerical model is outlined in Appendix B.




FIGURE 9.8: Simulated interferometer contrast for a cloud of ^{85}Rb at $1\ \mu\text{K}$ as a function of LMT order for different pulse sequences. The contrast is numerically averaged over Raman detunings due to a velocity distribution corresponding to a temperature of $1\ \mu\text{K}$ and a uniform distribution of static Rabi frequency errors of $\pm 10\%$ (a) or a random temporal variation of Rabi frequency errors of $\pm 10\%$ from pulse to pulse (b). Contrast values for bi-selective GRAPE pulses found with fidelities $\mathcal{F}_{\text{square}}$ and $\mathcal{F}_{\text{real}}$ (blue dotted line with diamonds and red solid line with diamonds respectively) are shown. The results for π augmentation pulses (black dotted line with rectangles) and the tanh/tan ARP pulse (dashed orange line with circles) are also shown.

The simulation results for an atomic temperature of $1\ \mu\text{K}$ are shown in Figure 9.8. To meaningfully compare the performance of the GRAPE bi-selective pulses with ARP, the total duration of the sequences is kept the same. GRAPE bi-selective pulses far outperform basic rectangular π pulses, and tanh/tan ARP pulses of equivalent duration. When no thermal expansion is modelled, the dynamic phase cancellation of the ARP pulses is perfect, meaning the reduction in contrast with increasing LMT order is purely due to the limited velocity acceptance of the pulses as the arms separate in detuning. For the ARP contrast to match that achievable with bi-selective pulses, the ARP pulse duration must be increased, thus increasing the susceptibility to spontaneous emission (not modelled currently) and dynamic phase dephasing [46]. This highlights one key advantage of our adaptive approach.

If the Rabi frequency varies from pulse to pulse for different atoms, the dynamic phase

imprinted on the atoms by ARP is not cancelled in the interferometer, meaning different atoms obtain different phases and the interference is washed out. Although ARP pulses have a velocity acceptance which increases with pulse duration, longer ARP pulses become more susceptible to errors in dynamic phase cancellation [46].

When the Rabi frequency is instead varied randomly between pulses in the simulation - within the $\pm 10\%$ interval for which the pulses are optimised - to emulate noise in the Raman beam intensity profile (Figure 9.8b), the bi-selective pulses found optimising the phase-sensitive fidelity measure $\mathcal{F}_{\text{real}}$ are able to maintain significantly higher contrast than other pulse sequences. This is because the phase variation of the wave-packets is minimised with respect to variations in the Rabi frequency. This is not the case with pulses found using $\mathcal{F}_{\text{square}}$ or ARP, and the contrast is significantly reduced in such cases. While we vary the Rabi frequency randomly from pulse to pulse, the resilience to the variations is quite general and would extend to systematic changes in Rabi frequency throughout the interferometer due to, for example, expansion of an atom cloud within a Gaussian beam.

We have repeated the bi-selective optimisation and contrast simulation for a hotter atomic temperature of $5 \mu\text{K}$ and for a longer pulse duration of $5t_\pi$. The results are summarised in Table 9.1. We also show the results for the WALTZ composite pulse and phase-modulated pulses obtained using GRAPE maximising $\mathcal{F}_{\text{square}}$ for a large range of detunings centred on resonance. These non-selective pulses were optimised following the procedure outlined in [1] and Chapter 7 - to compensate the off-resonance errors arising from a $40 \mu\text{K}$ atom cloud with pulse-length errors in the range $\pm 10\% \Omega_{\text{eff}}$. Longer pulses can achieve higher terminal fidelities, and using the phase-sensitive fidelity measure $\mathcal{F}_{\text{real}}$ requires a longer duration than $\mathcal{F}_{\text{square}}$ to reach an equivalent terminal fidelity. This is shown most clearly when optimising for the hotter cloud. There is no guarantee that we have found the global maximum for each choice of duration and temperature. Even so, pulse sequences were found which outperform the composite and ARP alternatives tested. Although for a given pulse duration, the phase-insensitive fidelity measure yields the highest expected interferometer contrast, phase-sensitive bi-selective pulses are expected to provide robustness to Rabi frequency variations throughout the interferometer, and therefore may still be the best choice in practice.

It is important to highlight the expected difference in performance between pulses found using the bi-selective approach and the non-selective robust mirror pulses introduced in Chapter 7. The non-selective GRAPE augmentation pulse of duration $5t_\pi$ found optimising $\mathcal{F}_{\text{square}}$ yields a contrast of 0.74 following an $n = 10$ LMT sequence at a temperature of $1 \mu\text{K}$. As we would expect from the results in Chapter 7, this contrast is significantly larger than that expected following a sequence comprised of ARP tanh/tan pulses of the same length. However, the contrast using this particular non-selective GRAPE pulse rapidly decays when the LMT order is extended beyond $n = 10$. For example, when simulating an $n = 16$ sequence, we observe that the contrast decays to

TABLE 9.1: Simulated contrast values at temperatures of 1 and 5 μK for different momentum separations with no Rabi frequency variation between pulses. The π , WALTZ, and $\mathcal{F}_{\text{square}}$ non-selective pulses are previous results with either defined (π , WALTZ) or previously chosen (non-selective) durations. The results for the ARP and bi-selective pulses are presented for two different durations. Bold values indicate the best performing pulse at each temperature.

Pulse	Length (T/t_π)	1 μK Contrast			5 μK Contrast		
		10 $\hbar k$	26 $\hbar k$	42 $\hbar k$	10 $\hbar k$	26 $\hbar k$	42 $\hbar k$
π	1	0.50	-	-	0.35	-	-
WALTZ	3	0.73	0.51	0.00	0.72	0.41	0.00
ARP tanh/tan	4	0.95	0.62	0.06	0.88	0.52	0.04
	5	0.96	0.79	0.19	0.90	0.69	0.15
$\mathcal{F}_{\text{square}}$ non-selective	4	0.92	0.72	0.38	0.85	0.67	0.34
	5	0.95	0.88	0.74	0.89	0.82	0.64
$\mathcal{F}_{\text{real}}$ bi-selective	4	0.91	0.78	0.69	0.79	0.43	0.22
	5	0.95	0.87	0.82	0.85	0.60	0.50
$\mathcal{F}_{\text{square}}$ bi-selective	4	0.97	0.92	0.90	0.83	0.64	0.57
	5	0.96	0.94	0.93	0.89	0.81	0.76

0.01. By comparison, we optimised $4t_\pi$ $\mathcal{F}_{\text{square}}$ bi-selective pulses up to $n = 16$ and observed that the contrast obtained at $n = 16$ was 0.88.

As the bi-selective pulses are the result of individual optimisations, there is no need to fix the duration of each pulse to be the same. In fact, a more adaptive approach would be to choose the minimum pulse duration that obtained a given fidelity for each pulse in the sequence ($n = 1, 2, \dots, N$). This tactic could reduce the total sequence duration, because optimising the pulses for higher n (larger momentum separation) generally requires a longer pulse to reach the same fidelity.

9.5 Suppression of off-resonant excitation

When the Doppler shift from gravitational acceleration is used to discriminate between retroreflected frequency pairs in vertically-orientated interferometers, off-resonant excitations from broadband atom-optics can lead to unwanted double-diffraction [44, 108]. For example, following a 10 ms drop time [42], the resonance conditions of the two frequency pairs will shift by $2gk_{\text{eff}} \times 10 \text{ ms} \approx 500 \text{ kHz}$ for ^{85}Rb , where g is the local gravitational acceleration. Conventional robust pulses can result in a non-zero transfer at comparable detunings (Figure 9.9), potentially leading to double-diffraction if the unwanted frequency pair has not been shifted far enough from resonance [42, 44].

We have optimised bi-selective pulses which suppress the transfer efficiency outside the detuning bands of interest using a suitable modification of our bi-selective fidelity measure (Equation 9.4). Figure 9.9 shows the preliminary results and depicts a pulse that




FIGURE 9.9: Panels (a) and (b) show the phase and amplitude profiles for two bi-selective pulses maximising $\mathcal{F}_{\text{square}}$. The dashed blue curve shows the case where no suppression of off-resonant excitation is included in the optimisation and the solid red curve shows the case where off-resonant excitation is suppressed. The transfer efficiency of these two pulses as a function of detuning is shown in panel (c), and panel (d) shows the transfer efficiency of the WALTZ and tanh/tan pulses. The diagonally hatched regions coinciding with the shaded detuning distributions show the optimisation bands of detuning during the 8th augmentation pulse. The outer hatched regions show the range of the detunings for which the state transfer is suppressed ($\delta_{\min}, \delta_{\max} = 500, 2500$ kHz). The response is unconstrained outside the hatched regions.

effectively suppresses far off-resonant excitation and still yields efficient state-transfer for the two frequency separated interferometer arms.

In order to achieve a good level of suppression at large detunings we have found it necessary to allow both the phase and amplitude of the pulse to vary in the optimisation procedure. If the amplitude of the pulse is fixed, noisy phase-profiles are obtained such as those in Figure 9.10. Although these pulses obtain good bi-selective efficiency and good suppression of excitation within the suppression band of detuning, significant unwanted excitation is “pushed” outside the suppression band. These ripples far off-resonance are




FIGURE 9.10: Panels (a) and (b) show the phase and amplitude profiles for $n=10$ bi-selective pulses maximising $\mathcal{F}_{\text{square}}$ where we have added a penalty (Equation 9.4) to suppress excitation far off-resonance. The dot-dashed blue curve shows a purely phase-modulated pulse, the dashed red curve shows a phase-modulated pulse with a smoothing constraint, and the solid orange curve shows an amplitude and phase-modulated pulse with a smoothing constraint. The transfer efficiency of these pulses as a function of detuning is shown in (c) and (d). The amplitude and phase-modulated pulse yields the best bi-selectivity and off-resonant suppression.

removed by applying the derivative-norm-square (DNS) smoothing penalty, proportional to the difference between adjacent phase steps. Smoothing the pulse waveform removes far off-resonant excitation, meaning the suppression band of detuning can consist of fewer detunings.³ However, the highest bi-selective fidelities and best off-resonant suppression are found when the amplitude is allowed to vary *and* a smoothing constraint is applied. We also applied the spill-out norm-square (SNS) penalty to limit the maximum intensity during a pulse [23].

³Alternatively, a low-pass filter could be applied to the waveform at each iteration of the optimisation. This was the approach taken by Janich *et al.* [187] when optimising slice-selective pulses for NMR applications.

9.6 Conclusions

We have presented an optimised Raman pulse scheme for the augmentation pulses in large-area atom interferometry, whereby individually tailored Raman pulses found using optimal control techniques maintain resonance with the diverging wave-packets as the Raman detuning increases between the interferometer arms. We optimise our bi-selective pulses to provide maximum state-transfer while minimising variation in the interferometer phase across the atomic ensemble. The pulses can be made robust to large spatial and temporal variations in the Rabi frequency, potentially allowing LMT interferometry in non-ideal experimental environments such as those with warmer atom clouds and inhomogeneous laser beam fronts. Our simulations show that our pulses can maintain contrast at significantly higher momentum splittings than interferometers that have employed the best augmentation pulses demonstrated to-date, including the WALTZ composite pulse and the tanh/tan adiabatic rapid passage pulse, whose finite velocity acceptances limit the LMT momentum range.

For large LMT orders, our pulses can be considerably shorter than ARP or composite equivalents of the same efficiency, reducing the susceptibility to spontaneous emission which - although assumed to cause negligible decoherence in our simulations - can be significant in some experiments [40, 46, 191] where, for instance, the single-photon detuning of the Raman lasers is kept lower than the ground state hyperfine splitting in order to null differential light-shifts [192]. In such cases the relative improvements of bi-selective pulses with respect to equal-duration ARP pulses are unaffected and, in fact, using pulses that optimise the phase-sensitive fidelity measure $\mathcal{F}_{\text{real}}$ could allow such interferometers to suppress spontaneous emission by operating further from single-photon resonance⁴. We have also shown how the technique may be adapted to suppress deleterious off-resonant excitation, which can limit the applicability of broadband pulses to atom interferometer configurations.

Future work will involve extending our analysis of the pulse shapes to gain further insight into how and why they work, and implementing the scheme experimentally to obtain a real-terms sensitivity enhancement in an LMT interferometer with bi-selective pulses.

⁴The light-shift manifests as a detuning error which, if it differs from pulse to pulse due to intensity variations, can result in an accumulated phase uncertainty. Optimising pulses that maintain a fixed interferometric phase over a range of off-resonance and pulse-length errors ensures a first-order robustness to these variations, thus negating the need to operate close to single-photon resonance to negate the light-shift.

Chapter 10

Conclusions and outlook

In this thesis we have adapted NMR optimal control methods to design robust, high-fidelity Raman pulses for atom interferometry by treating the laser phase and in some cases intensity as controllable parameters. The resulting pulses compensate variations in atomic detuning and Rabi frequency that limit the sensitivity of interferometers using conventional $\pi/2$ and π pulses with constant laser phase and intensity. We have explored different individual pulse fidelity measures, the role of pulse symmetry in interferometry, the resilience of optimised pulses to laser phase noise, and developed pulse sequences tailored to improve the sensitivity of large-momentum-transfer (LMT) interferometers. Optimised pulses yield significant improvements in the fidelity of interferometric beam-splitter and mirror operations compared with $\pi/2$ and π pulses, and also with the best composite [42, 45], and adiabatic rapid passage (ARP) [46] pulses that have been applied in atom interferometers to-date. Importantly, we have also verified many of these pulse designs experimentally, finding an immediate threefold increase in fringe visibility in a three-pulse Mach-Zehnder interferometer using an inhomogeneous cloud of ^{85}Rb atoms at $\sim 100\ \mu\text{K}$. In the following, we discuss the implications of this work and outline directions for future investigation.

In many interferometers designed for sensing applications steps are taken to restrict inhomogeneities in the atom-light interaction such as those due to a spread in atomic velocity and/or variations in laser intensity. Often this involves reducing the number of atoms by selecting those with the correct positions, states, and velocities, and discarding the rest. However, such filtering processes lower the signal-to-noise ratio (SNR) of the measurement. Our optimised pulses mitigate these inhomogeneities and should therefore enable interferometry to be performed with high fringe visibility using less homogeneous atomic samples, for example at higher temperatures, and enable the use of more complex interferometer sequences. We expect these techniques will also enhance the robustness and sensitivity of new compact and portable atom-based sensors, which must tolerate larger variations in the atom-light interactions than lab-based interferometers.

A spread in resonant frequency within the atomic sample cannot always be reduced by filtering the cloud. For example, a change in resonant frequency is integral to the technique of LMT interferometry, which is therefore an important application of the work in this thesis. In Raman pulse LMT interferometers, additional light pulses are used to impart extra photon momentum kicks to diffracting atoms and increase the measurement scale-factor. This technique requires that the extra “augmentation” pulses maintain fidelity as the two arms of the interferometer separate in resonant frequency. When π pulses are used, their finite velocity acceptance leads to a rapid decay in fringe visibility, limiting the increase in sensitivity [40]. Optimised augmentation pulses are more robust to off-resonance errors than π , composite [42], and ARP pulses [46], and the bi-selective pulse technique, introduced in Chapter 9, is a tailored solution for LMT. By optimising each pulse in the sequence to track the changing resonant frequency of each interferometer arm, significantly more momentum may be imparted before visibility is lost. Bi-selective pulses should therefore enable large increases in the sensitivity of Raman pulse LMT interferometers, and an important next step is to test this idea experimentally.

We have also discussed how modifying the bi-selective fidelity measure allows us to optimise pulses which suppress atomic state transfer within unwanted frequency bands, a modification necessary to suppress deleterious double diffraction in specific interferometer arrangements. However, in some interferometers, for example those operating in micro-gravity environments, double diffraction can be a useful technique [189, 190], and optimal control could be deployed to optimise the mirror and beamsplitters operations which, due to the double-diffraction geometry, lead to population loss into unwanted diffraction orders [193]. Our application of optimal control demonstrates that we can tailor the velocity-selectivity of a pulse. In future work, this feature may be used to optimise cooling schemes, either by tailoring the pulses within existing interferometric [105] and algorithmic [107] cooling sequences, or by designing new ones.

Central to this work has been the question of which fidelity measure to use in the optimisation. We have explored fidelity measures for individual pulses, including those that minimise unwanted variations in interferometer phase with errors in detuning and Rabi frequency, and we have also begun to investigate the role of pulse symmetry in sequence design, allowing one pulse to correct errors due to another. In all these cases, we have optimised the pulses within interferometer sequences individually. However, it is reasonable to expect that better performance could be reached and possibly in a shorter time by optimising the entire pulse sequence in one go. Such “cooperative” multi-pulse optimisation has recently been developed for Ramsey and spin-echo type sequences in NMR [184, 194] with the resulting pulses yielding a reduction in the total duration of a sequence needed to obtain a given fidelity. However, the results of these studies do not directly transfer to atom interferometry. The goals of the experiments and the errors involved are quite different; in interferometry we care about measurement sensitivity,

and both the laser intensity and resonant frequency can change from pulse to pulse due to the expansion of the atomic cloud. New measures of performance that incorporate measurement sensitivity, scale-factor stability, and application-dependent trade-offs will therefore be needed to realise optimisation of the entire interferometer, and perhaps this will also involve lifting the constraint of a three-pulse sequence, allowing the computer complete freedom to vary the order and timings of the individual operations.

Designing pulse sequences that suppress time-dependent noise is another important problem and the sensitivity of many interferometers is limited by phase and vibrational noise [182]. Although some studies have shown that certain NMR pulse shapes can suppress different frequencies of laser phase noise [96], and we have shown that these results also apply to optimised pulses with a fixed smooth amplitude profile, no attempt has yet been made in atom interferometry to incorporate phase noise suppression into the pulse optimisation itself. Here, an analysis of the fidelity landscape of the control problem is likely to be important, as properties of this landscape in the vicinity of local maxima provide information about the tolerance of a pulse to noise in the control parameters [157].

10.1 Physical interpretation of optimised pulses

A fundamental question raised by this work is: how do the pulses found by numerical optimisation compensate errors? Although it is satisfactory to find that optimised pulses provide significant improvements in the fidelity and robustness of interferometer operations, we would ideally like to explain why the pulses take the smooth and symmetrical shapes that they do. Understanding how individual optimal control pulses work is an important problem in quantum control more generally [11], and while intuitive geometric explanations are available for the simplest pulse shapes, such as the original Levitt pulse [5] and the adiabatic frequency sweeps [6], analysing more complicated pulses is not straightforward. Some approaches based on a Fourier analysis of NMR pulse shapes [195] may provide insight, and an analysis of noisy pulse shapes can reveal surprising structure in the state evolution [196]. In this section, we discuss some interpretations of the pulse shapes presented in this thesis and provide suggestions for future investigations.

Many solutions found by optimal control exhibit smooth features and possess distinct symmetries. In Chapter 7, we showed that different pulse shapes may emerge for a given choice of fidelity measure when the initial guess for the pulse is varied: some solutions are symmetric, some highly asymmetric, and some resemble adiabatic frequency sweeps. The optimised frequency-swept pulses appear to work using a similar mechanism to ARP but are of low adiabaticity at the pulse edges, and this departure from adiabatic behaviour allows for superior off-resonance error compensation.

The variety of pulse shapes found suggests that multiple error compensation strategies are being exploited by the algorithm. We have begun to explore how the pulse shapes depend upon parameters of the optimisation such as the number of time-steps. It is possible that more insight into the underlying pulse mechanisms may be found by exploring the dependence of the optimised pulse shapes upon the range of detuning and amplitude errors that are used. For example, how do the pulse shapes change as the detuning or amplitude error range is varied?

Pulse symmetry is an important property of optimised pulses: it determines how the overall Bloch sphere rotation described by a pulse behaves as a function of resonance offset. For example, inversion pulses with fully symmetric phase profiles have a transfer efficiency that is identical for positively and negatively detuned atoms¹. This could explain why these pulses arise naturally when a symmetric detuning range is supplied in the ensemble. Symmetry can also be exploited in the design of pulses and pulse sequences. In Chapter 8 we explored the role of antisymmetry, and found that anti-symmetric mirror pulses and beamsplitters combined in a “flip-reverse” configuration allowed phase errors in the interferometer sequence to cancel.

For many of the pulse shapes a simple explanation of the underlying mechanism remains elusive. One approach to tackle this problem may be to parameterise the distribution of atomic states on the Bloch sphere for an ensemble of atoms with different detunings and Rabi frequencies. We could, for example, compute statistics such as the spread and elliptical orientation of this distribution of points on the Bloch sphere. The next step would be to see how these parameters change throughout the pulse. Do different sections of the optimised waveform correspond to different transformations of the distribution? For example, certain parts of a pulse may rotate and stretch the distribution to minimise the spread in interferometer phase or excitation probability. Such an approach may offer insight into the role of each part of a pulse, and thereby improve our understanding of how optimised pulses work.

¹This fact follows from the symmetry relationships listed by Kobzar *et al.* [10] and the form of the pulse propagator given in Equation 8.16.

Part V

Appendices

Appendix A

Mach-Zehnder interferometer output

In this appendix, we show for a general sequence of 3 interferometer pulses, represented by the propagators $\hat{U}_{i=1,2,3}$, that the probability for a given atom to exit the interferometer in the internal state $|e\rangle$ is given by

$$P_e = \frac{1}{2}(\mathcal{A} - \mathcal{B} \cos(\Phi + \Delta\phi)), \quad (\text{A.1})$$

and derive expressions for the offset $\mathcal{A}$, contrast $\mathcal{B}$, and phase shift $\Delta\Phi$ in terms of elements of the pulse propagators. Φ represents the inertial phase shift, introduced in Chapter 3 and is given by Equation 3.5 in the case of a constant acceleration. In deriving Equation 3.11, we follow the procedure outlined by Stoner *et al.* [60], and assume that the change in detuning due to acceleration is negligible during the interferometer pulses.

We start by recalling the form of the propagator for a Raman pulse of constant laser phase, frequency, and amplitude (Equation 2.52):

$$\hat{U}_i = \begin{pmatrix} C_i^* & -iS_i^* \\ -iS_i & C_i \end{pmatrix}. \quad (\text{A.2})$$

In this expression, the terms C_i and S_i are defined by Equations 2.53 and 2.54. We emphasise that these definitions apply only for a pulse with constant amplitude, phase, and frequency. As long as a general pulse with time-dependent amplitude, phase, and frequency may be approximated by a piece-wise constant Hamiltonian, the propagator can still be written in the form of Equation 2.52, but in this case C_i and S_i are no longer defined by Equations 2.53 and 2.54 and must be computed numerically by concatenating the propagators for each time-step.

The propagators for the two periods of free-evolution, known as the dwell-times, take the form

$$\hat{U}_{T1} = \begin{pmatrix} \exp(-\frac{i}{2}\phi_{T1}) & 0 \\ 0 & \exp(+\frac{i}{2}\phi_{T1}) \end{pmatrix} \quad (\text{A.3})$$

$$\hat{U}_{T2} = \begin{pmatrix} \exp(-\frac{i}{2}\phi_{T2}) & 0 \\ 0 & \exp(+\frac{i}{2}\phi_{T2}) \end{pmatrix}, \quad (\text{A.4})$$

where $T1$ refers to the first dwell-time and $T2$ to the second. The phases $\phi_{T1,T2}$ are given by the following integrals:

$$\phi_{T1} = \int_{t_i}^{t_i+T1} \delta(t) dt \quad (\text{A.5})$$

$$\phi_{T2} = \int_{t_i+T1}^{t_i+T1+T2} \delta(t) dt. \quad (\text{A.6})$$

Here t_i is some arbitrary initial time and $\delta(t)$ is the time-dependent Doppler detuning during the dwell-times. The interferometer phase Φ is then defined to be $\phi_{T2} - \phi_{T1}$, the difference between the free-evolution phases.

The final excited state amplitude following the three-pulse Mach-Zehnder sequence is given, after a small number of matrix multiplications, by

$$\begin{aligned} \langle e | \hat{U}_3 \hat{U}_{T2} \hat{U}_2 \hat{U}_{T1} \hat{U}_1 | g \rangle = & i(-S_3 C_2^* C_1^* e^{-\frac{i}{2}(\phi_{T1} + \phi_{T2})} \\ & + S_3 S_2^* S_1^* e^{-\frac{i}{2}(\phi_{T2} - \phi_{T1})} \\ & - C_3 S_2 C_1^* e^{\frac{i}{2}(\phi_{T2} + \phi_{T1})} \\ & - C_3 C_2 S_1 e^{\frac{i}{2}(\phi_{T1} + \phi_{T2})}). \end{aligned} \quad (\text{A.7})$$

By computing the squared modulus of this amplitude and grouping together similar terms, we find the following expression for the excited state probability P_e :

$$\begin{aligned} P_e = & |S_3|^2 |C_2|^2 |C_1|^2 + |S_3|^2 |S_2|^2 |S_1|^2 \\ & + |C_3|^2 |S_2|^2 |C_1|^2 + |C_3|^2 |C_2|^2 |S_1|^2 \\ & - 2 \cos [\phi(C_1 S_1 (S_2^*)^2 S_3 C_3^*) - (\phi_{T2} - \phi_{T1})] \\ & + 2 \cos [\phi(C_1 S_1 C_2^2 C_3 S_3^*) + (\phi_{T2} + \phi_{T1})] \\ & + (|C_3|^2 - |S_3|^2) \times 2 \cos [\phi(S_1^* C_1^* S_2 C_2^*) + \frac{1}{2}(\phi_{T2} - \phi_{T1}) - \frac{1}{2}(\phi_{T2} + \phi_{T1})] \\ & + (|C_1|^2 - |S_1|^2) \times 2 \cos [\phi(S_2 C_2 C_3 S_3^*) + \frac{1}{2}(\phi_{T2} - \phi_{T1}) + \frac{1}{2}(\phi_{T2} + \phi_{T1})]. \end{aligned} \quad (\text{A.8})$$

In this expression, we have used the notation $\phi(A)$ to denote the argument of the complex number A . Following the argument made by Stoner *et al.* [60], we note that the phase combination $\phi_{T2} + \phi_{T1}$ exhibits a significant dependence on the initial atomic velocity because it contains terms proportional to $k_{\text{eff}} v (T_1 + T_2)$. If we let $v = \sigma_v \approx 10 \text{ mm s}^{-1}$

(the standard deviation of the Maxwell-Boltzmann velocity distribution for a cloud of ^{85}Rb at a temperature of $1\ \mu\text{K}$) and choose a combined dwell-time ($T_1 + T_2$) of $1\ \text{ms}$, we find that $\phi_{T_2} + \phi_{T_1} \approx 1.6 \times 10^2$ radians. Consequently, the terms in P_e which contain the phase combination $\phi_{T_2} + \phi_{T_1}$ ¹ oscillate rapidly and will average-out when the excited state probability is integrated over a cloud of atoms with a distribution of atomic velocities, and we therefore drop these terms from our expression. The phase combination $\Phi = \phi_{T_2} - \phi_{T_1}$ does not depend on the initial velocity providing the two dwell-times are exactly equal. This gives us our final expression for P_e in the form of Equation 3.11.

¹These terms are only present when the interferometer pulses are imperfect.

Appendix B

Numerical Simulations

In this appendix we provide more details of the numerical models used to simulate interferometer pulses and sequences in this thesis. Simulations were carried out in *Python* using the *NumPy* [173] and *SciPy* [197] libraries. To reduce the computation time of some functions we employed just-in-time compilation using the *Numba* library [198]. The *QuTiP* library [199] was used to obtain Bloch sphere plots. All the optimisations were carried out using modifications to the *Spinach* spin dynamics software suite [125] in MATLAB and the optimal control methods are explained in detail in Chapter 5.

B.1 Modelling individual Raman pulses

Throughout this thesis, we model ^{85}Rb atoms undergoing Raman transitions as two-level systems, described by the basis states $|g, \mathbf{p}\rangle$ and $|e, \mathbf{p} + \hbar\mathbf{k}_{\text{eff}}\rangle$ with corresponding time-dependent amplitudes $c_{g,e}(t)$. Pulses are described by propagators of the form given by Equation 2.52, and are given by time-ordered products of the propagators for individual slices in the case of shaped and composite pulses, where the amplitude, detuning (frequency), and phase may be varied for each step¹. Once the total propagator for a pulse is computed, we can find the final state of an atom subjected to that pulse for a given initial state and two-photon Rabi frequency Ω_R . We can then repeat this calculation for a range of detunings and/or Rabi frequencies to examine the effect of off-resonance and pulse-length errors on the final state. We can also compute the average fraction of atoms transferred to the excited state after a pulse within an ensemble of atoms with a distribution of sub-states, atomic velocities, and pulse-length errors.

¹To simulate adiabatic pulses we divide the adiabatic frequency sweeps and amplitude profiles into a large number of time-steps to ensure the piece-wise constant approximation is valid. When the tanh/tan pulse is simulated, we use the same sweep parameters as those used by Kotru *et al.* [46] and divide the waveforms into 2500 steps.

When simulations are compared against experimental data in Chapter 7, we use measured values of the Raman beam intensities and single-photon detuning to compute the two-photon Rabi frequency and AC Stark shift for each atom in an ensemble defined by a uniform distribution over the 5 Zeeman sub-levels and a Gaussian distribution of atomic velocities at an empirically determined temperature. The two-photon Rabi frequency and AC Stark shift for each Zeeman sub-level are found for the ‘lin-perp-lin’ polarisation arrangement using Equations 2.45 and 2.46 respectively, following the approach taken by Alex Dunning [56] using the dipole matrix elements for each Raman route given in [57]. We may then find the average fraction of atoms in the excited state throughout a pulse for a given laser detuning (Figure 7.18), or compute the average fraction of atoms in the excited state after a pulse for a range of laser detunings (Figure 7.19). In both cases we find good agreement between the data and the simulation results.

B.2 Modelling Mach-Zehnder sequences

To simulate the interferometer contrast following a Mach-Zehnder pulse sequence in Chapter 7, we first draw a random sample of atoms with different detunings and Rabi frequencies all in the magnetically insensitive $m_F = 0$ sub-level. The distribution of detunings is taken to come from the velocity distribution along the beam axis at a given temperature, and the distribution of pulse-length errors is typically taken from a uniform distribution with a given width. We then evolve the state amplitudes for an atom initially in the ground internal state $|g\rangle$ by applying the relevant pulse propagators in the correct order for each atom in the ensemble. This allows us to compute the average or expected population in the state $|e\rangle$ at the output of the interferometer.

We simulate interference fringes by repeating the evolution of the final beamsplitter pulse while varying a constant laser phase shift ϕ_{bs} between 0 and 2π . The ensemble average fringe is then fitted to the sinusoidal function $0.5(A + B \cos(\phi_{bs} + C))$, where A is an offset, B is the contrast, and C is a phase shift. This gives us the contrast of the average interferometer signal. Since this approach relies on averaging over a random sample of atoms, the sample size is increased until the results do not change appreciably from run to run.

B.3 Modelling LMT sequences

To simulate LMT sequences in Chapter 9 we adopt a similar approach to that outlined in Section B.2 for the Mach-Zehnder sequence but with some important distinctions.

Just as with the Mach-Zehnder simulations, we compute how the LMT interferometer contrast is affected by the atomic temperature and variations in the Rabi rate across




FIGURE B.1: Diagram of the first half of an $N=1$ LMT pulse sequence, shown up to the central mirror pulse, indicating how the momentum and atomic state vary in each arm of the interferometer during the sequence. The resonant Raman frequency depends upon the pair of states coupled at each stage. ‘A’ represents an augmentation pulse. Single and double arrows represent the internal states $|g\rangle$ and $|e\rangle$ respectively. Dotted arrows represent loss trajectories that are discarded in the simulation.

the atom cloud. In order to do this, we draw a sample of initial atomic velocities (and hence Raman detunings) from a Maxwell-Boltzmann distribution for ^{85}Rb at a given temperature and a uniform distribution of Rabi frequencies in the range $\pm 10\%$ of the intended Rabi frequency. Alternatively, the Rabi frequency may be randomly varied from pulse to pulse within $\pm 10\%$ of the intended Rabi frequency in order to explore the robustness of sequences to temporal variation throughout the interferometer.

In an LMT interferometer, the detuning and basis states change for atoms in each of the two primary arms of the interferometer. We therefore make sure that for a given LMT order the pulses are applied in the correct sequence and with the correct detunings for each arm. To reduce computation time, only the amplitudes following the two primary interferometer arms are included in the calculation, as shown in Figure B.1. This involves setting the C term (defined by Equation 2.53) in the propagators describing each augmentation pulse to zero, which is a good assumption when the pulses providing the large-momentum-transfer are efficient, or when there is no interference from atoms that have not followed the primary interferometer paths [44]. We simulate the state evolution for atoms within each interferometer arm separately, and coherently recombine the result before applying the final rectangular beamsplitter pulse. Repeating the calculation while varying the phase of the final beamsplitter allows us to generate a fringe and determine the contrast, as described in Section B.2.

The effect of spontaneous emission is not modelled at present but will ultimately limit the achievable contrast with extended pulse sequences in LMT interferometers.

Appendix C

Conference posters

This appendix contains a selection of two relevant posters presented at academic conferences.

Optimal pulses for enhanced interferometer sensitivity and contrast

Jack Saywell^{1,*}, Max Carey¹, David Elcock¹, Mohammad Belal¹, Ilya Kuprov², and Tim Freegarde¹

¹School of Physics and Astronomy, University of Southampton, SO17 1BJ, UK

²School of Chemistry, University of Southampton, SO17 1BJ, UK

*j.c.saywell@soton.ac.uk

UNIVERSITY OF
Southampton

Introduction

The sensitivity of atom interferometers is ultimately limited by the fidelity of the beamsplitter and mirror operations, which may be improved by the use of NMR composite pulses. We have used the **gradient ascent pulse engineering (GRAPE)** algorithm to design broadband mirror and beamsplitter pulses to increase the fringe contrast of a **Mach-Zehnder** interferometer in the presence of systematic inhomogeneities [1].

Model parameters

We drive **stimulated Raman transitions** between hyperfine ground levels $|5^2S_{1/2}, F=2\rangle$ and $|5^2S_{1/2}, F=3\rangle$ (states |1⟩ and |2⟩ respectively) in a cloud of $\sim 80\mu\text{K}$ ^{85}Rb . Our Hamiltonian may be expressed as:

$$\hat{H}(t) = \frac{\delta}{2}\sigma_z + \frac{1}{2}\Omega_R \left(\cos[\phi_L(t)]\sigma_x + \sin[\phi_L(t)]\sigma_y \right). \quad (1)$$

The effect of the laser pulse is to rotate the quantum state $|\psi\rangle$ on the surface of the Bloch sphere about a fixed axis, which we denote the **field vector**:

$$\Omega = \Omega_R \cos(\phi_L)\mathbf{x} + \Omega_R \sin(\phi_L)\mathbf{y} + (\delta)\mathbf{z}. \quad (2)$$

The field vector's position in the xy plane of the Bloch sphere is determined by the **relative phase** of the two lasers ϕ_L . Detuning from resonance δ causes a deflection of the atomic state, and variations in coupling strength or Rabi frequency Ω_R lead to incorrect rotation angles.

Optimal control theory

Optimal control theory optimises shaped pulses which are highly robust to the specific inhomogeneities present in an ensemble of systems. An interferometer sequence may be optimised through appropriate choices of pulse fidelity. We apply the gradient-based method, **GRAPE** [2], to our system. Our tailored pulses are characterised by smooth variations of the Raman laser phase $\phi_L(t)$.




Figure 1: Robust beamsplitter (left) and mirror (right) pulses designed using GRAPE.

References

- [1] Jack C. Saywell, Ilya Kuprov, David Goodwin, Max Carey, and Tim Freegarde. *Physical Review A*, 98(2):023625, 2018.
- [2] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbruggen, and Steffen J. Glaser. *Journal of Magnetic Resonance*, 172(2):296–305, 2005.
- [3] H. J. Hogben, M. Krzystyniak, G. T. P. Charnock, P. J. Hore, and Ilya Kuprov. *Journal of Magnetic Resonance*, 208(2):179–194, 2011.
- [4] J. R. Johansson, P. D. Nation, and Franco Nori. *Computer Physics Communications*, 184(4):1234–1240, 2013.
- [5] A. J. Shaka, James Keeler, Tom Frenkiel, and Ray Freeman. *Journal of Magnetic Resonance (1969)*, 52(2):335–338, 1983.

Acknowledgements

The *Spinach* spin dynamics software suite was used [3]. The author is grateful for the *baposter* LaTeX template by Brian Amberg. The QuTIP python toolbox was used [4]. This work was supported by the EPSRC under grants EP/M013294/1, and EP/L015382/, and by Dstl under grants DSTLX-1000091758 and DSTLX-1000097855.



Robust beamsplitter and mirror pulses

Inhomogeneities include: nonuniform beam intensity, variations in magnetic environment, Zeeman degeneracy, and broad velocity distributions. We characterise these pulse-errors as falling in two categories:

1. **Off-resonance errors** → Caused by non-zero detuning from resonance, δ
2. **Pulse-length errors** → Caused by variations in coupling strength or Rabi-frequency, $\Omega_R/\Omega_{\text{eff}}$




Figure 2: Final excited state probability for square π (left), composite WALTZ [5] (middle), and GRAPE (right) pulses.

Design Procedure: Pulse fidelities characterise how accurately our system is driven from an initial state $|\psi_0\rangle$ to a target state $|\psi_D\rangle$:

$$\mathcal{F} = f\left(\langle\psi_D|\hat{T}\exp\left[-\frac{i}{\hbar}\int_{t_1}^{t_2}\hat{H}(t)dt\right]|\psi_0\rangle\right) = f(\langle\psi_D|\hat{U}_N\hat{U}_{N-1}\dots\hat{U}_2\hat{U}_1|\psi_0\rangle). \quad (3)$$

GRAPE computes derivatives of the chosen fidelity w.r.t. the control pulse sequence efficiently, to permit fast optimisation. Pulses are discretised into N slices, and the choice of f , $|\psi_0\rangle$, and $|\psi_D\rangle$ determine which pulse is obtained. Robustness to errors is achieved by averaging over an ensemble of pulse-length and off-resonance errors in the fidelity calculation, maximising $\mathcal{F}$ for an ensemble of atoms.

Augmentation pulses for LMT interferometers

The use of extended pulse sequences to increase interferometer area requires robust high-fidelity augmentation pulses. We have demonstrated broadband pulses that optimise state transfer within a warm cloud of ^{85}Rb atoms in the presence of large off-resonance and pulse-length errors. We expect these pulses to:

- Improve the contrast for a given number of augmentation pulses (and equivalent interferometer area) compared with standard rectangular and composite pulses
- Allow for a greater interferometer area and sensitivity without the need to filter the atomic sample




Figure 3: Temporal evolution of atoms (left) and simulated Bloch sphere trajectories (right) for two GRAPE pulses showing optimised state transfer in a $\sim 30\mu\text{K}$ cloud. Our model has good agreement with the data, demonstrating optimal control theory works for our system.

GRAPE interferometer sequences

We obtain pulse sequences that optimise the contrast of interferometer fringes in warm, inhomogeneous clouds. We use **symmetry** in the design of pulse sequences to preserve the interferometer phase.




Figure 4: Mach-Zehnder fringes (left) with optimised GRAPE pulse sequence vs. rectangular pulses. The temperature was $\sim 100\mu\text{K}$. Contrast was improved by a factor of 3. A spectral comparison of pulses is also shown (right).

Outlook

GRAPE pulses prove more robust than established composite pulses and increases in pulse fidelity therefore allow for greater interferometric areas, increased contrast, warmer atom clouds, and thus improved interferometric sensitivity. We have demonstrated a three-fold improvement in fringe visibility for a $\sim 100\mu\text{K}$ cloud.

FIGURE C.1: Poster presented at the 2018 Frontiers of Matterwave Optics (FOMO) conference 17-21 September, Orthodox Academy of Crete, Greece. For this poster I was awarded a runner-up prize in the conference's poster competition.

Optimal Raman pulses for enhanced contrast and sensitivity of atom interferometers

Jack Saywell^{1,*}, Max Carey¹, Mohammad Belal¹, Ilya Kuprov², and Tim Freearge¹

¹School of Physics and Astronomy, University of Southampton, SO17 1BJ, UK

²School of Chemistry, University of Southampton, SO17 1BJ, UK

*j.c.saywell@soton.ac.uk



UNIVERSITY OF
Southampton

Introduction

The sensitivity of atom interferometers is limited by the beamsplitter and mirror fidelities in the presence of systematic inhomogeneities. We have designed and tested broadband Raman pulses that achieve **99.8(3)%** Raman population transfer in a $\sim 35\mu\text{K}$ cloud of ^{85}Rb and increase threefold the fringe contrast of an interferometer when compared with that obtained using rectangular pulses [1,2]. Our pulses will enhance measurement sensitivity by permitting larger atom clouds and greater path separations in Large Momentum Transfer (LMT) interferometers.

Model parameters

We drive **stimulated Raman transitions** between hyperfine ground levels $|5^2S_{1/2}, F=2\rangle$ and $|5^2S_{1/2}, F=3\rangle$ (states $|1\rangle$ and $|2\rangle$ respectively) in a cloud of $\sim 80\mu\text{K}$ ^{85}Rb released from a 3D MOT. Our Hamiltonian may be expressed as:

$$\hat{H}(t) = \frac{\delta}{2}\sigma_z + \frac{1}{2}\Omega_R \left(\cos[\phi_L(t)]\sigma_x + \sin[\phi_L(t)]\sigma_y \right). \quad (1)$$

The effect of a Raman pulse is to rotate the quantum state $|\psi\rangle$ on the surface of the Bloch sphere about a fixed axis, called the **field vector**:

$$\Omega = \Omega_R \cos(\phi_L)\mathbf{x} + \Omega_R \sin(\phi_L)\mathbf{y} + (\delta)\mathbf{z}. \quad (2)$$

The field vector's position in the xy plane of the Bloch sphere is determined by the **relative phase** of the two lasers ϕ_L . Detuning from resonance δ causes a deflection of the atomic state, and variations in coupling strength Ω_R lead to incorrect rotation angles and reduce interferometric contrast.

Optimal control theory

We optimise pulses which are highly robust to the specific inhomogeneities present in an atom interferometer. An interferometer sequence may be optimised through appropriate choices of pulse fidelity. We apply the gradient-based method, **GRAPE** [3], to our system. Our tailored pulses are characterised by sequences of the Raman laser phase $\phi_L(t) = \phi_1 \dots \phi_N$.




Figure 1: Robust augmentation (A), beamsplitter (B) and mirror (C) pulses designed using GRAPE.

References

- [1] J. Saywell et al., *Phys. Rev. A* **98**, 023625, (2018)
- [2] J. Saywell et al., *in preparation*, (2019)
- [3] N. Khaneja et al., *J. Magn. Reson.*, **172**, 296, (2005)
- [4] A. J. Shaka et al., *J. Magn. Reson.* **1969**, **52**, 335, (1983)
- [5] T. L. Hwang et al., *J. Magn. Reson.*, **133**, 200, (1998)
- [6] C. J. Hardy et al., *J. Magn. Reson.* **1969**, **66**, 470, (1986)
- [7] K. Kokru et al., *Phys. Rev. Lett.*, **115**, 103001, (2015)
- [8] H. J. Hogben et al., *J. Magn. Reson.*, **208**, 179, (2011)
- [9] J. R. Johansson et al., *Computer Physics Communications*, **184**, 1234, (2013)

Acknowledgements

The *Spinach* software suite was used [8]. The author is grateful for the *baposter* LaTeX template by Brian Amberg. The QuTiP toolbox was used [9]. This work was supported by the EPSRC (EP/M013294/1, EP/L015382/1) and by Dstl (DSTLX-1000091758, DSTLX-1000097855).

Robust beamsplitter and mirror pulses

Inhomogeneities include: nonuniform beam intensity, variations in magnetic environment, Zeeman degeneracy, and broad velocity distributions. We characterise these pulse-errors as falling in two categories:

1. **Off-resonance errors** $\rightarrow$ Caused by non-zero detuning from resonance, δ
2. **Pulse-length errors** $\rightarrow$ Caused by variations in coupling strength or Rabi-frequency, $\Omega_R/\Omega_{\text{eff}}$

Design Procedure: Pulses are split into N "slices", where the relative laser phase ϕ_L takes different values for each step, $\phi_L(t) = \phi_1 \dots \phi_N$. We choose a pulse fidelity, e.g. how accurately our system is driven from an initial state $|\psi_0\rangle$ to a target state $|\psi_D\rangle$:

$$\mathcal{F} = f\left(\langle\psi_D|\hat{T}\exp\left[-\frac{i}{\hbar}\int_{t_1}^{t_2}\hat{H}(t)dt\right]|\psi_0\rangle\right) = f(\langle\psi_D|\hat{U}_N\hat{U}_{N-1}\dots\hat{U}_2\hat{U}_1|\psi_0\rangle). \quad (3)$$

GRAPE computes derivatives of the chosen fidelity w.r.t. the control phases efficiently, to permit fast optimisation from an initial guess. f , $|\psi_0\rangle$, and $|\psi_D\rangle$ determine which pulse is obtained. Robustness is achieved by averaging over an ensemble of pulse-length and off-resonance errors in the fidelity calculation, maximising $\mathcal{F}$ for all atoms.

Augmentation pulses for LMT interferometers

Extended pulse sequences to increase interferometer area require robust high-fidelity augmentation pulses. We have demonstrated broadband pulses that optimise state transfer within a warm cloud of ^{85}Rb atoms in the presence of large off-resonance and pulse-length errors. We expect these pulses to:

- Improve the contrast for a given number of augmentation pulses (and equivalent interferometer area)
- Allow for a greater interferometer area and sensitivity without the need to filter the atomic sample




Figure 2: Measured temporal evolution of atoms (left) and simulated Bloch sphere trajectories (right) for two GRAPE pulses showing optimised state transfer in a $\sim 30\mu\text{K}$ cloud. Our model has good agreement with the data, demonstrating optimal control theory works for our system [2].




Figure 3: (left) Simulated average pulse efficiency in a $40\mu\text{K}$ atom cloud for different augmentation pulses as a function of pulse duration as a fraction of that for a rectangular π pulse. GRAPE offers more efficient and shorter pulses than composite [4] and Adiabatic Rapid Passage (ARP) alternatives [5-7]. (right) Simulated efficiency of GRAPE inversion pulse (A) and rectangular π pulse (B) computed for off-resonance and pulse-length errors.

Optimised 3-pulse Mach-Zehnder sequence

We obtain pulse sequences that maximise the contrast of interferometer fringes in warm, inhomogeneous clouds. We use **symmetry** in the design of pulse sequences to preserve the interferometer phase.




Figure 4: Mach-Zehnder fringes (left) with optimised GRAPE pulse sequence vs. rectangular pulses. The temperature was $\sim 100\mu\text{K}$. Contrast was improved by a factor of 3. A spectral comparison of pulses is also shown (right) [2].

Outlook

GRAPE pulses prove more robust than established composite pulses and increases in pulse fidelity therefore allow for greater interferometric areas, increased contrast, warmer atom clouds, and thus improved interferometric sensitivity. We have demonstrated a three-fold improvement in fringe visibility for a $\sim 100\mu\text{K}$ cloud [2].

FIGURE C.2: Poster presented at the 2019 International Conference on Quantum Error Correction (QEC) 29 July - 2 August, Senate House, London, UK.

Bibliography

- [1] 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).
- [2] M. Carey, J. Saywell, M. Belal, and T. Freegarde, “Velocimetry of cold atoms by matter-wave interferometry,” *Physical Review A* **99**, 023631 (2019).
- [3] 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).
- [4] J. Saywell, M. Carey, I. Kuprov, and T. Freegarde, “Biselective pulses for large-area atom interferometry,” *Physical Review A* **101**, 063625 (2020).
- [5] M. H. Levitt and R. Freeman, “NMR population inversion using a composite pulse,” *Journal of Magnetic Resonance* (1969) **33**, 473 (1979).
- [6] J. Baum, R. Tycko, and A. Pines, “Broadband and adiabatic inversion of a two-level system by phase-modulated pulses,” *Physical Review A* **32**, 3435 (1985).
- [7] T. E. Skinner, T. O. Reiss, B. Luy, N. Khaneja, and S. J. Glaser, “Application of optimal control theory to the design of broadband excitation pulses for high-resolution NMR,” *Journal of Magnetic Resonance* **163**, 8 (2003).
- [8] K. Kobzar, T. E. Skinner, N. Khaneja, S. J. Glaser, and B. Luy, “Exploring the limits of broadband excitation and inversion pulses,” *Journal of Magnetic Resonance* **170**, 236 (2004).
- [9] 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).
- [10] 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).

- [11] 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).
- [12] D. J. Tannor and S. A. Rice, “Control of selectivity of chemical reaction via control of wave packet evolution,” *The Journal of Chemical Physics* **83**, 5013 (1985).
- [13] R. Kosloff, S. Rice, P. Gaspard, S. Tersigni, and D. Tannor, “Wavepacket dancing: Achieving chemical selectivity by shaping light pulses,” *Chemical Physics* **139**, 201 (1989).
- [14] J. Somló, V. A. Kazakov, and D. J. Tannor, “Controlled dissociation of I_2 via optical transitions between the X and B electronic states,” *Chemical Physics* **172**, 85 (1993).
- [15] W. S. Warren, H. Rabitz, and M. Dahleh, “Coherent Control of Quantum Dynamics: The Dream Is Alive,” *Science* **259**, 1581 (1993).
- [16] S. Conolly, D. Nishimura, and A. Macovski, “Optimal Control Solutions to the Magnetic Resonance Selective Excitation Problem,” *IEEE Transactions on Medical Imaging* **5**, 106 (1986).
- [17] D. Rosenfeld and Y. Zur, “Design of adiabatic selective pulses using optimal control theory,” *Magnetic Resonance in Medicine* **36**, 401 (1996).
- [18] J. Mao, T. Mareci, K. Scott, and E. Andrew, “Selective inversion radiofrequency pulses by optimal control,” *Journal of Magnetic Resonance* (1969) **70**, 310 (1986).
- [19] P. de Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, “Second order gradient ascent pulse engineering,” *Journal of Magnetic Resonance* **212**, 412 (2011).
- [20] P. Doria, T. Calarco, and S. Montangero, “Optimal Control Technique for Many-Body Quantum Dynamics,” *Physical Review Letters* **106**, 190501 (2011).
- [21] T. Caneva, T. Calarco, and S. Montangero, “Chopped random-basis quantum optimization,” *Physical Review A* **84**, 022326 (2011).
- [22] T. E. Skinner and N. I. Gershenzon, “Optimal control design of pulse shapes as analytic functions,” *Journal of Magnetic Resonance* **204**, 248 (2010).
- [23] D. L. Goodwin and I. Kuprov, “Modified Newton-Raphson GRAPE methods for optimal control of spin systems,” *The Journal of Chemical Physics* **144**, 204107 (2016).
- [24] R. Eitan, M. Mundt, and D. J. Tannor, “Optimal control with accelerated convergence: Combining the Krotov and quasi-Newton methods,” *Physical Review A* **83**, 053426 (2011).

- [25] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, “Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits,” *Physical Review Letters* **120**, 150401 (2018).
- [26] B. Bartels and F. Mintert, “Smooth optimal control with Floquet theory,” *Physical Review A* **88**, 052315 (2013).
- [27] 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).
- [28] F. Poggiali, P. Cappellaro, and N. Fabbri, “Optimal Control for One-Qubit Quantum Sensing,” *Physical Review X* **8**, 021059 (2018).
- [29] U. V. Poulsen, S. Sklarz, D. Tannor, and T. Calarco, “Correcting errors in a quantum gate with pushed ions via optimal control,” *Physical Review A* **82**, 012339 (2010).
- [30] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, “Parallel entangling operations on a universal ion-trap quantum computer,” *Nature* **572**, 368 (2019).
- [31] A. Spörl, T. Schulte-Herbrüggen, S. J. Glaser, V. Bergholm, M. J. Storcz, J. Ferber, and F. K. Wilhelm, “Optimal control of coupled Josephson qubits,” *Physical Review A* **75**, 012302 (2007).
- [32] A. W. Cross and J. M. Gambetta, “Optimized pulse shapes for a resonator-induced phase gate,” *Physical Review A* **91**, 032325 (2015).
- [33] S. van Frank, A. Negretti, T. Berrada, R. Bückler, 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).
- [34] 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,” *Phys. Rev. A* **90**, 033628 (2014).
- [35] 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).
- [36] M. Kasevich and S. Chu, “Atomic interferometry using stimulated Raman transitions,” *Physical Review Letters* **67**, 181 (1991).

- [37] B. Young, M. Kasevich, and S. Chu, “Precision Atom Interferometry with Light Pulses,” in *Atom Interferometry*, edited by R. P. Berman (Academic Press, San Diego, 1997) pp. 363–406.
- [38] B. Barrett, R. Geiger, I. Dutta, M. Meunier, B. Canuel, A. Gauguier, P. Bouyer, and A. Landragin, “The Sagnac effect: 20 years of development in matter-wave interferometry,” *Comptes Rendus Physique* **15**, 875 (2014).
- [39] 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).
- [40] J. M. McGuirk, M. J. Snadden, and M. A. Kasevich, “Large Area Light-Pulse Atom Interferometry,” *Physical Review Letters* **85**, 4498 (2000).
- [41] 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**, 033608 (2002).
- [42] 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).
- [43] A. Shaka, J. Keeler, T. Frenkiel, and R. Freeman, “An improved sequence for broadband decoupling: WALTZ-16,” *Journal of Magnetic Resonance* (1969) **52**, 335 (1983).
- [44] K. Kotru, *Timekeeping and Accelerometry with Robust Light Pulse Atom Interferometers*, Ph.D. thesis, Massachusetts Institute of Technology (2015).
- [45] 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).
- [46] 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).
- [47] T. Kovachy, S.-w. Chiow, and M. A. Kasevich, “Adiabatic-rapid-passage multi-photon Bragg atom optics,” *Physical Review A* **86**, 011606(R) (2012).
- [48] 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).

- [49] 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).
- [50] M. Jaffe, *Atom Interferometry in an Optical Cavity*, Ph.D. thesis, University of California, Berkeley (2019).
- [51] T.-L. Hwang, P. C. van Zijl, and M. Garwood, “Fast Broadband Inversion by Adiabatic Pulses,” *Journal of Magnetic Resonance* **133**, 200 (1998).
- [52] K. Moler, D. S. Weiss, M. Kasevich, and S. Chu, “Theoretical analysis of velocity-selective Raman transitions,” *Physical Review A* **45**, 342 (1992).
- [53] C. J. Foot, *Atomic Physics* (Oxford University Press, Oxford, 2004).
- [54] C. Cohen-Tannoudji and D. Guéry-Odelin, *Advances in atomic physics: an overview* (World Scientific Publishing, Singapore, 2011).
- [55] S. M. Barnett, “On the recoil and Doppler shifts,” *Journal of Modern Optics* **57**, 1445 (2010).
- [56] A. Dunning, *Coherent atomic manipulation and cooling using composite optical pulse sequences*, Ph.D. thesis, University of Southampton (2014).
- [57] D. A. Steck, “Rubidium 85 D Line Data,” available online at <http://steck.us/alkalidata> (revision 2.2.1, 21 November 2019) .
- [58] J. Bateman, A. Xuereb, and T. Freegarde, “Stimulated Raman transitions via multiple atomic levels,” *Physical Review A* **81**, 043808 (2010).
- [59] M. Kasevich and S. Chu, “Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer,” *Applied Physics B* **54**, 321 (1992).
- [60] 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).
- [61] H. K. Cummins and J. A. Jones, “Use of composite rotations to correct systematic errors in NMR quantum computation,” *New Journal of Physics* **2**, 6 (2000).
- [62] 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).
- [63] M. H. Levitt, “Composite pulses,” *Progress in Nuclear Magnetic Resonance Spectroscopy* **18**, 61 (1986).
- [64] J. A. Jones, “Nested composite NOT gates for quantum computation,” *Physics Letters A* **377**, 2860 (2013).

- [65] L. Marton, J. A. Simpson, and J. A. Suddeth, “An Electron Interferometer,” *Review of Scientific Instruments* **25**, 1099 (1954).
- [66] H. Rauch, W. Treimer, and U. Bonse, “Test of a single crystal neutron interferometer,” *Physics Letters A* **47**, 369 (1974).
- [67] A. Ashkin, “Atomic-Beam Deflection by Resonance-Radiation Pressure,” *Physical Review Letters* **25**, 1321 (1970).
- [68] T. Hänsch and A. Schawlow, “Cooling of gases by laser radiation,” *Optics Communications* **13**, 68 (1975).
- [69] D. J. Wineland and W. M. Itano, “Laser cooling of atoms,” *Physical Review A* **20**, 1521 (1979).
- [70] 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).
- [71] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Graduate Texts in Contemporary Physics (Springer, New York, 1999).
- [72] A. Peters, K. Y. Chung, and S. Chu, “High-precision gravity measurements using atom interferometry,” *Metrologia* **38**, 25 (2001).
- [73] G. Rosi, F. Sorrentino, L. Cacciapuoti, M. Prevedelli, and G. M. Tino, “Precision measurement of the Newtonian gravitational constant using cold atoms,” *Nature* **510**, 518 (2014).
- [74] G. Rosi, G. D’Amico, L. Cacciapuoti, F. Sorrentino, M. Prevedelli, M. Zych, Č. Brukner, and G. M. Tino, “Quantum test of the equivalence principle for atoms in coherent superposition of internal energy states,” *Nature Communications* **8**, 15529 (2017).
- [75] A. Wicht, J. M. Hensley, E. Sarajlic, and S. Chu, “A Preliminary Measurement of the Fine Structure Constant Based on Atom Interferometry,” *Physica Scripta* **T102**, 82 (2002).
- [76] M. Cadoret, E. de Mirandes, P. Cladé, S. Guellati-Khélifa, C. Schwob, F. Nez, L. Julien, and F. Biraben, “Combination of Bloch Oscillations with a Ramsey-Bordé Interferometer: New Determination of the Fine Structure Constant,” *Physical Review Letters* **101**, 230801 (2008).
- [77] 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).

- [78] P. Cladé, F. Nez, F. Biraben, and S. Guellati-Khelifa, “State of the art in the determination of the fine-structure constant and the ratio h/m_u ,” *Comptes Rendus Physique* **20**, 77 (2019).
- [79] P. Hamilton, M. Jaffe, P. Haslinger, Q. Simmons, H. Müller, and J. Khoury, “Atom-interferometry constraints on dark energy,” *Science* **349**, 849 (2015).
- [80] C. Burrage, E. J. Copeland, and E. Hinds, “Probing dark energy with atom interferometry,” *Journal of Cosmology and Astroparticle Physics* **2015**, 042 (2015).
- [81] M. Jaffe, P. Haslinger, V. Xu, P. Hamilton, A. Upadhye, B. Elder, J. Khoury, and H. Müller, “Testing sub-gravitational forces on atoms from a miniature in-vacuum source mass,” *Nature Physics* **13**, 938 (2017).
- [82] 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).
- [83] P. W. Graham, J. M. Hogan, M. A. Kasevich, and S. Rajendran, “New Method for Gravitational Wave Detection with Atomic Sensors,” *Physical Review Letters* **110**, 171102 (2013).
- [84] T. L. Gustavson, P. Bouyer, and M. A. Kasevich, “Precision Rotation Measurements with an Atom Interferometer Gyroscope,” *Physical Review Letters* **78**, 2046 (1997).
- [85] B. Canuel, F. Leduc, D. Holleville, A. Gauguier, J. Fils, A. Virdis, A. Clairon, N. Dimarcq, C. J. Bordé, A. Landragin, and P. Bouyer, “Six-Axis Inertial Sensor Using Cold-Atom Interferometry,” *Physical Review Letters* **97**, 010402 (2006).
- [86] 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).
- [87] 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).
- [88] 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).
- [89] 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).

- [90] 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).
- [91] H. Müller, S.-w. Chiow, S. Herrmann, and S. Chu, “Atom Interferometers with Scalable Enclosed Area,” *Physical Review Letters* **102**, 240403 (2009).
- [92] 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).
- [93] S. Odedra and S. Wimperis, “Use of composite refocusing pulses to form spin echoes,” *Journal of Magnetic Resonance* **214**, 68 (2012).
- [94] P. Cheinet, B. Canuel, F. Pereira Dos Santos, A. Gauguier, F. Yver-Leduc, and A. Landragin, “Measurement of the Sensitivity Function in a Time-Domain Atomic Interferometer,” *IEEE Transactions on Instrumentation and Measurement* **57**, 1141 (2008).
- [95] A. Peters, *High Precision Gravity Measurements using Atom Interferometry*, Ph.D. thesis, Stanford University (1998).
- [96] 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).
- [97] X. Li, C.-G. Shao, and Z.-K. Hu, “Raman pulse duration effect in high-precision atom interferometry gravimeters,” *Journal of the Optical Society of America B* **32**, 248 (2015).
- [98] A. Bertoldi, F. Minardi, and M. Prevedelli, “Phase shift in atom interferometers: Corrections for nonquadratic potentials and finite-duration laser pulses,” *Physical Review A* **99**, 033619 (2019).
- [99] C. J. Bordé, “Atomic clocks and inertial sensors,” *Metrologia* **39**, 435 (2002).
- [100] S. Bize, P. Laurent, M. Abgrall, H. Marion, I. Maksimovic, L. Cacciapuoti, J. Grünert, C. Vian, F. P. dos Santos, P. Rosenbusch, P. Lemonde, G. Santarelli, P. Wolf, A. Clairon, A. Luiten, M. Tobar, and C. Salomon, “Cold atom clocks and applications,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **38**, S449 (2005).
- [101] 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).

- [102] 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).
- [103] M. S. Carey, *Velocimetry, trapping and optimal coherent manipulation of atomic rubidium*, Ph.D. thesis, University of Southampton (2019).
- [104] M. Weitz and T. W. Hänsch, “Frequency-independent laser cooling based on interferometry,” *Europhysics Letters* **49**, 302 (2000).
- [105] A. Dunning, R. Gregory, J. Bateman, M. Himsworth, and T. Freegarde, “Interferometric Laser Cooling of Atomic Rubidium,” *Physical Review Letters* **115**, 073004 (2015).
- [106] D. S. Weiss, B. C. Young, and S. Chu, “Precision measurement of the photon recoil of an atom using atomic interferometry,” *Physical Review Letters* **70**, 2706 (1993).
- [107] T. Freegarde and D. Segal, “Algorithmic Cooling in a Momentum State Quantum Computer,” *Physical Review Letters* **91**, 037904 (2003).
- [108] 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,” *European Physical Journal D* **70**, 262 (2016).
- [109] 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).
- [110] N. Mielec, M. Altorio, R. Sapam, D. Horville, D. Holleville, L. A. Sidorenkov, A. Landragin, and R. Geiger, “Atom interferometry with top-hat laser beams,” *Applied Physics Letters* **113**, 161108 (2018).
- [111] M. H. Levitt and R. Freeman, “Compensation for pulse imperfections in NMR spin-echo experiments,” *Journal of Magnetic Resonance* (1969) **43**, 65 (1981).
- [112] C. Kabytayev, T. J. Green, K. Khodjasteh, M. J. Biercuk, L. Viola, and K. R. Brown, “Robustness of composite pulses to time-dependent control noise,” *Physical Review A* **90**, 012316 (2014).
- [113] R. Tycko, A. Pines, and J. Guckenheimer, “Fixed point theory of iterative excitation schemes in NMR,” *The Journal of Chemical Physics* **83**, 2775 (1985).
- [114] C. Hardy, W. Edelstein, and D. Vatis, “Efficient adiabatic fast passage for NMR population inversion in the presence of radiofrequency field inhomogeneity and frequency offsets,” *Journal of Magnetic Resonance* (1969) **66**, 470 (1986).

- [115] A. Tannús and M. Garwood, “Improved Performance of Frequency-Swept Pulses Using Offset-Independent Adiabaticity,” *Journal of Magnetic Resonance, Series A* **120**, 133 (1996).
- [116] M. Garwood and L. DelaBarre, “The Return of the Frequency Sweep: Designing Adiabatic Pulses for Contemporary NMR,” *J. Magn. Reson.* **153**, 155 (2001).
- [117] L. S. Pontryagin, V. G. Bol’tanskii, R. S. Gamkrelidze, and E. F. Mischenko, *The Mathematical Theory of Optimal Processes* (Pergamon Press, New York, 1964).
- [118] R. Freeman, “Shaped radiofrequency pulses in high resolution NMR,” *Progress in Nuclear Magnetic Resonance Spectroscopy* **32**, 59 (1998).
- [119] J. A. Nelder and R. Mead, “A Simplex Method for Function Minimization,” *The Computer Journal* **7**, 308 (1965).
- [120] M. H. Levitt and R. Ernst, “Composite pulses constructed by a recursive expansion procedure,” *Journal of Magnetic Resonance* (1969) **55**, 247 (1983).
- [121] 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).
- [122] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell, “Multiparameter optimisation of a magneto-optical trap using deep learning,” *Nature Communications* **9**, 4360 (2018).
- [123] J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer, New York, 2006).
- [124] D. C. Liu and J. Nocedal, “On the limited memory BFGS method for large scale optimization,” *Mathematical Programming* **45**, 503 (1989).
- [125] 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).
- [126] D. L. Goodwin, *Advanced Optimal Control Methods for Spin Systems*, Ph.D. thesis, University of Southampton (2018).
- [127] C. G. Broyden, “The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations,” *IMA Journal of Applied Mathematics* **6**, 76 (1970).
- [128] C. G. Broyden, “The Convergence of a Class of Double-rank Minimization Algorithms: 2. The New Algorithm,” *IMA Journal of Applied Mathematics* **6**, 222 (1970).

- [129] R. Fletcher, “A new approach to variable metric algorithms,” *The Computer Journal* **13**, 317 (1970).
- [130] D. F. Shanno, “Conditioning of quasi-Newton methods for function minimization,” *Mathematics of Computation* **24**, 647 (1970).
- [131] D. Goldfarb, “A Family of Variable-Metric Methods Derived by Variational Means,” *Mathematics of Computation* **24**, 23 (1970).
- [132] V. Krotov, *Global Methods in Optimal Control* (Marcel Dekker, New York, 1996).
- [133] S. E. Sklarz and D. J. Tannor, “Loading a Bose-Einstein condensate onto an optical lattice: An application of optimal control theory to the nonlinear Schrödinger equation,” *Physical Review A* **66**, 053619 (2002).
- [134] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, “Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework,” *Physical Review A* **84**, 022305 (2011).
- [135] S. G. Schirmer and P. de Fouquieres, “Efficient algorithms for optimal control of quantum dynamics: the Krotov method unencumbered,” *New Journal of Physics* **13**, 073029 (2011).
- [136] D. B. Zax and S. Vega, “Broad-band excitation pulses of arbitrary flip angle,” *Physical Review Letters* **62**, 1840 (1989).
- [137] H. Geen and R. Freeman, “Band-selective radiofrequency pulses,” *Journal of Magnetic Resonance* (1969) **93**, 93 (1991).
- [138] D. J. Tannor, *Introduction to Quantum Mechanics: A Time-Dependent Perspective* (University Science Books, Sausalito, 2007).
- [139] D. L. Goodwin and I. Kuprov, “Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories,” *Journal of Chemical Physics* **143**, 084113 (2015).
- [140] M. Hochbruck and C. Lubich, “On Krylov Subspace Approximations to the Matrix Exponential Operator,” *SIAM Journal on Numerical Analysis* **34**, 1911 (1997).
- [141] I. Kuprov, “Polynomially scaling spin dynamics II: Further state-space compression using Krylov subspace techniques and zero track elimination,” *Journal of Magnetic Resonance* **195**, 45 (2008).
- [142] R. Fletcher, *Practical Methods of Optimization* (John Wiley & Sons, Chichester, 2000).
- [143] R. L. Kosut, M. D. Grace, and C. Brif, “Robust control of quantum gates via sequential convex programming,” *Physical Review A* **88**, 052326 (2013).

- [144] J. L. Allen, R. Kosut, J. Joo, P. Leek, and E. Ginossar, “Optimal control of two qubits via a single cavity drive in circuit quantum electrodynamics,” *Physical Review A* **95**, 042325 (2017).
- [145] H. Zhang and H. Rabitz, “Robust optimal control of quantum molecular systems in the presence of disturbances and uncertainties,” *Physical Review A* **49**, 2241 (1994).
- [146] J. G. B. Beumee and H. Rabitz, “Robust optimal control theory for selective vibrational excitation in molecules: A worst case analysis,” *The Journal of Chemical Physics* **97**, 1353 (1992).
- [147] P. Owrutsky and N. Khaneja, “Control of inhomogeneous ensembles on the Bloch sphere,” *Physical Review A* **86**, 022315 (2012).
- [148] J.-S. Li, J. Ruths, T.-Y. Yu, H. Arthanari, and G. Wagner, “Optimal pulse design in quantum control: A unified computational method,” *Proceedings of the National Academy of Sciences* **108**, 1879 (2011).
- [149] J.-S. Li and N. Khaneja, “Control of inhomogeneous quantum ensembles,” *Physical Review A* **73**, 030302(R) (2006).
- [150] S. Husain, M. Kawamura, and J. A. Jones, “Further analysis of some symmetric and antisymmetric composite pulses for tackling pulse strength errors,” *Journal of Magnetic Resonance* **230**, 145 (2013).
- [151] C. D. Schwieters, J. G. B. Beumee, and H. Rabitz, “Optical control of molecular motion with robustness and application to vinylidene fluoride,” *Journal of the Optical Society of America B* **7**, 1736 (1990).
- [152] H. Rabitz, “Optimal control of quantum systems: Origins of inherent robustness to control field fluctuations,” *Physical Review A* **66**, 063405 (2002).
- [153] H. Rabitz, T.-S. Ho, M. Hsieh, R. Kosut, and M. Demiralp, “Topology of optimally controlled quantum mechanical transition probability landscapes,” *Physical Review A* **74**, 012721 (2006).
- [154] M. Demiralp and H. Rabitz, “Assessing optimality and robustness of control over quantum dynamics,” *Physical Review A* **57**, 2420 (1998).
- [155] M. Demiralp and H. Rabitz, “Assessing optimality and robustness for the control of dynamical systems,” *Physical Review E* **61**, 2569 (2000).
- [156] T.-S. Ho, J. Dominy, and H. Rabitz, “Landscape of unitary transformations in controlled quantum dynamics,” *Physical Review A* **79**, 013422 (2009).

- [157] D. Hocker, C. Brif, M. D. Grace, A. Donovan, T.-S. Ho, K. M. Tibbetts, R. Wu, and H. Rabitz, “Characterization of control noise effects in optimal quantum unitary dynamics,” *Physical Review A* **90**, 062309 (2014).
- [158] 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).
- [159] M. A. Smith, H. Hu, and A. J. Shaka, “Improved broadband inversion performance for NMR in liquids,” *Journal of Magnetic Resonance* **151**, 269 (2001).
- [160] C. A. Ryan, J. S. Hodges, and D. G. Cory, “Robust Decoupling Techniques to Extend Quantum Coherence in Diamond,” *Physical Review Letters* **105**, 200402 (2010).
- [161] Z. Starčuk and V. Sklenář, “Composite pulse sequences with variable performance,” *Journal of Magnetic Resonance* (1969) **62**, 113 (1985).
- [162] R. Freeman, S. P. Kempell, and M. H. Levitt, “Radiofrequency pulse sequences which compensate their own imperfections,” *Journal of Magnetic Resonance* (1969) **38**, 453 (1980).
- [163] K. Kobzar, T. E. Skinner, N. Khaneja, S. J. Glaser, and B. Luy, “Exploring the limits of broadband excitation and inversion: II. Rf-power optimized pulses,” *Journal of Magnetic Resonance* **194**, 58 (2008).
- [164] T. E. Skinner, T. O. Reiss, B. Luy, N. Khaneja, and S. J. Glaser, “Reducing the duration of broadband excitation pulses using optimal control with limited RF amplitude,” *Journal of Magnetic Resonance* **167**, 68 (2004).
- [165] T. E. Skinner, T. O. Reiss, B. Luy, N. Khaneja, and S. J. Glaser, “Tailoring the optimal control cost function to a desired output: application to minimizing phase errors in short broadband excitation pulses,” *Journal of Magnetic Resonance* **172**, 17 (2005).
- [166] A. J. Shaka and A. Pines, “Symmetric phase-alternating composite pulses,” *Journal of Magnetic Resonance* (1969) **71**, 495 (1987).
- [167] S. Wimperis, “Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments,” *Journal of Magnetic Resonance, Series A* **109**, 221 (1994).
- [168] R. Freeman, J. Friedrich, and Wu Xi-li, “A pulse for all seasons. Fourier transform spectra without a phase gradient,” *Journal of Magnetic Resonance* (1969) **79**, 561 (1988).

- [169] N. I. Gershenzon, T. E. Skinner, B. Brutscher, N. Khaneja, M. Nimbalkar, B. Luy, and S. J. Glaser, “Linear phase slope in pulse design: Application to coherence transfer,” *Journal of Magnetic Resonance* **192**, 235 (2008).
- [170] 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).
- [171] P. Cheinet, *Conception et réalisation d’un gravimètre à atomes froids*, Ph.D. thesis, Université Pierre et Marie Curie (2006).
- [172] Z. Chen, J. G. Bohnet, J. M. Weiner, and J. K. Thompson, “General formalism for evaluating the impact of phase noise on Bloch vector rotations,” *Physical Review A* **86**, 032313 (2012).
- [173] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, “The NumPy array: A structure for efficient numerical computation,” *Computing in Science and Engineering* **13**, 22 (2011).
- [174] J. Rudolph, T. Wilkason, M. Nantel, H. Swan, C. M. Holland, Y. Jiang, B. E. Garber, S. P. Carman, and J. M. Hogan, “Large Momentum Transfer Clock Atom Interferometry on the 689 nm Intercombination Line of Strontium,” *Physical Review Letters* **124**, 083604 (2020).
- [175] S.-w. Chiow, T. Kovachy, H.-C. Chien, and M. A. Kasevich, “ $102\hbar k$ Large Area Atom Interferometers,” *Physical Review Letters* **107**, 130403 (2011).
- [176] B. Plotkin-Swing, D. Gochnauer, K. E. McAlpine, E. S. Cooper, A. O. Jamison, and S. Gupta, “Three-Path Atom Interferometry with Large Momentum Separation,” *Physical Review Letters* **121**, 133201 (2018).
- [177] P. Cladé, S. Guellati-Khélifa, F. Nez, and F. Biraben, “Large Momentum Beam Splitter Using Bloch Oscillations,” *Physical Review Letters* **102**, 240402 (2009).
- [178] 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).
- [179] Z. Pagel, W. Zhong, R. H. Parker, C. T. Olund, N. Y. Yao, and H. Müller, “Symmetric Bloch oscillations of matter waves,” *Physical Review A* **102**, 053312 (2020).
- [180] M. Gebbe, S. Abend, J.-N. Siemß, M. Gersemann, H. Ahlers, H. Müntinga, S. Herrmann, N. Gaaloul, C. Schubert, K. Hammerer, C. Lämmerzahl, W. Ertmer, and E. M. Rasel, “Twin-lattice atom interferometry,” (2019), [arXiv:1907.08416](https://arxiv.org/abs/1907.08416).

- [181] P. Cladé, T. Plisson, S. Guellati-Khélifa, F. Nez, and F. Biraben, “Theoretical analysis of a large momentum beamsplitter using Bloch oscillations,” *The European Physical Journal D* **59**, 349 (2010).
- [182] R. Geiger, A. Landragin, S. Merlet, and F. Pereira Dos Santos, “High-accuracy inertial measurements with cold-atom sensors,” *AVS Quantum Science* **2**, 024702 (2020).
- [183] K. Kobzar, B. Luy, N. Khaneja, and S. J. Glaser, “Pattern pulses: design of arbitrary excitation profiles as a function of pulse amplitude and offset,” *Journal of Magnetic Resonance* **173**, 229 (2005).
- [184] 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).
- [185] L. Emsley, I. Burghardt, and G. Bodenhausen, “Double selective inversion in NMR and multiple quantum effects in coupled spin systems,” *Journal of Magnetic Resonance* (1969) **90**, 214 (1990).
- [186] R. Freeman, “Shaped radiofrequency pulses in high resolution NMR,” *Progress in Nuclear Magnetic Resonance Spectroscopy* **32**, 59 (1998).
- [187] M. A. Janich, R. F. Schulte, M. Schwaiger, and S. J. Glaser, “Robust slice-selective broadband refocusing pulses,” *Journal of Magnetic Resonance* **213**, 126 (2011).
- [188] W. S. Warren, “Effects of arbitrary laser or NMR pulse shapes on population inversion and coherence,” *The Journal of Chemical Physics* **81**, 5437 (1984).
- [189] T. Lévêque, A. Gauguier, 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).
- [190] N. Malossi, Q. Bodart, S. Merlet, T. Lévêque, A. Landragin, and F. P. D. Santos, “Double diffraction in an atomic gravimeter,” *Physical Review A* **81**, 013617 (2010).
- [191] X. Wu, F. Zi, J. Dudley, R. J. Bilotta, P. Canoza, and H. Müller, “Multiaxis atom interferometry with a single-diode laser and a pyramidal magneto-optical trap,” *Optica* **4**, 1545 (2017).
- [192] T. L. Gustavson, *Precision rotation sensing using atom interferometry*, Ph.D. thesis, Stanford University (2000).
- [193] S. Hartmann, J. Jenewein, E. Giese, S. Abend, A. Roura, E. M. Rasel, and W. P. Schleich, “Regimes of atomic diffraction: Raman versus Bragg diffraction in retroreflective geometries,” *Physical Review A* **101**, 053610 (2020).

- [194] W. Kallies and S. J. Glaser, “Cooperative broadband spin echoes through optimal control,” *Journal of Magnetic Resonance* **286**, 115 (2018).
- [195] S. Köcher, T. Heydenreich, and S. Glaser, “Visualization and analysis of modulated pulses in magnetic resonance by joint time-frequency representations,” *Journal of Magnetic Resonance* **249**, 63 (2014).
- [196] I. Kuprov, “Spin system trajectory analysis under optimal control pulses,” *Journal of Magnetic Resonance* **233**, 107 (2013).
- [197] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt, “SciPy 1.0: fundamental algorithms for scientific computing in Python,” *Nature Methods* **17**, 261 (2020).
- [198] S. K. Lam, A. Pitrou, and S. Seibert, “Numba,” *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM '15*, 1 (2015).
- [199] J. Johansson, P. Nation, and F. Nori, “QuTiP 2: A Python framework for the dynamics of open quantum systems,” *Computer Physics Communications* **184**, 1234 (2013).