Spinor Bose-Einstein condensates in magnetic field gradients

Alexander A. Wood

BScAdv(Hons)

Submitted in total fulfilment of the requirements of the degree of
Doctor of Philosophy

Supervised by:
Dr Lincoln Turner
Dr Russell Anderson
Prof Kristian Helmerson

School of Physics and Astronomy, Monash University

June, 2015

Copyright notice

Abstract

Spinor Bose-Einstein condensates (BECs) are multi-component superfluids with magnetic interactions, formed from dilute gases of atoms in the Zeeman sublevels m_F of a single hyperfine ground state F. Spinor BECs are inherently *magnetic* quantum objects, and exhibit a multitude of features such as spin domains, magnon excitations, and distinct magnetic phases. At the heart of the magnetic nature of a spinor BEC is the spin-dependent collisional interaction c, which determines the overall magnetic behaviour of the superfluid. In conjunction with the quadratic Zeeman shift q, the magnitude and sign of c give rise to an intricate magnetic phase diagram. The magnetic phase of a spinor BEC is tunable, and c and q are the knobs.

In addition to realising quantum simulators of magnetism, spinor BECs are also precise magnetic sensors. Atomic magnetometry measures the Larmor precession frequency of atoms, which is proportional to the magnetic field strength. The sensitivity of an atomic magnetometer is principally dependent on the duration of the measurement (interrogation time) and number of atoms participating in the measurement. BECs are highly suitable as atomic magnetometers, allowing for long interrogation times and high atomic density, with microscopic sensor volumes. The latter is an important consideration, as measurement of small magnetic sources demand equally microscopic sensors. Established forms of magnetometry (such as warm-vapour magnetometers) use large sensing volumes and are not trivially miniaturised.

The exquisite sensitivity of a spinor BEC to magnetic fields becomes a liability when one studies low energy ($\sim 1\,\text{Hz}$) spinor physics over long evolution times ($100\,\text{ms}$) in a typical laboratory environment. To first order, the spin-collisional interaction is independent of the magnetic field. However, spatially inhomogeneous magnetic fields are problematic, exerting mechanical forces on the magnetically sensitive spin components, suppressing spin collisions and possibly introducing relaxation mechanisms that confound the observation of intrinsic equilibration. Precise characterisation of magnetic field gradients and their cancellation is thus of paramount importance.

In this thesis, the interaction of a spinor BEC with its environment is studied. The construction, operation and optimisation of a BEC machine is described in detail, as well as the quantum state preparation and control tech-

niques used. Spin-mixing collisions were studied by observing population oscillations, the characteristic hallmark of such coherent phenomena. The observed population dynamics showed strong evidence of the deleterious effects of magnetic field gradients. A series of experiments were performed using spin-echo pulse sequences to decouple the evolving spinor condensate from the inhomogeneous magnetic environment with encouraging results, but which posed further questions. The fidelity of spin-echo pulses in the strong radiative coupling regime was also studied.

Spinor condensates were then employed as sensitive magnetometers to characterise the magnetic landscape they inhabit. A new method of measuring magnetic field gradients was developed, using a pair of spinor BECs separated in space that are simultaneously addressed by a Ramsey interferometry sequence. Combined with a re-orientable magnetic bias field, the resulting gradiometer was used to measure the magnetic field gradient tensor *in vacuo*. This scheme can be extended to realise a high-precision vector co-magnetometer, with spatiotemporal sensitivity comparable to established forms of magnetometry.
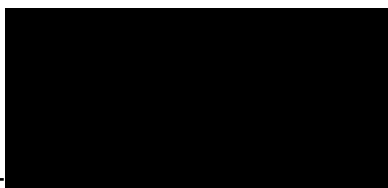
Vector light shifts (VLSs) are optically-induced energy shifts that result in effective magnetic fields, which in our case originate from the laser beams used to trap the atoms. The VLS vanishes for purely linearly polarised light, which is experimentally difficult to obtain *in vacuo*. Small elliptical polarisation imperfections originating from birefringent optics and the glass vacuum cell result in a non-zero shift that contaminates magnetometry measurements. The VLS is spatially varying due to the intensity gradient of the trapping light. Due to gravitational sag of the trapping potential, the spatial variation of the VLS is accentuated, resulting in an effective magnetic field gradient that contributes to dephasing and component separation. The simultaneous Ramsey interferometry experiments are then adapted to measure the VLS-induced Zeeman shift and diagnose its elimination. Up to 99.96% of the VLS is eliminated using this method, and extensions to the technique are described.

Declaration

This thesis contains no material that has been accepted for the award of any other degree or diploma in any university or other institution. To the best of my knowledge the thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis. For parts of this thesis that are based on joint research or publications, the relative contributions of the respective authors are detailed appropriately:

- *Apparatus construction.* My contributions to the design, construction and operation of the apparatus were extensive. The elements focused on in Chapter 3 reflect my primary contributions. Specific sub-projects led by other members of the group are also listed in Chapter 3.

- *Experimental spin-mixing dynamics.* These experiments were led by myself, with supervision from L. D. Turner and R. P. Anderson. Theory to describe beyond-RWA dynamics was developed in conjunction with R. P. Anderson.

- *Magnetic tensor gradiometry.* Experiments were devised by myself, L. D. Turner and R. P. Anderson. Experiments were performed and data collected by myself, L. M. Bennie, A. Duong, R. P. Anderson and with assistance from M. Jasperse. Data analysis was performed by R. P. Anderson and myself. Numerical simulations of magnetic bias coils were performed by R. P. Anderson.

- *Vector light shifts.* Experiments were devised, conducted and analysed by myself with supervision by L. D. Turner and R. P. Anderson.
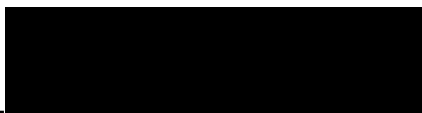
Alexander Wood

Dr Lincoln Turner

Dr Russell Anderson

Prof Kristian Helmerson

Acknowledgements

My life as a scientist has been shaped and influenced by many people, but none more so than my supervisors Lincoln Turner and Russell Anderson. Their genuine passion for science and communicating it, as well as their warm friendship and unceasing support has made my PhD a truly rewarding experience.

My associate supervisor Prof Kris Helmerson was also an invaluable source of advice and genuine wisdom, able to offer profound insight into any problem I sought him out for. Prof Robert Scholten provided great advice as my external supervisory panel member, and was always happy to provide help with our lasers when we needed it. I look forward to working with Rob in the future.

During my PhD candidature I had the great fortune to be a member of the Monash BEC group. In itself, this made the PhD experience all the more engaging and it has been a pleasure to work in such a remarkable group. I thank Dr Michael Egorov and Dr Tapio Simula for the advice and assistance they have provided me.

The other students I had the good fortune to work with in the group deserve special mention; Martijn, Lisa, Phil, Watkins, Chris B, Shaun, Andrew, Prasanna, Chris W and more recently Sam and Seb have been good friends and a pleasure to work with. Particular thanks go to Martijn, Chris, Phil and Shaun for their invaluable assistance in improving my coding proficiency. Although I never had a chance to work with him, I also thank Dr. Gary Ruben for his fantastic thesis template.

As well as doing their utmost to make my PhD a pleasant, stream-lined experience Jean Pettigrew, Dr Alexis Bishop, Prof David Paganin and Prof Michael Morgan were kind enough to provide helpful advice, of both academic and non-academic nature. I was fortunate enough to receive a post-graduate student travel grant to attend the DAMOP '13 conference in Quebec City, Canada, which would not have been possible without the assistance provided by the School of Physics.

I also thank the tireless efforts of Rob Seefeld in making our transition to the New Horizons centre as painless as possible, and Darren and the crew at Horne for their professional execution of the coordinate transform we undertook in March 2013. The New Horizons centre itself should also

be commended for its efforts to ensure our skills as experimentalists were always honed.

As an experimentalist, I am particularly indebted to the Electronics and Mechanical Workshops at Monash and the extraordinary staff therein: Brett Williams, Alan Holland, Roger Bucknell, Stephen Downing, Antonio Benci, David Zuidema and John Golja.

I would also like to acknowledge the assistance provided in the first year or so of my PhD by Vlad Negenevitsky. Vlad is a tremendous physicist, excellent programmer, and his attention to detail and supreme laser skills were an inspiration for me.

I have been lucky enough to share office space and the PhD experience with some truly fantastic people, and I thank them for their friendship, humour and tolerance of my fungus-encrusted coffee cups. I am particularly indebted to Jian Yen, whose friendship, sound advice and refreshing take on science were invaluable influences for me.

None of this would have been possible without the love and support of my family. Thanks, Mum and Dad.

Contents

List of Figures

Introduction

<div style="text-align: right">1</div>

This thesis describes experiments using spinor Bose-Einstein condensates (BECs) as sensitive probes of their environment. BECs are macroscopic quantum objects that allow a range of phenomena integral to the behaviour of myriad quantum systems to be observed in a simple, controllable form on modest spatio-temporal scales; from μm to mm and ms to seconds. BECs are isolated from thermal contact with their environment, as they exist in a vacuum cell. However, a number of environmental factors such as ambient magnetic field gradients and any residual spin-dependence of the trapping potential mean that dephasing mechanisms still exist, and can have deleterious consequences for observing fragile, low-energy spin interactions. This thesis investigates these effects, exploring the behaviour of spinor BECs in magnetic field gradients, and using spinor BECs as high-precision sensors to characterise the magnetic environment.

This thesis is divided into three parts. The construction and optimisation of experimental apparatus to produce and scientifically interrogate BEC is an elaborate process, which has become only marginally simplified over time. This forms the first part of this thesis. In the next part, I present an experimental study of spin-dependent collisional dynamics in magnetic field gradients. This work motivates the third part of the thesis, the precise characterisation of magnetic field gradients and trapping-potential induced spin-dependent light shifts. This introductory Chapter sets the scene for my work in the context of the achievements in spinor BEC physics to date.

1.1 Bose-Einstein condensates

The study of Bose-Einstein condensates is now a mature field; twenty years have passed since the first BECs were manufactured in 1995 by the JILA [1], MIT [2] and Rice groups [3]. A plethora of physics has issued forth, and BEC is being made in physics laboratories across the world. While none of the intrinsic interest in these fascinating macroscopic quantum objects has

waned, the uses to which BECs can be put is as relevant as ever.

There is significant interest in BEC and ultracold non-degenerate gases for realising quantum simulators [4]. A BEC is a clean, highly controllable system in which to study the fundamental quantum behaviour of analogue condensed matter systems, where constituent atoms are typically strongly interacting and coupled to a range of environmental factors. Computer simulations of these systems are generally infeasible or limited to a few special cases. Effects limited to extreme circumstances in real condensed matter systems, such as spin-orbit coupling, can be realised in ultracold atomic systems [5]. Ultracold atoms are essentially real life 'spherical cows' in vacuum chambers [6]: real quantum systems where real quantum behaviour is directly observable on comparatively macroscopic scales.

While insights into fundamental physics can be gained using BECs as quantum simulators, they also offer the prospect of advanced sensors. In this capacity cold atoms and BECs have been used to precisely probe the strength of gravitational [7,8] or magnetic fields [9–11].

1.2   Spinor BECs

Spinor BECs are multi-component condensates composed of different internal spin states of an atom. Our spinor condensates are formed from a superposition of Zeeman sublevels m_F of the $F = 1$ hyperfine ground state of ^{87}Rb, and are therefore three-component BECs, although five-component $F = 2$ [12] and seven-component $S = 3$ ^{53}Cr [13] have been produced. The spin degree of freedom is frozen in a magnetic trap, as only weak-field seeking states ($m_F = -1$ in $F = 1$ ^{87}Rb) can be trapped. Spinor BECs are produced in spin-independent trapping potentials, primarily optical dipole traps, which use the AC Stark shift of an atom in a high-intensity, off-resonant *linearly polarised* laser field to form a conservative potential that traps all m_F states equally.

We reserve Chapter 2 for an in-depth discussion of the myriad interesting features exhibited by spinor condensates, and summarise here a few important points. The key feature of spinor BECs is their fundamentally magnetic nature: they form magnetic spin domains [14] and exhibit ferro- and anti-ferromagnetic ground states [15] (without Néel order). Moreover, this quantum magnetism can be visualised and controlled directly: transitions between magnetic phases of BECs have been experimentally performed, showing the rich spatial dynamics at play [16,17]. At the heart of the magnetic nature of a spinor BEC is the spin-dependent collisional interaction: atoms in different spin states colliding coherently and reversibly, allowing the

population of spin states within the BEC to change without heating or trap losses [18] - something generally not possible with other multicomponent condensates, where collisions are typically associated with loss processes and dephasing. Mediated by these coherent collisions, the populations of the spin components of a condensate can then oscillate in time, in some cases for up to several seconds. The magnitude and sign of the spin-dependent interaction, c, determines the magnetic nature of the condensate, as well as the details of the population oscillations.

Despite being vivid quantum simulators of magnetism, the same benefits associated with a magnetic superfluid come with caveats. Spinor BECs are *very* sensitive to magnetic fields: this makes them fantastic magnetometers, as demonstrated in Ref. [10] and in Chapter 6 of this thesis. Spin-mixing interactions are immune to the linear Zeeman shift from a magnetic field, with only the weaker quadratic shift contributing. However, inhomogeneous linear Zeeman shifts introduce unwanted dephasing mechanisms that suppress the low-energy spin dynamics.

The first step to ensuring a truly clean environment for spinor quantum gases is precise characterisation of the magnetic landscape. It is not just magnetic field gradients that lead to dephasing in spinor experiments – the trapping potential itself is *assumed* to be spin-independent. For far off-resonant, high power laser beams used in optical dipole traps this is generally the case, provided the light is linearly polarised to an extremely high purity.

Any ellipticity results in an atomic vector light shift (VLS), a spin-state dependent energy shift that behaves like an effective, or 'fictitious', magnetic field. These fields may be fictitious, but that in no way makes them negligible, or any less problematic than real magnetic fields and gradients. Along with the overall VLS effective field is an effective *gradient*, proportional to the intensity gradient of the light. For tightly focused trapping beams, these gradients can be comparable to typical background gradients of order 10 mG/cm even for polarisation imperfections of order 1%. Precision characterisation and elimination of the VLS is of considerable interest to other areas in atomic physics: it contaminates optical magnetometry experiments [19, 20] and limits the attainable temperatures in single atom optical tweezer experiments [21, 22].

1.3 Spinor BEC – the state of the art

Spinor BEC research has matured. During the course of my PhD, two major review articles were published, one detailing primarily the theoretical advancements in the field [23] and the other strongly inclined towards

experimental progress [24]. Aside from the Les Houches Summer school notes [25] and several high-quality PhD theses, these review articles constitute the first major attempts to summarise the achievements and outstanding questions in the field of spinor condensates.

Spinor BECs are at the forefront of quantum simulation with neutral atoms. Recent progress has seen the experimental realisation of a Dirac monopole in a spinor BEC [26], an exquisite example of the capability of these systems for quantum simulation of exotic physics. A number of interesting proposals exist for further experiments, such as creating an interface between topologically distinct phases of a BEC [27].

The eminent controllability of these atomic systems can potentially be harnessed to change the magnetic nature of a spinor BEC. By manipulating the atomic scattering properties of the atoms, techniques such as optical Feshbach resonances could be used to change the sign of the spin-dependent interaction, effectively allowing a ferromagnetic BEC to be turned into a antiferromagnetic one, or *vice versa* [28, 29]. This remains an elusive goal due to the difficulty of implementing non-magnetic scattering resonances in the alkali atoms, but a worthwhile one nevertheless. It was this aim that initially formed the overall goal of my thesis.

Precision sensing with spinor BECs, however, has progressed rather slowly. Since the first results from the Stamper-Kurn group at Berkeley in 2007 [10], there has been little progress in realising the full capabilities of spinor BECs as precision magnetic sensors. There are reasons as to why this is so. Magnetometers are parametrised by both their sensitivity and sensing volume. Interesting, microscopic magnetic samples typically require highly sensitive, similarly microscale magnetometers. The most *sensitive* magnetometers measure polarisation rotations in optically pumped atomic vapour ('warm vapour' magnetometers) and attain sensitivities below $1\,\mathrm{fT}\,\mathrm{Hz}^{-1/2}$ [30]. These devices interrogate many spins, but necessarily have large sensing volumes: the useful sensing volume is typically the size of the atomic vapour cell, of order $\sim \mathrm{cm}^3$.

Obtaining the superlative sensitivities demonstrated by large atomic magnetometers in microscopic sensing regions is a profound challenge. Small volume atomic magnetometers necessarily interrogate fewer spins, increasing the effects of quantum noise. Recent advances have resulted in microfabricated glass cells with sub-femtotesla sensitivities [31] and sensor volumes of $\sim 1\,\mathrm{mm}^3$. Ultra-small volume magnetic sensors can be formed from SQUID devices [32] or nitrogen-vacancy (NV) centre detectors [33], with spatial resolutions ranging from μm to nm. The sensitivity of microscopic sensors is less than that of larger sensors that interrogate more spins, with NV centre probes typically achieving $\sim 100\,\mathrm{nT}\,\mathrm{Hz}^{-1/2}$ per active centre [34].

Where do BEC sensors fit in, and are condensed sensors *really* better than thermal atomic magnetometers? This parallels other discussions within the precision measurement community, such as the comparison of thermal and condensed sources for precision gravimetry [8]. In the first 'BEC magnetometer', that of Ref. [9], it is the coherent properties of the condensed state that confer sensitivity to magnetic fields, but this does not necessarily mean it is more sensitive than other atomic magnetometers of a similar size.

Harnessing the spin degree of freedom in atomic gases gives a much more direct measurement of a magnetic field, and in this case a spinor BEC allows the macroscopic coherent properties of the condensed state to be combined with the precision of atomic spin manipulation. Magnetometry with spinor condensates has advantages over thermal gases: compared to thermal clouds, condensed clouds have a higher density of spins to interrogate in a microscopic volume, which reduces the effects of atomic shot noise. A sensor formed from a thermal cloud can have many more atoms than a condensed cloud, but is necessarily larger; a small thermal cloud of the same spatial extent as a condensate is correspondingly less dense. Condensation can be seen as an inevitable consequence of increasing the density of the atomic sample within microscopic volumes, making BECs perfect for high-precision small-volume magnetometry. The nature of collisions in a condensed cloud is also different to that of a thermal cloud – spin-mixing collisions conserve the condensate magnetisation – and allow for lifetime-limited measurement times [35] with concomitantly higher precision.

Consider the parameters of the most prominent example to date of a spinor BEC-based magnetometer, that of Ref. [10], which has a sensitivity of $8.3\,\mathrm{pT}\,\mathrm{Hz}^{-1/2}$ and a sensing region of $\sim 100\,\mathrm{\mu m}^2$, *i.e.* an effective spatial resolution of $10\,\mathrm{\mu m}$. This is roughly in-between the two extreme specifications mentioned earlier. The applicability of BECs to useful magnetic measurement is often criticised due to the requirements of in-vacuum measurement with vacuum-compatible samples, and generally speaking, a BEC apparatus is not particularly field-deployable, yet. Nevertheless, there are prospects for practical magnetic sensing with such a scheme [36].

1.4 What's new in this thesis

Initially, my project was to devise methods to coherently control collisions in spinor BECs. In order to implement and quantitatively evaluate these schemes, we require an unambiguous signature of spin-dependent collisional dynamics – namely spin population oscillations. The magnitude and sign of the spin-dependent interaction c can be determined by observing population

oscillations. If we want to change c, this is a good place to start.

But first, it was necessary to build the apparatus to produce BEC. Although the BEC recipe is a well known one, it is still complex and well worth documenting for the sake of posterity. Every BEC apparatus is unique, improving on previous designs or utilising schemes from other apparatuses. Every step towards simplifying and optimising the BEC production process is a step closer to realising practical technologies with quantum gases. Our apparatus routinely produces spinor condensates of $\sim 5 \times 10^5$ atoms every 25 s, and utilises a hybrid magnetic quadrupole crossed-beam optical dipole trap method, which is being adopted in many labs since the initial realisation [37].

Having built a spinor BEC factory, we immediately put it to work. However, we were unable to observe sustained, large amplitude population oscillations. Inhomogeneous magnetic fields – both real and fictitious – meant that the delicate spin-mixing oscillations were strongly suppressed, and maximum evolution times limited to 50 ms. We conducted a series of experiments to decouple a spinor BEC from this treacherous magnetic environment. We used spin-echo pulses to reverse the deleterious effects of magnetic field gradients, switching the populations of the magnetically sensitive spin components before the gradients could pull the BEC apart. These experiments met with limited success. Much investigation eventually led to the understanding that this was primarily due to the time-varying nature of vector light shift-induced effective field gradients manifesting as irreproducibility. Nevertheless, we were able to demonstrate successful rephasing of spin-mixing oscillations in a noisy, inhomogeneous magnetic environment, with coherent population oscillations preserved to almost 200 ms for particular choices of rephasing times. We also studied the robustness of radiofrequency π-pulses beyond the rotating-wave approximation, and the consequences this has for demonstrating improved spin-echo rephasing schemes in noisy magnetic environments.

We then set out to perform a series of measurements to precisely characterise the magnetic landscape of a spinor BEC. Using simple modifications to the optical dipole trap, we are able to create a pair of spinor BECs, in two separate dipole traps spatially separated by just less than 1 mm. We developed a new measurement technique, differential Ramsey interferometry (DRI), that allows us to use the pair of spinor BECs as a high-precision, interferometric magnetic gradiometer. In this scheme, radiofrequency (rf) spin-manipulation pulses are used to drive transitions between all Zeeman states m_F of each BEC. DRI simultaneously addresses the pair of BECs: each BEC is an independent spin-1 interferometer, and the relative phase of Ramsey fringes is proportional to the magnetic field difference between the two BECs.

Shot-to-shot magnetic field fluctuations of order 0.1 mG scramble the phase of Ramsey fringes from each interferometer for interrogation times longer than 1 ms, which impedes high-precision interferometry at long interrogation times. However, parametrically plotting the output from one interferometer against the other yields an ellipse: the *relative* phase is unaffected by common-mode magnetic field noise. Differential interferometry exhibits common-mode rejection of 50 dB, making high-precision gradiometry and magnetometry possible in noisy environments.

The dipole trap beam positions are easily reconfigurable: by changing the position of the BECs and the orientation of the magnetic bias field, we are able to measure the entire magnetic field gradient tensor. This is the near-definitive measurement of the magnetic field spatial variation on the scale of a trapped quantum gas, a significant result. Our magnetic gradiometer is a substantial achievement in its own right, demonstrating high sensitivity at microscale spatial resolution and has applications as potential high-precision *in vacuo* co-magnetometer. Conservative enhancements to the technique could result in spatiotemporal sensitivities comparable to or even surpassing established, mature forms of magnetometry.

The magnetic gradiometry experiments were performed in freefall to circumvent the spurious contribution of vector light shifts to the measured gradients. This means that the long interrogation times necessary for superlative sensitivity could not be reached, as the cloud will have long since left the microscopic interrogation region: the longest freefall interrogation time used was 3 ms. The vector light shift was also found to change shot-to-shot due to thermal polarisation transients in acousto-optic modulators used for intensity and position control of the trapping beams. High-precision in-trap gradiometry is possible only when the vector shift is eliminated; this motivated a systematic characterisation of the VLS, adapting the techniques developed for measuring magnetic field gradients to realise high-precision measurement of optically-induced 'fictitious' magnetic fields.

The origin of the vector shift is a polarisation imperfection in the dipole trapping light, due to passage through birefringent optics (such as the glass vacuum cell) making the near linear input light slightly elliptical at the atoms. Carefully placed Glan-laser polarising prisms were ultimately used to eliminate the polarisation drifts that wreaked havoc with spin-mixing experiments and differential interferometry was used to measure the trapping light polarisation *in vacuo*. A single quarter-wave plate was then used to counter the birefringence of the downstream optics and glass cell, which in theory is sufficient to attain perfect linear polarisation in the vacuum cell. Using this method we were able to eliminate up to 99.96 % of the vector light shift.

1.5   Outline of this thesis

This thesis is divided into eight Chapters. In the next Chapter, we focus on the theoretical description of spinor BECs, deriving the single-mode theory from the second-quantised many-body Hamiltonian. Chapter 2 also contains a more detailed summary of the major results in experimental spinor BEC physics. The key results of this Chapter are the mathematical formulation of single-mode spin-mixing dynamics, which is explored in more detail experimentally in Chapter 5.

The construction and optimisation of the BEC apparatus is described in Chapter 3. The most important specifications of the experiment such as atom numbers and temperatures at various points of the experiment are included here, and serve as a benchmark for future generations in the Spinor BEC lab at Monash. Chapter 4 continues the discussion about experimental apparatus but focuses on the techniques used for state preparation and manipulation.

Our experimental studies of spin-mixing dynamics are presented in Chapter 5. In this Chapter, we describe how the presence of magnetic field gradients suppresses spin-mixing dynamics, leading to substantially reduced coherent evolution times. We implemented a spin-echo rf pulse sequence in an attempt to mitigate the effects of gradients, with modest success, but confounded by time varying vector light shifts which resulted in irreproducibility. We describe these effects, including the fidelity of refocusing pulses outside the domain of validity of the rotating-wave approximation in this Chapter as well.

Our magnetic tensor gradiometry experiments are discussed in Chapter 6 and the work on characterising and eliminating vector light shifts is presented in Chapter 7. These Chapters form the basis of two upcoming publications:

- A. A. Wood, L. M. Bennie, A. Duong, M. Jasperse, L. D. Turner and R. P. Anderson *Magnetic tensor gradiometry using Ramsey interferometry of spinor condensates.* Accepted, to appear in *Physical Review A.*

- A. A. Wood, L. D. Turner and R. P. Anderson *Measurement of vector light shifts using differential Ramsey interferometry of spin-1 Bose-Einstein condensates.* In preparation for submission to *Physical Review A.*

In Chapter 8, we make concluding remarks and outline extensions of the results of this thesis as well as future directions.

Spinor Bose-Einstein condensates

<div style="text-align: right; font-size: 2em;">2</div>

The wealth of physics that comes from unlocking the spin degree of freedom in a BEC cannot be understated. Simultaneous confinement of BECs in different internal states allows the study of multicomponent superfluidity, with intercomponent interactions and multicomponent ground states a few of the intrinsically interesting facets. A spinor condensate is but one example of such fascinating quantum matter, an example of a *magnetic* superfluid which also has practical applications: a high precision sensor for magnetic fields or a quantum amplifier, for instance [10, 38].

It is the purpose of this Chapter to explain what a spinor BEC is and why it is interesting. The first part is a brief review of scientific literature, outlining the major achievements, from the creation of the first spinor BECs in 1998 to the first major reviews which appeared in 2012-13. In Sections 2.3 and 2.4 the theoretical formulation of a spin-1 BEC is looked at in more detail. As a magnetic condensate, the effect of external magnetic fields and the spin-mixing interaction gives rise to magnetic ground states, this is covered in Section 2.5. The final Sections comprise details about the aspect of a spinor BEC most studied in this thesis, that of coherent spin-mixing collisions.

2.1 Introduction

The first spinor BECs were produced by the Ketterle group at MIT in 1998 [39] from a BEC of sodium atoms loaded into an optical dipole trap from a magnetic trap. Only one Zeeman state in the hyperfine ground state of ^{23}Na can be magnetically trapped; the optical dipole trap confines all three Zeeman sublevels of the $F = 1$ manifold equally, and radiofrequency or microwave fields can be used to create various superpositions of Zeeman states once loaded into the dipole trap. Several seminal results followed, the most notable of which was the exploration of the ground state magnetic phase diagram of a sodium spinor BEC [14], which highlighted the role of both the quadratic Zeeman shift and conservation of magnetisation in

determining the ground state phase diagram. These experiments confirmed the antiferromagnetic nature of the $F = 1$ sodium spinor BEC ground state.

Early theoretical works by Ho [15] and Ohmi and Machida [40] gave the first in-depth theoretical treatment of the zero magnetic field properties of an optically trapped spinor gas. Law et al. had identified the process of spin collisional mixing [41], however without consideration of the effect of magnetic fields. Later works described the single-mode approximation (SMA) [42] and the effect of the quadratic Zeeman shift arising from magnetic fields [43].

The next period of substantial experimental progress occurred in 2005, from the groups of Sengstock and Bongs (Institut für Laserphysik, Hamburg) and Chapman (Georgia Tech). The population of Zeeman states in a spinor BEC is generally not constant in time: coherent, reversible collisions between atoms result in coherent population oscillations, called *spin-mixing oscillations*. The Hamburg group observed spin-mixing oscillations in an $F = 1$ [44] and $F = 2$ [12] ^{87}Rb spinor BEC, albeit in the presence of a thermal component that lead to dissipation and short evolution times. The Chapman group made the first unambiguous observations of long-timescale spin population evolution in $F = 1$ ^{87}Rb [18, 45], confirming the predictions of their previous theoretical formulation of spinor condensates in the SMA [46].

Spin-mixing oscillations in optical lattices were investigated in the group of Bloch (Mainz), who observed population dynamics between pairs of ^{87}Rb atoms in single sites of an optical lattice [47]. They also demonstrated coherent control of spin-mixing collisions by applying a detuned microwave field to create an effective quadratic Zeeman shift [48].

Non-destructive imaging of spinor BECs was developed [35] by the group of Stamper-Kurn (UC Berkeley), and in a remarkable experiment they showed the formation of spin domains in an $F = 1$ ^{87}Rb condensate that was quenched across a phase transition from a polar state into a ferromagnetic state [16]. Another experiment from this group in 2007 used spinor BECs as small volume, high-precision magnetometers by observing spatially resolved Larmor precession across the condensate due to an imposed optical vector light shift [10]. The year 2007 also saw the first observation of spin-mixing oscillations in ^{23}Na, in the Lett group at NIST Gaithersberg [49]. Later results from this group demonstrated a minimally destructive dispersive measurement of *in situ* spinor evolution using Faraday rotation of an off-resonant probe laser [50] as well as investigations into long-term dissipative relaxation [51].

Considerable theoretical interest was also devoted to dipolar effects in spinor condensates [52–55] as well as experimental probing of the weak dipolar effect in ^{87}Rb [56]. Dipolar effects significantly enrich the physics

of spin-1 condensates, but are weak and elusive in alkali metals. Strong dipolar effects are a motivation for producing spin-3 BECs of chromium [57], a system where – due to strong dipolar interactions – magnetisation is no longer conserved [58], collisions can change the magnetisation of the BEC. Experimental efforts to study dipolar effects in alkali metal spinor BECs continues [59], with more magnetic atoms such as erbium [60] and dysprosium [61] offering interesting prospects.

Spinor condensates as sensors for precision metrology have also been explored. Chief among the benefits a spinor BEC brings to atomic magnetometry is the ultra-small volume of the sensor, offering precise, calibration free measurement of magnetic fields over spatial regions of $\sim (10\,\mu m)^3$. Squeezing too has been of interest, with spinor BECs providing a highly suitable system to explore squeezing of macroscopic atomic ensembles [62]. In a spinor BEC, the squeezing arises from the atom-optical analogue of parametric down conversion: spin-mixing collisions produce a entangled pairs of $m_F = \pm 1$ atoms from a pure $m_F = 0$ condensate, with vacuum fluctuations seeding pair creation [63]. Squeezing of up to 10 dB in the nematic spin space of an $F = 1$ ^{87}Rb BEC was demonstrated in condensates of 10^4 atoms [64].

As a magnetic superfluid with an accessible, intricate phase diagram, quantum phase transitions of spinor condensates have received substantial experimental attention in both ^{87}Rb [16] and ^{23}Na [17, 50, 65]. Particular importance is attached to the equilibration behaviour of spinor condensates over long evolution times: as an isolated, essentially dissipationless quantum system, a spinor BEC prepared out of equilibrium takes on the order of seconds to relax to an equilibrium state. The precise mechanism for this equilibration remains unclear [24]. Experiments in ^{87}Rb [66] and ^{23}Na [51] have explored long term equilibration, and others have studied the mean-field phase diagram of ^{23}Na [67, 68].

The study of topological defects and structures in spinor condensates has matured recently, with several notable experimental results [26, 69–72]. Creating an interface between BECs of distinct order parameter symmetries has been proposed, which gives rise to a range of novel topological features at the interface [27].

2.2 General features of a BEC

Before plunging into the details of a spinor BEC, we note here a few general properties of single-component 'scalar' BECs that are of importance. Our experiments typically begin with single-component condensates, and most of what we discuss here carries over to spinor condensates in the single-mode

approximation, such as the characteristic spatial profile of the condensate wavefunction. The detailed theory of BEC is well documented in numerous sources, such as Ref. [73] or the monograph [74]. In this Section, we present only the key theory that aids the understanding of a spinor BEC, as well as other results that are important for the experimental procedure discussed in Chapter 3.

In the mean-field approximation, a single-component ground state condensate described by the order parameter $\psi(\mathbf{r}, t) = \phi(\mathbf{r})e^{-i\mu t/\hbar}$ obeys the Gross-Pitaevskii equation [73]:

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g|\phi(\mathbf{r})|^2 \right)\phi(\mathbf{r}) = \mu\phi(\mathbf{r}), \tag{2.1}$$

with μ the chemical potential, $g = 4\pi\hbar^2 a/m$ (where a is the single-component s-wave scattering length and m the atomic mass) and $\int |\phi(\mathbf{r})|^2 d\mathbf{r} = N$, the number of atoms. The trapping potential $V(\mathbf{r})$ is assumed to be harmonic, $V(\mathbf{r}) = \sum_i \frac{1}{2}m\omega_i^2 x_i^2$, with ω_i the harmonic oscillator frequency of the trap along the spatial axes $x_i = x, y, z$.

To determine the ground-state structure of the BEC, we make the *Thomas-Fermi* approximation: for large N, the nonlinear interaction $g|\phi(\mathbf{r})|^2$ is much larger than the kinetic energy, which we neglect [75]. Solving the GP equation (Eq. 2.1) in the Thomas-Fermi approximation, we arrive at

$$|\phi(\mathbf{r})|^2 = \frac{\mu - V(\mathbf{r})}{g}. \tag{2.2}$$

Since the number density $|\phi(\mathbf{r})|^2 = n(\mathbf{r}) > 0$, we can write the density profile as

$$n(\mathbf{r}) = n_0 \max\left(1 - \left(\frac{x}{R_x}\right)^2 - \left(\frac{y}{R_y}\right)^2 - \left(\frac{z}{R_z}\right)^2, 0 \right) \tag{2.3}$$

where the Thomas-Fermi radii R_i are given by

$$\frac{1}{2}m\omega_i^2 R_i^2 = \mu \tag{2.4}$$

and the peak density $n_0 = \mu/g$. The spatial profile of the condensate is therefore parabolic, with the density going to zero at the Thomas-Fermi radii, as shown in Figure 2.1. In general, the interaction between different components of a multicomponent spinor condensate give rise to complicated time-dependent spatial dynamics. In the discussions that follow, we will introduce the concept of the single-mode approximation, which states that under certain conditions the spatial and spin degrees of freedom of a spinor BEC can be decoupled. In particular, we will consider the case where all three components of an $F = 1$ BEC share the same *spatial* wavefunction,
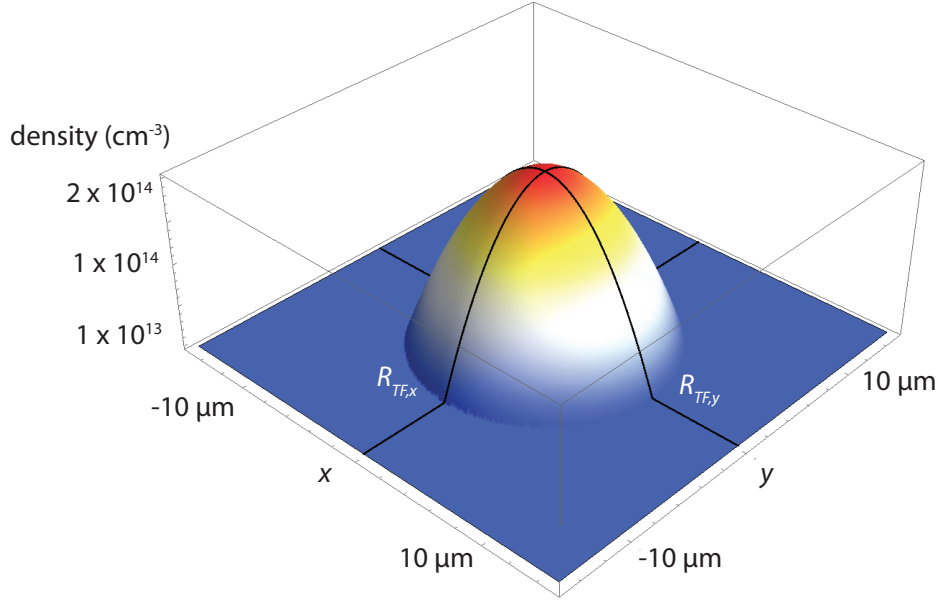
Figure 2.1: Radial Thomas-Fermi density profile ($z = 0$ in Eq. 2.3) of a
3×10^5 atom ^{87}Rb condensate, elongated along z in a harmonic trap with
$\omega_x, \omega_y = 2\pi \times 70\,\text{Hz}$, the Thomas-Fermi radii are $\sim 8.5\,\mu\text{m}$ and the peak
density is $n_0 = 2 \times 10^{14}\text{cm}^{-3}$.

which satisfies Eq. 2.1: the spin-dependent physics are then isolated to the
time-dependent evolution of the relative populations. Another key feature of
a single-component BEC that is important is the condensate healing length

$$\zeta = \sqrt{\frac{\hbar^2}{2mgn_0}},$$ (2.5)

which describes the minimum distance over which density perturbations
exist in a stationary condensate wavefunction. For spinor condensates, there
exists an analogous quantity for the spatial variation of the condensate spin,
called the spin healing length (described in Section 2.4). We will refer back
to some of the results of this Section where useful analogies and pertinent
results can be understood with reference to the case of a single-component
BEC.

2.3 Theory of a spin-1 BEC

In this Section, the theory of a spin-1 BEC is derived. In the interests of
completeness, I feel it proper to derive the primary aspect of a spinor BEC
studied during this thesis (in particular Chapter 5), coherent spin-mixing,
from the theoretical formulation of the spin-dependent collisional interaction.
In doing so, I collate the results of several other works and echo their

discussions, making small corrections and changes to notation for consistency. Among the works I have drawn from are primarily the PhD theses of Ming-Shien Chang and Wenxian Zhang from the Georgia Tech group [76, 77] and the major review articles [23, 24] as well as the 2001 Les Houches summer school notes by Ketterle and Stamper-Kurn [25]. M-S. Chang's thesis was the first document to clearly illuminate the theory of spinor BECs to me, and I have attempted to parallel the clear, concise formulation presented therein.

A spinor BEC is a multi-component BEC composed of a superposition of atoms in different Zeeman sublevels of a single hyperfine manifold F. For an $F = 1$ BEC, the order parameter is a three-component spinor,

$$\boldsymbol{\psi} = \begin{pmatrix} \psi_{-1} \\ \psi_0 \\ \psi_{+1} \end{pmatrix}. \tag{2.6}$$

The mean-field order parameter of a spinor condensate transforms as a vector under rotations in spin space, but in contrast to other multicomponent condensates comes with a rotationally-invariant interaction between the different components. This has important consequences when we consider the collisional interaction that underlies much of the physics of a spinor BEC, and results in a considerable simplification of the interaction Hamiltonian.

2.3.1 Hyperfine and Zeeman interactions

Alkali metal atoms such as ^{87}Rb have hyperfine structure. The nuclear angular momentum $\mathbf{I}$ couples to the electronic angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$, itself the sum of the spin ($\mathbf{S}$) and orbital ($\mathbf{L}$) electron angular momentum. The resulting total angular momentum is $\mathbf{F} = \mathbf{I} + \mathbf{J}$ and the hyperfine interaction Hamiltonian is

$$H_{\mathrm{hfs}} = \frac{A_{\mathrm{hfs}}}{2}(\mathbf{F}^2 - \mathbf{I}^2 - \mathbf{J}^2). \tag{2.7}$$

The hyperfine interaction splits the energy levels of the $5S_{1/2}$ ground state of ^{87}Rb into states $F = 1, 2$ separated by an energy $E_{\mathrm{hfs}} = h \times 6.834\,\mathrm{GHz}$. In the presence of a magnetic field $\mathbf{B}$, the degeneracy of the F hyperfine manifold is lifted by the Zeeman interaction $H_Z = -\boldsymbol{\mu} \cdot \mathbf{B}$, and the eigenstates of $H_{\mathrm{hfs}} + H_Z$ at low fields are $|F, m_F\rangle \equiv |I, J, F, m_F\rangle$, with m_F the projection of $\mathbf{F}$ on the quantisation axis z, normally taken as the direction of the magnetic field.

The energies of the Zeeman states $|F, m_F\rangle$ are dependent on the strength $B = |\mathbf{B}|$ of the magnetic field, and are described by the Breit-Rabi equation, an analytic expression that applies to ground state alkali metal atoms [78–80]:

$$E_{|F=1, m_F\rangle} = -\frac{E_{\mathrm{hfs}}}{2(2I+1)} + g_I \mu_{\mathrm{B}} m_F B - \frac{E_{\mathrm{hfs}}}{2}\sqrt{1 + \frac{4m_F}{2I+1}x + x^2} \tag{2.8}$$

with $x = \frac{(g_J - g_I)\mu_B B}{E_{\text{hfs}}}$. Retaining only m_F-dependent terms to second order in a series expansion of Eq. 2.8 for small B gives a total Zeeman energy E_Z for an atom in the state $|F = 1, m_F\rangle$ as

$$E_Z/\hbar = p\, m_F + q\, m_F^2, \qquad (2.9)$$

with

$$
\begin{aligned}
p &= \frac{E_{+1} - E_{-1}}{2\hbar} \\
&= -\frac{g_J - 2(I+1)g_I}{\hbar(2I+1)} B\mu_B \\
&= -2\pi\, B \times 702.37\,\text{kHz/G} \qquad (2.10a)
\end{aligned}
$$

and

$$
\begin{aligned}
q &= \frac{E_{+1} + E_{-1} - 2E_0}{2\hbar} \\
&= \frac{(g_J - g_I)^2 \mu_B^2}{\hbar\, E_{\text{hfs}}(2I+1)^2} B^2 \\
&= 2\pi\, B^2 \times 71.89\,\text{Hz/G}^2 \qquad (2.10b)
\end{aligned}
$$

the *linear* and *quadratic* Zeeman shifts, respectively for B in Gauss. At typical laboratory fields of $B \sim 1\,\text{G}$, the linear Zeeman shift is four orders of magnitude larger than the quadratic shift, but as we shall see in Section 2.4.1, the quadratic shift primarily determines the magnetic interaction of a spinor BEC. The $F = 1$ and $F = 2$ hyperfine levels, Zeeman states and energy shifts from the linear and quadratic Zeeman shifts are depicted in Figure 2.2.

2.3.2 Collisional interactions

Interatomic interactions determine the physics of all ultracold gases. For spinor BECs, the internal state of each colliding atom is fundamental to the outcome of the collision, and underlies the nature of the spin-dependent interaction. Coherent population oscillations are driven by spin-exchange collisions, where a colliding $|m_F = \pm 1\rangle$ pair can result in the production of an $|m_F = 0\rangle$ pair and vice versa.

In the case of atoms in eigenstates $|F_1, m_{F_1}\rangle$ and $|F_2, m_{F_2}\rangle$ colliding, the collision couples the individual angular momenta F_1, F_2 of each atom to form a total angular momentum $f = F_1 + F_2$. To describe atomic collisions theoretically, we make the assumption that the temperature and density of the gas are low enough so that two-body, rotationally symmetric s-wave collisions dominate, and that the interaction can be described as a contact interaction. The interaction between two atoms with angular momenta F_1, F_2 with position vectors r_1, r_2 can be written in the form of a Fermi
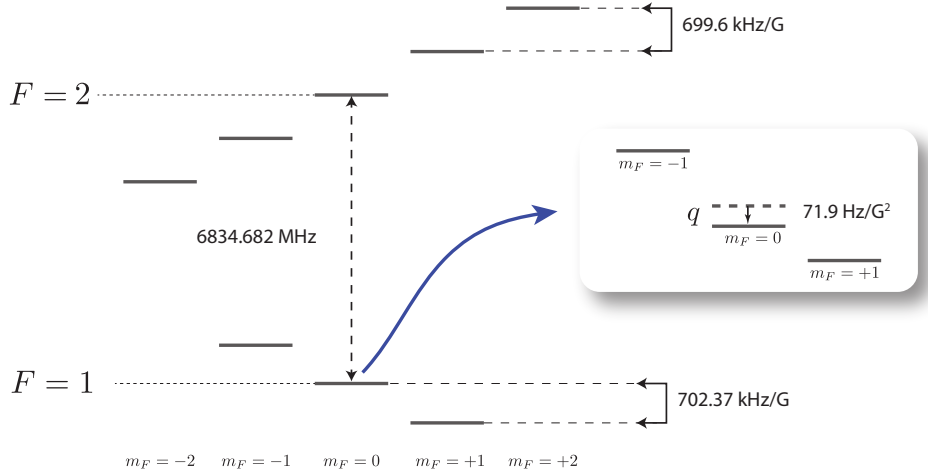
Figure 2.2: Hyperfine splitting, Zeeman states and energy shifts from magnetic fields for the ^{87}Rb $5S_{1/2}$ ground state. Hyperfine levels $F = 1$ and $F = 2$ are split into Zeeman sublevels by a magnetic field; the magnitude of the linear Zeeman splitting is given by $|p|$ in Eq. 2.10a for $F = 1$. The small quadratic shift q (Eq. 2.10b) of the $m_F = 0$ state relative to the other two Zeeman states is shown.

pseudopotential that projects the spin state of the individual atoms $|F_i, m_{F_i}\rangle$ onto the total spin $|f, m_f\rangle$, with $f = |F_1 - F_2|, ..., F_1 + F_2$ [76]:

$$V(\boldsymbol{r}_1 - \boldsymbol{r}_2) = \frac{4\pi\hbar^2}{m}\delta(\boldsymbol{r}_1 - \boldsymbol{r}_2)\sum_{f=0}^{F_1+F_2} a_f \mathbf{P}_f, \qquad (2.11)$$

where $\mathbf{P}_f$ is the projection operator onto the total spin f and the s-wave scattering length for the relevant f collision channel is denoted by a_f: bosonic symmetry restricts the total spin $f = 0, 2, ...$ to even values. For spin-1 bosons, the sum evaluates to:

$$V(\boldsymbol{r}_1 - \boldsymbol{r}_2) = \delta(\boldsymbol{r}_1 - \boldsymbol{r}_2)(c_0 + c_2 \boldsymbol{F}_1 \cdot \boldsymbol{F}_2), \qquad (2.12)$$

with:

$$c_0 = \frac{4\pi\hbar^2}{3m}(a_{f=0} + 2a_{f=2}) \qquad c_2 = \frac{4\pi\hbar^2}{3m}(a_{f=2} - a_{f=0}). \qquad (2.13)$$

The constants $c_0 = 5.16 \times 10^{-51}$ J m^3 and $c_2 = -2.39 \times 10^{-53}$ J m^3 are referred to as the spin-independent and spin-dependent interactions strengths, respectively, with $|c_0| \gg |c_2|$ due to the near equality of the singlet and triplet scattering lengths. The fact that all inter- and intra-component collisional effects can be parametrised by only two parameters is a direct consequence of the rotational symmetry of the system: different states are related to each other by rotations in spin space. The collisional interaction depends only

on the total spin of the colliding atom pair, not on the orientation of the constituent spins, which allows the interaction Hamiltonian to be expressed in the form of Eq. 2.11 [15, 25].

The two-atom picture outlined above can be generalised to the many-body atomic ensemble using the formalism of second quantisation. We introduce the field creation and annihilation operators $\Psi_m^\dagger$ and Ψ_m, with the index $m \in \{+1, 0, -1\}$ corresponding to the Zeeman sublevels m_F. The second-quantised field operators $\Psi_m(\mathbf{r}, t)$ satisfy the commutation relations:

$$\left[\Psi_i, \Psi_j^\dagger \right] = \delta_{ij}\delta(\mathbf{r} - \mathbf{r}'), \tag{2.14a}$$

$$\left[\Psi_i, \Psi_j \right] = \left[\Psi_i^\dagger, \Psi_j^\dagger \right] = 0. \tag{2.14b}$$

The spin-1 Hamiltonian in second-quantised notation (summation over repeated indices is implied) is given by [15]:

$$H = \int d\mathbf{r}\, \Psi_i^\dagger \left(-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r}) + \hat{E}_Z \right) \Psi_i$$
$$+ \frac{c_0}{2}\Psi_i^\dagger \Psi_j^\dagger \Psi_j \Psi_i + \frac{c_2}{2}\Psi_i^\dagger \Psi_k^\dagger (F_\alpha)_{ij}(F_\alpha)_{kl}\Psi_l \Psi_j, \tag{2.15}$$

where we have ignored dipole-dipole interactions. The Zeeman energy operator is $\hat{E}_Z = pF_z + qF_z^2$, we also assume uniform magnetic fields. The matrices F_α are the spin-1 Pauli matrices:

$$F_x = \frac{1}{\sqrt{2}}\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},\; F_y = \frac{i}{\sqrt{2}}\begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix},\; F_z = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}. \tag{2.16}$$

The trapping potential $U(\mathbf{r})$ is state-independent and is assumed to not couple different hyperfine states m_F. The spin-independent and spin-dependent interaction terms c_0 and c_2 are given by Equation 2.13.

The final two terms in the lower line of Eq. 2.15 can then be expanded, and we obtain the spin interaction Hamiltonian:

$$H_{\text{spin}} = \frac{1}{2}\int d\mathbf{r}\left[(c_0 + c_2)\Psi_+^\dagger \Psi_+^\dagger \Psi_+ \Psi_+ + c_0\Psi_0^\dagger \Psi_0^\dagger \Psi_0 \Psi_0 \right.$$
$$+ (c_0 + c_2)\Psi_-^\dagger \Psi_-^\dagger \Psi_- \Psi_- + 2(c_0 + c_2)\Psi_+^\dagger \Psi_0^\dagger \Psi_0 \Psi_+$$
$$+ 2(c_0 - c_2)\Psi_+^\dagger \Psi_-^\dagger \Psi_- \Psi_+ + 2(c_0 + c_2)\Psi_0^\dagger \Psi_-^\dagger \Psi_- \Psi_0$$
$$\left. + 2c_2\Psi_-^\dagger \Psi_+^\dagger \Psi_0 \Psi_0 + 2c_2\Psi_0^\dagger \Psi_0^\dagger \Psi_- \Psi_+ \right]. \tag{2.17}$$

The last two terms in Eq. 2.17 describe the spin-mixing collisions and have a particularly simple intuitive explanation, *i.e.* the annihilation of an $m_F = 0$

pair to create a $m_F = \pm 1$ pair and vice versa. The equations of motion for the field operators are found by solving the Heisenberg equation of motion, $\frac{\partial \Psi}{\partial t} = \frac{-i}{\hbar}[H, \Psi]$. Using the commutation relations Eq. 2.14, this evaluates to three coupled partial differential equations:

$$i\hbar \frac{\partial \Psi_0}{\partial t} = (H_S + E_{Z,0})\Psi_0 + c_0 \left(\Psi_+^\dagger \Psi_+ + \Psi_-^\dagger \Psi_- + \Psi_0^\dagger \Psi_0 \right) \Psi_0$$
$$+ c_2 \left(\Psi_+^\dagger \Psi_+ + \Psi_-^\dagger \Psi_- \right) \Psi_0 + c_2 \Psi_0^\dagger \Psi_+ \Psi_- \tag{2.18a}$$

$$i\hbar \frac{\partial \Psi_\pm}{\partial t} = (H_S + E_{Z,\pm})\Psi_\pm + c_0 \left(\Psi_\pm^\dagger \Psi_\pm + \Psi_\mp^\dagger \Psi_\mp + \Psi_0^\dagger \Psi_0 \right) \Psi_\pm$$
$$+ c_2 \left(\Psi_\pm^\dagger \Psi_\pm - \Psi_\mp^\dagger \Psi_\mp + \Psi_0^\dagger \Psi_0 \right) \Psi_\pm + c_2 \Psi_\mp^\dagger \Psi_0 \Psi_0. \tag{2.18b}$$

with $H_S = -\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r})$. The above equations encapsulate all non-dipole physics of an $F = 1$ spinor BEC.

2.3.3 Mean-field theory

To construct a mean-field theory from the results of the previous Section, we make the assumption of large atom numbers, such that quantum fluctuations may be neglected. With the substitution $\Psi_m \to \psi_m \equiv \langle \Psi_m \rangle$ into the Hamiltonian (Eq. 2.15), we obtain the mean-field energy functional $\langle H_S \rangle + \langle H_{\text{int}} \rangle$ [23]:

$$\langle H_{\text{int}} \rangle = \int d\mathbf{r} \sum_{m=-1}^{1} E_Z n_m + \frac{c_0}{2} n^2 + \frac{c_2}{2} |\langle \mathbf{F} \rangle|^2, \tag{2.19}$$

where the density $n_m = \psi_m^\dagger \psi_m$, $n = \sum_m n_m$ is the total density, and the final term

$$|\langle \mathbf{F} \rangle|^2 = \psi_i^\dagger \psi_k^\dagger (F_\alpha)_{ij} (F_\alpha)_{kl} \psi_l \psi_j \tag{2.20}$$

is interpreted as a magnetisation density. The bare Hamiltonian contains the kinetic energy terms and spin-independent potential

$$\langle H_S \rangle = \int d\mathbf{r} \sum_{i=-1}^{1} \psi_i^\dagger \left(-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r}) \right) \psi_i. \tag{2.21}$$

The mean-field equations of motion are:

$$i\hbar \frac{\partial \psi_0}{\partial t} = (H_S + E_{Z,0})\psi_0 + c_0 \left(n_+ + n_- + n_0 \right) \psi_0$$
$$+ c_2 \left(n_+ + n_- \right) \psi_0 + c_2 \psi_0^\dagger \psi_+ \psi_-, \tag{2.22a}$$

$$i\hbar \frac{\partial \psi_\pm}{\partial t} = (H_S + E_{Z,\pm})\psi_\pm + c_0 \left(n_\pm + n_\mp + n_0 \right) \psi_\pm$$
$$+ c_2 \left(n_\pm - n_\mp + n_0 \right) \psi_\pm + c_2 \psi_\mp^\dagger \psi_0 \psi_0. \tag{2.22b}$$

Equations 2.22 are the coupled Gross-Pitaevskii (GP) equations for a spin-1 BEC.

2.4 Single-mode approximation

The equations of motion (Eq. 2.22) may be simplified further by making the single mode approximation (SMA) [42]. In the SMA, each spin component is assumed to have the same spatial wavefunction, with the spin dependence isolated to the time-dependent relative population amplitudes. The validity of the SMA hinges on the energetic suppression of spatial structures, or spin domains. The spin healing length ζ_S, in analogy with the healing length of the condensate wavefunction (Eq. 2.5), describes the minimum length scale over which the spin projection of the BEC may spatially vary [24]:

$$\zeta_S = \sqrt{\frac{\hbar^2}{2m|c_2|n}}.$$ (2.23)

For typical densities of $n \sim 10^{14}\,\text{cm}^{-3}$, the spin healing length of ^{87}Rb condensates is $\zeta_S \sim 4\,\mu\text{m}$, which is generally smaller than many trapped clouds; the SMA is only strictly valid for low atomic densities. Provided that the overall spatial extent of the cloud[1] is less than ζ_S, spin population dynamics are well described by SMA theory. Magnetic field gradients inherently violate the SMA by driving spatial separation of the magnetically sensitive $m_F = \pm 1$ components, an effect discussed in Chapter 5 of this thesis. For now, we will neglect this effect and proceed with simplifying Eqs. 2.22 under the SMA. The SMA wavefunction is written as

$$\psi_m(\mathbf{r}, t) = \sqrt{N}\Phi_{\text{SMA}}(\mathbf{r})e^{-i\mu t/\hbar}\zeta_m(t)$$ (2.24)

with N the total atom number and μ the chemical potential. The time dependent spin dynamics are wholly described by spinor $\zeta_m(t)$. All spin components share the same time-independent spatial wavefunction $\Phi_{\text{SMA}}(\mathbf{r})$, which satisfies [23, 81]:

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r}) + c_0 n \right)\Phi_{\text{SMA}}(\mathbf{r}) = \mu\Phi_{\text{SMA}}(\mathbf{r}).$$ (2.25)

This is essentially Eq. 2.1, the time-independent GP equation for the case of a single-component BEC, with $c_0 n$ playing the role of the spin-independent self-interaction term. We now substitute the SMA wavefunction Eq. 2.24 into the GP equations Eq. 2.22. Integrating over space and using the fact that $\int d\mathbf{r}\,|\Phi_{\text{SMA}}(\mathbf{r})|^2 = 1$, we obtain:

$$i\hbar\frac{\partial\zeta_0}{\partial t} = E_{Z,0}\zeta_0 + c\left(|\zeta_+|^2 + |\zeta_-|^2\right)\zeta_0 + c\zeta_0^{\dagger}\zeta_+\zeta_-$$ (2.26)

$$i\hbar\frac{\partial\zeta_\pm}{\partial t} = E_{Z,\pm}\zeta_\pm + c\left(|\zeta_\pm|^2 - |\zeta_\mp|^2 + |\zeta_0|^2\right)\zeta_\pm + c\zeta_\mp^{\dagger}\zeta_0\zeta_0$$ (2.27)

---

[1]Elongated clouds – much longer than ζ_S – where the tight radial Thomas-Fermi radii are less than ζ_S are still well described by SMA theory for reasonably long evolution times.

with $c = c_2 \langle n(\mathbf{r}) \rangle = c_2 N \int d\mathbf{r} |\Phi_{\text{SMA}}(\mathbf{r})|^4$. This final relation is important when we consider violation of the SMA, due to magnetic field gradients. Since the spin interaction strength c is effectively an overlap integral between the spatial wavefunctions of the three Zeeman components, when $\psi_{\pm 1,0}(\mathbf{r}) \neq \Phi_{\text{SMA}}(\mathbf{r})$, the effective spin-mixing interaction strength is reduced.

2.4.1 Conservation of magnetisation

Before proceeding with further simplification of Equations 2.27, we must first consider the conservation of the overall condensate spin. This is an important result, as it allows spinor physics to be observed in realistic laboratory conditions. Two of the energy scales we have considered to this point are roughly comparable at low magnetic fields: the quadratic shift at 250 mG is $q \sim 2\pi \times 4.5$ Hz, the spin interaction energy is $|c| \sim 2\pi \times 4$ Hz for an ^{87}Rb condensate with $n \sim 10^{14}$, but the linear Zeeman shift is $p = 2\pi \times 176$ kHz. For a linear Zeeman shift to be comparable to c, a magnetic field environment of $\sim 1\,\mu$G is required, which would be a singularly demanding experimental task. However, conservation of magnetisation eliminates the linear Zeeman shift from the process of spin-mixing collisions.

We use the terms 'spin projection' and 'magnetisation' interchangeably, as both correspond to the expectation value of the longitudinal spin projection operator.[2] The spin of an alkali metal spinor condensate is a conserved quantity as the relaxation mechanisms are extremely weak – dipolar interactions being the primary means of such relaxation [24]. This means that a condensate prepared in a superposition of Zeeman states with spin $F_z \in \{-1, 1\}$ will persist in that state for a long time. A spinor condensate composed of atoms with strong dipole-dipole interactions, such as ^{52}Cr, does not exhibit conservation of magnetisation [55].

The early theoretical treatments of spinor ground states and dynamics did not consider conservation of magnetisation [15, 40–42]. Conservation of magnetisation is also essential for minimising the mean-field energy to determine the spinor ground state [14]. The mean-field ground state of $m_F = +1$ for $F = 1$ ^{87}Rb is thus not the ground state with the requirement of spin conservation; there is no obvious mechanism to relax to this state [25]. We will discuss this in more detail in Section 2.5.

Conservation of magnetisation has very important consequences for the linear Zeeman shift, as described in Ref. [25]. The mean-field energy of a spinor BEC in the presence of a uniform magnetic field is $E = E_0 + $

---

[2]We will define this as $F_z = \langle \hat{F}_z \rangle$ throughout this thesis (particularly in the context of spin-mixing observables), except where the context requires explicit representation as an expectation value, such as in the following discussion.

$p \int d\mathbf{r}\, n\langle F_z\rangle$, with E_0 the energy of all other interaction terms. The linear Zeeman energy[3] is minimised by making $\langle F_z\rangle = +1$, that is, all atoms in $m_F = +1$. Since magnetisation is a conserved quantity, the condensate with an arbitrary configuration of spins cannot relax to a pure $m_F = +1$ state, E is minimised subject to the constraint $\int d\mathbf{r}\, n\langle F_z\rangle$ is fixed. This is done with the insertion of a Lagrange multiplier, p_0, which depends on the net magnetisation of the condensate. The energy is then minimised subject to this constraint: *i.e.* $E = E_0 + (p - p_0) \int d\mathbf{r}\, n\langle F_z\rangle$. For the case of $\langle F_z\rangle = 0$, $p_0 = p$ and so the linear Zeeman shift vanishes.

An equivalent picture to describe how conservation of magnetisation affects spinor dynamics is to transform to a frame rotating at the Larmor frequency. The magnetic energy of a colliding $m_F = \pm 1$ atom pair is the same as that of a colliding $m_F = 0$ pair, up to the quadratic shift. In the presence of a spatially varying magnetic field the magnetisation is only locally conserved, as the linear Zeeman shift $p(z)$ varies across the length z of the condensate (the associated spatial variation of q is negligible).

Whilst a uniform linear Zeeman shift has no contribution to spin-mixing collisions, the magnetic energy per atom is still dependent on the value of p. Even a small magnetic field gradient results in a variation of p across the extent of a condensate comparable to c and q, resulting in spin domains, as described in the next Section.

2.5 Spinor ground states

The magnetic nature of spinor BECs is most apparent when we consider the ground state properties. For clarity, we will cast Eq. 2.19 in a simpler form, where we consider only the spin-dependent terms of the mean-field interaction energy within the SMA [24]:

$$E = \frac{c}{2}\langle\mathbf{F}\rangle^2 + p\langle F_z\rangle + q\langle F_z^2\rangle. \qquad (2.28)$$

The sign of c_2, and hence c, determines the nature of the mean-field spinor ground state: for $c_2 > 0$, the spinor that minimises Eq. 2.28 is one with $|\mathbf{F}| = 0$. Such an unmagnetised state is called antiferromagnetic if the condensate is composed primarily of $m_F = \pm 1$ states[4] or polar if composed of $m_F = 0$ atoms. For $c_2 < 0$, the condensate favours a magnetised state with $|\mathbf{F}| = 1$, termed a ferromagnetic state. For ^{87}Rb, $c_2 < 0$ and for ^{23}Na, $c_2 > 0$, with the result that these two atomic species are said to have ferromagnetic and antiferromagnetic interactions, respectively.

---

[3] Recall the definition of p in Eq. 2.10a, for positive B, $p < 0$.
[4] There is no spatial correlation of anti-aligned spins, *i.e.* there is no Neél order.
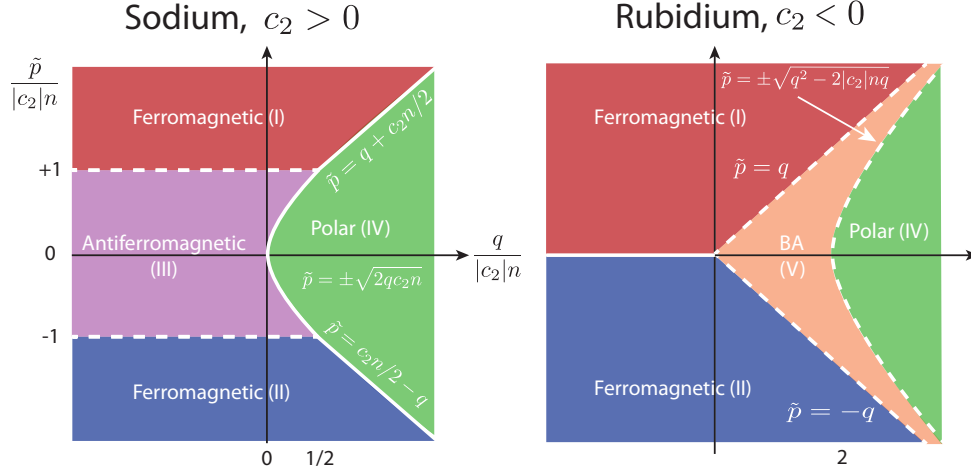
Figure 2.3: Spin-domain diagrams for antiferromagnetic ($c > 0$) ^{23}Na and ferromagnetic ($c < 0$) ^{87}Rb. The mean-field ground state consists of different population distributions depending on the quadratic Zeeman energy q and the linear Zeeman shift/magnetisation $\tilde{p}$, resulting in distinct phases (discussed in body text). In the antiferromagnetic (III) and broken-axisymmetric (V) phases, multiple Zeeman states overlap, whereas the I, II and IV states correspond to pure occupation of $m_F = +1, -1$ and $m_F = 0$, respectively; dashed or solid lines correspond to smooth or sudden changes of population distributions [14]. Figure is the author's work, based on examples in Refs. [14,23].

Consideration of magnetic fields, including conservation of magnetisation discussed previously, leads to a rather more detailed phase diagram. Some of the first high-profile experiments with spinor BECs concerned the elucidation of the ground state phase diagram [14,82]. Figure 2.3 illustrates the different magnetic phases of antiferromagnetic (left) and ferromagnetic (right) spin-1 BECs. The full derivation of such a figure may be found in the original work [14] or in the most recent reviews [23,24]. The mean-field ground state for a given value of c, q and $\tilde{p} = p_0 - p$ can be deduced from such spin-domain diagrams. The $\tilde{p}$-axis is effectively the variation of the linear Zeeman shift from the Larmor frequency: a BEC at a magnetic bias field B_0 with $\langle F_z \rangle = 0$ occupies a single point, $\tilde{p} = 0, q = q(B_0)$. A magnetic field gradient samples a range of linear Zeeman shifts, it is represented on Figure 2.3 as a line. If the overall condensate spin is zero, this line is symmetric about $\tilde{p} = 0$, for positive magnetisations it is shifted to $\tilde{p} > 0$ and for $\langle F_z \rangle < 0$ to $\tilde{p} < 0$.

The boundaries separating different phases correspond to sharp (solid) or gradual (dashed) changes of population across the boundary (a domain wall). The important role of magnetic field gradients is immediately apparent: gradients result in spatial structure in a spinor BEC.

Five magnetic phases [14] are represented in Fig 2.3:

- **I, II Ferromagnetic:** Longitudinally magnetised states $|m_F = \pm 1\rangle$ respectively.

- **III Antiferromagnetic:** Spatially overlapping superposition of $|\pm 1\rangle$ states.

- **IV Polar:** longitudinally polar state, purely $|m_F = 0\rangle$.

- **V Broken-axisymmetric (BA) phase:** superposition of overlapping magnetically-polarised and polar states predominantly consisting of $m_F = 0$. The overall condensate spin is tilted away slightly from quantisation axis, hence the name.

Experimental investigations of quantum phase transitions can be well understood using the map provided by Figure 2.3. One notable example is the study of the polar-antiferromagnetic transition in ^{23}Na by the Raman group at Georgia Tech [17,65]. In their experiments, an off-resonant microwave field induced a detuning-dependent quadratic Zeeman shift that could be swept from positive to negative, the phase transition from polar to antiferromagnetic was studied in terms of spatial structure and population instabilities.

Studies of spinor ground states are linked to studies of relaxation and equilibrium behaviour. Early experiments looked at metastable spin structures and populations [82] over long evolution times of several seconds, and more recently with thermal components [66]. The precise mechanism for such equilibration is not apparent. We will discuss this in more detail later when discussing spin population dynamics in Chapter 5.

2.6 Spin-mixing oscillations

As well as long-timescale equilibration physics, there are many interesting dynamical features of spinor condensates that happen on timescales on the order of 10 ms to greater than 1 s. As mentioned previously, coherent spin-changing collisions result in population interconversion amongst Zeeman sublevels. The process can be represented by

$$|m_F = +1\rangle + |m_F = -1\rangle \rightleftharpoons |m_F = 0\rangle + |m_F = 0\rangle. \qquad (2.29)$$

We can utilise conservation of magnetisation to simplify Equations (2.22) and derive the equations of motion that describe spin-mixing collisions. We define the time-dependent fractional populations $\rho_m(t) = |\zeta_m|^2$ with $\zeta_m(t)$ in Eq. 2.24, which satisfy $\sum_m \rho_m = 1$; the population fractions can equivalently

be represented by N_m/N, the fractional atom numbers in each spin state. We then make the transformation to the rotating frame [46,76]:

$$\zeta_{\pm 1} \rightarrow \zeta_{\pm 1} e^{-i(E_{z,0}\pm p)t/\hbar}; \qquad \zeta_0 \rightarrow \zeta_0 e^{-iE_{z,0}t/\hbar}, \qquad (2.30)$$

and take $\zeta_m(t) = \sqrt{\rho_m(t)}e^{-i\theta_m(t)}$, with $\theta_m(t)$ the phase of each spin component. The single-mode dynamics of a spin-1 BEC can then be described by only two variables: ρ_0, the fractional population in the $m_F = 0$ state and $\theta = \theta_+ + \theta_- - 2\theta_0$, the relative phase:

$$\dot{\rho}_0 = \frac{2c}{\hbar}\rho_0\sqrt{(1-\rho_0)^2 - m^2}\sin\theta, \qquad (2.31a)$$

$$\dot{\theta} = -2q + \frac{2c(1-2\rho_0)}{\hbar} + \frac{2c}{\hbar}\left(\frac{(1-\rho_0)(1-2\rho_0)-m^2}{\sqrt{(1-\rho_0)^2-m^2}}\right)\cos\theta, \qquad (2.31b)$$

with $m = \langle F_z \rangle = \rho_+ - \rho_-$ the magnetisation. Equations (2.31) can also be derived from the classical energy functional [46]:

$$E = (1-\rho_0)q + c\rho_0\left((1-\rho_0) + \sqrt{(1-\rho_0)^2 - m^2}\cos\theta\right), \qquad (2.32)$$

with

$$\dot{\rho}_0 = -\frac{2}{\hbar}\frac{\partial E}{\partial \theta}, \quad \dot{\theta} = \frac{2}{\hbar}\frac{\partial E}{\partial \rho_0}. \qquad (2.33)$$

Equation 2.32 is the energy of a classical non-rigid pendulum, with the dynamical variables (ρ_0, θ) in place of the length and rotational phase of the pendulum. The equal energy contours of Eq. 2.32 are plotted in Figure 2.4 for two different magnetic fields in terms of ρ_0 and θ. Overlaid with the energy landscape are solutions of Equations 2.31 for a state initialised with $m = 0$ and $\rho_0(t = 0) = 0.5$. These are the spin-mixing oscillations, as shown in Figure 2.5. The two particular solutions in Figure 2.4 are differentiated from each other by the phase behaviour: for low fields, the phase oscillates in time (closed contours) whereas at higher fields, the term linear in q in Eq. 2.31b dominates and the phase increases with time. These two distinct regions – *oscillating phase* and *running phase* – are visible in the plot on the left. The boundary known as the *separatrix* defines the crossover point, a contour in the phase space with $E(\rho_0, \theta) = 0$ for a given quadratic shift. For some configuration of ρ_0 and θ, the system will evolve along the separatrix at a quadratic shift $q(B_{sep})$, which for ^{87}Rb, with $c < 0$, is given by [46]

$$q(B_{sep}) = |c|\rho_0(1 + \cos\theta). \qquad (2.34)$$

At the separatrix, the period of the spin-mixing oscillations is longest, and the specific behaviour for different magnetisations differs between ferro- and
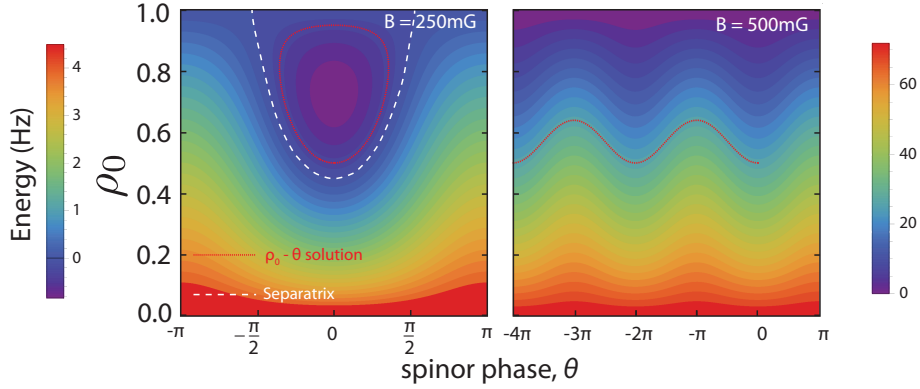
Figure 2.4: Equal energy contours of the classical spinor energy E in terms of the conjugate variables ρ_0 and θ for two different magnetic bias fields (different quadratic shifts), with $m = 0$ and $c = -2\pi \times 5\,\mathrm{Hz}$. The dashed red lines are parametric plots of solutions to the Eq. 2.31, which correspond to spin-mixing oscillations, and the white line on the left plot the separatrix. The example on the left corresponds to an oscillating phase solution, with $B < B_{\mathrm{sep}}$, whereas a running phase solution is shown on the right. The separatrix can be seen to divide the phase space between oscillating and running phase solutions. Note the different energy scales used.

antiferromagnetic atomic species [46]. Returning to the nonlinear pendulum analogy, the oscillating and running phase regimes as well as the separatrix can be given a satisfying physical interpretation. Oscillating phase solutions correspond to familiar libration modes of the pendulum, and running phase solutions to that of rotation around the pivot. The separatrix corresponds to the pendulum perfectly inverted, neither librating nor rotating.

The interplay between the quadratic Zeeman shift and the spin interaction strength determine the period and amplitude of spin-mixing oscillations. The period of population oscillations is not as easily shown as the amplitude in the phase space trajectories of Figure 2.4, in general the period varies considerably with each trajectory. Figure 2.5 shows population oscillations in the time domain for ^{87}Rb, initialised with $m = 0$ and $\rho_0 = 0.5$ for different magnetic bias fields.

Single-mode spinor dynamics are a fascinating manifestation of the spin collisional interaction in BECs. Observing the amplitude and period of spin-mixing oscillations informs the magnitude and sign of c, providing a useful means of characterising the magnetic nature of a spinor BEC. In particular, when the aim of an experiment is to *change c* (by tuning the scattering lengths, for example), the observation of the resulting spin-mixing oscillations allows the imparted change to c to be measured directly.

Spin-mixing dynamics are intrinsically compelling, a manifestation of
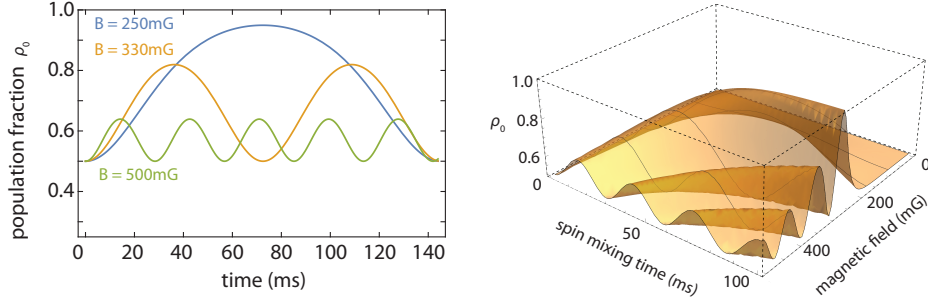
Figure 2.5: Spin-mixing oscillations. The amplitude and period of spin-mixing oscillations is determined by the value of the quadratic Zeeman shift $q = 2\pi B^2 \times 72\,\mathrm{Hz/G^2}$ and the spin-mixing interaction, which in this case is $c = -2\pi \times 5\,\mathrm{Hz}$.

coherent, reversible atomic collisions in a many body ensemble. The analogue dynamics of spin-mixing collisions and oscillations of a non-rigid pendulum have been explored in a recent work, where a condensate prepared in the unstable equilibrium state (on the separatrix) is subject to quantum fluctuations that drive evolution [83]. Application of periodic microwave pulses that modulate q can also be used to prevent fluctuations driving spin population evolution, dynamically stabilising the pendulum [84]. Spin-mixing oscillations are also not limited to condensed clouds, as recent experiments have discovered for both ^{23}Na and ^{87}Rb [85,86].

While conservation of magnetisation results in reasonable experimental requirements for magnetic fields (for $B = 500\,\mathrm{mG}$, $q \sim c$), single-mode spin-mixing oscillations are difficult to observe in the presence of magnetic field gradients: spatial dynamics are incompatible with the SMA, and generally lead to loss of collisional coherence; this can be readily observed in damping of the population oscillations [18]. Magnetic field gradients are often present in labs, originating from magnetic sources such as ion pumps or magnetic vacuum hardware. Whilst the exact phenomenology exhibited in the presence of an inhomogeneous magnetic field depends on factors such as the size and density of the condensate, one can readily estimate the magnitude of gradients that affect spinor condensates from simple energy considerations, and show that small gradients create energy shifts on the order of c and q. Consider a condensate with a spatial extent of $10\,\mu\mathrm{m}$, with $|c| = 2\pi \times 4\,\mathrm{Hz}$ in a magnetic field gradient $B' = 10\,\mathrm{mG/cm}$. This gradient imposes a $2\pi \times 7\,\mathrm{Hz}$ frequency shift across the condensate, almost double c. Gradients of this strength are common in laboratories, and must be cancelled or circumvented in order to observe long-timescale spinor physics. In Chapter 5, we focus on the experimental observations of coherent spin-mixing, in the presence

of these stray magnetic field gradients. In that Chapter, the motivations behind studying spin-mixing oscillations are also discussed, as well as the implications of the experimental results.

2.7 Chapter summary

In this Chapter, we have described how atomic collisions give rise to the myriad rich physics of spinor condensates, such as magnetic ground states and coherent spin-mixing oscillations. The latter is studied experimentally in Chapter 5. The results discussed herein describe the essential theory necessary to understand single-mode spin-mixing oscillations.

Apparatus

<div style="text-align: right; font-size: 2em; color: #8B0000;">3</div>

The first Bose-Einstein condensates of dilute gases were produced in 1995 [1–3]. It is a testament to the complexity of the experimental techniques that it took 70 years before theoretical predictions made by Bose and Einstein could be realised in the lab. A completely new device (the laser) and an entire field of physics (laser cooling) had to be invented along the way. Making a BEC was worthy of a Nobel prize in 1995, whereas today there are about 50 active groups worldwide who routinely make BECs of 12 different atomic species.

Lasers have become cheaper and more powerful, the complicated optics and electronics to stabilise and control them have been engineered into self contained off-the-shelf systems and ultracold atoms have become common-place in atomic physics laboratories. The first ever commercial BEC machine was released in 2013 by ColdQuanta Inc., a company born from the JILA group that made the first BECs. Although still in their infancy, it may be not long before commercial BEC machines are being used in undergraduate teaching laboratories.

The drive to make better BEC machines has largely been influenced by prospective applications, such as metrology and quantum simulation. The future of ultracold science lies in what BECs *do*, and this requires exceptionally well characterised, robust experimental apparatus that mass produce ultracold atoms in the same way that lasers produce coherent light.

After a little more than twenty years since making the first BEC of ^{87}Rb, the procedure may be well-known and highly optimised for the alkali metal atoms, with only a few major variations, but it is still technically challenging. Just as every atom 'marches in lockstep' in a BEC [87], so must every laser, current driver, shutter, camera, radiofrequency antenna and acousto-optic modulator in a BEC experiment. The laboratory environment must be temperature and humidity controlled to combat the nemesis of drift; stray magnetic fields must be actively cancelled or screened; atomic beam sources must be monitored and interlocked to prevent their catastrophic self-annihilation

and the entire vacuum system must be kept at a pressure less than that at
the surface of the moon. Only a few of these requirements have become
simplified over time.

3.1 Getting to BEC

Bose-Einstein condensation in alkali gases is a mature field. The in-depth
theoretical formulation of BEC can be found in any of the notable reviews on
the subject [73,74]. Here I provide a brief summary, sufficient to understand
and motivate the experimental requirements.

A BEC is an ensemble of bosonic atoms that have all condensed into a
single-particle quantum state. This macroscopic occupation takes place under
specific conditions for dilute gases of atoms: it is restricted to certain atomic
species and environmental parameters; the most important of which is low
temperature, a direct consequence of the low-density atomic samples used. It
is generally regarded that BECs are the coldest known objects in the universe,
with temperatures of only 100 nK or less.

Low temperature is the essential requirement for making a BEC out of
an alkali gas. Consider an ensemble of bosonic atoms with mass m and
temperature T. The *thermal de Broglie wavelength* of each atom is given by [88]:

$$\lambda_{\mathrm{dB}} = \sqrt{\frac{2\pi\hbar^2}{mk_BT}},$$ (3.1)

which increases with decreasing temperature. When λ_{dB} becomes comparable
to the interatomic separation, the individual atomic wavepackets begin to
overlap. For a critical temperature T_c, the atoms begin to condense into the
lowest available energy level, which depends on the confining potential.

For gas of N atoms, T_c is such that the *phase-space density* (PSD), $n\lambda_{\mathrm{dB}}^3 = 2.61$,
where n is the peak number density. The fraction of condensed atoms N_0/N
near the critical temperature for a 3D harmonic potential is well-approximated
by [74]

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c}\right)^3.$$ (3.2)

The PSD of a thermal atomic vapour of ^{87}Rb, our starting point, is about
10^{-21} for a temperature of 80°C. To get to BEC we thus need to increase
the PSD by over 20 orders of magnitude. A BEC can only exist when it
is thermally isolated from the ambient environment: for this reason, it is
suspended by electromagnetic forces in an ultra-high vacuum vessel.

The remainder of this Chapter is devoted to the experimental methods
used to manufacture BECs of ^{87}Rb. In Chapter 4 we consider the primary

tool we use to prepare and manipulate spinor BECs: radiofrequency state manipulation.

3.2 Overview

What we now call the spinor BEC machine was a pair of empty optics tables in Room G07a, Building 27, when I first started in the group in Summer 2009. Over the course of four undergraduate projects, three honours projects and the best part of three PhD projects it came into being. It is now a reliable, well understood conglomeration of equipment that can produce over 2800 BECs a day, with a 'service delivery rate' in excess of 90%. Like any other accumulation of occasionally somewhat reluctant equipment it has its days, where it reminds us with every failed experiment run "I'm doing *exactly* what you told me to do!".

In this Chapter I compile the construction and optimisation of the apparatus. The process was inherently a team effort, with various components or processes the focus of one or more lab members. That said, everyone had at least some input into the development of most aspects of the machine, whether from design, development and construction up to using, optimising and of course, debugging.

The most natural way to introduce all the components of the apparatus is to summarise what actually happens with each experimental run. In each 25 s iteration (or 'shot') of the experiment, a BEC of several hundred thousand ^{87}Rb atoms is created by reducing the temperature of a thermal atomic beam to the point it can be trapped, and then performing a series of cooling stages on the trapped gas until the BEC transition is reached. At the conclusion of this process, the temperature has been reduced by nine orders of magnitude and the phase-space density increased by over 20 orders of magnitude. The atom number is reduced similarly; of the 10^{15} atoms that leave the oven during the loading cycle, we trap only $\sim 10^9$ and from that condense only 10^5.

Here I briefly summarise the function of each component and design work I and the rest of the group participated in. Each of stage of the experiment is discussed in more detail later in this Chapter.

Oven

The oven is the atomic source. It forms an atomic beam by heating a 5 g sample of rubidium metal to 80°C and collimating the resultant atomic vapour through a thin tube heated to 120°C. The atomic beam then passes through an aperture in a copper cup kept at −30°C that catches the off-axis atomic

emission from the tube. An in-vacuum shutter allows the atomic beam to be gated, and a pneumatically actuated gate valve forms an isolating vacuum seal between the oven vacuum and downstream UHV system. The oven is controlled by a microcontroller system that ensures stable temperatures and robustness to component failures.

I was particularly involved in the development of the oven, and it constituted my undergraduate Summer (2009) and Honours (2010) projects [89,90]. Development and testing of the microcontroller and associated electronics were carried out in conjunction with M. Jasperse, and are discussed further in Section 3.3.

Zeeman slower

The atomic beam is at a temperature of 353 K, corresponding to an average atomic velocity of > 300 m/s. The atomic beam is used to load the magneto-optical trap (MOT), which has a capture velocity of ~ 30 m/s for typical MOT parameters[1]. It is therefore necessary to slow the atoms down, which is accomplished using the Zeeman slower. Our Zeeman slower is a length of vacuum tubing, surrounded by a variable pitch solenoid [37,91]. A red detuned laser propagates against the atomic beam. Atoms in the beam absorb photons and experience a momentum kick opposite to the direction of motion, and after many scattering events the atoms experience an overall deceleration and slow down. The changing velocity results in a decreasing Doppler shift of the oncoming photons. The varying pitch solenoid produces a magnetic field that adds a compensating Zeeman shift that varies along the length of the slower to keep atoms resonant with the slowing laser.

The design and construction of the Zeeman slower was the third-year undergraduate project of fellow PhD student L. Bennie [92]. Installing it on the apparatus and getting it operational was a whole-group effort. More details about the specifics of the Zeeman slower are given in Section 3.3.3.

MOT

The MOT is the first trapping stage of the experiment. It is formed from three orthogonal pairs of counter-propagating beams and a magnetic quadrupole field. The laser beams alone do not form a trap, as there is no centre-directed restoring force associated with the optical potential. This 'optical molasses' results in a velocity-dependent damping force. The addition of the quadrupole field results in a spatially varying Zeeman shift and so atoms

---

[1]The capture velocity can be deduced from simple kinematic arguments to be $v_c \approx \sqrt{2 F_{sc} d/m}$ the maximum atomic speed that can be brought to rest across the spatial extent (diameter d) of the MOT beam, assuming $F_{sc} = \hbar k \Gamma / 2$.

in a MOT feel a net laser cooling force that results in preferential scattering towards the quadrupole zero.

The MOT was a whole group effort. The design, purchasing and manufacturing of vacuum components and supports was conducted by all group members. Assembly and vacuum bakeout was conducted in parallel with Helmerson group's efforts on their own apparatus, and mutual exchange of ideas and effort was of great importance. The optics layout and design was completed by L. Bennie, while the associated laser design, locking and assembly was conducted by all group members. The magnetic quadrupole and bias coil assembly (and current driver) was designed by R. P. Anderson with assistance from A. Benci.

Magnetic trap and evaporative cooling

The MOT cannot be used to cool ^{87}Rb atoms to quantum degeneracy. The Doppler limit of 290μK is the minimum attainable temperature in a MOT, and so we transfer to a purely magnetic trap (MT). Before loading the MT, the 15 G/cm quadrupole gradient is relaxed and polarisation gradient cooling takes place, resulting in an 'optical molasses' with temperature of between 10-30 μK. The increase in density and reduction in temperature this process confers vastly improves the MT load efficiency. The molasses stage additionally optically depumps atoms into the $F = 1$ state, from which we 'catch' 10^9 atoms in the $|F = 1, m + F = -1\rangle$ state.

Unlike the MOT, there is no active dissipative cooling process at work in a magnetic trap (there is in fact a heating process). In order to cool the atoms further, we perform evaporative cooling by effectively reducing the trap depth. To accomplish this, a radiofrequency (rf) magnetic field with a frequency corresponding to the Zeeman splitting at a given magnetic equipotential is applied. The Zeeman shift is highest at the greatest spatial extent of the cloud, and so at the edge of the cloud, warm atoms are transferred to anti-trapped states and lost. Reducing the frequency of the rf field continuously removes the warmest atoms, the atoms rethermalise via collisions, reducing the overall temperature.

In a pure quadrupole trap non-adiabatic spin flips at the quadrupole zero lose atoms from the trap and increase temperature. A red detuned far off resonant laser beam positioned near the quadrupole zero forms a hybrid optical dipole-magnetic quadrupole trap [37]. The spin-flip losses are suppressed and the loading into the purely dipole trap is enhanced.

There is minimal technical overhead in adapting the apparatus used in the MOT stage for the MT setup. The coil drivers for the MOT coils designed by R. P. Anderson were capable of fast switching of high currents, and a simple

high power rf amplifier and coil is sufficient to attain high Rabi frequencies for optimum evaporative cooling.

Optical dipole trap

The final trap is a purely optical potential, which is essential for making spinor condensates. As discussed in Chapter 2, large magnetic fields and gradients present in the magnetic trap would strongly suppress spinor collisions, but more importantly a magnetic trap can only trap a single Zeeman state in the $F = 1$ hyperfine level.

The scalar AC Stark shift is independent of spin state[2], and so a high power, far off resonant red-detuned laser beam can be used to trap atoms at an intensity maximum, such as a focused waist or the intersection between two beams. The crossed-beam optical dipole trap we use is switched on during the magnetic trapping stage to suppress spin flip losses, forming the *hybrid trap*. The quadrupole gradient is then relaxed to zero after rf evaporative cooling, resulting in a purely optical trap. Like the magnetic trap, we perform evaporative cooling in the dipole trap by reducing the trap depth, which is accomplished by simply reducing the laser power. The power in our crossed beam trap is reduced from 4 W in each beam to about 500 mW in each beam over 5 s after which we are left with a pure Bose-Einstein condensate of typically 3×10^5 atoms.

The design, construction, characterisation and optimisation of the dipole trap was principally my responsibility. This included design of the optical layout, selection and purchase of suitable optics and the fibre laser as well as testing and optimisation.

3.3 Oven

We now discuss the operation and theory of each stage of the apparatus in more detail. The following in-depth discussions of each stage of the apparatus additionally provide a benchmark for future generations of lab users. There is therefore particular attention drawn to technical details; more theoretically inclined readers may wish to peruse the following Sections, examine the figures and otherwise proceed unhindered by the realisation that very rarely things go to plan in experimental physics.

The atom beam source, which of the experimental stages mentioned earlier in Section 3.2 comprises the oven and Zeeman slower, was a focus area of my PhD. I was responsible for the design and construction of the

---

[2]For linearly polarised light, this is true, but elliptically polarised trapping light results in a slight spin-dependence to the trapping potential, the vector light shift.
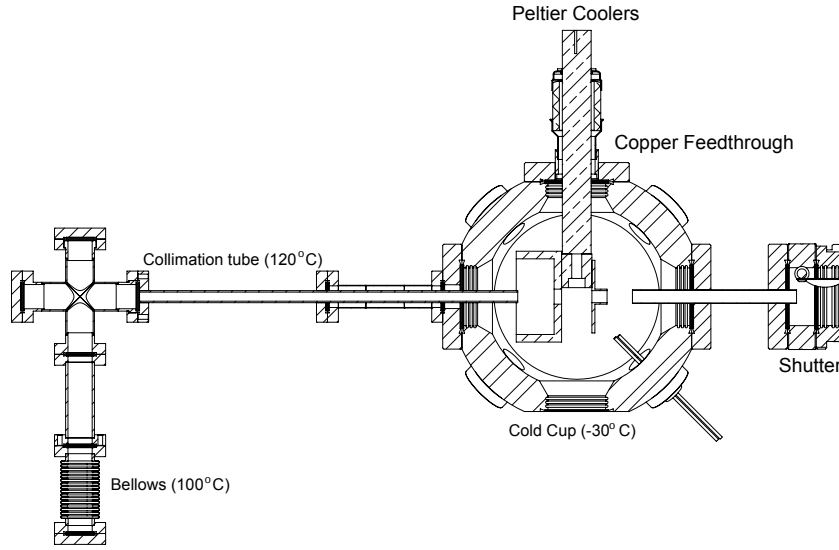
Figure 3.1: Schematic of the oven and associated vacuum system. Molten rubidium in the bellows is collimated into an atomic beam inside the 6 mm diameter, 250 mm long collimation tube. Off-axis emission from the tube is collected by the copper cold cup ($-30°$C). A differential pumping tube isolates the oven from the remainder of the apparatus, and a shutter allows the beam to be blocked when loading of the MOT is complete. A pneumatic gate valve (not shown) isolates the oven from the remainder of the apparatus. The ion pump is connected to the rear of the main chamber. Figure reproduced from Ref. [90].

oven, which is detailed in my Honours thesis [90]. The atomic beam became fully operational in the first year of my PhD.

3.3.1   Theory

The purpose of the atomic beam is simple: to load the MOT with as many atoms as possible in a short time. A high-flux source of atoms is combined with a cooling mechanism to achieve this; our experiment uses an effusive oven combined with a Zeeman slower. The oven contains a 5 g ampoule of rubidium atoms, of which 27.8% (natural isotopic abundance) are ^{87}Rb. It is contained within an ultra-high vacuum (UHV) chamber at a pressure of 1×10^{-9} torr, maintained by a large ion pump. Figure 3.1 schematically shows the oven.

Rubidium melts at 39°C; heating the sample of metal to 80°C creates a vapour pressure of 2×10^{-4} torr in the reservoir section (the 'bellows'). The atomic beam is formed by channelling the vapour through a 6 mm

diameter, 250 mm long collimation tube. A common problem with alkali metal sources is lifetime – high flux sources typically expend the metal source very quickly without some form of recirculation mechanism, which adds additional complexity [93, 94]. In the molecular flow regime, the mean free path of the atoms at UHV pressure is larger than the dimensions of the apparatus: atomic trajectories follow straight lines. The collimation tube effectively limits the bulk of the output atomic flux to *usable* atomic beam [78], with atoms striking the walls of the tube returning to the reservoir or exiting the tube off-axis.

The off-axis atomic emission from the collimation tube presents a problem. The lifetime of trapped ultracold atoms is highly dependent on losses due to one-body inelastic collisions from background atoms, the overall motivation for achieving ultra-high vacuum. In addition to background losses, alkali metal vapour is detrimental to the operation of ion pumps. If an ion pump is exposed to rubidium vapour for an extended period, a process referred to as 'alkali poisoning' deteriorates the pump cathode. Off-axis atomic emission in our oven is prevented from entering the ion pumps by a copper baffle, the 'cold cup', which is kept sufficiently cold that the rubidium atoms that strike it stick. The cold cup is connected to a thick copper vacuum feedthrough, which is cooled by two two-stage thermoelectric coolers (TECs, Melcor MS2-11). The TECs are housed in an airtight vessel mounted to the exterior of the vacuum system, which is pumped down to ~ 1 torr. The vessel completely eliminates condensation of water vapour onto the TECs, improving system reliability. Feedthroughs to the vessel supply the electrical power and recirculated cooling water for the TECs. Ensuring the cold cup remains at temperature and does not overheat in a cooling failure is the primary motivation for the control and interlock system discussed in Section 3.3.2.

3.3.2   Interlocking and Control

Of all the elements in the apparatus, the oven is capable of the most severe failures that would result in extensive downtime. The most pressing failure mode is loss of the cooling water supply to the TECs. If there is no cooling water to remove the heat pumped by the TECs they will simply dissipate the 140 W of electrical power into the cold cup, which will subsequently heat up, re-emitting all of the rubidium metal that has accumulated on its surface. Such a failure event will likely destroy the ion pump in a few hours, and would require disassembly, removal of the contaminating rubidium, replacement of the ion pump, reassembly and re-baking. An approximate down time of several months would be a conservative estimate.

As well as an interlock, there are several other aspects of the oven that demand automation and control. The temperature stability of the reservoir and collimation tube are good examples, as they determine atomic flux; a variable, drifting MOT load rate is highly undesirable. The malfunction of a simple heater may also result in premature expenditure of the metal sample.

Motivated by these concerns, we implemented a robust interlock and control system using a Galil RIO-47100 programmable logic controller (PLC). Sensors for temperature, water flow and vacuum pressure are interfaced with the PLC's eight analog inputs, and the digital outputs of the PLC are used to control relays that actuate the gate valve or deactivate the TEC power supply. A digital state-machine was implemented by Martijn Jasperse using the PLC. I will summarise here only several salient features of the system, which is shown in Figure 3.2.

- *Interlock.* In response to a failure of cooling water, which is detected using a flowmeter in the return line, the PLC responds by entering a fail-safe state: the TEC power supply is turned off, the oven heaters are deactivated and the gate valve shuts (if open). This mode is entered regardless of the previous state of the controller, resulting in a well-defined response to failures if the system is either running or idle.

- *Control.* The oven heaters are each connected to programmable power supplies (Manson HCS-3042) that are driven by two analog output channels from the PLC. The temperatures of the bellows and collimation tube are each measured by a thermistor, which is part of a proportional-integral-derivative (PID) control loop. This confers the benefits of rapid heat-up and stable atomic flux.

- Logging: the PLC is interfaced with the network, and can output the values of critical sensors (such as the temperature of the cold cup or the oven) for long term monitoring and remote diagnosis.

More details can be found in the PhD thesis of Martijn Jasperse [95]. The oven interlock, as well as similar interlocking setups for the water cooling of the Zeeman slower and quadrupole coils are indispensable elements of the apparatus, and have saved us from many a catastrophe.

3.3.3 Zeeman slower

The Zeeman slower is a zero-crossing design [37,91], with the axial magnetic field changing direction along the length of the slower. This allows the use of more modest magnetic field strengths, and thus currents. Two varying pitch coils surround the vacuum tubing of the Zeeman slower, as shown in
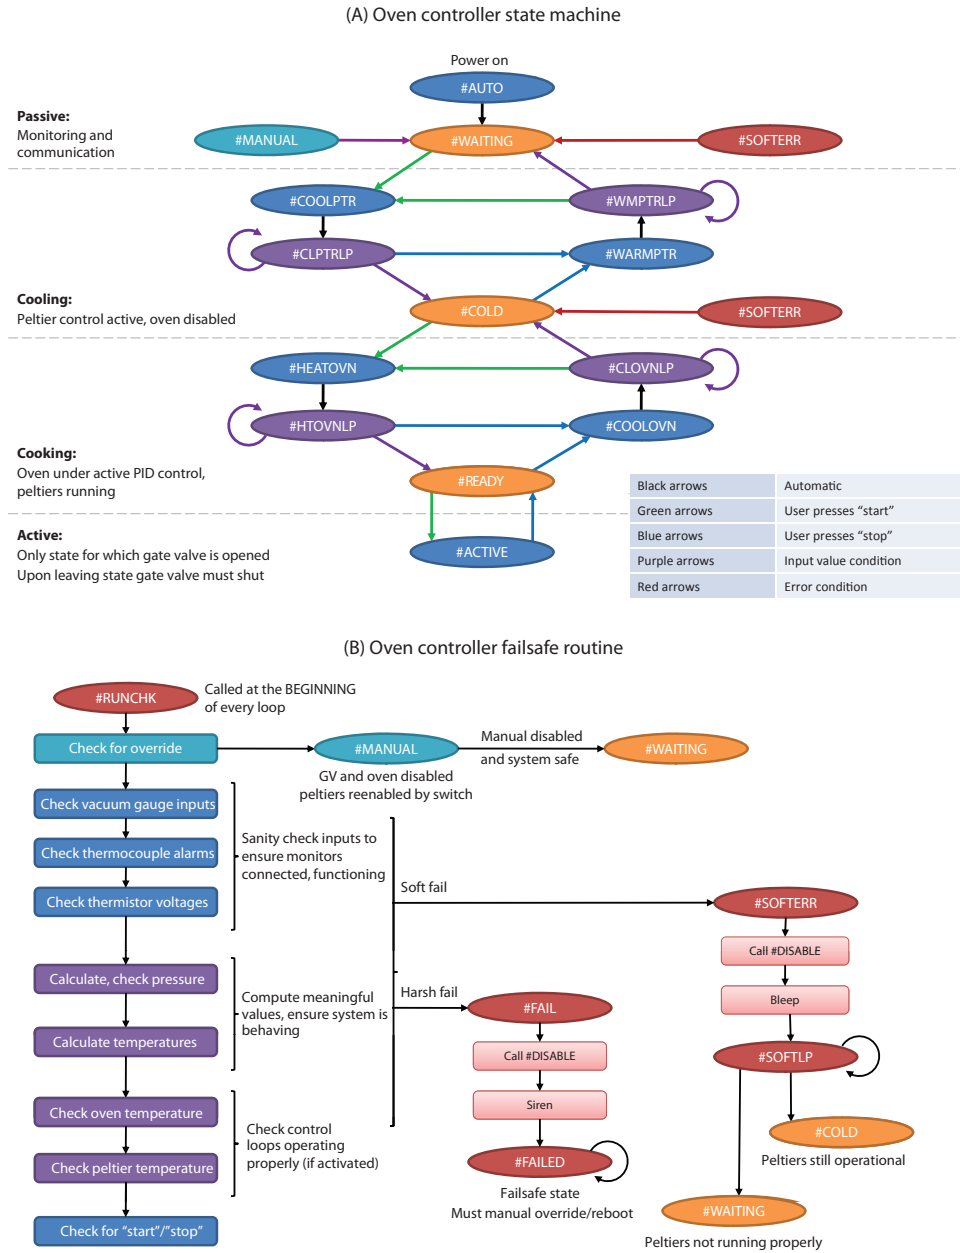
Figure 3.2: The oven control and interlock state machine. A) The operational states of the machine, with transitions between states determined by user input (pressing control buttons, green and blue), loops converging on setpoints (purple) or errors (red). B) Typical sequence of operations for failsaifing the experiment, where monitor inputs are checked before execution of stages. Harsh errors deactivate the TECs, close the gate valve and require manual intervention, whereas soft errors return the experiment to a 'cold' state with the TECs operational. Image from Martijn Jasperse's PhD thesis [95], where further details can be found.

Figure 3.3. The first coil carries a current of 95 A and the second 46 A. Both coils are hollow, and allow recirculated cooling water to be pumped through to prevent overheating. Two lasers propagate counter to the atomic beam in the slower; a 13 mW circularly polarised cooling beam detuned -162 MHz from the $F = 2 \rightarrow F' = 3$ transition and an overlapping 22 mW repump beam, detuned -183 MHz from the $F = 1 \rightarrow F' = 1$ transition. We will discuss the optimisation of the Zeeman slower in more detail in Section 3.8.1.

3.4   Ultra high vacuum system and bakeout

It is natural to delineate at some point the technical exploits of construction from the more interesting (though no less challenging) process of working with atoms and lasers. In this Section we discuss the final, gruelling, and most demanding aspect of vacuum system construction: bakeout.

All UHV apparatus regardless of supplier, cleanliness or cost, must be *baked* to attain ultra high vacuum. When assembled and pumped down, the lowest attainable pressure is unremarkable, typically between 10^{-6} and 10^{-9} torr, even with high capacity pumps operating. This is because *outgassing* from the steel, glass or copper surfaces in the system vastly outpaces the pumping speed. Water vapour is the primary contributor, although gases like hydrogen become important at lower pressures once the water is eliminated. This outgassing can be accelerated, and the resident contaminants thus depleted, by increasing the temperature of the outgassing components. A 'vacuum oven' is a machine built specifically for the task, and allows individual components to be heated to $\sim 300°$C under low vacuum. It is however demanding, or rather impossible, to assemble the system under vacuum or to put the entire ~ 1.5 m apparatus in a vacuum oven. As soon as the baked components are exposed to atmosphere a microlayer of water forms and must be baked again when fully assembled.

Bakeout of the assembled system is rather hazardous. The whole vacuum system must be heated to 200-300°C with minimal temperature gradients, particularly across sensitive areas like the glass-to-metal seal of the quartz science cell. Once heated, it must be kept in this precarious state for 2 weeks. The heat-up and cool-down phases are of particular difficulty. Such was the story of our bakeout.

Figure 3.3 shows the constituent vacuum components. Two turbomolecular pumps (Pfeiffer TMU-064, backed by a Varian Triscroll roughing pump), one on either end of the system were used to pump from atmospheric pressure to $\sim 10^{-6}$ torr. Below this pressure the three ion pumps were active: a Perkin-Elmer D-I 50 L/s on the oven, a smaller 10 L/s Perkin-Elmer pump
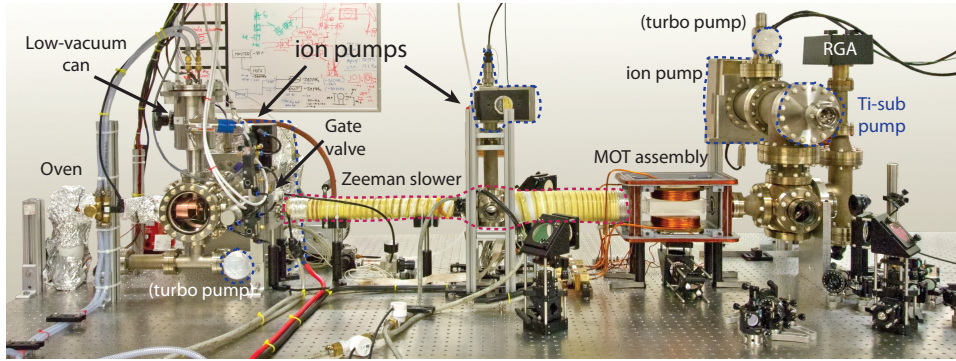
Figure 3.3: Profile of the apparatus following assembly (late 2010). Vacuum pumps are shown outlined in blue – including the blocked-off flanges where the turbo pumps were connected for pumpdown and bakeout. During bakeout, the MOT assembly was absent whereas the Zeeman slower coils (outlined in magenta) could not be removed. Absent from this image are the dipole trapping optics, imaging setup and cameras which would otherwise obscure the view of the science cell.

in the Zeeman slower and a 75 L/s Gamma pump. During bakeout, all pumps were operational, whereas the turbo pumps are isolated by conflat angle valves and removed during normal operation. A titanium sublimation pump (TSP) is used after bakeout to reach pressures $< 10^{-11}$ torr. Pressure is monitored with a UHV-grade ion gauge at the oven and UHV ends, and a residual gas analyser (RGA) provides a readout of the constituent elemental partial pressures in the vacuum system down to $\sim 10^{-12}$ torr.

3.4.1 Procedure

Our cleaning and bakeout preparation loosely followed that described in Refs. [96, 97]. To begin with, all vacuum parts were cleaned using solvents – acetone for bulky items, methanol for glassware – before pre-baking in a small vacuum oven. Small parts were treated in an ultrasonic bath of acetone or methanol. It was often necessary to use detergents on parts that had come from the mechanical workshop, or in several cases, fresh from the suppliers. The Pyrex rubidium ampoule was treated with a similar cleaning process,[3] albeit with a mild pre-bake in the vacuum oven.

The entire system was then assembled over the course of several weeks. The oven was actually assembled earlier than the rest of the apparatus due to delays with the arrival of the quartz science cell.[4] Pre-assembly and testing of the oven was done until the cell arrived. I will only discriminate

---

[3]A safety precautions sticker on the ampoule proved difficult to remove. A citrus-based cleaner was used to lift it prior to solvent treatment.

[4]Complications associated with 'disconnection' of the glass-metal seals, we were informed.

between the two bakeouts where necessary, as the procedures for each were essentially identical. The only major difference between the two bakes was the temperature reached. The oven contains numerous complicated vacuum components, such as valves, shutters, feedthroughs and the rubidium ampoule which limited the bakeout to 180°C.

After assembly and initial pumpdown, resistive heater tapes (AC, 0-250 V) were wrapped around the steel components and covered in aluminium foil to distribute the heat. Figure 3.4 shows the vacuum system being prepared for a bake. Glassware would be thermally stressed if in direct contact with the heater tape or aluminium foil, so foil cages incorporating tape were assembled around windows. Thermocouples were connected to monitor temperature of each section of the system. The Zeeman slower coils could not be removed for baking, and so a heater tape was wound around the vacuum tubing inside, with water flowing in the coils to prevent the adhesive tape holding the coil forms together from melting. A current of $\sim 20\,\text{A}$ was maintained in the TSP filaments during the bakeout.

The glass science cell is the most challenging aspect of the bake. An oven was built around the 'locus of anxiety' from fire bricks and thermally insulating fibreglass (shown in Figure 3.4). A single heater element was placed inside the oven, with several heater tape assemblies on each glass-metal seal to minimise the temperature gradient to $< 10°\text{C}$. Another concern was the temperature of the optics table – the bonding adhesive inside the table limited surface temperatures to $\sim 70°\text{C}$. This required some degree of thermal isolation of the baking apparatus from the table, although heat loss from the exposed aluminium support mounts was sufficient to maintain a safe temperature. This effect would become critically important to the structural integrity of the system during cool-down phases.

Each heater tape and the heater element was connected to a 0-250 V variable transformer (Variac). The voltage was gradually increased to ensure the temperature would increase by more than 20°C/hour. The temperature of the bakeout was limited to 300°C by the glass-to-metal seal on the cell, and we aimed for higher temperatures in other parts of the system, primarily to accelerate the outgassing rate of hydrogen from large metal surfaces. The RGA allowed us to determine the presence of leaks, contaminants and additionally analyse the efficacy of the bake on certain partial pressures. We also monitored total pressure with an ion gauge, which was the sole source of pressure information above 100°C, where the RGA could no longer be operated.

The *first* bake proceeded well initially, reaching an average temperature of 300°C. Upon cooldown however, a massive leak opened up in the 4.5" chamber underneath the titanium sublimation pump (Figure 3.3) while the
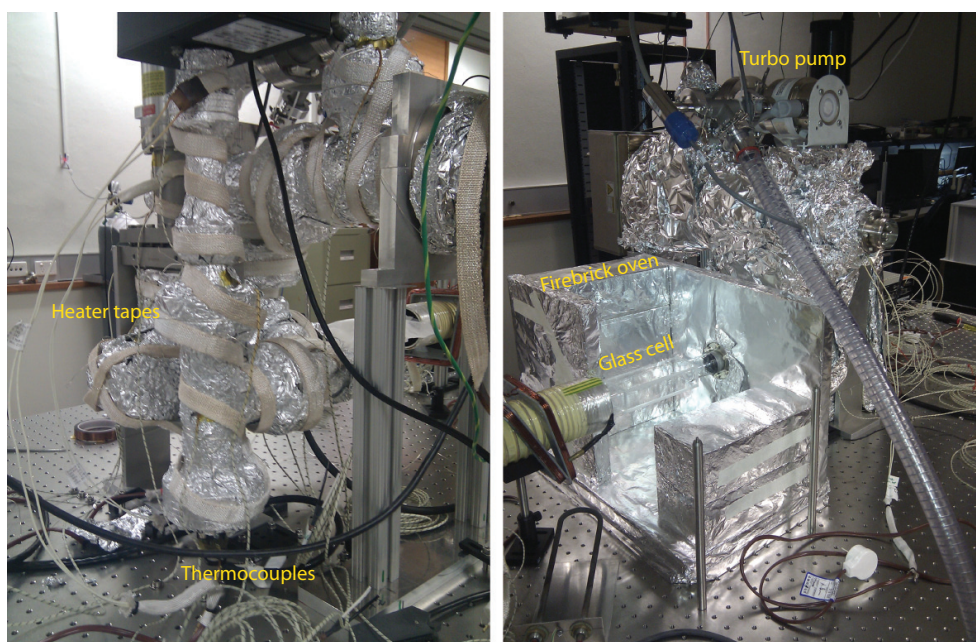
Figure 3.4: Bakeout of the UHV system. Left: Heater tapes wrapped around the aluminium foil-coated vacuum system in preparation for a bake. Right: The glass cell inside the fire brick baking oven prior to the roof being put on.

system was at 150°C. This leak was sufficiently large to limit the system pressure to a few torr under roughing vacuum, which allowed the leak to be detected by spraying methanol on the system joins.[5] The cause of the leak was suspected to be an excessive temperature rate-of-change at the time but we now suspect that unequal thermal expansion of the vacuum system support structures, which tightly constrained the system, may have led to shearing forces between the vacuum joins which opened under the strain.

Leaks opened up during subsequent bakes, possibly due to the copper gaskets having suffered damage upon sudden exposure to oxygen at high temperature during first leak. With each leak, we had to cool down the system (usually over the course of a day or two if the leaks opened at high temperature), analyse the RGA spectrum to confirm the presence of a leak (a 4:1 N_2:O_2 ratio was the first sign), unwrap the layers of heater tape and aluminium foil covering the system and finally locate the leak by using the RGA to detect helium gas sprayed over suspect joins. From here we could either attempt an *in-situ* repair, or in most cases undergo a full re-pressurisation and replacement of the offending part. After four such aborted bakes, removal of an all-metal gate valve and close-coupler from the system,

---

[5]The leak could be detected by the 'slurping' sound it made when solvent was applied.
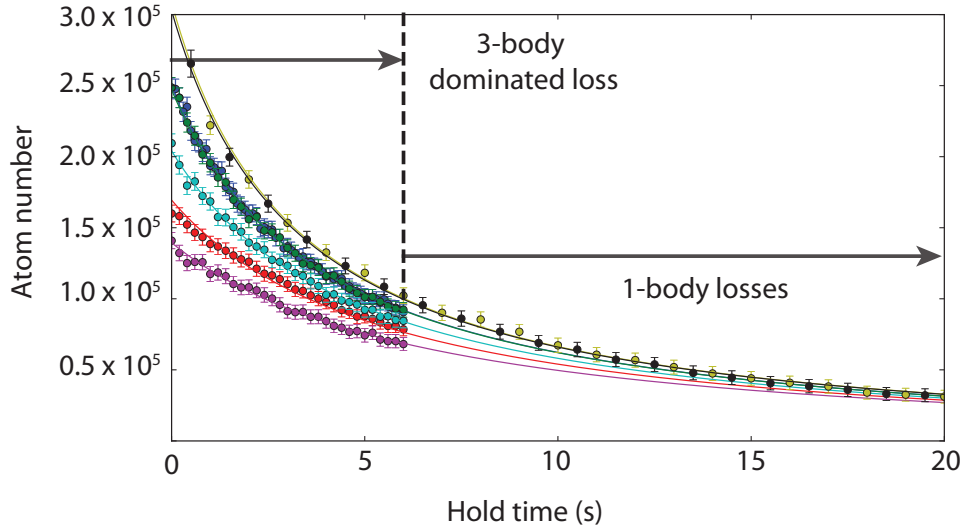
Figure 3.5: Combined loss rate curves for BECs in the crossed dipole trap; different coloured curves correspond to different initial atom numbers. Accounting for density-dependent losses due to three-body recombination, the vacuum limited lifetime due to one-body collisions is 39 ± 7 s. Figure courtesy M. Jasperse [95]

and the replacement of a damaged turbo pump, we achieved a post-bake pressure of 10^{-10} torr. The final, successful bake was undertaken with only minimal support structures, suggesting the cause of the leaks was likely due to over-constraining the system during bakeout. We fired the titanium sublimation pump filaments and achieved a pressure of 3×10^{-11} torr, as measured by an ion gauge in the UHV system. This pressure gradually fell, and when the ion gauge and RGA were disconnected it had reached $1\text{-}2 \times 10^{-11}$ torr. From RGA readings, only residual amounts of hydrogen remained with water being well into the noise floor of the device. The presence of methane, comparable to hydrogen partial pressures, was observed – methane is actually outgassed from TSP filaments and is not readily pumped [98, 99].

The most definitive metric for successful vacuum preparation is not necessarily pressure measured by the vacuum gauges. There is no ion gauge in the glass cell; the only measurement of background pressure there is the lifetime of the magnetically trapped thermal cloud or a BEC. Figure 3.5 shows the number of atoms remaining in a BEC as a function of hold time. The measured lifetime of 39 ± 7 s indicates the preparation of high quality vacuum, despite the setbacks.

3.5   Lasers

Lasers are a major component of the apparatus. We require a minimum of five laser frequencies to laser cool atoms: a red detuned cooling ($F = 2 \rightarrow F' = 3$) and resonant repump ($F = 1 \rightarrow F' = 2$) for each of the MOT and Zeeman slower as well as a resonant imaging probe. We source our cooling light from a single MOGlabs external-cavity diode laser (ECDL) with an Eagleyard RWE-810 150 mW diode driving a tapered amplifier (m2k Laser TA-0785-2000-DHP), in a master oscillator-power amplifier (MOPA) configuration. The MOPA produces over 800 mW of light, which is then split into beamlines for the MOT, Zeeman slower and imaging. Acousto-optic modulators (AOMs) are used to both adjust the frequency of light in each beam line and as high-speed gates. Repump light for the MOT and Zeeman slower is sourced from two independent MOGlabs ECDLs: an Axcel Photonics M9-780-0150-S5S for the Zeeman slower repump and a 30 mW Sharp diode on the MOT repump. All of the lasers are located on a different optics table; polarisation-maintaining (PM) fibres are used to transport light to the vacuum system.

3.5.1   Locking

Precise stabilisation of laser frequency is essential for laser cooling. A laser locked to a single frequency must additionally be relatively immune to ambient perturbations for robust day-to-day operation. For the purposes of laser cooling, the requirements for laser linewidths are modest; typically several hundred kilohertz is acceptable. All our ECDLs are driven by MOGlabs diode laser controllers ('MOGboxes', models DLC-202 Revs. 5 and 7) that provide current and frequency control as well as locking functions.

The master cooling laser has the most stringent requirements for frequency stability and linewidth. It is locked to an absolute atomic reference, the $F = 2 \rightarrow F' = 3$ cooling transition using modulation transfer spectroscopy (MTS) [100]. MTS is notable among laser locking schemes in that it allows for high signal-to-noise dispersive error signal lineshapes primarily localised to cycling atomic transitions. Our MTS setup, designed and built by Vlad Negnevitsky, uses AOMs rather than EOMs for frequency modulation of the pump beam [101]. Additionally, the demodulation electronics for synthesising the error signal are implemented digitally on a commercial software radio (Ettus Research USRP N210).

The Zeeman repump laser is locked 183 MHz below the $F = 1 \rightarrow F' = 1$ transition. This frequency was optimised based on operation of the Zeeman slower, the reasoning behind this is discussed in Section 3.8. We use an offset locking scheme [102] to lock the Zeeman repump laser: light from the master

laser is mixed with a repump beam on a high speed photodiode (Hamamatsu G4176). The resulting beatnote is then electronically processed to create an error signal, with the laser controller closing the feedback loop.

The final laser is the dedicated MOT repump laser. It is locked to the $F = 1 \rightarrow F' = 0, 1$ crossover transition using a simple saturated absorption spectroscopy setup and current modulation provided by the MOGbox. A double-passed 80 MHz AOM in the locking setup and single-pass 110 MHz AOM before the fibre couple are used to bring the light to resonance with the $F = 1 \rightarrow F' = 1$ transition.

The detunings used for the various laser cooling stages of the experiment as well as that of the individual ECDLs and the MOPA are depicted in Figure 3.6.

3.5.2 Optical layout

The complete laser setup is shown in Figure 3.7. Frequency control of each beamline is achieved with AOMs in either single- or double-pass configurations [103]. Double-pass AOMs are useful when frequency changes are required *during* experiments: if the frequency of a single-pass AOM system were changed, the pointing direction of the beam would also change and substantially reduce the fibre coupling efficiency. Double-pass AOMs alleviate this difficulty by retro-reflecting the diffracted order through the AOM. Higher diffraction bandwidths are achievable, which is an essential requirement for reaching large red detunings (> 60 MHz) for polarisation gradient cooling. The AOMs also serve as very high speed (< 1 μs) beam shutters. Mechanical shutters (Stanford Research Systems SR-474 as well as shutters built in-house based on Ref. [104]) are positioned after each AOM to totally isolate optical leakage from the atoms.

Polarisation-maintaining fibres connect the optics table to the vacuum table. We typically achieve fibre coupling efficiencies of 50-85% or more, subject primarily to the spatial mode of various lasers and the tapered amplifier output.

3.6 Absorption imaging

We use absorption imaging to characterise cold atomic clouds and BECs. There are many references that cover absorption imaging in depth, and Ref. [105] is perhaps the most definitive source. Here, we will briefly summarise the main features of the technique.

The atomic cloud is exposed to a resonant probe laser that is directed onto a camera via magnifying (or demagnifying) optics. The image we obtain
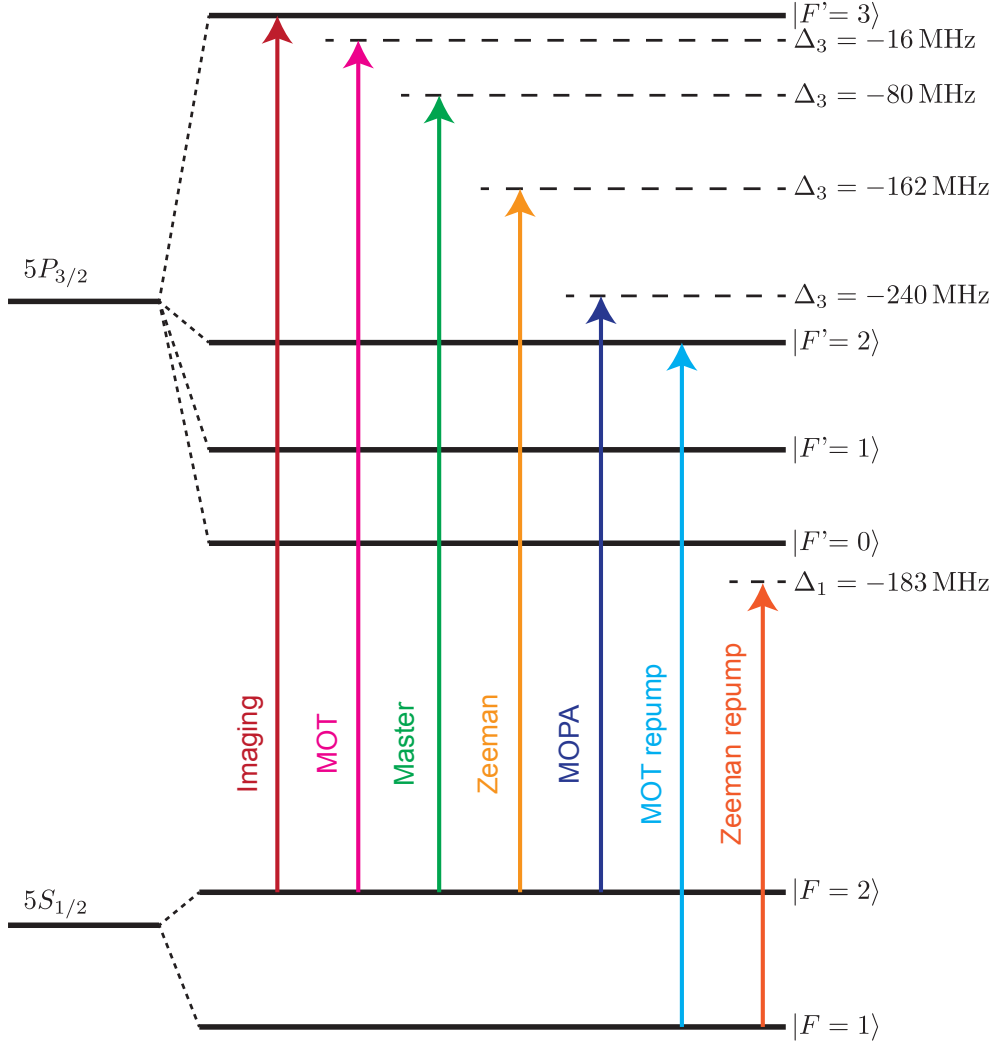
Figure 3.6: Laser detunings for ECDLs, MOPA and various laser cooling frequencies used in experiments. The master, Zeeman repump and resonant MOT repump are independent ECDLs; the master laser locked $-80\,\text{MHz}$ below resonance and is down-shifted $-160\,\text{MHz}$ before seeding the MOPA ($-240\,\text{MHz}$), and the Zeeman repump is offset-locked $6305\,\text{MHz}$ to the red of the master ($-183\,\text{MHz}$ detuned from the $F = 1 \rightarrow F' = 1$ transition). The MOT ($-16\,\text{MHz}$), Zeeman slower ($-162\,\text{MHz}$) and resonant imaging beam are all up-shifted by AOMs in independent beamlines after the MOPA (Figure 3.7). Figure reproduced with permission from the PhD thesis of M. Jasperse [95].
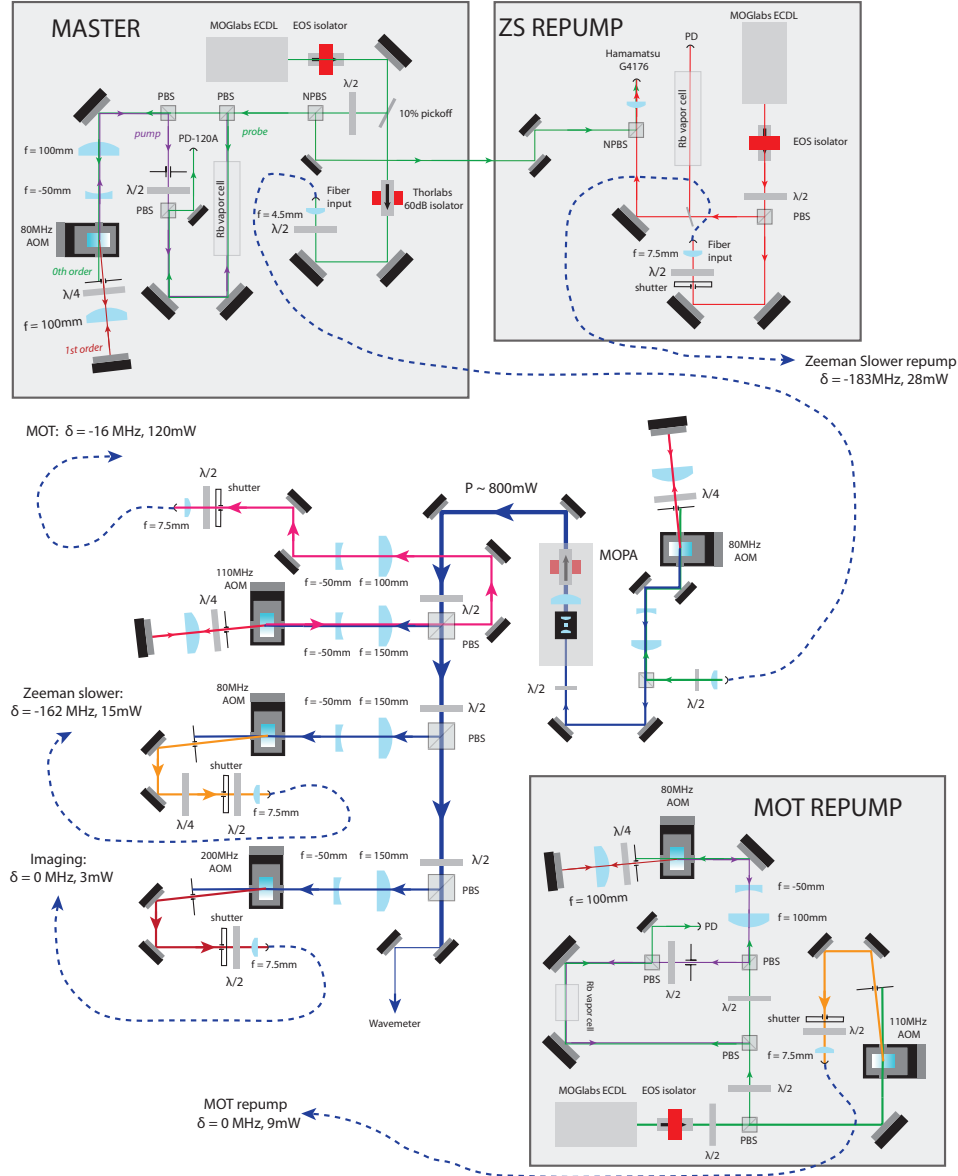
Figure 3.7: The complete laser setup. A master laser ECDL locked 80 MHz to the red of the ^{87}Rb $5S_{1/2}F = 2 \rightarrow 5P_{3/2}F' = 3$ transition using MTS drives a tapered amplifier, which provides all the cooling light for the MOT, Zeeman slower (ZS) and absorption imaging. Independent lasers are used for the Zeeman slower repump beam (offset locked to the master) and MOT repump (independent spectroscopic lock).
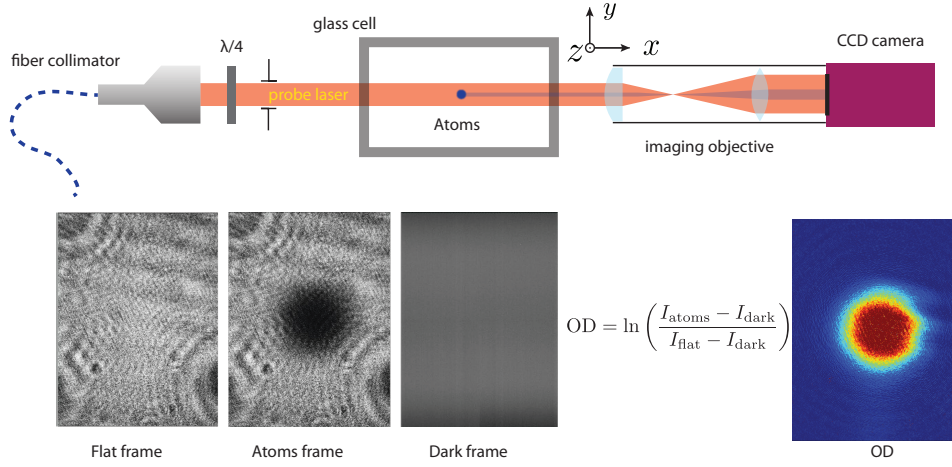
Figure 3.8: Schematic example of absorption imaging setup. Atoms absorb light from a resonant probe laser, casting a shadow. The image of the cloud is focused onto a CCD camera. Three images are taken to produce an image of the optical depth (OD): atoms (laser beam + atom signal), flat (laser beam alone) and dark (no laser, only dark counts and stray light).

is that of the shadow cast by the atoms; from this image we can deduce the atom number, temperature and the spatial structure of the cloud. A basic schematic of the imaging setup is shown in Figure 3.8, with the imaging probe beam propagating along the x-axis, as it does in our apparatus.

The intensity of the probe beam depends on x, and thus so to does the optical absorption cross-section when saturation of the atomic transition is significant. Beer's law relates the transmitted intensity to the optical absorption cross-section and atomic density along the direction of propagation:

$$\frac{dI(x)}{dx} = -I(x)\frac{\sigma_0}{1 + \delta^2 + I(x)/I_{\text{sat}}}n(x),\tag{3.3}$$

where the on-resonant absorption cross section $\sigma_0 = 3\lambda^2/2\pi$ for circularly polarised light on the cycling transition, I/I_{sat} the ratio of the intensity to the saturation intensity I_{sat} and Δ the detuning, expressed in half-linewidths ($\delta = \Delta/\Gamma/2$).

In practice, a single 'absorption image' is actually composed of at least two raw frames: the *flat* frame, which is essentially an image of the spatial profile of the imaging light and the *atoms* frame, which actually contains the shadow. A third frame is also taken to account for the effects of dark counts on the CCD. These three frames are then combined on a computer into a final image of the optical depth (OD) as shown in Figure 3.8. The OD is given by

$$\text{OD} = \ln\left(\frac{I_{\text{atoms}} - I_{\text{dark}}}{I_{\text{flat}} - I_{\text{dark}}}\right)\tag{3.4}$$

with I_{atoms}, I_{flat} and I_{dark} corresponding to the intensity in the atomic, flat and dark exposures. The fraction of light transmitted through the atoms can be written as $e^{-\text{OD}(y,z)}$.

3.6.1 Saturation correction

As is evident from Eq. 3.3, the effects of saturation must be considered to accurately determine the atomic density. High optical intensities saturate the atomic transition, and if not accounted for the atomic density will be underestimated. It can be shown that Eqs. 3.3 and 3.4 can be combined and solved to determine the atomic column density $\tilde{n}(y,z) = \int n(x,y,z)\, dx$ from an OD image [106]:

$$\tilde{n}(y,z) = \frac{1}{\sigma_0} \left[(1 + \delta^2)\text{OD}(y,z) + \frac{I_{\text{flat}}(y,z) - I_{\text{dark}}(y,z)}{I_{\text{sat}}} \left(1 - e^{-\text{OD}(y,z)} \right) \right],$$

(3.5)

and the total atom number given by

$$N_{\text{atoms}} = d_{\text{pix}}^2 \sum_{\text{pixels}} \tilde{n}(y,z).$$

(3.6)

with d_{pix} the effective pixel size (accounting for the magnification of imaging optics). To correct for saturation, I_{sat} must be determined. We follow the procedure outlined in Ref. [106], whereby the saturation intensity is determined (in units of camera analog digital units). The atomic column density calculated at each pixel is then based on the local saturation correction parameter $(I_{\text{flat}}(y,z) - I_{\text{dark}}(y,z))/I_{\text{sat}}$.

3.6.2 Imaging orientations and hardware

We can image the atoms from two different directions. 'Side imaging', where we can see the atoms fall along gravity, is our workhorse imaging diagnostic. The imaging probe beam and objective lens are parallel to the x-axis of the experiment (as shown in Figure 3.10). The objective lens used in side images has a magnification of 2.2 and is composed of two achromatic lenses of focal lengths $f = 50\,\text{mm}$ and $f = 100\,\text{mm}$. We use an interline CCD camera (Allied Vision Technologies AVT-GX1920) for side imaging, which is capable of very fast interframe times.

We can also image clouds from above. 'Top imaging' is highly beneficial – when combined with side imaging, it gives an almost complete description of the atomic cloud and is useful for diagnosing the alignment of the dipole trap. Our top imaging setup uses the MOT beams as the resonant probe and a flipper mirror to deflect the MOT beams either onto the atoms or into the

camera. As such, the top imaging setup is only able to measure stages of the experiment post-molasses; the flipper mirror must be positioned to allow the MOT probe beam into the camera, which takes $\sim 100\,$ms, the molasses cloud will be destroyed within several ms of the flipper mirror moving. The probe beam is focused onto a camera (Andor Neo sCMOS) by a pair of achromats with $f = 200\,$mm and $f = 400\,$mm as well as a single plano-concave lens with $f = -100\,$mm, the overall magnification is 2.04.

3.6.3   Back-to-back images

The flat-field frame is most useful when the spatial profile of the imaging beam does not change significantly between flat and atomic frames. Imaging with coherent light inevitably produces a multitude of interference fringes from the glass cell[6] as well as other optical components, as can be seen in the flat frame shown in Figure 3.8. Any vibration of the apparatus would result in these fringes changing position between the atomic and flat exposures, and as such will be present as spatial fringes appearing in the resulting OD image. This motivates reducing the time between frames: the shorter inter-frame time reduces the translation of fringes between frames. The AVT GX1920 has a minimum interframe time of 30 µs, but there is one caveat: the atoms cannot be in both atomic *and* flat frames. Acquiring the atoms frame before the flat frame requires that the cloud fall from the interrogation region between frames, but this limits the interframe time to tens of ms, and if vibrations are particularly intense[7] then the flat-field subtraction will be poor.

   Since we are principally interested in imaging $F = 1$ atoms, we are able perform the flat frame first while the atoms are still in the $F = 1$ state, $6.8\,$GHz detuned from the $F = 2 \rightarrow F' = 3$ probe laser. We then repump the atoms into the $F = 2$ state with the MOT repump beams for 40 µs between the atomic and flat frame, immediately before performing the 100 µs atomic exposure. With *back-to-back* imaging we can reach interframe times of 40 µs, with substantially better subtraction of fringes.

3.6.4   Time of flight imaging and Stern-Gerlach analysis

Absorption images of trapped atomic clouds and BEC are performed after freefall, so that the gas can expand. A trapped BEC would have an OD of over

---

[6]Our glass cell is not anti-reflection coated – the additional cost for this was prohibitive. Multiple reflections of the imaging beam (at normal incidence to the cell) within the cell result in light passing multiple times through the atoms; creating spatial modulation of the column density (present in the absorption image in Figure 3.9, left), this effect is not trivially corrected.

[7]As is the case of our experiment. We attribute this to the water cooling of the TECs in the oven.

100 and is typically only 10 μm in size. Extracting atom number, for example, is therefore very difficult. Extinguishing the trapping forces and allowing for time of flight, the condensate can expand enough so that the optical depth is within the realm of the imaging system (which is limited by imperfect subtraction in Eq. 3.4 and the bit depth of the camera to $OD_{max} \sim 5$). Time-of-flight absorption imaging is a powerful technique, albeit destructive; each image requires a new atomic sample.

As well as calculating BEC atom number accurately, freefall allows the temperature of a thermal cloud can be deduced by monitoring the expansion of the cloud. This is an important diagnostic for the success of various cooling stages of the experiment. The time-dependent root mean square widths ($\sigma_x(t)$ and $\sigma_y(t)$, unrelated to laboratory coordinates) of an expanding thermal cloud is related to the temperature by

$$\sigma(t)^2 = \frac{k_B T}{m} t^2 + \sigma_0^2, \tag{3.7}$$

where $\sigma(t)$ is the width, t is the time of flight and σ_0 the initial width, which can be used to determine the phase space density, from

$$\mathrm{PSD} = \frac{N_{\mathrm{atoms}}}{(2\pi)^{\frac{3}{2}}} \frac{1}{\sigma_{0,x}^2 \, \sigma_{0,y}} \left(\frac{h}{\sqrt{2\pi m k_B T_{\mathrm{ave}}}} \right)^3 \tag{3.8}$$

where $T_{\mathrm{ave}} = (T_x + T_y)/2$ and T_x, T_y are the effective temperatures determined by applying Eq. 3.7 to expansion data $\sigma_x(t)$, $\sigma_y(t)$ and we have assumed cylindrical symmetry about the y-axis of the cloud, and harmonic confinement.

For long enough drop times, the thermal wings of even a cloud on the threshold of condensation can be reliably extracted from a Gaussian fit to an absorption image. A sample thermometry experiment is shown in Figure 3.9, for two clouds: a partially evaporated magnetically trapped cloud and a cloud nearing the threshold of condensation in the crossed dipole trap.

For optically trapped clouds (where multiple Zeeman sublevels can be trapped simultaneously) we can image the population in each magnetic sublevel by briefly applying a sufficiently strong magnetic field gradient during time of flight to spatially separate the spin components. The reduction in density effectively switches off the interatomic interactions, the gradient thus exerts only a mechanical force on the atoms. In analogy with the seminal experiment [107] this technique is called *Stern-Gerlach* absorption imaging. The population as well as the spatial structure of each sublevel is then mapped to spatially distinct regions.

For our Stern-Gerlach separation, we apply a 60 G/cm gradient with the quadrupole coils over 3 ms, 11 ms after the atoms are released from the
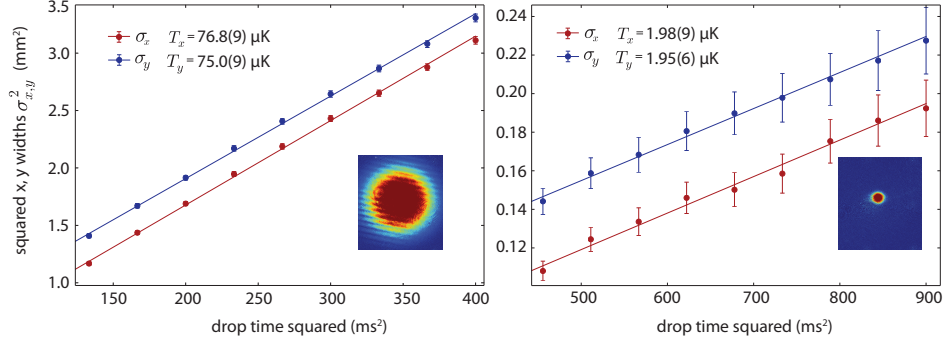
Figure 3.9: Thermometry performed using time of flight. Plotting the squared Gaussian widths of a thermal cloud in a magnetic trap (left) and approaching the threshold of condensation in the dipole trap (right) against time of flight (drop time) yields a straight line. Shown in inset are absorption images of each cloud at 20 ms drop time, scale in comparison with each other.

trap. These parameters are chosen so that the Stern-Gerlach kick distinctly separates all three Zeeman components after typical 20 ms drop times.

3.7   Controlling experiments

Experimental hardware is controlled by computer. A single control computer interfaces with a SpinCore PulseBlaster DDS-II-300 USB, which serves as the master clock for experimental timing. Three National Instruments data acquisition and control cards: $2\times$NI-PCIe6363 (digital outputs mostly) and NI-PCIe6733 (analog outputs) provide analog and digital I/O with which we control individual coil currents, beam amplitudes and shutters, for example.

3.7.1   Radiofrequency devices

We use two different types of direct-digital synthesiser (DDS) as sources for radiofrequency. The commercial SpinCore PulseBlaster DDS is used for low-grade applications where pulse phase stability is not important, such as driving AOMs and rf evaporation. An in-house designed and manufactured DDS[8], the *RFBlaster*, is used in demanding applications, such as rf pulse synthesis for state preparation. The superior phase stability and flexibility of the RFBlaster over the PulseBlaster was demonstrated by the ANU Atom Laser group in high precision gravity measurements [108]. The RFBlaster is programmable with up to $\sim 10^6$ amplitude/frequency/phase instructions, in contrast to the PulseBlaster, which is limited to $\sim 10^3$ in our application.

---

[8]Designed and assembled by Vlad Negnevitsky.

This allows the RFBlaster to implement many sophisticated pulse sequences, ramps, and sweeps without recourse to additional devices.

3.7.2 Control system

The control system we use was created in-house by members of the Monash BEC group. It differs from more widespread schemes by allowing experiments to be written as Python code and then compiled into hardware instructions for the various devices. The full details of the experimental control system can be found in Ref. [109].

3.8 MOT and Magnetic trap

Our MOT is formed from six laser beams arranged as three orthogonal counter-propagating pairs, as shown in Figure 3.10. An anti-Helmholtz coil pair provides a 14.8 G/cm quadrupole magnetic field gradient. The 6 mm $1/e^2$ diameter beam output from the fibre collimator is increased to 20 mm in diameter by a beam expander for each MOT beam, the full details of which can be found in Ref. [110]. The MOT beams pairs are $\sigma^+ \sigma^-$ polarised and have about 14 mW of optical power in each beam, each with a 1 mW overlapping resonant repump beam. The cooling light in the MOT is detuned -16 MHz from resonance. We typically load $\sim 3 \times 10^9$ atoms in 5 s, corresponding to a peak load rate of 7×10^8 atoms/s. The temperature of the atoms in the MOT is about 1-2 mK.

3.8.1 Optimising the MOT load

Optimising the MOT is a complicated procedure, as it is intrinsically linked to the performance of the Zeeman slower. The goal is simple – to achieve maximum loaded atom number. Temperature is less of a concern, as sub-Doppler temperatures can be routinely achieved after an optical molasses stage (Section 3.8.2).

We tried a number of methods to independently diagnose the performance of the Zeeman slower: atomic absorption of a resonant probe beam was used to measure atomic flux, and varying the angle of the probe allows for velocimetry. The absorption signal was very weak, and difficult to interpret in most cases. The perturbative effect of the Zeeman slowing beam on the MOT (which depending on beam intensity and detuning can result in poor MOT loading) could not be considered without a MOT to load. We thus resolved to optimise the performance of the MOT and Zeeman slower with but one metric, loaded atom number.
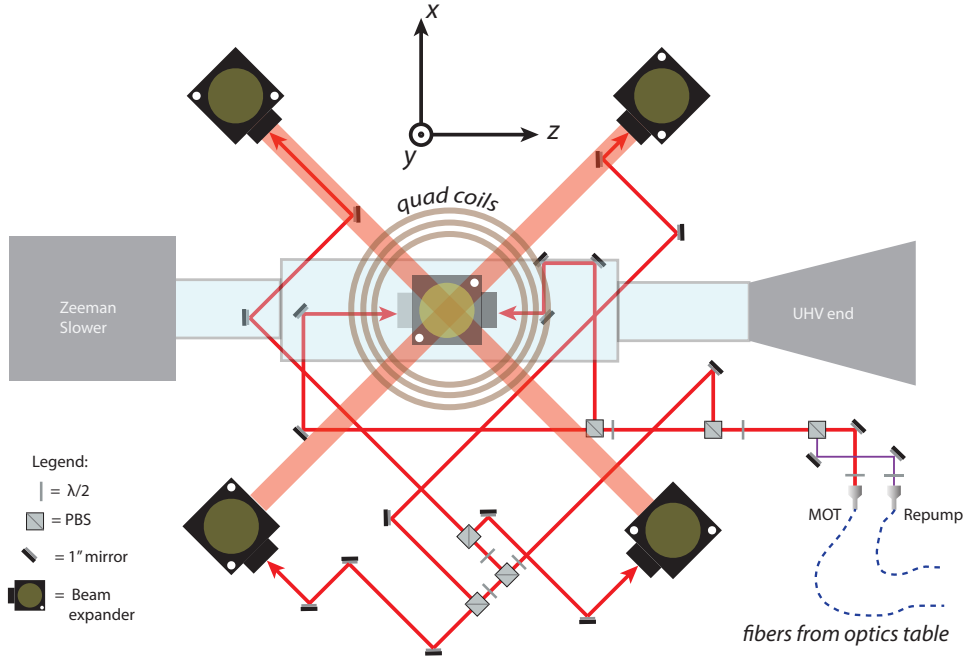
Figure 3.10: Beam geometry of the MOT. Each MOT beam enters a beam expander that produces a 20 mm collimated beam. Approximately 1 mW of resonant repump light is mixed with each MOT beam, for a combined total of 14 mW in each beam. The quadrupole magnetic field is produced by two anti-Helmholtz coils, not shown are the bias magnetic field coils used to shim stray magnetic fields.

The Zeeman slower laser is focused to a waist of $\sim 300\,\mu m$ at the oven collimation tube output, from an initial diameter of 30 mm. This maximises scattering of photons at the start of the slower while minimising the perturbative effect on the MOT, which was found to be quite severe for smaller beam widths at the position of the MOT. The wide control bandwidth of the microwave offset lock allows the Zeeman repump detuning to be changed over a substantial range. The optimum detuning is found to be $-183\,\mathrm{MHz}$ detuned from the $F = 1 \rightarrow F' = 1$ transition, the corresponding beatnote frequency is 6305 MHz.[9]

Power drifts in the MOT cooling light cause the load rate and total atom number to drift, which can be problematic for subsequent stages of the experiment. We monitor the MOT atom number by collecting a fraction of the MOT fluorescence with a photodiode each experimental run. The photodiode signal is calibrated into atom number, from which we extract the MOT load rate and total atom number from each run. Figure 3.11 shows the

---

[9]The double-passed MTS setup results in the lock point being 80 MHz to the red of the actual cooling transition.
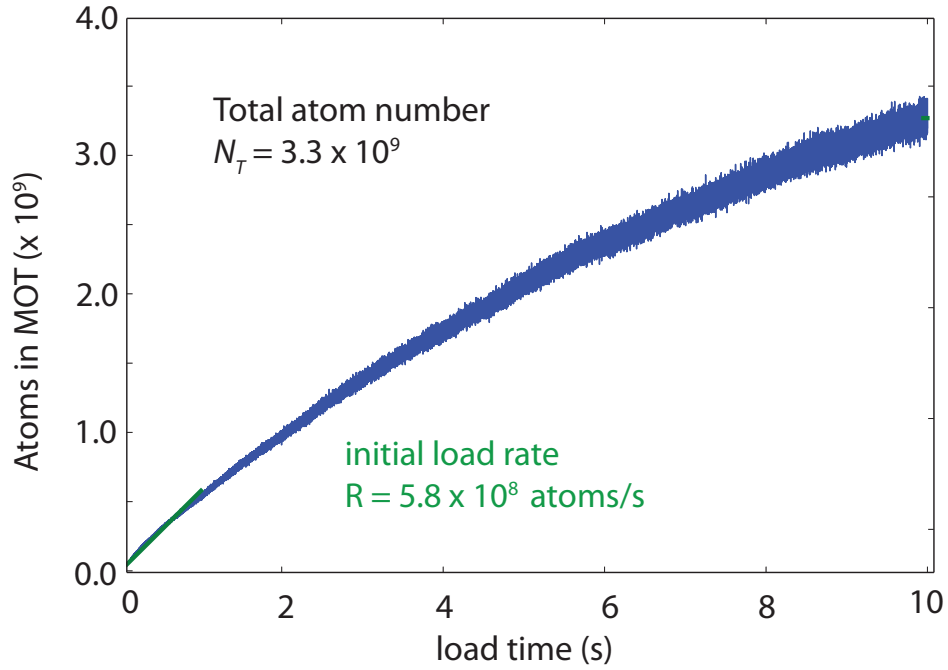
Figure 3.11: Loading the MOT. A photodiode collects flourescence from the loading MOT, which is then calibrated to absolute atom number. Each experimental run, we acquire a load rate curve to monitor the performance of the MOT.

result of such a measurement. The predominant cause of MOT beam power drift is thermal stabilisation of the MOPA over the course of several hours, a design failure of particular irritance. In future, the photodiode signal will be used to measure the atom number load and feedback to the control system in a servo, alleviating these difficulties.

3.8.2 Molasses

Each stage of the production process is designed to optimise the number of atoms transferred to the next stage. To optimise the transfer of a spatially extended and comparatively warm MOT into the magnetic trap, we perform an optical molasses stage. The Doppler temperature $T_{\mathrm{Doppler}} = \hbar\Gamma/2k_B = 290\mu K$ for ^{87}Rb is the lowest attainable temperature predicted by a simple two-level analysis of laser cooling [75]. The transfer of atoms into the magnetic trap is optimum for cold, dense clouds with consequently less spatial extent. Attaining sub-Doppler temperatures is possible by performing optical molasses, which in turn greatly improves the load into the magnetic trap.

Optical molasses circumvents the Doppler limit due to additional, multi-level cooling processes that take place when the quadrupole magnetic field of

the MOT is switched off. The resulting purely optical potential is no longer a trap, an atom transiting the optical molasses thus behaves like a particle moving through a viscous fluid, hence the name.

In the absence of the quadrupole field, the magnetic states of the atoms are approximately degenerate. Each pair of MOT beams consists of an opposing pair of $\sigma^+\sigma^-$ polarised beams, which form a spatially periodic polarisation grating. This in turn gives rise to a spatially varying vector AC Stark shift of the magnetic sublevels, which drives polarisation gradient cooling (PGC) [75, 111]. Although we omit here the details of PGC, we make note of an important experimental requirement for its success: low magnetic fields, where the Zeeman shift of magnetic sublevels is small compared with the optical pumping rate [112]. We minimise the background magnetic fields using bias magnetic fields that cancel the contribution from Earth's field as well as that of stray fields from magnetic lab components. A bias field coil pair is aligned along each axis (x, y and z) and each coil is driven by a separate transconductance current driver.

The minimum temperature attainable with PGC scales with the amplitude of the spatial polarisation modulation, which is determined by the magnitude of the Stark shift, $T_{\min} \sim \frac{1}{k_B} I / |\Delta|$ [75]. Coldest temperatures are thus achieved by increasing the detuning and reducing the power, which in turn reduces the rate of the cooling process. In practice one must limit the time over which PGC takes place, as the absence of a restoring force in the molasses potential can lead to the cloud falling substantially, making loading into the magnetic trap inefficient. Before initiating pure PGC, the quadrupole gradient is reduced to 2 G/cm and the MOT beam detuning increased to -30 MHz for 5 ms. We found that this precursor stage improved the atom number compared to when the final molasses (0 G/cm, -60 MHz) parameters were immediately switched to after the MOT stage, which may be attributable to an issue with the switch-off time of the coils.

We then perform PGC for 8 ms by setting the quadrupole gradient to zero, reducing the repump amplitude to $\sim 100\,\mu$W and detuning the MOT beams -60 MHz from resonance. The repump beam is then completely extinguished for the final 2.5 ms of PGC, which transfers most of the atoms to the $F = 1$ hyperfine state via passive optical depumping. At the conclusion of optical molasses, the temperature of the $\sim 10^9$ atoms[10] is 31(3) μK.

---

[10]Attaining accurate atom numbers for the molasses stage with absorption imaging is very difficult, due to the high optical depth. Ideally the molasses number is N_{MOT}.

3.8.3 Magnetic trap

After concluding the laser cooling processes with molasses, the atoms are transferred to a purely magnetic trap. Each ^{87}Rb atom in the state $|F = 1, m_F = \pm 1\rangle$ carries a magnetic moment of $|\mu| = \mu_B/2$; $m_F = -1$ atoms, which are attracted to regions of low magnetic field strength, can be magnetically trapped in a local magnetic field minimum. The magnetic potential derives from the Zeeman Hamiltonian:

$$H_Z = -\mu \cdot \mathbf{B}. \tag{3.9}$$

The simplest magnetic trap is a quadrupole potential, which we already discussed in the context of providing a spatially varying Zeeman detuning in a MOT. A quadrupole magnetic field can be produced by a pair of anti-Helmholtz coils aligned along the $y-$axis:

$$\mathbf{B} = \frac{B'_q}{2} \left(x\,\hat{\mathbf{x}} - 2y\,\hat{\mathbf{y}} + z\,\hat{\mathbf{z}} \right), \tag{3.10}$$

with the field gradient strength B'_q. The Zeeman potential for a quadrupole magnetic field is thus

$$U(x,y,z) = m_F g_F \mu_B B'_q \sqrt{\frac{x^2}{2} + y^2 + \frac{z^2}{2}}. \tag{3.11}$$

It is straightforward to show that a quadrupole gradient of $B'_q = 30.45\,\text{G/cm}$ is sufficient to support $F = 1, m_F = -1$ atoms against gravity (along y).

Quadrupole coils

Our quadrupole coils are formed from a pair of water-cooled, square cross-section copper wire coils with 42 turns per coil. Current is driven through the coils from by a high-current driver circuit, which takes a 0-5 V analog control signal and sources up to 100 A from an Agilent 6690A power supply. The approximate current-to-gradient calibration is $1.89\,\text{G}\,\text{cm}^{-1}\,\text{A}^{-1}$.

The water flow through the quadrupole coils is maintained parallel through each coil by a magnetically-geared pump (Micropump GJ-120), at a rate of 0.26 l/min. The Micropump also drives cooling water through the coils of the Zeeman slower, with a Julabo HE-FP60 chiller keeping the coolant at 20°C. The water cooling system is interlocked using a Galil RIO-47100 PLC in a similar manner to the oven system: water cooling failures trigger the PLC to disable both Zeeman supplies as well as the quadrupole power supply.

Bias coils

Magnetic bias fields are integral elements of the apparatus. We use three sets of rectangular coil pairs to produce magnetic fields that cancel stray fields or provide bias fields for state preparation. The application of bias fields also allows the quadrupole centre to be moved, which is important for optimising the transfer of the atomic cloud from one stage of the experiment to another, such as molasses to magnetic trap and the particularly sensitive decompression of the magnetic trap into the dipole trap. Figure 3.12 depicts the bias coil geometry; each coil of the y and z oriented pair is driven independently, allowing gradients to be applied along these directions, whereas the x coils are driven in series. Current is sourced from a 40 A power supply by five independent transconductance circuits that drive current through each coil in proportion to a 0-5 V control voltage from the NI-PCIe6733 analog outputs. Double-pole double throw relays are installed on the y-bias coils to allow the field direction to be inverted, as the current driver circuits are not able to provide negative currents in the interest of maintaining linear response near zero current. The driver circuits were designed and built by R. P. Anderson, and are encased in a rack-unit box affectionately called the *Magneato*.

Operation

The process of transferring the laser cooled cloud into the magnetic trap is known as the *catch*. The name is apt, as during the molasses stage the cloud is essentially in freefall experiencing no restoring force from the optical potential. We catch the atoms by ramping the quadrupole gradient from zero to 40 G/cm in 3 ms. The bias coils are again configured to fields that optimise this process; any severe discrepancy in the position of the quadrupole zero (which is assumed to be at the geometric centre of the coil pair) and the molasses (assumed to be close to the centre of the laser beam intersection) results in sloshing of the cloud. We find that a 5 G y-bias field is needed to ensure the quadrupole zero matches the molasses – this is may be due to a manufacturing fault that renders one of the quadrupole coils with more turns than the other.

A very important metric for the success of the catch is the number of transferred atoms. Although temperature should be kept as close to that of the molasses, some heating during the transfer is inevitable. We can determine the caught atom number in two ways: comparing the number of atoms in the MOT (determined via fluorescence) to the integrated atom number from an absorption image, or by performing a *recapture* measurement. The former is straightforward but is not a like-for-like comparison, as there are inherent ambiguities in comparing the fluorescence measurement to an
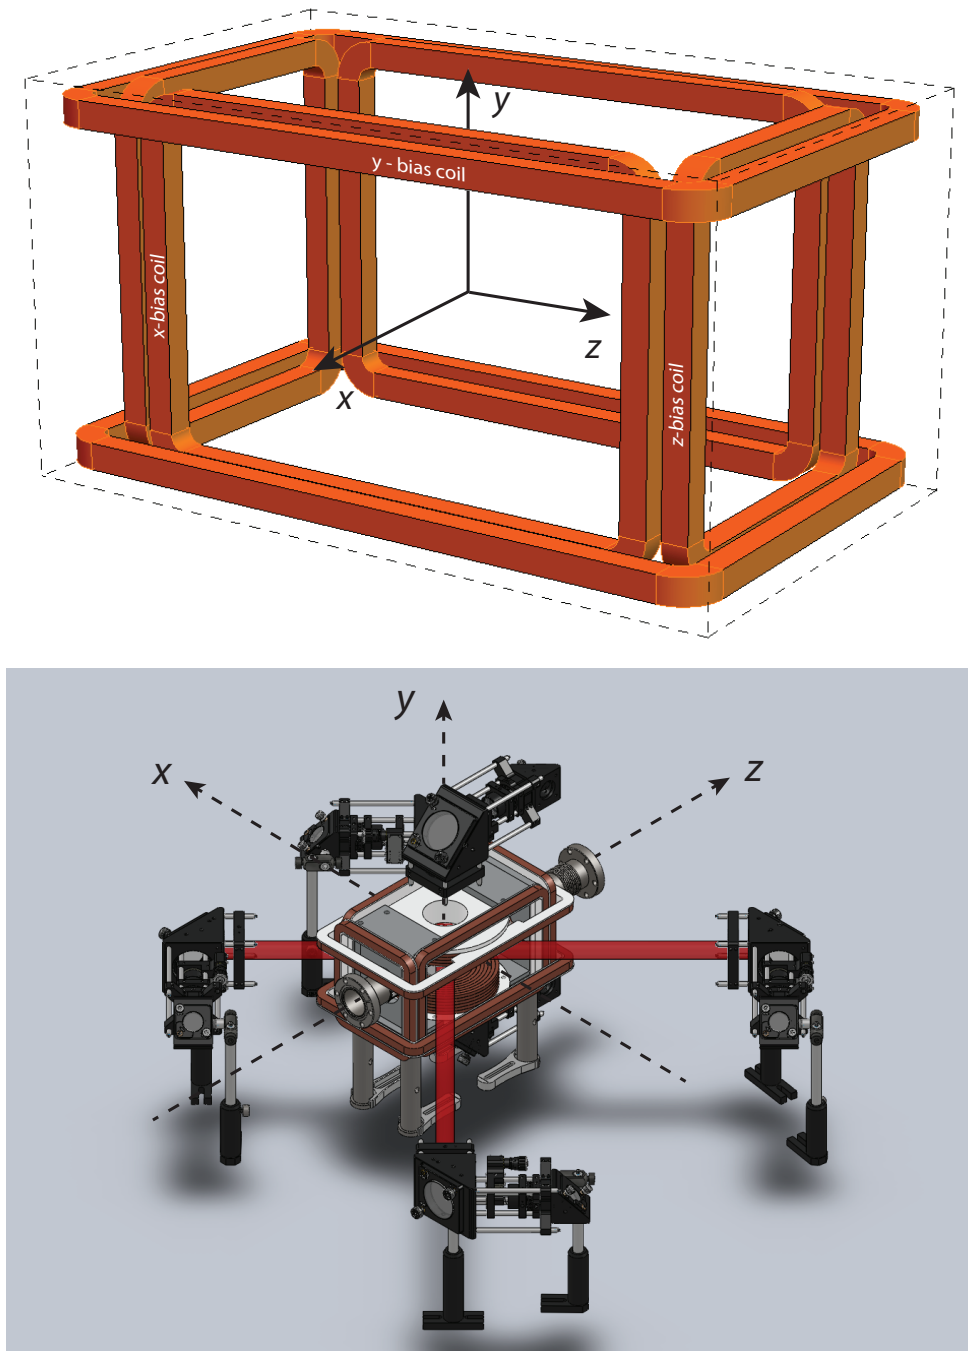
Figure 3.12: Magnetic bias coil assembly. Top: Design schematic of the bias coils: each coil has a rectangular shape, with 16 turns per coil. Bottom: Bias coil assembly set up around glass cell, showing MOT beams and bias field axes. Images courtesy R. P. Anderson.
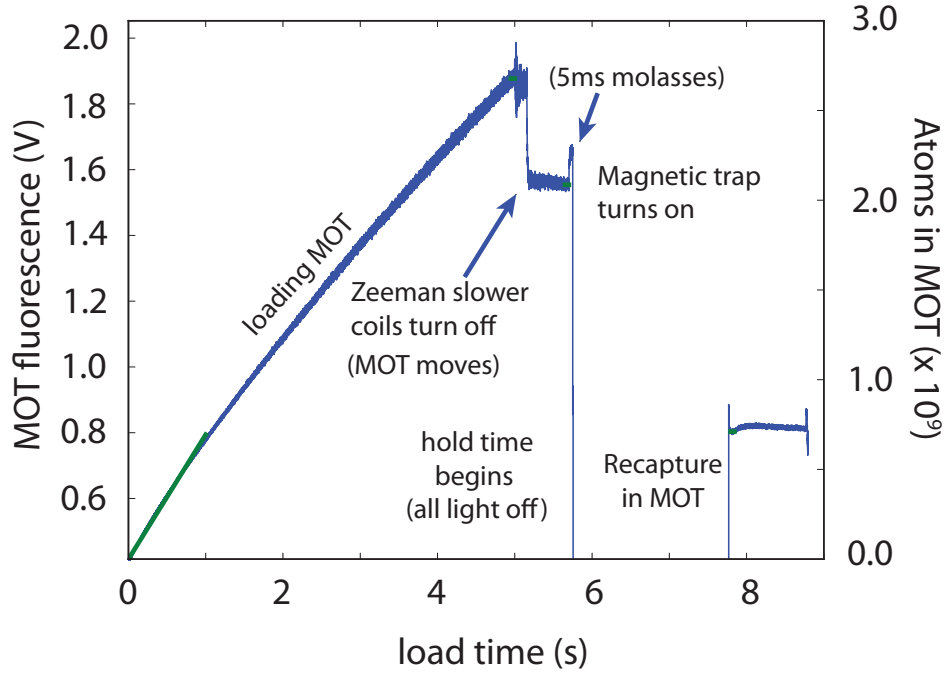
Figure 3.13: Recapture measurement. Fluorescence signal from the MOT is used to determine the fraction of atoms transferred to the magnetic trap by loading a MOT, performing a molasses stage and then transferring the laser cooled atoms into the magnetic trap. After a hold time, the magnetic trap is turned off and the MOT switched back on, while the atom beam is shuttered. The fluorescence signal is thus a measure of the atoms transferred to the magnetic trap. The capture fraction of 34% is consistent with only passive optical pumping to the $|F = 1, m_F = -1\rangle$ magnetically trappable state.

absorption image; the molasses is too dense to reliably extract atom numbers from an absorption image directly. A recapture measurement in contrast is like-for-like. We can compare the fluorescence of the loaded MOT to the magnetically trapped cloud by 'recapturing' the latter with the MOT, so that the fluorescence of the recaptured atoms may be compared to the loaded MOT fluorescence. We show the results of a recapture measurement in Figure 3.13. The atom number and temperature of the cloud loaded into the magnetic trap is $N \approx 1 \times 10^9$ and $T = 46(4)\,\mu$K, corresponding to a phase space density of $7(2) \times 10^{-5}$.

3.8.4 Evaporative cooling

Until only recently [113], evaporative cooling was used at some point in the creation of every BEC. Evaporative cooling works by selectively removing warm, energetic atoms from a thermal distribution of energies. When these

atoms are removed from the sample, the ensemble rethermalises via elastic collisions to a lower overall temperature, and atoms repopulate the high-energy tail of the lower-temperature thermal distribution. Continuously removing these hot atoms allows lower temperatures and high phase space densities to be reached. We use evaporative cooling in the magnetic trap to cool atoms for loading into the hybrid trap as well as during final cooling in the pure dipole trap to quantum degeneracy. In this Section, we give a brief overview of the process and refer the reader to several notable references (such as Ref. [114]) for a more detailed treatment.

Forced evaporative cooling in the magnetic trap

Forced evaporative cooling is performed in the magnetic trap by using radiofrequency (rf) transitions to selectively outcouple atoms from the outermost region of the cloud. For the quadrupole potential, the magnetic field increases linearly from the centre of the trap, thus the atoms with the highest potential energies populate the edges of the trap. Applying an rf field resonant at a point in space such that (along the x-axis, for instance) $\hbar\omega_{\mathrm{rf}} = m_F g_F \mu_B B'_q x$ will selectively drive rf transitions between the Zeeman sublevels. As $m_F = 0$ atoms are not trappable and $m_F = +1$ atoms are repelled from the trap, the atoms interacting with the rf are expelled. Changing the rf frequency to be resonant with other, lower energy equipotential shells allows us to progressively remove the warmest atoms from the quadrupole trap.

The same effect could also be achieved by reducing coil current, which would 'spill' hot atoms from the trap edges. However, the crucial factor in sustaining successful evaporative cooling is the elastic collision rate, which rethermalises the atoms to a lower temperature. As the trap depth is reduced, the confinement strength in turn drops, and cooling stagnates when the ratio of 'good' elastic collisions to 'bad' collisions (such as inelastic losses from background atoms) becomes comparable. Forced evaporative cooling in the dipole trap can *only* be done by reducing the trap depth at the expense of the elastic collision rate.

After the catch stage, we adiabatically ramp the quadrupole gradient up to 155 G/cm over 500 ms. Increasing the strength of the trapping potential by a factor of k increases the collision rate by $k^{\frac{4}{3}}$, the density by k and the temperature by $k^{\frac{2}{3}}$ [114]. We then apply an rf field at 25 MHz and gradually reduce the frequency linearly to 6 MHz over 5 s. The atom number, temperature and phase-space density of the cloud in the compressed MT prior to evaporation is $N = 8.2 \times 10^8$, $T = 128(2)\,\mu$K and $3.7(2) \times 10^{-6}$, and at the conclusion of the rf evaporation sequence we are left with 1.5×10^8

atoms with $T = 36.1(6)\,\mu\text{K}$ with a phase-space density of $3.42(5) \times 10^{-4}$.

3.9 Dipole trap

At the instigation of magnetic trapping, we also switch on the dipole beams, forming a hybrid magnetic quadrupole – optical dipole trap. We now consider the operation of the dipole trap in more detail.

3.9.1 Theory

The optical forces on an atom in the dipole trap are distinct from those at work in laser cooling, in which near-resonant photons impart momentum onto atoms with each scattering event. In a dipole trap, the light is typically many nanometres off-resonant from electric dipole transitions, where photon scattering is minimal. The dipole force is due to the induced dipole moment of the atom due to the electric field of the laser interacting with the intensity gradient of the light. The review by Grimm et al. [115] is the definitive reference for optical trapping using far off-resonant light, and this Section summarises a few key results.

The strength of the induced dipole moment is related to the electric field amplitude by $p = \alpha E$, where the complex polarisability α depends on the atomic eigenstate, polarisation of the light and detuning from atomic transition lines. The corresponding potential is given by

$$U_{\text{dip}}(\boldsymbol{r}) = -\frac{1}{2\epsilon_0 c}\text{Re}(\alpha)I(\boldsymbol{r}), \qquad (3.12)$$

with ϵ_0 the free space permittivity, c the speed of light, and I the optical intensity. A spatially varying electric field, such as that of a tightly focused laser beam can be used to form a trap with the force imparted on the atoms proportional to the intensity gradient of the light field. The dipole potential can also be explained as the AC Stark shift of the atomic eigenstate, with the detuning-dependent energy shift giving rise to a attractive potential for red-detuned light ($\Delta = \omega - \omega_0 < 0$) and repulsive for blue ($\Delta > 0$). The residual photon scattering rate $\gamma \propto I/\Delta^2$ is reduced for large detuning, and trap depths can be made appreciable by using higher intensities.

In the case of an alkali metal atom in the state $|F, m_F\rangle$, Eq. 19 of Ref. [115] gives the form of the dipole potential for light with $\Delta \gg \Delta'_{\text{hfs}}$, the excited state hyperfine splitting ($\sim 500\,\text{MHz}$ for ^{87}Rb):

$$U_{\text{dip}}(\boldsymbol{r}) = \frac{\pi c^2 \Gamma}{2\omega_0^3}\left(\frac{2 + Pg_F m_F}{\Delta_{\text{D2},F}} + \frac{1 - Pg_F m_F}{\Delta_{\text{D1},F}}\right)I(\boldsymbol{r}), \qquad (3.13)$$

with $\Gamma = 2\pi \times 6.066\,\text{MHz}$ and the detunings $\Delta_{\text{D1,D2}}$ are from line-centre of the D1 $(^2\text{S}_{1/2} \rightarrow^2 \text{P}_{1/2})$ and D2 $(^2\text{S}_{1/2} \rightarrow^2 \text{P}_{3/2})$ lines respectively. The polarisation of the light $P = 0, \pm1$ for linear and $\sigma^\pm$ circularly polarised light respectively lifts the degeneracy of the m_F states for $P \neq 0$, yielding a potential that depends on the spin state via m_F. This is a manifestation of the vector atomic polarisability, discussed at length in Chapter 7.

When designing a dipole trap, the principal concerns are that the trap can support the atoms against gravity and that the scattering rate is low. High power near-infra red lasers are eminently suited for the latter purpose, with powers of tens of watts readily available at Nd:YAG (1064 nm) or telecommunications (1565 nm) wavelengths. Larger beams result in higher volume traps with more atoms, but require more power than beams focused to a smaller waist.

In cylindrical coordinates, a beam with wavelength λ_0 and total optical power P focused to a waist w_0 has an intensity profile given by

$$I(r,z) = \frac{2P}{\pi w(z)^2}\exp\left(-\frac{2r^2}{w(z)^2}\right),\qquad(3.14)$$

with the $1/e^2$ width $w(z)$ as a function of the distance along the axial direction from the focal point ($z = 0$) given by $w(z) = w_0\sqrt{1 + (\lambda_0 z/\pi w_0)^2}$. The dipole potential (Equation 3.13) for various optical powers is plotted in Figure 3.14.

In order to form BEC in a dipole trap, forced evaporative cooling is initiated by reducing the trap depth. In a dipole trap, this is done by reducing the power in the beam, with a caveat: reducing the power by k also reduces the confinement strength by $\sqrt{k}$, and thus makes efficient rethermalisation difficult to achieve. Figure 3.14 shows the reduction in depth and accompanying reduction in harmonic confinement. Numerous schemes have been conceived to circumvent this issue [116–118], but the added complexity is often avoided by using a crossed-beam trap geometry or axial quadrupole gradient, or both [119]. In all cases the confinement relaxes along with the power reduction, but the initial (and thus final) trapping frequencies are higher than the single-beam case due to the added axial confinement. Evaporation is thus faster and more efficient in a crossed-beam optical dipole trap, allowing for large BECs.

3.9.2 Optical layout

Our dipole trap is formed from the intersection of two beams originating from the same 1064 nm 20 W fibre laser (Keopsys CYFL-20W-MEGA) with a nominal linewidth of 5 MHz. Figure 3.15 shows the layout of the optical components. We based our optical design off the cross-beam trap used by
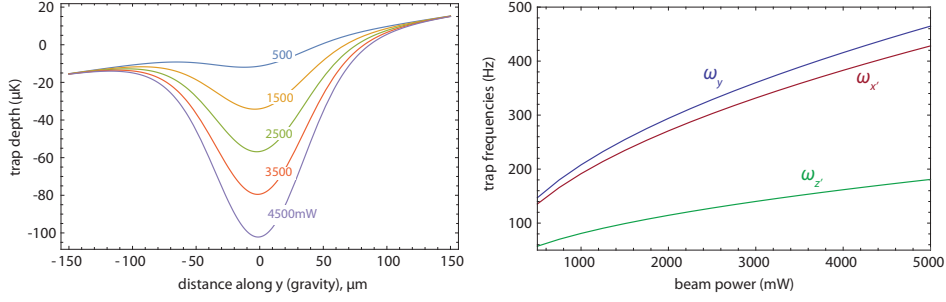
Figure 3.14: Evaporation in the crossed-beam dipole trap. Plotting the dipole trap potential (Eq. 3.13, with gravity) for different beam powers (left) and computing trap frequencies as a function of power (right). The powers are equal in each beam, whereas the beam waists are $w_{z'} = 65\,\mu m$ and $w_{x'} = 100\,\mu m$ for the beams travelling along the z' and x' directions, as described in Figure 3.15.

the group of David Hall at Amherst College [120]. Notably, we use dichroic mirrors to steer the dipole beams into the glass cell along the direction of the MOT beams. The 780 nm MOT beams are transmitted by the dichroics while 1064 nm is reflected. This configuration has the advantage of maximising the optical access for probing the condensate without obstruction by trapping light optics.

We refer to the dipole beams by the approximately orthogonal axes x' and z', pointing in the propagation directions of each beam (Figure 3.15). The two beams are quite different, the z' beam has a design waist[11] of $w_{z'} = 70\,\mu m$ and the x' beam waist is $w_{x'} \sim 100\,\mu m$. The z' beam is the principal trapping beam, with the x' beam providing axial trapping along the z' axis, as shown in Figure 3.15, inset.

The amplitude and position of the dipole beams are controlled by a 110 MHz AOM in each beam path. The power in each beam is controlled by altering the amplitude of the AOM rf drive, and the position may be varied by $\sim 500\,\mu m$ by changing the rf frequency. Each AOM is operated at a centre frequency of 101.1 MHz and use opposite diffraction orders to attain a relative frequency shift of 202.2 MHz between the two beams, which eliminates standing wave effects.

3.9.3 Hybrid trapping

The addition of the dipole beams circumvent a severe difficulty of quadrupole magnetic traps, Majorana spin-flip losses. Atoms cannot be cooled to quantum degeneracy in a quadrupole trap. An atom traversing the quadrupole

---

[11]Due to aberrations, the actual waist is different. When the dipole beams are split and translated, as in Chapters 6 and 7, the waist is as low as 65 μm.
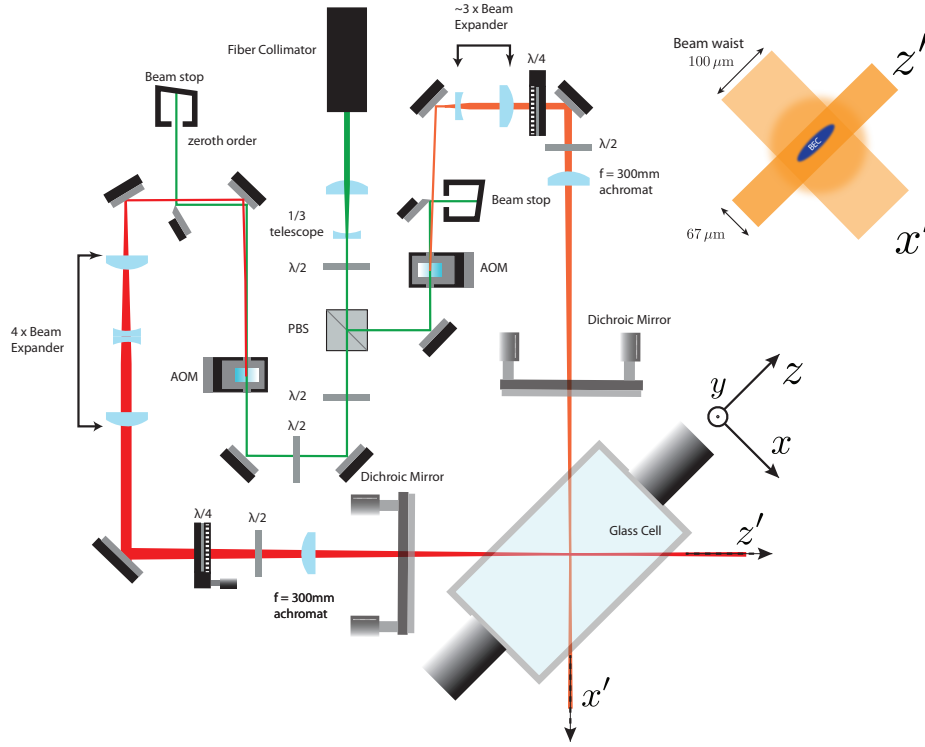
Figure 3.15: Schematic of the dipole trap optical setup. A fibre collimator with an in-built optical isolator supplies 13 W of 1064 nm light, which is split into two beam paths using a polarising beam-splitter cube (PBS). Each beam is then focused to a waist of 1 mm and passes through an AOM. The first diffracted order from each AOM is then split off (the zeroth order enters a beam dump) and increased to ~ 4 mm in diameter. An $f = 300$ mm achromatic doublet lens in each beam path then focuses each beam to a waist in the glass cell. Each dipole beam is steered by a dichroic mirror in the path of a MOT beam. The dichroic mirrors reflect the 1064 nm dipole trapping light into the glass cell; the dipole beams define near orthogonal axis x' and z'; the z' beam provides the primary trapping force, with the x' giving radial confinement (inset top right). See Chapters 6 and 7 for additional changes to the optical setup.

zero has a probability to undergo a non-adiabatic spin-flip, which transfers it to an untrapped state. As this process is strongest around the quadrupole zero, the atoms with the lowest potential energy are lost, which heats the gas. The rate of Majorana losses can be estimated as [37]

$$\Gamma_{\text{Maj}} = 1.85 \frac{\hbar}{m} \left(\frac{m_F g_F \mu_B B_q'}{k_B T} \right)^2 , \tag{3.15}$$

with B_q' the quadrupole gradient and T the temperature of the gas. Majorana losses are thus most severe at low temperatures and high field gradients. The quadrupole trap 'leaks' atoms with lowest potential energy, making runaway evaporation at low temperatures impossible [114].

There are several methods that overcome Majorana losses in quadrupole traps: a blue-detuned plug beam located at the centre of the trap, for instance, actively repels atoms from the quadrupole zero [2, 121]. Additional static or time-varying magnetic bias fields may be added to make Ioffe-Pritchard traps or time-orbiting potential traps, which formed the basis of the first generation of BEC experiments.

In the *hybrid* trapping scheme [37], the quadrupole trap is used to enhance the load into a red detuned dipole trap. Loading a dipole trap from a purely laser cooled cloud is highly inefficient, as the spatial extent of the cloud is far greater than the trapping volume of the tightly focused dipole beam and this leads to a poor transfer fraction. Evaporative cooling in the quadrupole trap results in very cold, dense clouds that can be transferred into the dipole trap with considerably greater efficiency. The hybrid trap is also a considerably simpler system to implement than Ioffe-Pritchard magnetic traps, requiring lower coil currents and simpler coil geometries. Although our hybrid trap serves as merely a precursor stage to the purely optical trap, which is discussed in depth in the following Sections, the addition of the quadrupole field provides axial trapping along the dipole beam, which is essential for sustaining evaporative cooling in a large volume single-beam trap.

Loading the hybrid trap

To load the hybrid trap, we reduce the quadrupole gradient from $155\,\text{G/cm}$ to $29\,\text{G/cm}$ over $3\,\text{s}$, while simultaneously sweeping the rf knife from $6\,\text{MHz}$ to $3.75\,\text{MHz}$ during the first half of this decompression. The decompression phase results in a substantial temperature reduction, from $36.1(6)\,\mu\text{K}$ to $6.3(1)\,\mu\text{K}$. The positioning of the dipole beam is crucial to this stage, as this temperature reduction increases the spin flip loss rate when the dipole beam is misaligned. We were able to ensure the correct positioning of the dipole

beams by observing heating and losses of the optically trapped atoms with reference to the position of the quadrupole zero, using methods described in Ref. [122]. The loaded hybrid trap contains 2.3×10^7 atoms at a temperature of $6.3(1)\,\mu K$, with a corresponding phase-space density of $9(1) \times 10^{-3}$.

3.9.4 Formation of BEC

When evaporating to BEC in the crossed-beam trap, the quadrupole gradient is further reduced from $29\,G/cm$ to zero over half the duration of the optical evaporation, meaning that for the last half of the sequence the confinement is purely optical. Forced evaporation is then instigated in the dipole trap by reducing the optical power in each beam from $\sim 4\,W$ initially to about $600\,mW$ over $5\,s$ in the crossed-beam trap. Figure 3.16 shows the optical trap evaporation sequence after various truncated evaporation sequences. The formation of BEC is heralded by the presence of a bi-modal (Thomas-Fermi + Gaussian) density distribution in absorption images. We found that the crossed-beam trap can yield BEC in as little as $3\,s$ of optical evaporation, whereas at least $15\,s$ was required in the single beam hybrid trap. This difference may be due to the sensitivity of the evaporation process lower trapping frequencies in the single-beam hybrid trap and possibly the proximity of the quadrupole zero. Also, we find that the alignment of the crossed-beam trap is very robust, remaining aligned and operational over the course of several months.

The entire trajectory from MOT to BEC is shown in Figure 3.17. There are some minor areas for improvement - particularly in hybrid load stage and dipole evaporation stages. We suspect the excessive loss of atoms in these stages is related to the population of atoms in the 'wings' of the crossed dipole potential [117], thermal atoms that contribute to the measured atom number but are not loaded into the beam intersection. This effect is absent for single beam hybrid traps. With minor optimisation work we anticipate the atom number can be improved significantly, so that pure BECs of 1×10^6 atoms could be created in as little as $13\,s$ ($2\,s$ MOT load + $5\,s$ rf evaporation + $3\,s$ hybrid load + $3\,s$ optical evaporation).

Trap varieties: The semi-isotropic trap

The simplest trap we can form consists of roughly equal powers in each dipole beam, so that the relative confinement is dictated primarily by the beam waists. To form BEC in this trap, we lower the beam powers until a pure condensate is formed; any further reduction of the trap depth cuts into the condensed fraction. We refer to this trap as the 'semi-isotropic' (SI) trap
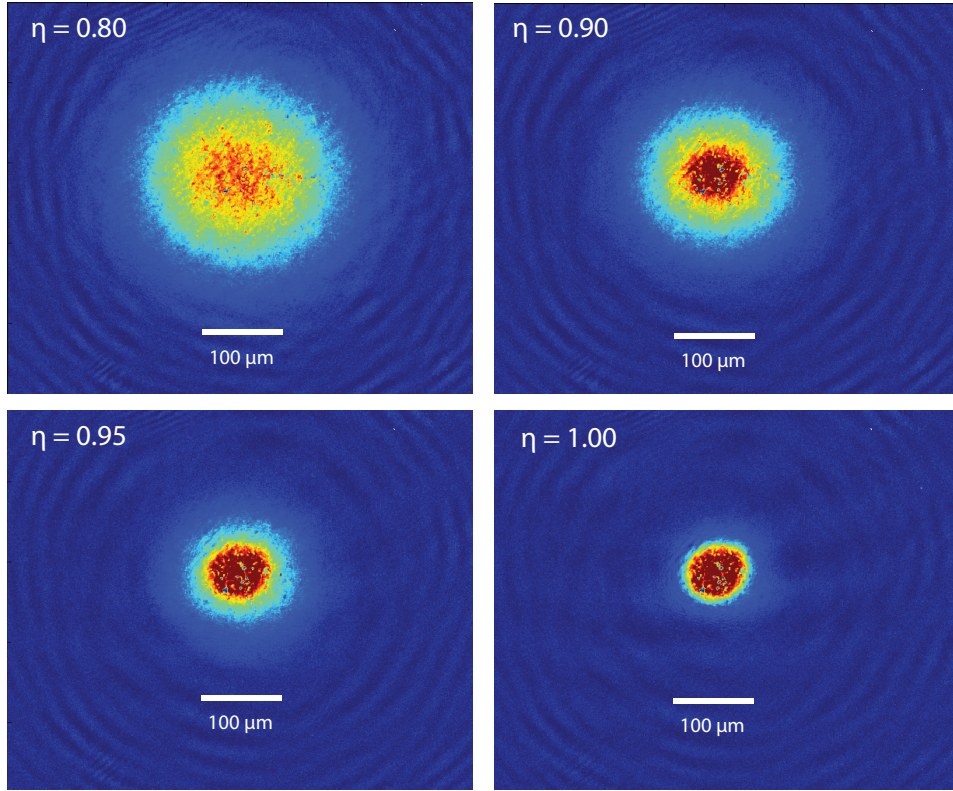
Figure 3.16: Onset of BEC in the crossed dipole trap. The truncation parameter η characterises the progress of the 5 s evaporation ramp: $\eta = 0.8$ corresponds to the dipole beam amplitudes stopping at 80% of their final amplitude, over 4 s. The final frame shows an almost pure BEC of 3×10^5 atoms. The maximum optical depth shown is OD 3. The time-of-flight is 20 ms.

due to the near equality[12] of the trapping frequencies. The SI trap is our workhorse trap, due to its considerable simplicity and robustness, although the low atomic densities make observing spinor physics difficult.

Trap varieties: The elongated trap

The trap geometry in which we can evaporatively cool a thermal cloud to BEC is limited almost exclusively to the semi-isotropic trap. As the trap depth must be reduced to force evaporation, the scope to alter the geometry by increasing the power in one of the beams is limited. However, once BEC is formed, we can then modify the trapping potential by changing the relative intensities of the two beams. We can form an elongated trap by increasing the intensity of one dipole beam and reducing that of the other, so that the BEC

---

[12]In contrast to the elongated trap!
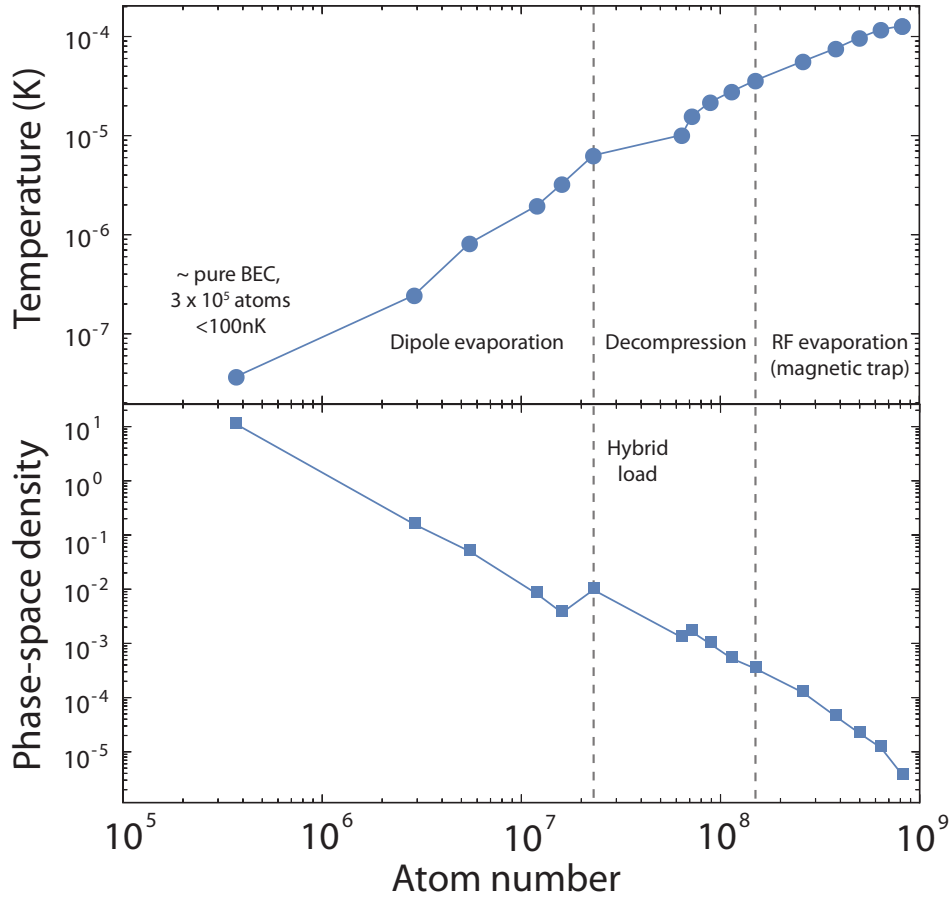
Figure 3.17: The path to BEC. Atom number vs. temperature (top) and phase-space density (bottom) during three stages of the experiment - RF evaporation in the magnetic trap, decompression and transfer into the hybrid trap followed by pure optical evaporation in the crossed dipole trap.

is spatially extended along the weaker beam. We can never totally extinguish one of the beams, as the condensate will simply expand to no end along the single beam[13] or, in our observations, slide off to one side due to slight tilting of the beams. To transfer the BEC to the elongated trap, the z' beam power is increased from $500\,\mathrm{mW}$ to $\sim 3\,\mathrm{W}$ and the x' beam power reduced from $500\,\mathrm{mW}$ to $50\,\mathrm{mW}$ over $2\,\mathrm{s}$, the trap is left to stabilise over a further $100\,\mathrm{ms}$.

The elongated trap is particularly useful for exploring spinor physics, as the higher densities available (discussed in the next Section) give rise to a spin interaction strength almost double that of the SI trap. This comes at the cost of several irksome features, the most serious of which is the presence of a thermal component that appears during the transfer process. Due to the

---

[13]Axial trapping frequency from a single beam is $< 1\,\mathrm{Hz}$.

higher trap depth, this thermal component persists.

3.9.5   Trap frequency measurements

The trapping strength of the dipole trap and the size of the condensate may be characterised by approximating the potential as harmonic and determining the trapping frequencies. In the simplest case, exciting bulk motion (such as a sudden translation) along each direction of symmetry x', y, z' will result in centre of mass oscillation ('sloshing') at an angular frequencies $\omega_{x',y,z'}$. Both our crossed-beam dipole traps can be described by three trap frequencies: a tight radial frequency $\omega_{x'}$ from the beam with a smaller waist, $\omega_{z'}$, the weaker radial confinement of the larger dipole beam that serves as the elongation axis of the condensate and ω_y, the trapping frequency along the direction of gravity (see Figure 3.15). Although the dipole beam intensity profile is roughly symmetric in the tight radial direction, gravity displaces the trap minimum below the peak intensity point by $\sim 10\,\mu\text{m}$; this gravitational sag typically results in $\omega_y < \omega_{x'}$.[14]

In practice it is often not possible to excite single oscillation modes of the trap without exciting others, particularly in the SI trap. The resultant motion is difficult to study in 2D absorption images due to mode interconversion between motion along and orthogonal to the imaging axis. Figure 3.18 shows sloshing excited in the SI trap, with strong amplitude modulation due to precession in the trap and sloshing excited along the elongated (z') axis of the trap, showing simpler behaviour.

An alternative to sloshing measurements is parametric heating. Modulating the strength of a trap with trapping frequencies ω_i at frequencies $2\omega_i/n$ for integer n results in parametric resonance (PR), with atoms being rapidly heated and expelled from the trap [123, 124]. To excite PR, we program the RFBlaster that controls the dipole trap beams to amplitude-modulate the rf power (by between 5-10%) sent to the AOMs at a frequency ω for between 1-2 s. With each iteration of the experiment we vary ω and measure the number of atoms remaining in the trap after time of flight. Figure 3.19 shows the parametric resonance 'spectrum' for the two different trap varieties.

Determining the trap frequencies from a PR spectrum requires some attention. Although the strongest resonances occur at $\omega_{\text{mod}} = 2\omega_0$, they are not the only feature we observe. We use sloshing measurements to estimate the fundamental modes $\omega_{x'}, \omega_y, \omega_{z'}$, and thus identify features in the PR spectrum. Trap frequency measurements in the elongated trap showed strong resonances at $2\omega_r, 2\omega_y$ as well as the fundamentals, whereas only

---

[14]Any misalignment of the crossing beams along y strongly effects the overall trapping strength along that direction.
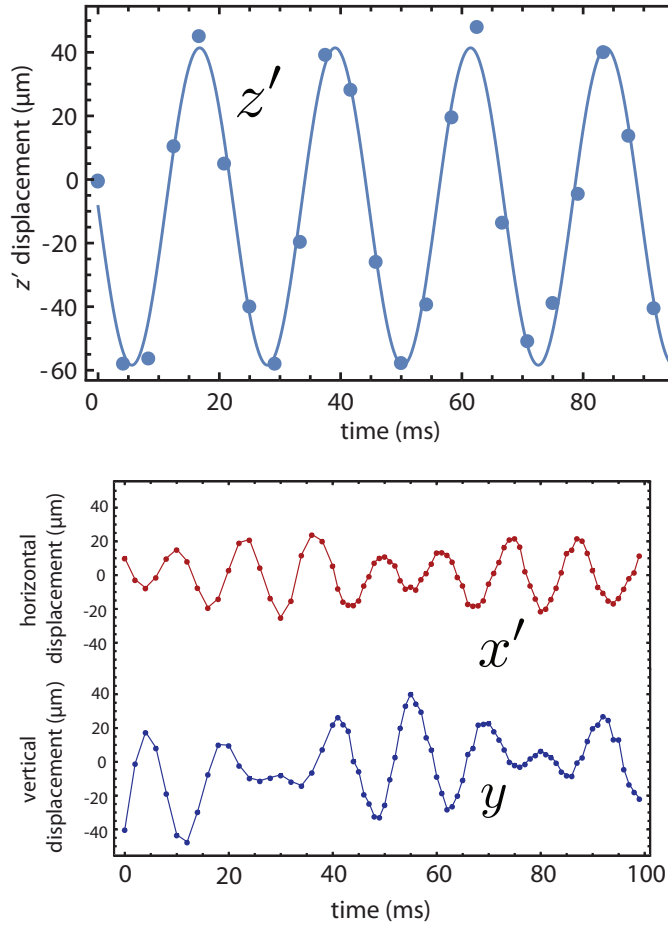
Figure 3.18: Sloshing measurements in the SI trap. Left: sloshing along the long (z') axis of the trap. Right: mode interconversion between sloshing in the x' and y directions. Observation of this behaviour informs the assignment of trap frequencies in parametric resonance experiments.

resonances at the fundamental modes were observed in the weaker, SI trap. The underlying reason for this is unclear and possibly related to the near equality of two of the trap frequencies ($\omega_{x'}$ and ω_y), although we can confirm from sloshing measurements that the features observed in the PR spectrum correspond to the fundamental modes.

3.9.6 Thomas-Fermi radii

The characteristic size of the condensate is parametrised by the Thomas-Fermi (TF) radii, which are dependent on the trap frequencies and number of atoms. Referring to Eq. 2.4, the TF radii expressed in terms of the chemical
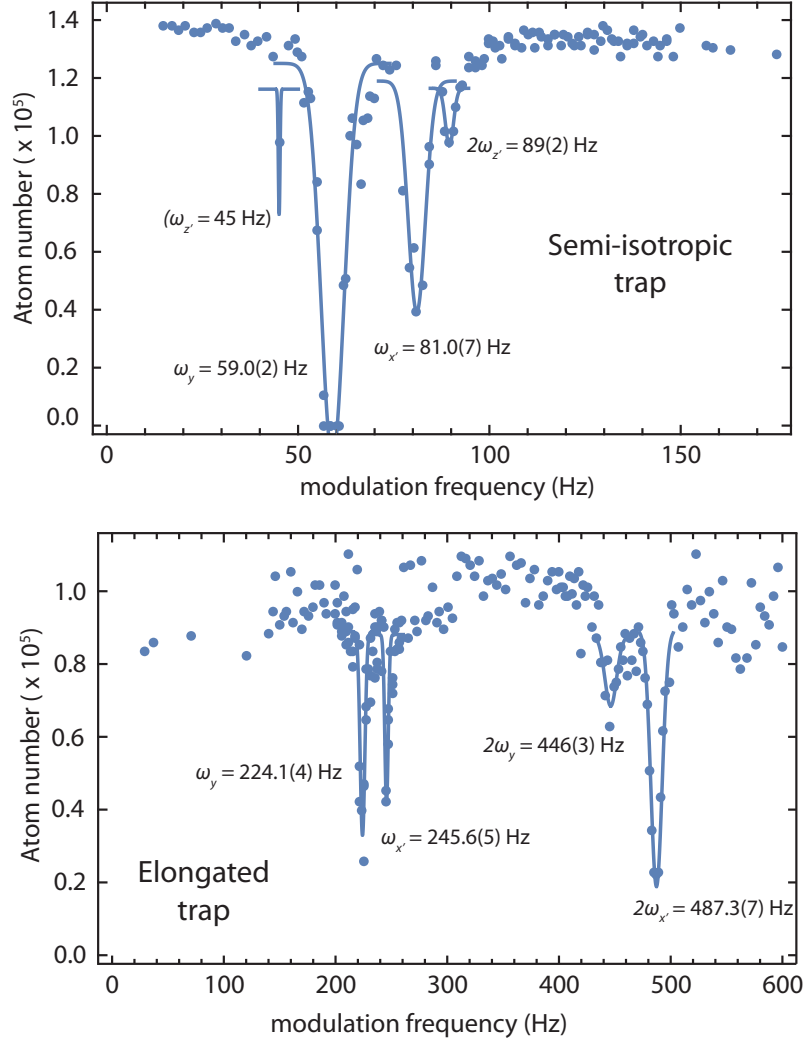
Figure 3.19: Parametric resonance spectra for the semi-isotropic trap (left) and the elongated trap (right). PR is excited by modulating the optical power of one or the other dipole beams in the crossed-beam trap. Resonances are detected by measuring the number of atoms remaining in the trap after time of flight. Modulation of the AOM rf drive power is over 1 s at depths 1% (left) and 5% (right).

potential are

$$R_i = \sqrt{\frac{2\mu}{m\omega_i^2}}. \tag{3.16}$$

The chemical potential is given by [73]

$$\mu = \frac{\hbar\bar{\omega}}{2}\left(\frac{15Na}{\bar{a}}\right)^{\frac{2}{5}} \tag{3.17}$$

with $\bar{a} = \sqrt{\hbar/(m\bar{\omega})}$ the harmonic oscillator length and $\bar{\omega} = (\omega_{x'}\omega_y\omega_{z'})^{1/3}$ the geometric mean of the trapping frequencies. We can determine the TF radii from the trapping frequencies for both the semi-isotropic trap and the elongated trap, as shown in Table 3.1. Comparing the geometric mean of the trapping frequencies for each trap, we see that $\bar{\omega}_{\text{SI trap}} = 2\pi \times 59.8\,\text{Hz}$ whereas $\bar{\omega}_{\text{E trap}} = 2\pi \times 98.6\,\text{Hz}$. The average densities for the SI trap and elongated trap are $8.8(4) \times 10^{13}\,\text{cm}^{-3}$ and $1.43(3) \times 10^{14}\,\text{cm}^{-3}$.

Trap	$(\omega_{x'},\omega_y,\omega_{z'})/2\pi$ (Hz)	$N_T(\times 10^5)$	$R_{x'}, R_y, R_{z'}$ (μm)
SI trap	$81.0(7), 59.0(2), 44.8(1)^S$	$1.8(2)$	$6.5(2), 9.0(2), 11.8(3)$
Elongated trap	$245.6(5), 224.1(4), 17.4(1)^S$	$1.35(7)$	$2.75(3), 3.01(3), 38.8(5)$

Table 3.1: Trapping frequencies, total atom numbers N_T and calculated Thomas-Fermi radii for the two different trapping geometries. Note: S indicates trap frequency derived from sloshing measurements; all other trap frequencies determined from parametric resonance.

3.10 Chapter summary

In this Chapter we have described our experimental apparatus, used for producing BECs. Our apparatus produces condensates of $\sim 3 \times 10^5$ atoms in about 25 s, and is routinely capable of producing over 2800 condensates a day, a satisfying outcome given our initial design philosophy of industrial production of BEC. The apparatus was constructed and commissioned during the first two years of my candidature, before relocating to the New Horizons centre across campus in early 2013. The gruelling task of moving the apparatus was performed whilst under vacuum to minimise the associated down-time upon resettlement in the new facility; nevertheless it still absorbed at least nine months of time in preparation, delays and re-commissioning. In the next Chapter, we discuss radiofrequency and microwave state preparation: the principal tools we use to make spinor condensates from single-component BECs and extract information from them.

State preparation techniques 4

In this Chapter we continue the description of our experimental apparatus and procedure, but this time focus on the steps used to produce well characterised superpositions of Zeeman states – *state preparation*. Although the state of the BEC throughout the production process is important, only passive state preparation techniques are used. For instance, depumping the molasses with a weak repump beam transfers at best 33% of the atoms to the $|F = 1, m_F = -1\rangle$ state for magnetic trapping. Experiments with a spinor condensate require that the initial state be a highly reproducible combination of Zeeman sublevels with a given phase.

4.1   Radiofrequency and microwave transitions

We prepare BEC in one or a superposition of several quantum states by driving magnetic dipole transitions between angular momentum states. The frequency and pulse duration of the applied field can be controlled with high precision, allowing preparation of states with well defined relative population fractions and phases. Resonant transitions between Zeeman sublevels $|F, m_F\rangle$ are dependent on the amplitude of the bias magnetic field, with the requisite frequency typically in the radiofrequency (rf) region of the electromagnetic spectrum (Equation 2.8). Transitions between states of different F are centred around 6.8 GHz, requiring microwave (mw) frequency photons to drive population transfer. We routinely use magnetic dipole transitions to prepare certain superpositions and measure the magnitude of magnetic fields.

   In this Section, the theory of rf and mw transitions is discussed, before describing the various experimental apparatus designed and built over the course of this project. Several concepts introduced in this Section will be discussed further in Chapter 5, such as rf coupling effects outside the rotating wave approximation.

4.1.1   Magnetic dipole transitions

The simplest example we can consider to formulate a mathematical under-
standing of how an atom interacts with a coupling field is that of the *two-level
system*. There is considerable motivation to do so, as the physics is universally
similar for most multi-level quantum systems – indeed a three-level spin-1
system can be represented using two spin-1/2 two level systems – and much
insight can be gained from such a simple model.

We consider a two-level system with states $|a\rangle$ and $|b\rangle$, separated by
an energy $\hbar\omega_0$. In the matrix representation, with the convention that the
energy splitting $\hbar\omega_0$ is symmetric about zero, the individual eigenstates and
Hamiltonian can be written as

$$a = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad b = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad H_0 = \frac{\hbar\omega_0}{2}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (4.1)$$

We now introduce the coupling field, oscillating at $\omega = \omega_0 + \Delta$ with am-
plitude $\hbar\Omega/2$: we will later identify Ω as the the Rabi frequency, a more
familiar measure of coupling strength. The coupling field can be written as
$G(t) = (\hbar\Omega/2)\exp(i(\omega t + \phi)) + c.c$, with ϕ representing the phase of the
field. The Hamiltonian describing magnetic dipole coupling is a matrix that
has only non-zero elements in the off-diagonal entries: $V_{ii} = 0, V_{ij} = G$. The
combined time-dependent Hamiltonian $H = H_0 + V$ is then expressed in
terms of only the driving frequency ω and detuning Δ:

$$H = \frac{\hbar}{2}\begin{pmatrix} \Delta - \omega & G \\ G & \omega - \Delta \end{pmatrix} = \frac{\hbar}{2}\begin{pmatrix} -\omega & 0 \\ 0 & \omega \end{pmatrix} + \frac{\hbar}{2}\begin{pmatrix} \Delta & G \\ G & -\Delta \end{pmatrix}, \quad (4.2)$$

where in the final step we have decomposed $H = H' + V$ into a sum of
a Hamiltonian rotating at the driving frequency H' and a coupling poten-
tial V. We can simplify the coupling potential by transforming into the
frame rotating at the driving frequency ω. Defining the unitary opera-
tor $U = \exp(-iH't/\hbar)$, we transform the Hamiltonian H into the inter-
action picture, $H_{\text{int}} = U^\dagger H U$. Invoking the rotating-wave approximation
(RWA) [80, 125] allows us to ignore terms rotating at twice the driving fre-
quency $(\exp(\pm 2i(\omega t - \phi)) \to 0)$, and we are left with

$$H_{\text{int}} = \frac{\hbar}{2}\begin{pmatrix} \Delta & \Omega e^{-i\phi} \\ \Omega e^{i\phi} & -\Delta \end{pmatrix}. \quad (4.3)$$

Now let us solve the Schrodinger equation $i\hbar\frac{\partial}{\partial t}|\psi\rangle = H_{\text{int}}|\psi\rangle$, for $|\psi\rangle =
c_a|a\rangle + c_b|b\rangle$ with H_{int} and the initial condition $|\psi(t=0)\rangle = |a\rangle$. Computing
the squared amplitude of the coefficients $c_i(t)$ allows us to determine the
probability that the system will be found in states $|a\rangle$ or $|b\rangle$ at time t. For a

large ensemble of two-level atoms, this is equivalently the normalised number of atoms in each state after a projective measurement onto the $\{|a\rangle, |b\rangle\}$ basis. For the case of resonant coupling $(\Delta = 0)$ we have the well known result of 'Rabi flopping' between the two states:

$$|c_a(t)|^2 = \cos^2\frac{\Omega t}{2}; \quad |c_b(t)|^2 = \sin^2\frac{\Omega t}{2}. \tag{4.4}$$

This simple model accurately describes one of the most profound and ubiquitous properties of quantum systems, the oscillation of probability amplitude between coupled states.

Extending the description of a two-level system to the $F = 1$ hyperfine system, the coupling field specifically takes the form of an oscillating magnetic field,

$$H_{\text{rf}} = -\boldsymbol{\mu} \cdot \mathbf{B}_{\text{rf}} \cos(\omega_{\text{rf}} t), \tag{4.5}$$

with $\boldsymbol{\mu}$ the atomic magnetic dipole moment.

In Section 2.3 of Chapter 2 we discussed how the hyperfine and Zeeman interaction give rise to hyperfine ground states denoted by total angular momentum F with associated Zeeman sublevels m_F in alkali metal atoms (Figure 2.2). The Hamiltonian H_{rf} couples different hyperfine levels or Zeeman sublevels, depending on the magnetic field-induced Zeeman splitting determined by Eq. 2.8. Coupling two Zeeman states in different hyperfine levels with microwaves can be treated as a spin-1/2 two-level system, but radiofrequency resonant with the Zeeman splitting *within* the hyperfine manifold drives transitions between all three Zeeman sublevels due to the near degeneracy of the Zeeman splitting: for low fields the quadratic Zeeman shift is much smaller than the linear Zeeman shift and in most cases $\Omega_{\text{rf}} \gg q$. The interaction Hamiltonian of $F = 1$ magnetic dipole coupling within the RWA, in analogy with Eq. 4.3 is

$$H_{\text{int}} = \hbar \begin{pmatrix} -\Delta & \frac{1}{\sqrt{2}}\Omega_{\text{rf}} & 0 \\ \frac{1}{\sqrt{2}}\Omega_{\text{rf}} & q & \frac{1}{\sqrt{2}}\Omega_{\text{rf}} \\ 0 & \frac{1}{\sqrt{2}}\Omega_{\text{rf}} & \Delta \end{pmatrix}, \tag{4.6}$$

where we have identified $\Omega_{\text{rf}} = \mu_B g_F B_{\text{rf}}/2\hbar$ as the rf Rabi frequency [80]. We can compute the probability amplitudes directly by acting on a state[1] $\boldsymbol{\psi} = (1, 0, 0)^T$ with the operator $\exp(-iH_{\text{int}}t/\hbar)$, and computing the population amplitudes. Varying the time traces out characteristic Rabi oscillations amongst the three Zeeman sublevels.

---

[1]In this thesis, when writing spinors in-line the order of magnetic sublevels is $\boldsymbol{\psi} = (\psi_{-1}, \psi_0, \psi_{+1})$.
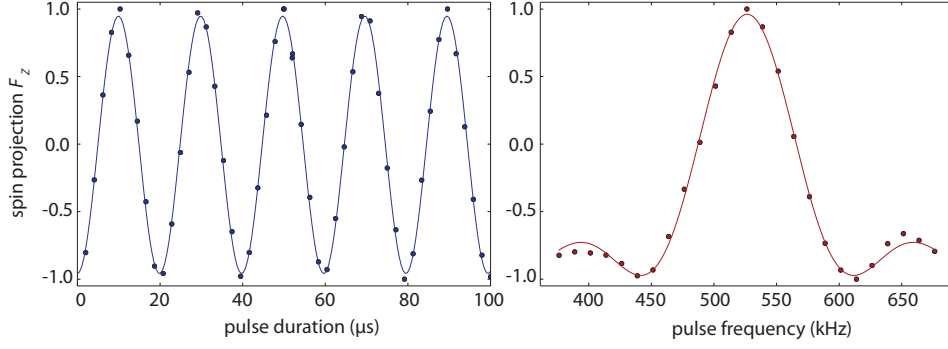
Figure 4.1: Left: Time domain Rabi flopping of the pseudospin F_z, as a function of the applied rf pulse duration. Each point corresponds to a run of the experiment, with F_z computed from absorption images of Stern-Gerlach separated clouds (Section 3.6). Right: Frequency domain 'Rabi spectrum', centred at the resonant transition frequency for a magnetic field of 752 mG.

It is often convenient to consider the Rabi oscillations of the expectation value $F_z = \langle \hat{F}_z \rangle$ rather than individual states. We thus define

$$F_z \quad = \quad \sum m_F |c_{m_F}(t)|^2 \tag{4.7}$$

$$= \quad \frac{\sum m_F N_{m_F}}{\sum N_{m_F}} \tag{4.8}$$

where the lower line corresponds to a measurement on a large ensemble of atoms, with N_{m_F} the total number of atoms in each spin state. We can now look at the behaviour of the spin projection as the detuning is varied for constant pulse duration (Figure 4.1, right). Such a *Rabi spectrum* is an indispensable tool for measuring magnetic fields: the centre frequency gives a direct measurement of the magnetic field via the Larmor frequency $\omega_L = \gamma B$. For $F = 1$ ^{87}Rb, γ is related to the linear Zeeman shift, Equation 2.10a by $\gamma = |p|/B$.

As shown in Figure 4.1, F_z oscillates in the presence of a resonant driving field at the Rabi frequency. Off-resonant Rabi oscillations are faster, and at lower amplitude, with the oscillation frequency given by the generalised Rabi frequency, $\Omega = \sqrt{\Omega_{\text{rf}}^2 + \Delta^2}$. For the sake of brevity, we will hereafter denote the resonant Rabi frequency $\Omega(\Delta \to 0)$ simply by Ω.

Figure 4.2 shows the spin projection as a function of detuning and pulse duration, with the experimentally measured plots corresponding to the resonant case, *i.e* lines bisecting the centre of each axis. We can create a superposition of states with spin projection F_z between -1 and 1 by varying the pulse duration of a resonant pulse, for which we have two special cases that are particularly important:
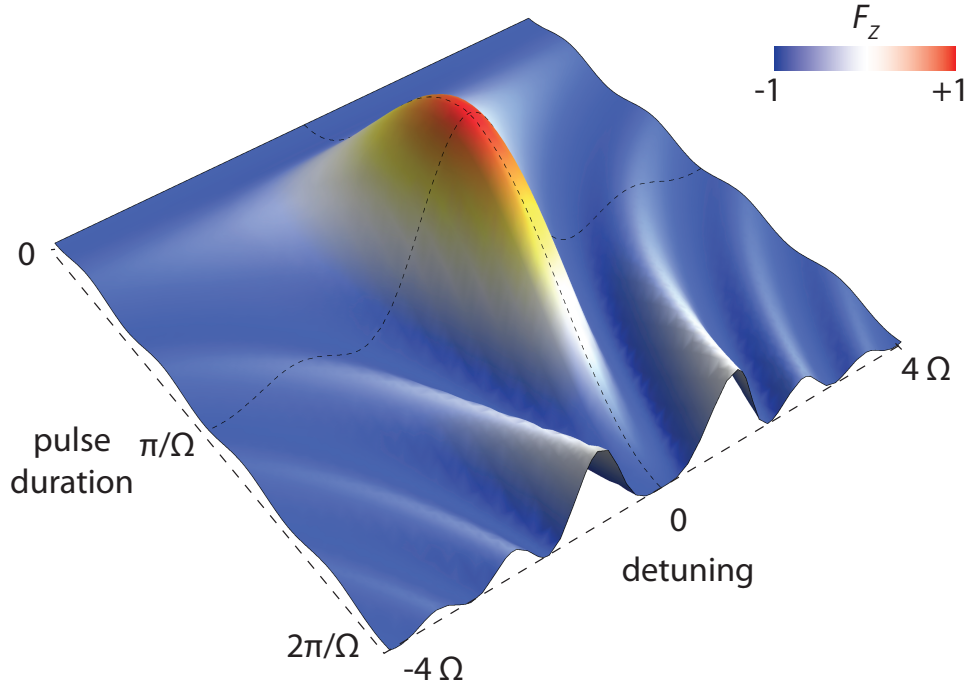
Figure 4.2: The 'Rabi landscape' from applying an rf pulse of variable detuning and duration to a state with $F_z = -1$. Resonant time oscillations (Figure 4.1, left) and a frequency spectrum with a pulse duration π/Ω (Figure 4.1, right) are overlaid. Only a resonant π-pulse results in full transfer to the $m_F = +1$ state.

- $\pi/2$-pulse: A resonant pulse with duration $\pi/2\Omega$ that creates a state with $F_z = 0$. After acting on a large ensemble of atoms initially all in state $m_F = -1$ with a $\pi/2$-pulse, a projective measurement of a spin-1 system would find the population distributed amongst the Zeeman states with $(N_{-1}, N_0, N_{+1})/N = (1/4, 1/2, 1/4)$. This corresponds to a $90°$ rotation of the spin vector on the Bloch sphere, as discussed below.

- π-pulse: a resonant pulse twice as long as a $\pi/2$-pulse that is equivalent to a $180°$ spin rotation. It maps F_z to $-F_z$, so for an atomic sample initially all in $m_F = -1$, a π-pulse transfers all population to $m_F = +1$.

A useful tool for understanding these spin rotations is the *Bloch sphere*. The reader is referred to any of the references listed in this Section, particularly [80] for a more thorough treatment. We can represent a two-level quantum system in the absence of decoherence by the unit-normalised state vector

$$|\psi\rangle = \cos(\theta/2)|a\rangle + e^{i\phi}\sin(\theta/2)|b\rangle, \tag{4.9}$$
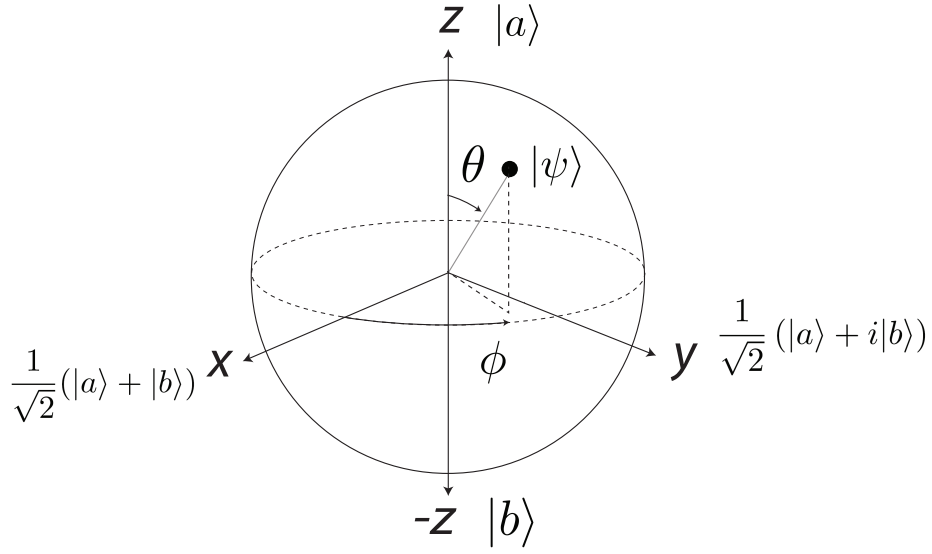
Figure 4.3: The Bloch sphere provides a very satisfying depiction of spin rotation pulses. A general spin-1/2 state $|\psi\rangle = \cos(\theta/2)|a\rangle + e^{i\phi}\sin(\theta/2)|b\rangle$ can be represented (in the absence of decoherence) by a vector of unit length with polar angle θ and azimuthal angle ϕ. The poles of the Bloch sphere correspond to pure states, the angle θ the degree of state mixture and ϕ the relative phase.

where $\theta \in [0, \pi]$ is an angle representing the degree of population admixture and $\phi \in [0, 2\pi]$ the relative phase. Such a state can be diagrammatically represented as vector on a sphere with unit radius, azimuthal coordinate ϕ and polar angle θ, as shown in Figure 4.3. The poles of the Bloch sphere correspond to pure states, and the equator an equal superposition with some relative phase. Although the Bloch sphere is a fundamentally two-level construct, many of the state preparation and interferometric procedures we perform on a spin-1 system can be understood in an analogous manner. For example, a spin-1 state vector can be represented as two spin-1/2 state vectors on the Bloch sphere, in the *Majorana representation* (Ref. [23] and references therein). However, for understanding spin rotation pulses and Ramsey interferometry in a spin-1 system, the conceptual picture provided by a single spin vector on the Bloch sphere is very illuminating and encapsulates much of the essential physics.

4.2   Ramsey interferometry

The principal magnetic resonance technique we use for precision measurement is Ramsey interferometry. Originally devised as a way of measuring
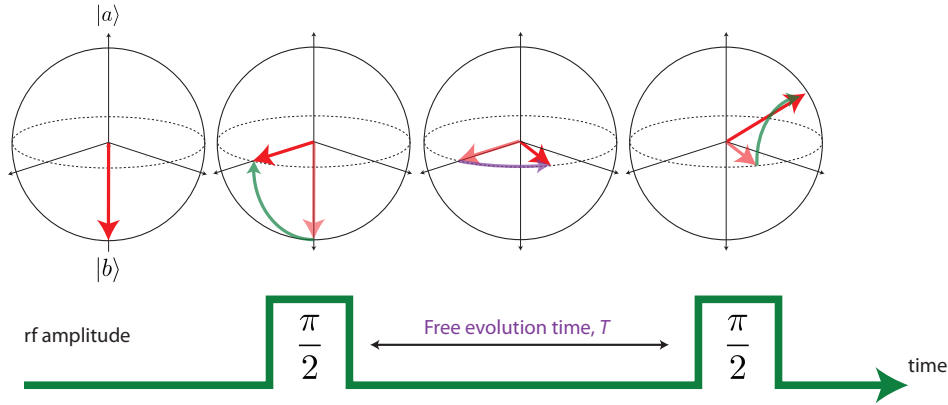
Figure 4.4: Ramsey interferometry explained on the Bloch sphere. A $\pi/2$-pulse is applied to the pure state $|\psi\rangle = |b\rangle$. In the frame rotating at the pulse frequency ω, the state vector $|\psi'\rangle$ rotates around the equator at the detuning Δ. After the free evolution time T, a second $\pi/2$-pulse is applied, which converts the relative phase into relative population.

nuclear magnetic moments in a molecular beam using two oscillating fields at different points along the beam [126], Ramsey interferometry has evolved into a widely applicable metrological technique: atomic clocks, which define the SI second, use Ramsey interferometry to attain extremely precise measurements of atomic transition frequencies [127].

We can utilise our simple model of two-level systems developed in Section 4.1.1 to explain how Ramsey interferometry works. The original method called for 'separated oscillatory fields', in the sense that a molecule in an atomic beam would briefly sample each a region of oscillating field during transit along the apparatus. Our fields are separated not in space, but in time.

The interferometry sequence begins by acting on an atomic sample in an eigenstate of the bare atomic Hamiltonian with a $\pi/2$-pulse. The $\pi/2$-pulse puts the ensemble of atoms into an equal superposition of states, after which the sample is then left to freely evolve. In this free evolution time, which may be called the 'interrogation time' or 'Ramsey time', the two states develop a relative phase; if the two energy states have opposite sign magnetic moments, the phase difference is due to the different linear Zeeman shifts experienced. A second, phase-coherent $\pi/2$-pulse is then applied, which converts the phase difference into population difference. Schematically this is represented on the Bloch sphere in Figure 4.4.

The action of a $\pi/2$-pulse on the state $|\psi_i\rangle = (0, 1)^{\mathrm{T}}$ can be described by applying the operator $U = e^{-iH_{\mathrm{int}}t/\hbar}$ for $t = \pi/2\Omega$, with H_{int} for the spin-1/2 case defined in Eq. 4.3. The $\pi/2$-pulse operator is then $U_{\pi/2} = U(t =$

$\pi/2\Omega, \Delta$). Rabi broadening of the $\pi/2$-pulse when $\Omega \gg \Delta$ allows us to neglect the detuning *during* each pulse: the detuning of the centre frequency of the pulse from resonance is what is measured in a Ramsey experiment. The Ramsey interferometry sequence can thus be written succinctly as

$$|\psi_f\rangle = U_{\pi/2}\, U_{\text{free}}\, U_{\pi/2}|\psi_i\rangle, \qquad (4.10)$$

with $U_{\text{free}} = U(t = T, \Omega \to 0)$ the free evolution operator over the interrogation time T. The final $\pi/2$-pulse converts the phase acquired in the free evolution time to a population difference: we then compute $|\langle\psi_f|\psi_f\rangle|^2$ and find for the probability amplitudes c_a, c_b and spin projection $F_z = \langle\hat{F}_z\rangle = |c_a|^2 - |c_b|^2$

$$|c_a|^2 = \sin^2(T\Delta/2); \quad |c_b|^2 = \cos^2(T\Delta/2); \quad F_z = \cos(T\Delta). \qquad (4.11)$$

These characteristic oscillations are the output from the interferometer: *Ramsey fringes*. The measurement utility is immediately apparent: the period of the Ramsey fringes corresponds to the detuning of the pulse from resonance. If the oscillator frequency (ω_{rf}) is well known the resonance frequency can then be determined with extremely high precision from the slow Ramsey fringes.

Time-domain fringes may be acquired by increasing the Ramsey time, which increases the precision of a detuning measurement. A complimentary technique is to vary the phase of the second pulse, which then traces out Ramsey fringes as a function of pulse phase: this is equivalent to changing azimuthal direction of the final Bloch vector rotation depicted in Figure 4.4.

For the case of a spin-1 system, the operators in Eq. 4.10 are given by $U = e^{-iH_{\text{int}}t/\hbar}$, with H_{int} now the spin-1 interaction Hamiltonian, Eq. 4.6. Evaluating Eq. 4.10, we find

$$F_z = \cos(qT)\cos(T\Delta), \qquad (4.12)$$

i.e. Ramsey fringes at the detuning from resonance Δ as per the spin-1/2 case, amplitude modulated at the quadratic Zeeman frequency. The physical interpretation of this can be garnered by considering two spin-$\frac{1}{2}$ interferometers, formed from the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 1, m_F = 0\rangle$ and $|F = 1, m_F = 0\rangle \leftrightarrow |F = 1, m_F = +1\rangle$ spin-1/2 transitions being simultaneously addressed in the Ramsey sequence. The quadratic shift gives rise to a slightly different detuning for each interferometer; the Ramsey fringes from each interferometer beat together at the angular frequency q.

Having now elucidated the physical basis of how oscillating magnetic fields couple quantum states in alkali atoms, we turn now to the experimental realisation.

4.3 Radiofrequency apparatus

We use separate DDSs to produce low-amplitude (typically less than 5 dBm) rf signals in the range 50 kHz to 2 MHz for state preparation (RFBlaster) and 3 to 25 MHz for rf evaporation (PulseBlaster). For rf evaporation, the PulseBlaster drives a 40 dB Delta RF amplifier (Delta RF LA10-1-525-40) that produces 10 W of rf power. This is then connected to a single-loop antenna oriented in the *x-z* plane of the apparatus, concentric with the quadrupole coils. Antennas must be designed for maximum transmitted power in a given frequency band. As such, the single-turn loop antenna is well suited to the range of rf frequencies used in rf evaporation, but is unsuitable for attaining high Rabi frequencies[2] for low field state preparation, where the requisite frequencies are of order 0.2 to 1 MHz.

Our low field state preparation setup consists of an RFblaster driving a high power 30 W amplifier (MiniCircuits LZY-22-+), which is connected to one of two multi-turn coils oriented along the *y*- and *x*-axes. Depending on the magnetic bias field orientation[3] we switch between the two coils.[4] The coils are near identical and optimised for superior response at frequencies < 1 MHz: both are 60 mm in diameter and have 20 turns. The two coils are located different distances from the atoms, the 'side' coil, along the *x*-axis is approximately 100 mm from the atoms whereas the 'top' coil (along the *y*-axis) is ∼ 30 mm distant. We can achieve typical Rabi frequencies of 20-50 kHz with the side coil, and over 250 kHz with the top coil. Since the 'top' state preparation coil is concentric with the single-turn evaporation antenna, cross coupling between the two antennae results in shielding currents that disrupt evaporation unless the high power amplifier is gated during rf evaporation.

4.3.1 Measuring magnetic fields

Radiofrequency is the workhorse state preparation tool in our experiments. One of the most common experiments we perform is to measure the magnetic field $|B|$, which determines the quadratic shift for a spinor BEC. Accurate knowledge of the magnetic field is necessary in order to optimise the fidelity of state preparation pulses. Resonant rf transitions occur when $\gamma B = \omega_{\mathrm{rf}}$; varying the rf frequency for constant pulse duration and magnetic field traces out a *Rabi spectrum*, as described in Section 4.1.1. The symmetrical

---

[2]The Rabi frequency should ideally be high enough to confer the benefits of power broadening on a driven transition, making state preparation relatively insensitive to magnetic field noise. Typically $\Omega \sim 2\pi \times 10$ kHz is required in our applications.

[3]As there are no linearly polarised rf transitions, rf parallel to the magnetic bias field will not drive any transitions.

[4]The evolution of our rf state preparation hardware was largely driven by the results discussed in Section 5.3.

centre point of the Rabi spectrum defines the resonant centre frequency, *i.e.* the Larmor frequency for a given magnetic field. If we then vary the pulse duration for rf fixed at the resonant centre frequency we trace out time domain Rabi oscillations. From these two measurements, shown in Figure 4.1, we have all the parameters required for a well characterised resonant state preparation pulse. Moreover, acquiring Rabi spectra at various applied external field currents allows the magnetic bias fields to be calibrated.

4.4   Microwaves

To drive magnetic dipole transitions between hyperfine states, which are split by 6.834 GHz, we require microwave photons. Unlike rf coupling within the hyperfine ground state, hyperfine transitions can be driven with both linear and circularly polarised microwave photons. Microwaves thus bring a substantial amount of state preparation controllability: using a combination of microwave and rf pulses and sweeps, any combination of fractional populations in either (or both) hyperfine manifolds can be created. In addition to state preparation, microwave fields are used to create effective quadratic Zeeman shifts (Section 4.5).

Fine control of microwave coupling translates to similarly high levels of state preparation control. To this end we have designed an built a flexible microwave setup with exquisite frequency, phase and amplitude control. Our scheme is centred around the concept of single-sideband quadrature amplitude modulation, which in contrast to traditional mixing of carrier and low-frequency rf signal results in only one dominant frequency component. Tuning the low-frequency rf signal thus controls a single-sideband output at the sum or difference frequency $f_{\text{carrier}} \pm f_{\text{rf}}$, allowing the precision of rf-band electronics to be translated to microwave frequency domains. This is important, as obtaining comparable levels of frequency control with a single commercial microwave source requires substantial investment.

In quadrature amplitude modulation, a high frequency local oscillator (LO) is mixed with two low frequency signals, called the *in-phase* (I) and *quadrature* (Q) signals, which differ in phase by $90°$. The output from a quadrature amplitude modulator, or IQ mixer, consists of a single sideband and suppressed carrier (typically up to 30 dB from the sideband).

Figure 4.5 schematically shows the quadrature amplitude modulation hardware we use for microwaves. We drive a Polyphase Microwave AM4080N quadrature modulator with two channels of an RFBlaster DDS and a local oscillator (GPS-locked PhaseMatrix FSW-0010). The LO frequency is set to 6934.682 MHz, and the DDS channels centred at 100 MHz. The relative phase
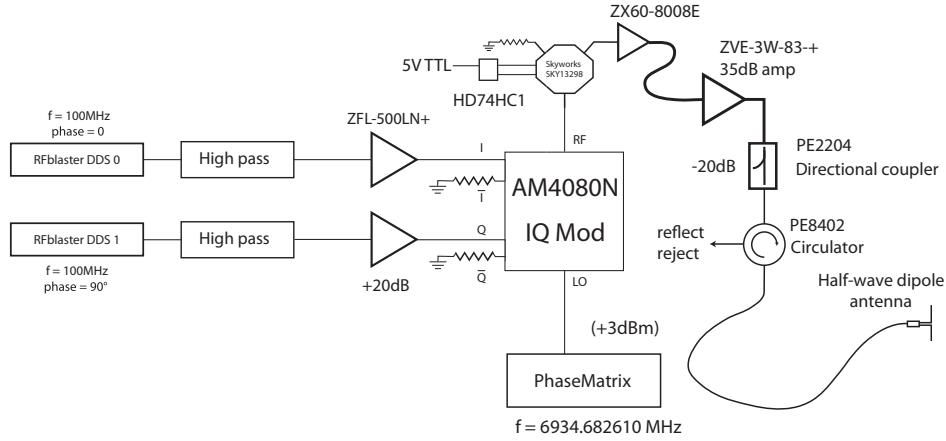
Figure 4.5: Microwave setup. A RFBlaster DDS drives a quadrature modulator with two 100 MHz rf signals 90° out of phase, which imparts single-sideband modulation onto the microwave frequency carrier (6.934 GHz) at the difference frequency (6.834 GHz). The microwave signal is then amplified and applied to the atoms using a half-wave dipole antenna.

of the two IQ signals is adjusted to utilise the lower sideband, we observe carrier suppression by 20 dB and up to 40 dB upper sideband suppression.[5] Changing the frequency of the DDS outputs to $f_{DDS} + \delta f$ changes the output frequency to $f_{LO} - f_{DDS} - \delta f$. The RFBlaster outputs are essentially always phase coherent, so sweeping the DDS frequency allows the microwave frequency to be swept.

The output from the IQ mixer is preamplified before being fed into a high power (3 W) 35 dB amplifier (MiniCircuits ZVE-3W-83+). The output from the amplifier is used to drive a single half-wave dipole antenna. Between the amplifier and antenna are a directional coupler (Pasternack PE2204), to monitor output power and a circulator (Pasternack PE8402), which prevents back reflections from the antenna returning to the amplifier. The dipole antenna is a simple half-wave dipole antenna formed from an SMA connector and non-magnetic coaxial cable core, with the long axis approximately 2.4 cm – this is optimised by using the circulator to measure back reflected power with a constant power 6.834 GHz signal input to the amplifier. When the dimensions of the antenna are optimal for 6.834 GHz radiation to be broadcast, the reflected power is minimised. Our optimised antenna reflects about 1% of input power back.

Maximising the input power and broadcast power are essential for a dipole antenna due to the poor directionality – the microwave field is radiated

---

[5]The suppression *should* be up to 35 dB, but damaged SMA pins on the LO input lead to increased carrier leakage.

near isotropically. Since our target, the BEC, is suspended in an entirely glass vacuum chamber, there is not steel tubing to act as a convenient reflector/waveguide and so the dipole antenna is positioned as close as possible to the cell. The output from the antenna is primarily linearly polarised, but inevitable defects associated with such a simplistic design result in a preference for either right or left circularly polarised radiation, as can be seen by measuring on-resonant Rabi frequencies for $\sigma^+\sigma^-$ microwave transitions and accounting for the transition strength factor. This becomes an important factor when we consider the contribution of these transitions to microwave-induced quadratic Zeeman shifts, discussed in Section 4.5.

4.4.1   AC line synchronisation

Parasitic time-varying magnetic fields are also present in the lab. These fields are predominantly at 50 Hz and odd harmonics which originate from lab equipment power supplies. Although rudimentary measures may be taken to minimise these fields, such as maximising the distance of the noise source from the atoms, there is typically a substantial component that remains, which contributes as background magnetic noise. Since the experimental duty cycle is $\gg$ AC line period and subject to minor variations shot-to-shot, a state preparation pulse samples a variable phase of AC magnetic noise every run, which we measure as shot-to-shot noise. For microwave transitions, where the maximum attainable Rabi frequencies are only $\sim$ 5 kHz due to drive power and distance from the atoms, AC line noise is particularly detrimental to the fidelity of mw state preparation pulses.

Like the stray DC fields, cancellation coils may be added that create opposing AC fields [128]. However, a simpler solution is to ensure the experiment performs magnetically sensitive steps at well-defined phases of the AC line, which can be assumed to be fairly stable over 100 ms. We can synchronise the experiment to a zero-crossing of the AC power line using a straightforward circuit based on an lm339 comparator [129] and a monostable multivibrator ('one-shot', 74hc423) [130], as shown in Figure 4.6.

In this circuit, the comparator produces an edge whenever the 50 Hz AC input (18 V rms from an AC power supply) crosses zero. Without the additional one-shot board, the free-running comparator circuit produces an output rising or falling edge at the respective zero-crossing of the AC line. The one-shot allows for integration into the control system. As shown in Figure 4.6, when the line-sync sequence is initiated, a TTL pulse from the control hardware arms the one-shot, which is configured so that it outputs a falling edge upon receiving a falling edge from the comparator circuit.
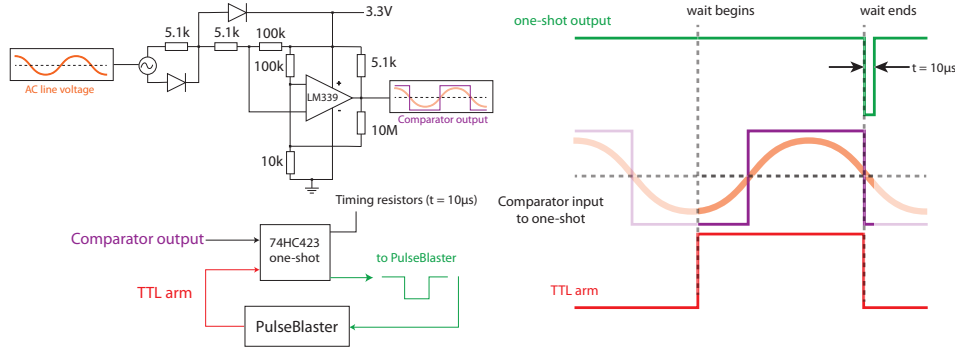
Figure 4.6: Schematic of the AC line synchronisation circuit (left) and operation (right). The comparator circuit, based on the LM339 comparator, outputs rising and falling edges on the zero-crossings of the 50 Hz AC mains voltage. The comparator output is directed into a one-shot circuit, using the 74HC423 metastable monovibrator. When armed by the control system, the one-shot outputs a falling edge on the falling-edge zero-crossings of the AC line, allowing the experimental procedure to be phase synchronised with the same point of the AC line every run, improving repeatability of magnetically sensitive state preparation.

The control system then *waits*[6] for the time between activating the arming pulse and receiving the output edge from the one-shot. This output pulse is fed to the master clock (PulseBlaster, Section 3.7); upon receipt of this pulse the control hardware deactivates the arming input, and proceeds to execute subsequent stages of the experiment. Since the experiment only proceeds from the moment it receives an edge synchronous with the same part of the AC line every shot, the sequence of experimental procedure takes place at well-defined phases of the AC line magnetic noise profile.[7] The circuit is armed 85 ms before the state preparation pulse. Synchronising the experiment to the AC line effectively removes a source of shot-to-shot noise and improves repeatability of state preparation. We can directly quantify the AC noise by performing Rabi spectroscopy. We alter the detuning of a microwave π-pulse on the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition so that the spin projection is zero, creating a state maximally sensitive to magnetic field fluctuations (Figure 4.7, left). Around $F_z \approx 0$, the deviation of Fz due to small detuning perturbations from AC magnetic fields is approximately linear, allowing direct mapping of δF_z to B_{AC}. Varying the time between acquiring AC synchronisation and applying the π-pulse samples the time variation of the AC field, as shown in Figure 4.7, right. The fidelity of microwave

---

[6]All experimental procedure proceeds with initial settings for a variable time run-to-run.

[7]The measured jitter of the line-synchronisation pulse is on the order of 100 µs about the zero-crossing of the AC line.
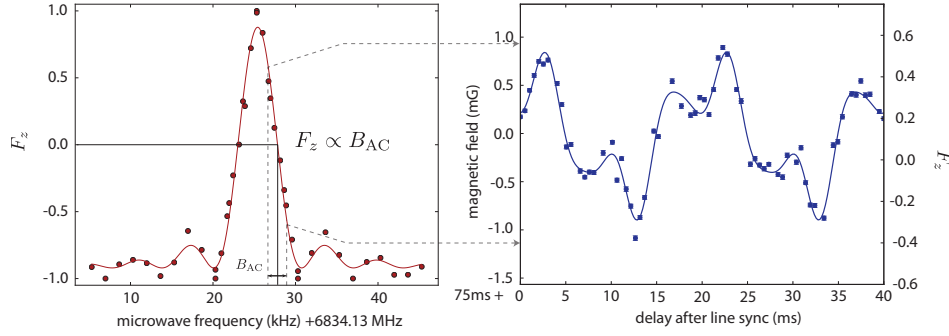
Figure 4.7: A measurement of the background AC field variation. Left: frequency-domain Rabi spectrum of the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition, with $\Omega_{mw} = 2\pi \times 2.8\,\text{kHz}$. Around $F_z \approx 0$, the deviation of Fz due to small detuning perturbations from AC magnetic fields is approximately linear, allowing direct mapping of δF_z to B_{AC}. We vary the time at which the state preparation pulse is performed after the experiment is line-synchronised and measure the variation in F_z due to the AC field magnetic field B_{AC} (right). The curve corresponds to a five-harmonic sinusoidal fit, yielding $B_{AC}(\text{rms}) = 0.925\,\text{mG}$.

and rf state preparation pulses is greatly improved with the AC line sync active. Microwave transitions with low Rabi frequencies ($< 3\,\text{kHz}$) are only able to be driven reproducibly with the line sync active (Figure 4.7, left, where $\Omega_{mw} = 2\pi \times 2.8\,\text{kHz}$). Radiofrequency transitions with higher Rabi frequencies also benefit; the standard deviation of F_z from 15 repeated rf $\pi/2$-pulses is found to be typically $\delta F_z = 0.007$ and as low as $\delta F_z = 0.004$. The former corresponds to a value ~ 3 times that of the spin-projection noise $(\delta F_z)_{SQL} = 1/\sqrt{2N}$ for typical BEC atom numbers of $N \approx 1 \times 10^5$ spin-1 atoms.

4.5 Microwave-induced quadratic shift

In addition to state preparation, off-resonant microwaves can also be used to impart an effective quadratic Zeeman shift, and as such are valuable for investigating spinor physics. This effective shift can be negative depending on detuning and can be thus used to reduce the overall shift $q_{tot} = q(B) + q_{mw}(\Omega_{mw}, \Delta_{mw})$ to low values at large background magnetic fields of 1 Gauss or more. This is in contrast to actually reducing the magnetic field to $< 100\,\text{mG}$, where AC field noise would otherwise cause spurious spin-flip transitions [131]. A microwave-induced quadratic shift can also be used to achieve $q < 0$, in contrast to $q > 0$ in Eq. 2.10b from magnetic fields only.

We can write the microwave induced quadratic shift from a microwave
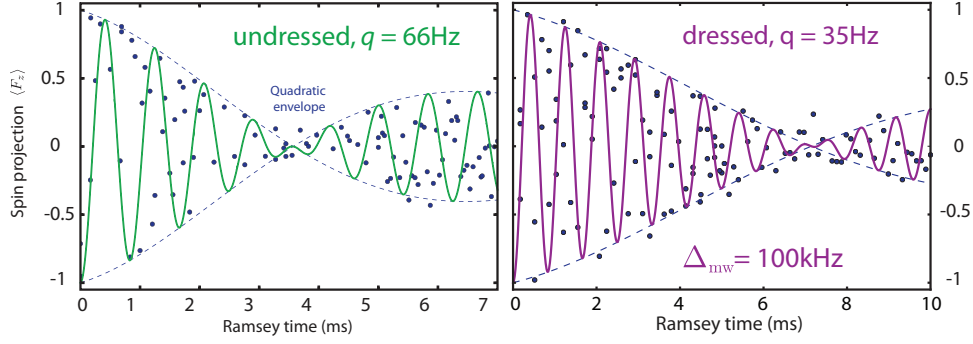
Figure 4.8: Tuning the quadratic Zeeman shift with microwaves. The contrast of Ramsey fringes oscillating at $\sim 1\,\mathrm{kHz}$ is amplitude modulated by the quadratic Zeeman shift, with an overall dephasing envelope. The fringe phase is quickly scrambled after $1\,\mathrm{ms}$. Applying a microwave field detuned $100\,\mathrm{kHz}$ from the clock transition imposes an effective quadratic Zeeman shift, which changes the position of the contrast nulls.

field detuned by Δ_mw from a transition with resonant Rabi frequency Ω_mw as:

$$q_\mathrm{mw} = -\frac{\Omega_\mathrm{mw}^2}{4\Delta_\mathrm{mw}}. \tag{4.13}$$

For $\Omega_\mathrm{mw} = 2\pi \times 5\,\mathrm{kHz}$ microwaves detuned $90\,\mathrm{kHz}$ from the $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ clock transition, the microwave induced shift approximately cancels the quadratic shift $q(B)$ of the state $|F = 1, m_F = 0\rangle$ from a $1\,\mathrm{G}$ bias field.

We developed a method to precisely measure the induced quadratic shift using Ramsey interferometry. As described in Section 4.2, Ramsey fringes from a spin-1 interferometer initiated with a $\pi/2$-pulse with detuning Δ take a simple form for short interrogation times, Eq. 4.12. The faster fringes, oscillating at Δ are amplitude modulated by the overall quadratic shift. We can measure the quadratic shift by determining the position of nulls or maxima in the fringe contrast as a function of microwave detuning. Figure 4.8 shows such a measurement, where the position of the first contrast null in time-domain Ramsey fringes has been pushed out to $\sim 7\,\mathrm{ms}$ from $3\,\mathrm{ms}$ by applying microwaves blue-detuned $72\,\mathrm{kHz}$ from the clock transition. The absolute phase of the fast Ramsey fringes is quickly scrambled due to AC field fluctuations[8] and shot-to-shot drifts of the magnetic bias field strength, but the amplitude modulation contrast nulls are relatively unaffected. A dephasing envelope due to magnetic field gradients damps the overall interference contrast (discussed further in Section 5.3.3).

---

[8]These experiments pre-date the implementation of the line synchronisation circuit.

We can refine the measurement in several ways. First, we perform phase domain Ramsey interferometry by varying the phase of the second $\pi/2$-pulse: this allows phase-domain Ramsey fringes to be acquired at fixed evolution time. We then measure the contrast, or fringe visibility, by determining the maximum and minimum values from 20 runs of the experiment. The contrast of Ramsey fringes vanishes past about 15 ms Ramsey time due to the overall dephasing envelope. We can access longer interrogation times by interposing a single π-pulse in the interferometry sequence, which mitigates the gradient induced dephasing. There is no loss of *contrast* information by performing a spin-echo pulse, as shown in Figure 4.9, the phase of the fast fringes is random and essentially irrelevant, with the contrast the only quantity of interest.

We can then precisely determine the microwave detuning that nulls the quadratic shift by measuring the fringe contrast for different detunings at a fixed Ramsey time. When $q_{tot} = 0$, the amplitude modulation disappears. Fixing the Ramsey time and varying the total quadratic shift would thus result in maximum fringe contrast when $q_{tot} = 0$. The results of exactly such a measurement are shown in Figure 4.9, at $T = 18.5$ ms. Each data point corresponds to the fringe contrast determined from 20 iterations of the experiment sampling fringe phase fluctuations from magnetic noise. The peak of the curve determines the microwave detuning where $q_{tot} = 0$.

The overlaid red curve in Figure 4.9 corresponds to a fit of the fringe contrast, $|\cos(q_{tot}T)|$, where $T = 18.5$ ms and q_{tot} the total microwave-induced shift:

$$q_{tot} = q(B) + \sum_i -\alpha_i \frac{\Omega_i^2}{4\delta_i}, \tag{4.14}$$

where the sum is taken over the three possible transitions from the state $|1,0\rangle = |F = 1, m_F = 0\rangle$ due to off-resonant dressing (at detuning δ_i) of the nearby $\sigma^\pm$ transitions, and α_i parametrises the relevant transition strengths (Clebsch-Gordan coefficients). The model contains free parameters for the Rabi frequencies for the $|1,0\rangle \leftrightarrow |2,0\rangle$ clock transition as well as off-resonant dressing from the $|1,0\rangle \leftrightarrow |2,\pm1\rangle$ transitions due to ellipticity of the microwave polarisation; these fitted parameters are close to the experimentally measured Rabi frequencies.[9]

As well as elimination of the quadratic shift, we can equivalently use this method to determine a precise calibration of induced quadratic shift. Such a measurement may be of interest to experiments seeking to precisely access the spinor condensate phase diagram [132]. We also point out that the

---

[9]Accurate measurement of $\sigma^\pm$ Rabi frequencies without the line sync in these experiments was difficult, with only the clock transition Rabi frequency able to be measured with any degree of confidence.

technique of differential Ramsey interferometry, expounded in Chapter 6 of this thesis, could be used in a similar manner to measure small differences in quadratic shifts from microwaves, or even optically induced effective shifts. The reader is also directed to the PhD thesis of Martijn Jasperse [95], where related measurements show how the amplitude of a Faraday rotation signal depends on the detuning of applied off-resonant microwaves.

4.6 Chapter summary

In this Chapter we have introduced the relevant theory for understanding radiofrequency and microwave state preparation and manipulation of a spinor BEC. We have discussed the state preparation apparatus, including the techniques used to measure magnetic fields. Ramsey interferometry was also introduced, which underlies several results in Chapter 5 and forms the basis of Chapters 6 and 7.
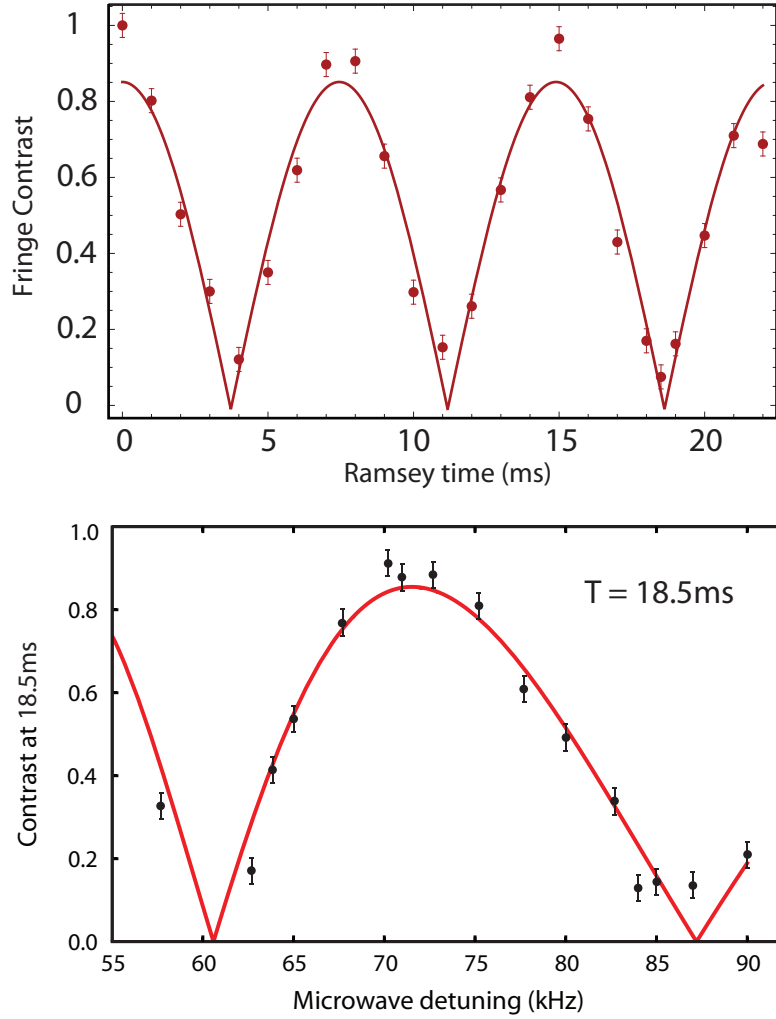
Figure 4.9: Top: Ramsey fringe contrasts from 20 experimental shots at free evolution times up to 25 ms. Adding a spin-echo pulse halfway through the free evolution time reverses the effects of gradient dephasing and recovers the fringe contrast. The addition of the spin echo pulse does not effect the fringe contrast; and allows for higher precision through longer interrogation times. Bottom: High precision nulling of the quadratic shift. The interrogation time is now fixed and an off-resonant microwave field is applied. The contrast of fringes is maximum when the quadratic shift is minimised. Each point corresponds to a measurement of the contrast at 18.5 ms Ramsey time from 20 realisations of the Ramsey experiment, as in the plot on the left. Plotting the measured contrasts against microwave detuning allows us to determine the nulling detuning. The solid line corresponds to a model that includes off resonant contributions from dressing nearby states and $\sigma^{\pm}$ transitions (described in text).

Experimental studies of spin-mixing dynamics

<div style="text-align: right">5</div>

An initial aim of this thesis was to develop new methods for controlling collisional interactions in spinor condensates. Controlling the magnitude and sign of both the quadratic Zeeman shift and the spin collisional interaction allows full access to the phase diagram (Figure 2.3) of a magnetic superfluid. A plethora of interesting experiments await the realisation of this aim: phase transitions between ferromagnetic and antiferromagnetic states, improved prospects for metrologically useful squeezing and the creation of junctions and interfaces between magnetic superfluids of differing nature.

Controlling spinor collisions amounts to experimentally manipulating the collisional energetics: the quadratic shift q and the collisional strength c. A range of techniques exist to control the quadratic shift, such as magnetic bias fields, microwave dressing (as discussed in Chapter 4) as well as more exotic schemes using optical fields that are yet to be demonstrated experimentally [133]. While there are difficulties associated with controlling q with lasers or microwaves (such as losses to other hyperfine levels or photon scattering), they pale in comparison to the difficulties associated with controlling c – which is a fundamental property of the scattering wavefunction.

Controlling c involves changing the scattering properties of a colliding atom pair, which requires the use of scattering resonances in one form or another. While tremendously successful in changing the scattering properties of single-component BECs and degenerate Fermi gases [134], magnetic Feshbach resonances are inflexible, occurring at high magnetic fields that preclude use in spinor BEC experiments, since they result in $q \gg c$. Schemes based on photoassocation (optical Feshbach resonances [135–137]) or microwave coupling of free scattering atoms to bound states (microwave Feshbach resonances [138]) are in theory are highly applicable to spinor condensates, although optical Feshbach resonances have been given less attention since the experiments of Ref. [29].

In order to study and quantify any scheme that purports to alter c or q, we require an atomic metric to diagnose its effect. As discussed in Chapter

4, we developed a scheme for measuring and controlling the quadratic shift imposed by off-resonant microwave dressing. To study schemes that change c, we look at the period and amplitude of spin-mixing oscillations, which provide a direct means of determining the magnitude and sign of c. We require a well characterised spinor condensate with a long coherence time to study population oscillations, which in turn requires a well characterised magnetic environment.

Ultimately this thesis did not explore changing c experimentally. Instead, this Chapter details our experimental efforts to realise spin-mixing oscillations, and to understand why our observations differed from the expected behaviour, primarily because of magnetic field gradients and vector light shifts. Chapters 6 and 7 detail our efforts to measure and minimise these systematics. In this Chapter, we also discuss a scheme we developed that uses spin-echoes to eliminate the effects of gradient-induced dephasing.

5.1 Introduction

Spin-mixing oscillations are initiated by preparing an out of equilibrium spin superposition. In the simplest case, we consider the dynamics of the state created by applying a $\pi/2$-pulse to a single-component $m_F = -1$ condensate. The resulting fractional population distribution is $(\rho_{-1}, \rho_0, \rho_{+1}) = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$, i.e. $\rho_0 = 0.5$, $F_z = 0$.

There are several reasons why this initial superposition is chosen. Following on from the reasoning in [139], we summarise the two salient benefits for us as

- Simple, highly reproducible state initialisation using rf pulses applied to a single-component BEC, that produce a state with $\theta = 0$;

- The single $\pi/2$ initialisation pulse is the first step in other magnetic resonance experiments, such as radiative spin-echo and Ramsey interferometry.

In an archetypical spin mixing experiment, we apply a single $\pi/2$-pulse to a pure $m_F = -1$ BEC at some magnetic bias field, with or without a microwave dressing field (Section 4.5) to attain lower (or even negative) values of q. We then vary the evolution time after the state preparation pulse in subsequent iterations ('shots') of the experiment. In each shot we use Stern-Gerlach absorption imaging to measure the population in each magnetic sublevel, and compute the normalised fractional populations $\rho_i = N_i/N$. For the bulk of our spin mixing experiments we used the elongated trap described in Section 3.9.4. It was not until we developed the elongated trap
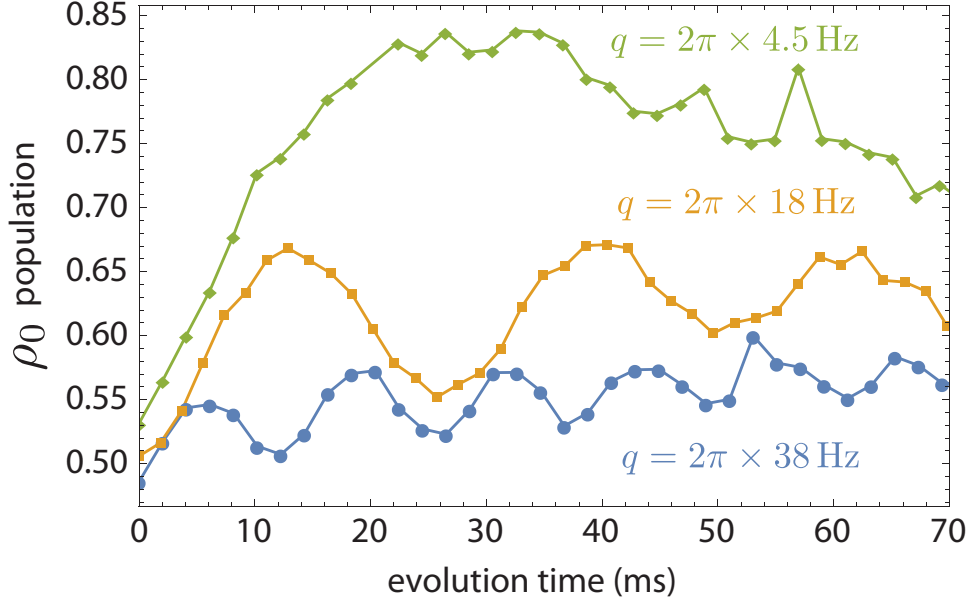
Figure 5.1: Spin mixing oscillations at different quadratic shifts, from bottom: $q/2\pi = 38\,\text{Hz}$, from a $B_x = 727\,\text{mG}$ bias field; $q/2\pi = 18\,\text{Hz}$ achieved by applying an off-resonant microwave dressing field (described in Chapter 4) with the same 727 mG bias field; and $q/2\pi = 4.5\,\text{Hz}$ reached by ramping the field from 727 mG to 251 mG immediately after the state preparation pulse. All results exhibit damped population oscillations after $\sim 50\,\text{ms}$, as well as a general trend towards higher ρ_0 population.

that we saw any evidence of spin mixing, which appeared to be totally absent in the semi-isotropic trap and were instead dominated by gradient-driven component separation. The results for several spin-mixing experiments at different quadratic shifts are shown in Figure 5.1.

Two observations that deviate from theoretical expectation are immediately apparent. Firstly, the amplitude of spin mixing oscillations is damped in all cases, and exhibit a general drift towards increasing ρ_0. Both effects – but most notably the latter – are present at all magnetic field strengths and when shots have been shuffled, *i.e.* spin-mixing times are randomly sampled are not increased sequentially with each run (to account for drift of atom number or trapping beam optical power). We also observe component separation: spatial separation of the constituent spin components of the condensates. We observe the centre-of-mass motion of the spin components after time of flight, which gives an indication of the component separation occurring in-trap. Figure 5.2 shows the centre-of-mass positions of the magnetically sensitive components after time-of-flight, normalised to the coordinates of the $m_F = 0$ component, as well as the spin-mixing oscillations at 727 mG. Unlike the semi-isotropic trap, where initial experiments showed strong component
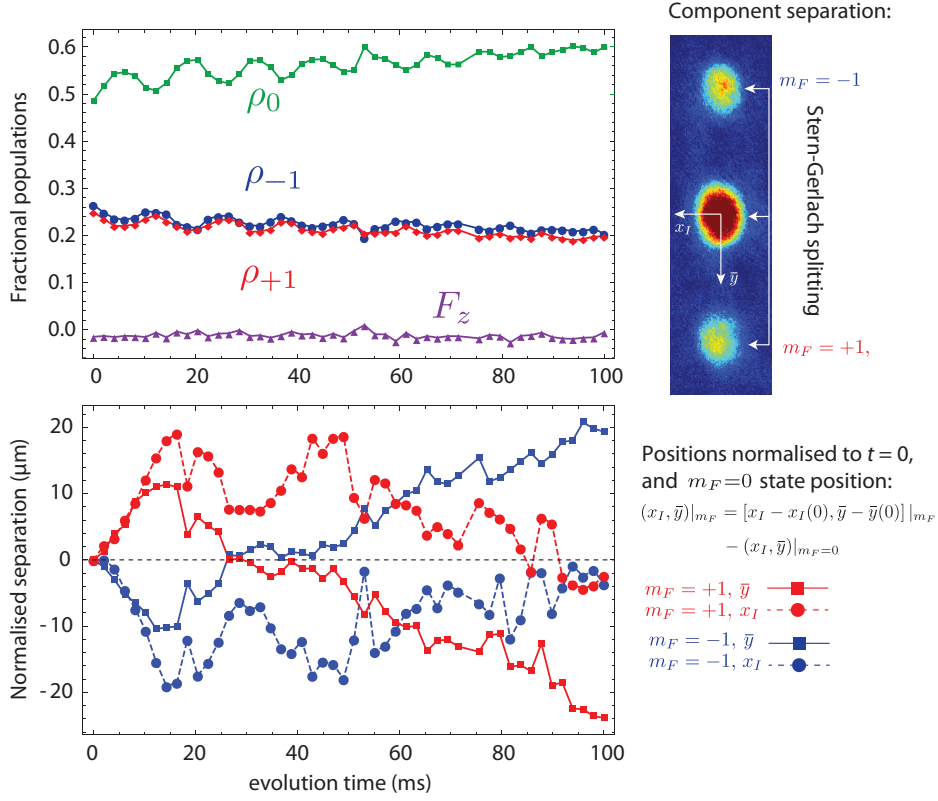
Figure 5.2: Spin mixing oscillations and spin component separation at $B_x = 727\,\mathrm{mG}$ ($q/2\pi = 38\,\mathrm{Hz}$). We define the separation directions as $\bar{y}$, along gravity and x_I, perpendicular to the imaging plane (*i.e.* a combination of x and z), and normalise the fitted centroid position to the initial position at $t = 0$ as well as the position of the $m_F = 0$ component. This ensures that we detect only relative, state dependent motion. Such m_F-dependent behaviour is indicative of magnetic gradient-driven separation. The normalised separation is after $23\,\mathrm{ms}$ time-of-flight, and not the in-trap separation.

separation into distinct spin domains, there remains some overlap between all components in the elongated trap.

The observation of component separation is indicative of the presence of magnetic field gradients, which account for the damping of spin-mixing oscillations and possibly for the observed ρ_0 drift as well. Following Ref. [76], in the Thomas-Fermi approximation (Section 2.2) the spin interaction strength can be determined using $c = c_2 \langle n(\mathbf{r}) \rangle = c_2 n_{\mathrm{ave}}$, with $n_{\mathrm{ave}} = 4 n_0 / 7 = 4\mu/7g$. Using the measured trap frequencies $\bar{\omega} = (\omega_{x'} \omega_y \omega_{z'})^{\frac{1}{3}}$ and atom number N, we compute the chemical potential in the Thomas-Fermi approximation

$$\mu = \frac{1}{2} \left(15 \hbar^2 m^{1/2} N \bar{\omega}^3 \bar{a} \right)^{\frac{2}{5}} \tag{5.1}$$

with $\bar{a} = (2 a_2 + a_0)/3$. For the elongated trap (Table 3.1), $n_{\mathrm{ave}} = 1.43(3) \times$

$10^{14}\,\mathrm{cm}^{-3}$ and so $c = -2\pi \times 5.2(1)\,\mathrm{Hz}$. This is consistent with what we observe experimentally, fits to the first period of oscillation suggesting $c \sim -2\pi \times 5\,\mathrm{Hz}$. Precise determination of c from experimental results is difficult due to the ρ_0 drift, which is significant even during the first period of oscillations. We can compute the corresponding spin healing length, $\zeta_S = 3.35\,\mu\mathrm{m}$ for the elongated trap (using Eq. 2.23), and we observe that since ζ_S is greater than the tight Thomas-Fermi radii (2.7μm and 3 μm), the system is effectively 1D along the elongation axis. A BEC in the semi-isotropic trap has $c = -2\pi \times 3.2\,\mathrm{Hz}$ and $\zeta_S = 4.3\,\mu\mathrm{m}$, and so is fully three-dimensional (since the smallest Thomas-Fermi radius is 6.5 μm). This is likely why we cannot see spin mixing in the semi-isotropic trap in the presence of gradients, as the spatially dependent spinor dynamics are more prevalent.

5.1.1 Dissipative relaxation?

There are a number of features of the observed ρ_0 drift that warrant investigation. Using the phenomenological dissipation model of Ref. [51], the equations of motion Eq. 2.31 are modified to include a phenomenological phase-damping term, β. With E given by the energy functional in Eq. 2.32, we have

$$\dot{\rho}_0 = -\frac{2}{\hbar}\frac{\partial E}{\partial \theta}, \quad \dot{\theta} = \frac{2}{\hbar}\frac{\partial E}{\partial \rho_0} + \beta\dot{\rho}_0. \tag{5.2}$$

The addition of the parameter β is equivalent to dissipative energy loss from the system,

$$\dot{E} = -\frac{\hbar}{2}\beta(\dot{\rho}_0)^2. \tag{5.3}$$

Trajectories in the (ρ_0, θ) phase space are altered with dissipation added. The spinor loses energy with time, and so the (ρ_0, θ) solutions to Eqs. 5.2 dissipatively relax to the lowest energy configuration, as shown in Figure 5.3. For ^{87}Rb with $q > q_{\mathrm{sep}}$ (with $q_{\mathrm{sep}} = q(B_{\mathrm{sep}})$ the quadratic shift corresponding to the separatrix magnetic field defined in Eq. 2.34), the ρ_0 population drifts upwards with evolution time. The addition of dissipation to the SMA model was first used by the Lett group to explain the long-term behaviour of spin mixing in sodium condensates [51]. Condensates prepared on the low-field (high energy in ^{23}Na) side of the separatrix were observed to lose enough energy to cross the separatrix after evolution times of between 0.1 to 1 s. The underlying cause for this dissipation was not determined, although thermal effects are highlighted as a possible cause [24, 140]. Ref. [51] did not report significant magnetic field gradients, or speculate if they were responsible for observed phenomena.

Several aspects of our observations are consistent with the predictions of the dissipative model for ^{87}Rb, such as the characteristic relaxation to
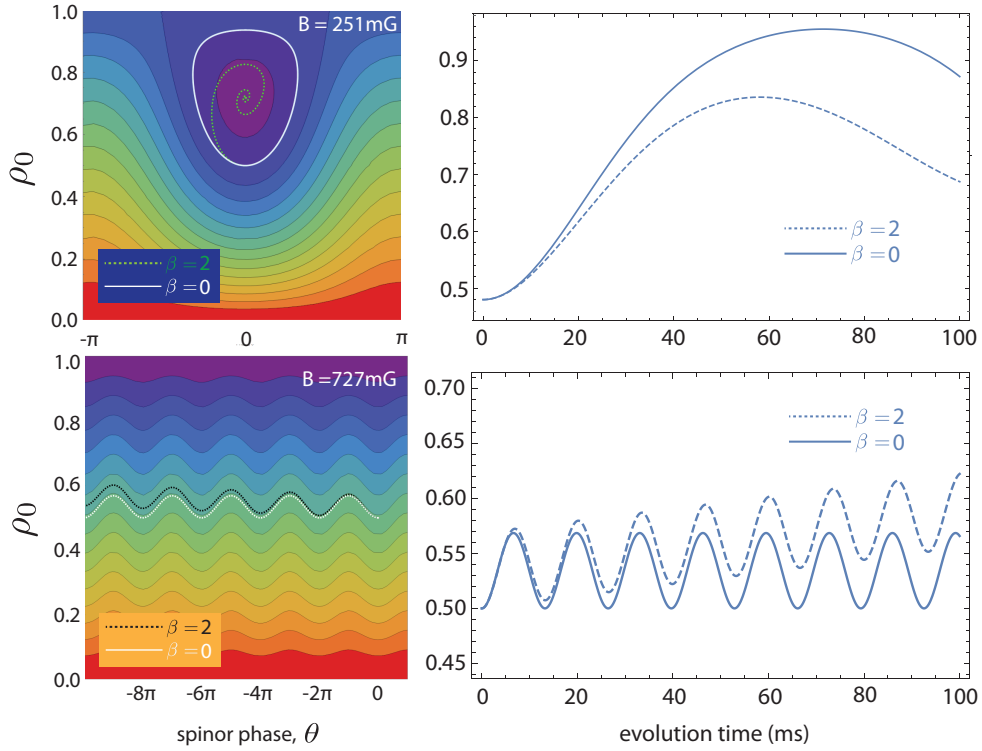
Figure 5.3: (ρ_0, θ) trajectories (left) and spin-mixing oscillations with and without the simple dissipative relaxation term β. Two example fields are shown, corresponding to oscillating phase and running phase solutions: for magnetic fields below the separatrix ($B = 251$ mG) the spinor relaxes to the ground state $\rho_0 \sim 0.7$ (top), whereas for higher fields dissipation manifests as a gradual drift towards higher ρ_0 population (bottom).

$\rho_0 \sim 0.7$ for q below the separatrix, as exhibited by spin-mixing data with $q = 2\pi \times 4.5$ Hz, $\rho_0(0) = 0.5$ and $\theta = 0$, $q_{\text{sep}} = |c| = 2\pi \times 5.2$ Hz from Eq. 2.34. For $q > q_{\text{sep}}$ the drift is towards still higher ρ_0. In general, the increasing ρ_0 population is evidence of energy loss, and relaxation to the mean-field ground state $\rho_0 = 1$ for ferromagnetic ^{87}Rb (The 'polar' state in Figure 2.3, Chapter 2). While the simple dissipative model and our data share qualitative similarities, our observations suggest that the relaxation tapers off, or has a more complicated phase dependence, with no value of β satisfactorily fitting the data. A concrete, physical explanation of relaxation behaviour is warranted, and constitutes an outstanding problem in spinor condensate physics.

A number of other experimental works have encountered similar relaxation/equilibration effects, albeit over longer timescales (up to a second or more) in both ^{23}Na and ^{87}Rb [51, 66, 68]. Our observations are distinguished from others by the timescale on which relaxation occurs; namely, relaxation

is obvious after ~ 10 ms and essentially over by 100 ms. No two experiments are identical, with trapping geometry, atom number and magnetic environments differing considerably between the experiments of other groups and our own. For both sodium and rubidium the effect is characterised by relaxation to the mean-field ground state, and thus indicative of an environmental factor, rather than some intrinsic characteristic of the atomic species such as dipolar effects or ferro/antiferromagnetic interactions. Perhaps the most distinguishing feature of our system are the comparatively high ambient field gradients, 10-20 mG/cm: in Ref. [141] the gradient is claimed to be minimised to less than 0.1 mG/cm, with relaxation occurring over several seconds. The results discussed later in this Chapter also further implicate magnetic field gradients as agents contributing to relaxation behaviour. In general, the behaviour of spinor BECs in magnetic field gradients remains a somewhat experimentally unexplored research area with scant documentation, and thorough, concerted efforts in this direction may elucidate the mechanism for equilibrium relaxation.

In addition to the relaxation behaviour described above, the oscillation *amplitude* is observed to decay. For fields far away from the separatrix, the amplitude of spin-mixing oscillations depends primarily on the magnitude of c and the period on q. Increasing q reduces the period of oscillations and also reduces the amplitude, whereas increasing (reducing) c increases (reduces) the amplitude of oscillations, with little effect on the period (Figure 2.5). Since c is weakly dependent on atom number ($\propto N^{\frac{2}{5}}$) and the trapping potential is assumed to remain constant over the course of an experiment (we observe insignificant atom loss over typical spin-mixing times), any reduction in spin-mixing oscillation amplitude is due to either

- Increasing q; or

- Reduced spatial overlap due to component separation.

The latter can be understood by noting that $c \propto \int d\mathbf{r} \, |\Phi_{\text{SMA}}(\mathbf{r})|^4$, with the SMA assuming $\phi_{m_F}(\mathbf{r}) = \Phi_{\text{SMA}}(\mathbf{r})$. Component separation inherently invalidates this assumption. In the case of a changing q, we note that the period appears to be roughly constant, perhaps slightly *increasing* with time, indicating that q is either constant or decreasing, which is incongruent with the observed amplitude decay.

Another possible explanation may lie in the presence of thermal (non-condensed) atoms. Among the early experiments on spinor BEC performed by the group of Sengstock and Bongs, their first observations of spin-mixing oscillations exhibited both damping and drift towards increasing ρ_0 [44]. Their experiments were performed using a cloud with a significant non-

condensed fraction; bi-modal fitting of the density distributions showed the thermal component does not exhibit spin-mixing oscillations, but instead drifts towards an equipartition distribution [44]. Typically only one period of oscillation could be resolved. Although the results later described in Sections 5.2.1 and 5.2.2 are highly indicative that the cause of relaxation and damping is not due to thermal effects, we occasionally detect the presence of a small thermal component in our spinor BECs that may have deleterious effects. Typically, this thermal fraction is very small, and experiments studying relaxation with varied thermal fractions show no clear dependence on the fraction of thermal atoms.

5.2 Eliminating gradient dephasing

Magnetic field gradients are the prime candidates for the underlying cause of both drift and damping. The most obvious effects are the mechanical forces on the $|\pm 1\rangle$ components that drive component separation. Gradients also result in a spatially varying magnetisation across the BEC, so that the mean-field ground states include spin domains [14]. We hypothesise that the spatial separation is responsible for damping of the spin-mixing oscillations. The drifting $m_F = 0$ population may be due to gradient-driven relaxation towards the mean-field ground state.

Whatever the mechanisms at work, it is reasonable to posit field gradients as the cause of the problems. In order to eliminate them, we need to apply cancelling gradients with magnetic field coils or current-carrying bars. This is no simple undertaking, as the coil geometry required depends on the gradients to be cancelled: only diagonal gradients $\partial B_{x_i}/\partial x_i$ could be cancelled by imbalancing the currents in the existing bias coils. Component separation provides a useful metric for determining the magnitude and direction of the predominant gradient, but can be misleading due to the imaging integration direction being oblique to the condensate's elongated axis.

A brute force approach could be used, whereby we apply different bias coil current imbalances and use an atomic metric (spin-mixing amplitude or component separation at $\sim 50\,\text{ms}$, for example, or width of a microwave frequency-domain Rabi spectrum) to diagnose the nulling procedure, but this is highly indirect and adding diagonal gradients does not affect off-diagonal gradient terms $\partial B_i/\partial x_j$, with $i \neq j$. In the case of Ref. [141], three anti-Helmholtz coils combined with their dispersive optical probing technique (that directly determines the spatial variation of the Larmor frequency) were able to reduce the background field inhomogeneity to $< 0.1\,\text{mG/cm}$. Our early attempts to minimise component separation with applied gradients

proved largely fruitless, but the findings of Section 5.4 may shed some light on why.

In Chapter 6 of this thesis, we describe a high precision interferometric measurement – using spinor BECs as magnetometers – that measures the entire magnetic field gradient tensor. Once the whole gradient tensor is known, the total field inhomogeneity can be reduced substantially by changing the direction of the bias field alone, which effectively picks out different terms of the gradient tensor, as discussed in Section 6.7.1 of Chapter 6. Before setting out to measure the gradient tensor, we developed a promising scheme that uses radiofrequency spin-echo pulses to counteract deleterious gradient dephasing.

The scheme is essentially a simple spin-echo pulse sequence used in nuclear magnetic resonance experiments [142]: a $\pi/2$-pulse initiates spin-mixing, and after a given time we apply an rf π-pulse that swaps the populations and phases in the $|\pm 1\rangle$ states. Prior to the π-pulse, gradients drive component separation; applying the π-pulse effectively reverses the sign of the gradient and hence the component separation is also reversed, as shown in Figure 5.4. For this reason, we also refer to a π-pulse used in this manner as a *refocusing* pulse. Since the magnetisation, ρ_0, and θ before and after the π-pulse are assumed to be the same, the coherent collisional process (Eq. 2.29) and thus spin-mixing dynamics are unaffected by the π-pulse. Earlier work described the use of spin-echo pulses to revive coherence in thermal mixtures of atoms in different hyperfine states [143, 144]. In these works, spatial inhomogeneity of the differential scalar light shift results in broadening of microwave hyperfine transitions, an effect discussed in more detail in Chapter 7. Spin-echo pulses have also been used to rephase interferometric contrast in a pseudospin-1/2 condensed Ramsey interferometer [145], where mean-field dephasing driven by interactions has deleterious effects similar to that of gradients in our system.

In theory, we could prolong the gradient-mitigating effect of this scheme by applying a sequence of refocusing pulses, reminiscent of Carr-Purcell (CP) [146] and Carr-Purcell-Meiboom-Gill (CPMG) [147] pulse sequences. Provided the cumulative error of refocusing pulses is kept small, it is foreseeable that such a scheme could mitigate the effects of magnetic field gradients entirely.
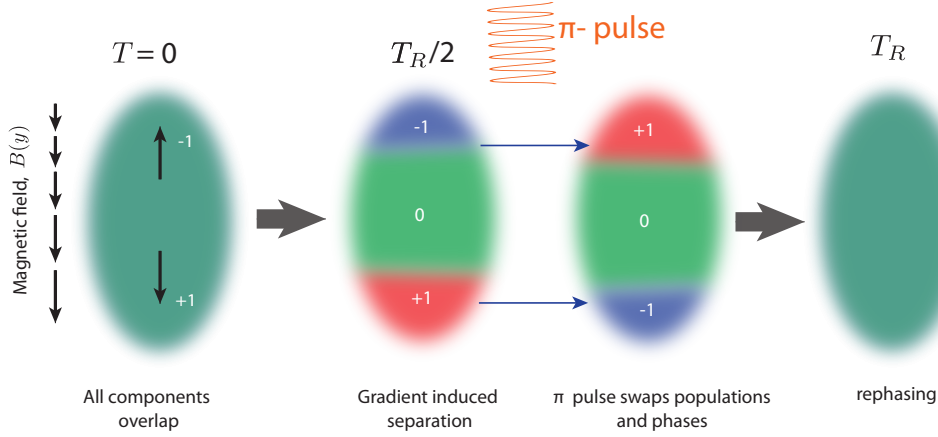
Figure 5.4: Spin-echo pulses to eliminate gradient dephasing. Immediately following preparation of the $\rho_i = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4}), \theta = 0$ state, gradient magnetic fields exert opposite forces on the magnetically sensitive states, causing spatial separation. Defining the rephasing time T_R, a π-pulse is applied at $T_R/2$, effectively swapping the populations in the $m_F = \pm 1$ states. The gradient force now acts to "reassemble" the components, leading to rephasing of the overall spin. Reducing component separation is expected to eliminate the observed damping of spin-mixing oscillations.

5.2.1 Single π-pulses

To begin with, we apply a single π-pulse after a time $T_R/2$ of evolution.[1] Several factors must be considered that affect the fidelity of the refocusing pulse, which is assumed to map $F_z(\mathbf{r}) \to -F_z(\mathbf{r})$ and leave ρ_0 and the overall spinor phase θ unchanged. Any imperfection in the pulse amplitude, frequency or phase will result imperfect refocusing and different population distributions, confounding the observation of spin-mixing oscillations.

The most important contributing factor is the magnetic field strength at the time of the pulse: this determines the instantaneous detuning of the pulse, and is determined primarily by the 50 Hz mains power (AC line) variation. Synchronising to the AC line (as described in Section 4.4.1) only partially alleviates this, as the π-pulse is then restricted to be applied when $B_{AC}(t = T_R) = B_{AC}(t = 0)$, *i.e.* at odd integer multiples of the AC line period. The effect of magnetic field fluctuations can be reduced by using stronger rf coupling. Increasing the rf Rabi frequency reduces the requisite time for the pulses and in turn makes the pulses less susceptible to small detuning fluctuations due to power broadening.

Initial experiments had mixed results. These experiments utilised the

---

[1]By convention, we define the time at which the first pulse is applied as $T_R/2$, with refocusing occurring at T_R, subsequent pulses are spaced in time by T_R.

semi-isotropic (SI) trap, where component separation due to gradients is unambiguously present, and thus suppression of these effects provides a useful metric for determining the success of the scheme. Refocusing pulses were applied at evolution times $T_R/2 < 20\,\text{ms}$ at several different quadratic shifts. We *did* observe the desired preservation of a common spatial mode structure of the different spin components, but also observed the onset of 'noise' on both F_z and ρ_0, presumably due to pulse imperfections.

We refer to the $\rho_0(t)$ and $F_z(t)$ data as noisy in the sense that the values of these quantities after the refocusing pulses were highly variable shot-to-shot, even for fixed evolution times, indicating that some random process is at work. This is of some importance, and contrary to expectations commensurate with our efforts to ensure pulse reproducibility. The observation of this noise led to much investigation into possible causes of the irreproducibility. The results of this suggested that shot-to-shot pulse amplitude, frequency, and phase variation were not significant, and thus unable to account for the post-refocusing onset of noise. Issues with the phase behaviour of the PulseBlaster DDS (used as the primary rf source for these earlier experiments) were identified, but did not appear to correlate to the observed (ρ_0, F_z) fluctuations.

Without any clear spin-mixing metric to refer to in the semi-isotropic trap we began implementing refocusing pulses in the elongated trap, where spin-mixing oscillations are higher amplitude. We observed much of the same 'noise' for certain times when π-pulses were applied in the SI trap. Figure 5.5 shows the onset of noise on both ρ_0 and F_z immediately after the application of a π-pulse at $10\,\text{ms}$ free evolution time. Before uncovering the source of this noise, we experimented with increasing the refocusing time. The first hint of success came from a single refocusing pulse applied at $T_R/2 = 30\,\text{ms}$ to a spinor BEC in the elongated trap, with a microwave dressing field (Section 4.5) applied to give $q/2\pi \sim -17\,\text{Hz}$, as shown in Figure 5.6. Such a case is interesting: for $q < 0$ the observed drift is towards *lower* ρ_0, consistent with relaxation to the mean-field ground state for $q < 0$. As well as resolving two clear oscillations not present in the data without refocusing pulses, the drift towards lower ρ_0 is arrested, strongly indicating that gradients are the underlying cause of this behaviour.

5.2.2 Multiple spin echo pulses

As is evident from Figure 5.6, a single refocusing pulse is shown to reduce damping and relaxation of the spin-mixing oscillations. However, a single pulse is not sufficient to ensure gradient-free evolution over several hundreds of ms. We then turned our attention to applying multiple pulses, using a fixed bias field of $B_x = 727\,\text{mG}$ for all experiments to this end. Although
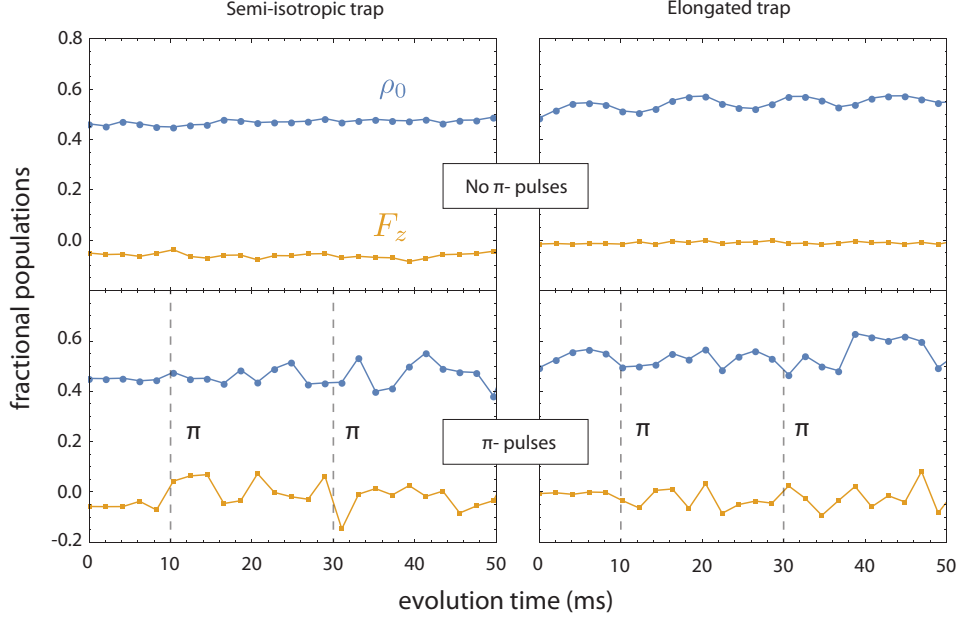
Figure 5.5: Noise on spin-mixing observables ρ_0 and F_z after the application of a refocusing pulse. Shown are ρ_0 and the magnetisation (F_z) during free evolution at a bias field of 727 mG for the semi-isotropic trap (left, top) and the elongated trap (right, top). Applying a refocusing pulse at $T_R/2 = 10$ ms immediately induces strong shot-to-shot variations ('noise') on both ρ_0 and F_z for both traps (bottom). The origin of this effect is discussed in Section 5.3. In the bottom two datasets, a second refocusing pulse was applied at $T = 30$ ms (as part of other experiments). It is clear however that the onset of noise after the first pulse is severe, and unrelated to pulse errors, which are negligible in comparison.

the previous example utilised microwave dressing fields to attain $q < 0$, we stress that there is no reason to expect general behaviour in this regime to be dissimilar from a positive q from a bias field alone. Gradients still act to drive component separation and relaxation. Initially, it was intended that a large (~ 1 G) bias field would be applied, for optimum response of the rf amplifier (originally the 10 W Delta RF LA10-1-525-40) and hence highest possible Rabi frequency. Microwaves would then be used to dress the system, and achieve values of q where spin-mixing oscillations were high-amplitude. Dressing to $q/2\pi = -17$ Hz from $q/2\pi = 38$ Hz is essentially a step in this direction. The behaviour of spinor condensates at negative quadratic shifts has also recently been characterised [68].

Improvements to the trap load and state-preparation allowed us to focus on the $q/2\pi = 38$ Hz case, without applied microwave dressing fields to increase the amplitude of oscillations. The PulseBlaster rf source was replaced with the RFBlaster DDS and the amplifier with the MiniCircuits LZY-22+, a
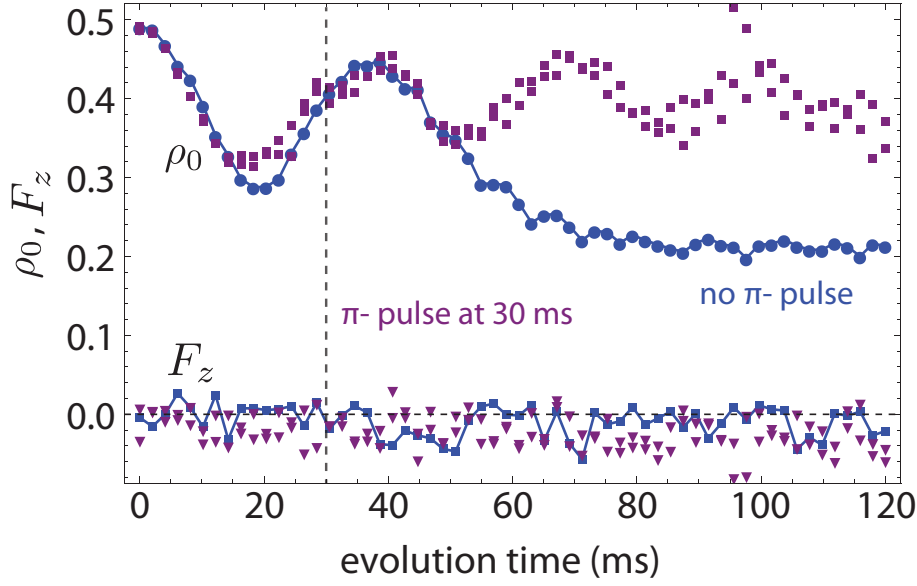
Figure 5.6: Rephasing of spin-mixing oscillations using a spin-echo pulse. A microwave dressing field is applied that creates an effective $q/2\pi = -17$ Hz, the observed spin-mixing oscillations are reminiscent of earlier results: damping of oscillation amplitude and relaxation (towards lower ρ_0 in this case, consistent with mean-field theory for $q < 0$). Applying a π-pulse at $T_R/2 = 30$ ms reverses the deleterious gradient-induced dephasing, and resolves at least two further oscillations as well as arresting the relaxation towards lower ρ_0 population.

low-frequency optimised multi-turn rf coil was added specifically for low field state preparation, as described in Chapter 4. The source of noise on F_z and ρ_0 at short refocusing times encountered previously (Figure 5.5) was assumed to be from detuning noise; the new rf setup was capable of rf Rabi frequencies of greater than 100 kHz. However, the noise remained. Even with large, 100 kHz Rabi frequencies we still observed significant variation of F_z and ρ_0 after a refocusing pulse applied 5-10 ms after the initial $\pi/2$-pulse. Indeed, at such high Rabi frequencies the fidelity of $\pi/2$-pulses seemed to exhibit fluctuation[2], ultimately we reduced the rf drive power to achieve Rabi frequencies on the order of $\Omega_{\mathrm{rf}} \sim 2\pi \times 50$ kHz, a modest improvement over the original configuration, which at best yielded $\Omega_{\mathrm{rf}} \sim 2\pi \times 30$ kHz. The source of the noise on refocusing pulses at short times remained a mystery. Before determining the source of this noise, we were able to find a regime where for certain refocusing times, the pulse fidelity was satisfactory.

An example of applying multiple refocusing pulses is shown in Figure

---

[2]This was no doubt due to coarse sampling of Rabi oscillations and effects beyond the RWA, as discussed in Section 5.3.
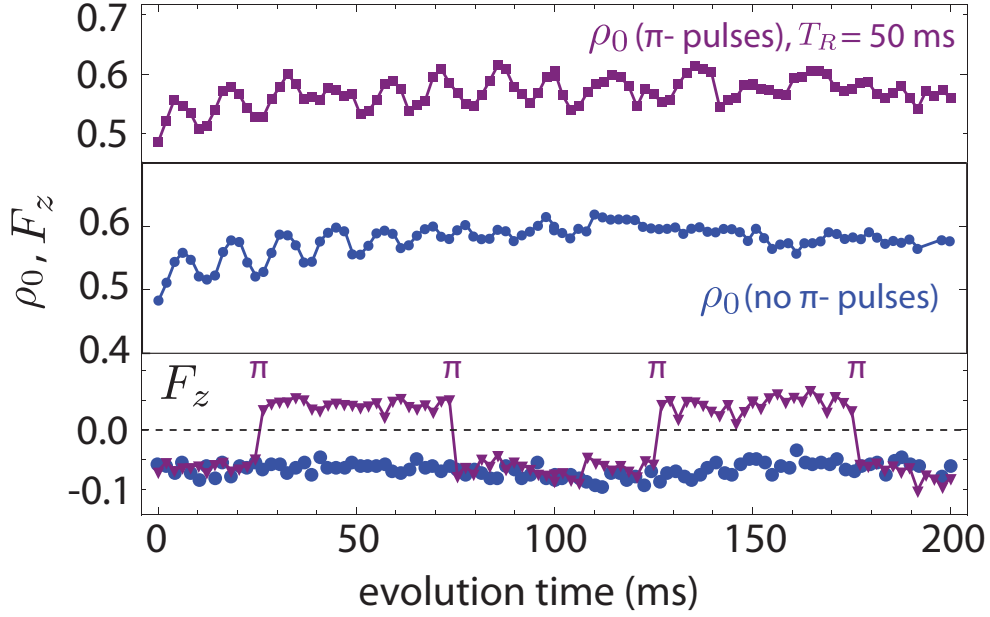
Figure 5.7: Rephasing of spin-mixing oscillations with multiple π-pulses. Pulses are applied every 50 ms, with the first pulse at $T_R/2 = 25$ ms. At least six further oscillation maxima are unambiguously resolved, and the relaxation towards higher ρ_0 evident in the bare data (centre panel) is arrested. The lower panel shows the behaviour of the magnetisation, F_z: a small residual negative magnetisation is due to imperfect state preparation, refocusing pulses flip this to positive and back again. The onset of noise is evident after subsequent pulses, which originates from cumulative pulse errors.

5.7. The first π-pulse is applied at $T_R/2 = 25$ ms, with subsequent pulses separated by $T_R = 50$ ms. Spin-mixing oscillations are resolved well past 100 ms, in contrast to the case where no π-pulses are applied. The small negative magnetisation in the data is due to miscalibration of the $\pi/2$-pulse duration: as discussed in Section 5.3, for our pulse parameters, the rotating wave approximation is not valid, and so the pulse duration for a $\pi/2$-pulse, $T_{\pi/2} \neq T_\pi/2$. Since we had used a preparation pulse duration of half the measured π-pulse duration T_π, we were consistently failing to prepare $F_z = 0$, but very reproducibly so. Here the residual negative magnetisation provides a convenient metric for the reproducibility of our refocusing pulses; the π-pulse inverts it to a small positive magnetisation, confirming that the pulse duration is accurate, and quite reproducible. A small increase in the spread of points is apparent by the second and third pulses, evidence of cumulative noise from pulse errors.

We also performed identical experiments on BECs with $F_z \approx 0$, by empirically changing the pulse duration to achieve $F_z = 0$, as shown in Figure
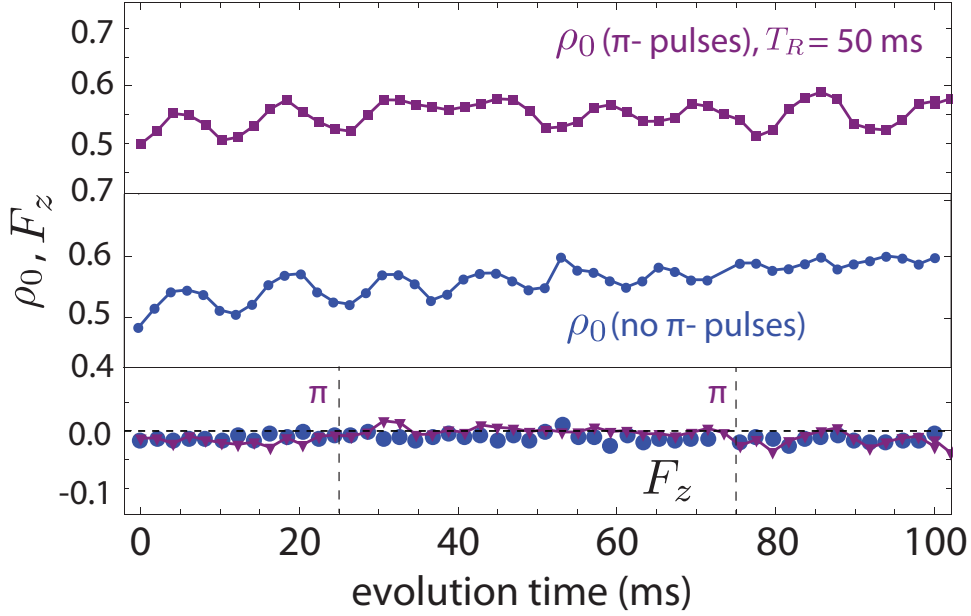
Figure 5.8: Multiple-pulse refocusing with zero magnetisation. Although the magnetisation is optimised to be zero, no major differences are observed between this the data shown in Figure 5.7. Around $T = 40\,\mathrm{ms}$, distortion to the spin-mixing oscillations is clearly apparent after refocusing pulses have been applied.

5.8, with very similar results. Although encouraging, our results also pose some questions. For instance, why does an inter-pulse time of $T_R = 50\,\mathrm{ms}$ 'work', when shorter inter-pulse times resulted in the onset of noise (Figure 5.5), which seemed inconsistent with detuning fluctuations or pulse errors? A number of reproducibility issues also arose, with the quality of rephasing varying for repeated shots as well as day-to-day.

We therefore return to considering component separation, and the effects refocusing pulses have on the centre-of-mass motion of each Zeeman state. Figure 5.9 shows component separation data (fitted centroids exhibiting relative motion) for the corresponding spin-mixing results (with and without π-pulses) shown in Figure 5.7. A number of features are immediately apparent; the most obvious of which is the centre-of-mass separation *before* the application of the first π-pulse. The initial centre-of-mass motion is vastly different for the two datasets; indeed the separation directions are opposite. This suggests a form of irreproducibility, and is consistent with a time-dependent gradient that drifts over the course of one or two experimental runs (10-15 minutes).

Before we consider the consequences of this rather disappointing finding, we address a few general observations. For example, the maximum sepa-

ration reached after the application of the first refocusing pulse is $\sim 30\,\mu\text{m}$, and during the period where spin mixing is unambiguously preserved by refocusing pulses, the two components appear to periodically move apart and rephase. For the bare case with no refocusing pulses, the centroids diverge away from each other. The former observation somewhat confirms our expectation that spatial component overlap of the spin components is necessary to preserve coherent spin-mixing, but also points to a more detailed mechanism at work – the component overlap is by no means constant in time. Also, shown in Figure 5.10 are Stern-Gerlach absorption images of the BEC immediately before and after the application of the π-pulse at 75 ms for the data shown in Figure 5.7, showing exactly the behaviour depicted in the simple schematic of Figure 5.4.

5.3   Refocusing pulse fidelity

To determine how the irreproducible component separation illustrated in Figure 5.9 affects the preservation of spin-mixing oscillations with refocusing pulses, we first need to uncover why the fidelity of the refocusing pulses is dependent on the time the pulse is applied. This is important: a time varying gradient means that the optimum rephasing time will vary shot-to-shot.

It is not immediately clear why applying π-pulses with $T_R = 50\,\text{ms}$ seems to work. The fidelity of refocusing pulses was highly dependent on T_R: at short inter-pulse times, such as $T_R = 10\,\text{ms}$ (first pulse at 5 ms) we observed significant shot-to-shot variation of F_z and ρ_0, whereas at longer T_R we observed successful rephasing and no noise on either quantity. For multiple refocusing pulses such as that depicted in Figure 5.7, we observed very little cumulative pulse error. For π-pulses applied after $T_R/2 = 30\,\text{ms}$, we observed no rephasing of spin-mixing oscillations. The observation of noise at short T_R is particularly concerning: as encouraging as the results discussed earlier are, the success of the scheme hinges on being able to mitigate gradient dephasing at the earliest possible time, before the components have a chance to separate, *i.e.* we desire $T_R/2 \ll 2\pi/\omega_{\text{trap}}$.

5.3.1   Radiofrequency coupling beyond the rotating-wave approximation

Such observations may not be able to be explained with the simplistic consideration of the rf coupling scheme described in Chapter 4. Given the tenuous validity of the RWA for the refocusing pulse experiments, we investigated the possibility that the observations of ρ_0 and F_z variability were due to effects beyond the RWA. In order to reduce the effect of detuning errors on pulses, we
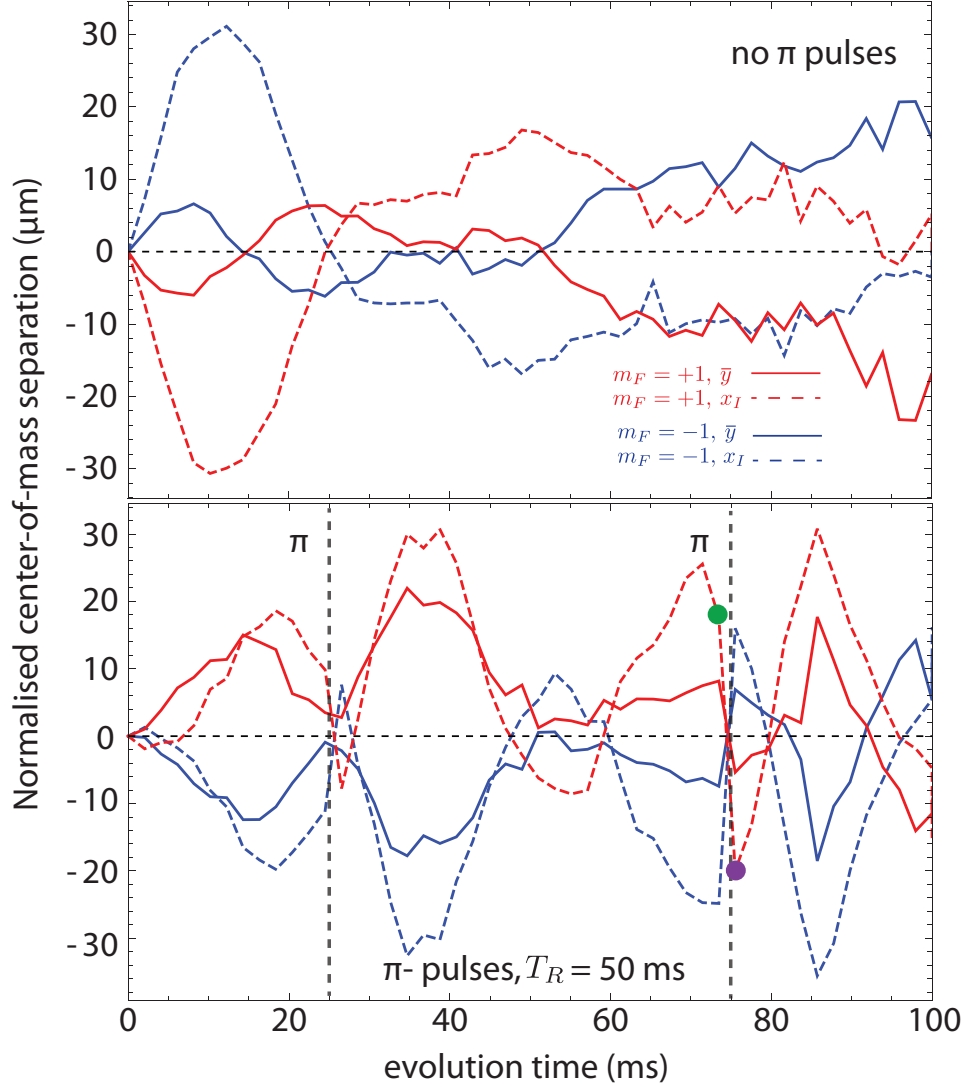
Figure 5.9: Component separation (after time of flight) as measured by centre-of-mass dynamics of each spin component; without (top) and with (bottom) applied refocusing pulses. The irreproducibility between the two experiments is clearly apparent, as the fitted centroid displacement *before* the application of the π-pulse is markedly different, the components actually separate in different directions. At the time refocusing pulses are applied, we can clearly see that they swap the positions of the $m_F = \pm 1$ components. The coloured points around the 75 ms mark refer to the absorption images shown in Figure 5.10.
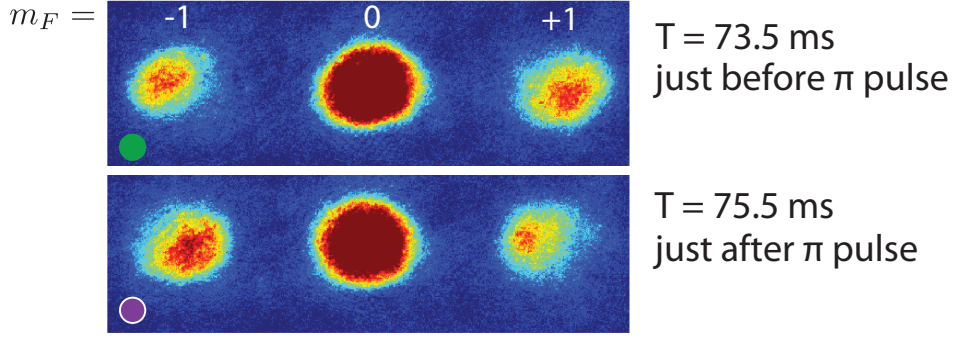
Figure 5.10: Stern-Gerlach absorption images of the BEC before (top) and after (bottom) the application of a π-pulse. It is clearly seen that the pulses effectively swap the $m_F = \pm 1$ populations, as described in the simple schematic of Figure 5.4.

use high rf powers with Rabi frequencies ranging from between 50-100 kHz at Larmor frequencies of ~ 500 kHz, so that we typically have $\omega_L/\Omega_{\mathrm{rf}} \sim 5$-10. The RWA is only valid when $\Omega_{\mathrm{rf}} \ll \omega_L$; in which case rapidly oscillating terms $\omega_{\mathrm{rf}} + \omega_L$ average to zero over the shortest timescale of the system, *i.e.* $(\Omega_{\mathrm{rf}})^{-1}$. Any effects of rf coupling not explained by the simple RWA model in Chapter 4 may be contributing to the observed refocusing pulse variability.

We first consider the solution to the Schrödinger equation for the three component spinor $\boldsymbol{\psi}(t) = (\psi_{-1}(t), \psi_0(t), \psi_{+1}(t))^{\mathsf{T}}$ for a $\pi/2$-pulse at $t = 0$. Spin collisions and spatial dynamics are neglected in this analysis, where we focus purely on the rf coupling scheme, for which the dynamics are on a much faster timescale. We write:

$$i\hbar \boldsymbol{\psi}'(t) = H_{\mathrm{rf}} \boldsymbol{\psi}(t) \tag{5.4}$$

With the coupling matrix

$$H_{\mathrm{rf}} = \hbar \begin{pmatrix} \omega_L & \sqrt{2}\Omega_{\mathrm{rf}} \sin(\omega_{\mathrm{rf}}t + \phi_{\mathrm{rf}}) & 0 \\ \sqrt{2}\Omega_{\mathrm{rf}} \sin(\omega_{\mathrm{rf}}t + \phi_{\mathrm{rf}}) & -q & \sqrt{2}\Omega_{\mathrm{rf}} \sin(\omega_{\mathrm{rf}}t + \phi_{\mathrm{rf}}) \\ 0 & \sqrt{2}\Omega_{\mathrm{rf}} \sin(\omega_{\mathrm{rf}}t + \phi_{\mathrm{rf}}) & -\omega_L \end{pmatrix} \tag{5.5}$$

the Hamiltonian for spin-1 rf coupling in the stationary laboratory frame *without* the RWA applied[3]; ω_L is the Larmor precession frequency, q the quadratic shift, ω_{rf} the rf frequency, Ω_{rf} the Rabi frequency and ϕ_{rf} the pulse phase. The operators for a π- and $\pi/2$-pulse applied at time t_0 are denoted by $U_\pi, U_{\pi/2}$ and correspond to application of H_{rf} a duration $T_\pi = \pi/\Omega_{\mathrm{rf}}$ and $T_{\pi/2} = \pi/2\Omega_{\mathrm{rf}}$ respectively: $U_{T_i} = \exp\left(-i \int_{t_0}^{t_0+T_i} H_{\mathrm{rf}}(t)/\hbar\right)$.

---

[3]The spin-1 rf coupling Hamiltonian within the RWA is Eq. 4.6.

We see evidence of the breakdown of the RWA in simple Rabi flopping experiments. On top of the familiar temporal oscillations of F_z from which we determine the Rabi frequency, there is a lower amplitude, high-frequency modulation at $\sim 2\,\omega_L$. When $\omega_L/\Omega_{\mathrm{rf}} \sim 10$, as in our experiments, this effect becomes substantial. A Rabi oscillation sampled coarsely in time appears noisy, but when densely sampled reveals a higher frequency component, as shown in Figure 5.11. Overlaid with the data in Figure 5.11 is the solution to Equation 5.4, with $\boldsymbol{\psi}(t = 0) = (1,0,0)^{\mathsf{T}}$. The solution uses our known experimental parameters $\omega_L = 2\pi \times 527.4\,\mathrm{kHz}$ and $\Omega_{\mathrm{rf}} = 2\pi \times 51.3\,\mathrm{kHz}$ – the only free parameter is due to saturation effects associated with measuring optically thick clouds, which causes $|F_z|$ near unity to be underestimated. It pays to mention that $F_z(T_{\pi/2}) \neq 0$, where the pulse duration (Rabi frequency) is measured by fitting to multiple complete oscillation periods. Radiofrequency coupling beyond the RWA is accurately described by such theory, and we can thus consider what effects violating the RWA has on the fidelity of refocusing pulses.

5.3.2 Sensitivity of refocusing pulses to Larmor phase

We consider the full sequence for a refocusing pulse experiment. The free evolution operator from the time immediately after the $\pi/2$-pulse for a time $T_R/2$ when π-pulse is applied is written as $U_{\mathrm{free}} = e^{-iH_{\mathrm{rf}}(\Omega_{\mathrm{rf}} \to 0)T_R/2\hbar}$, so that the state immediately after the refocusing pulse can be written as

$$\boldsymbol{\psi}_f = U_\pi\, U_{\mathrm{free}}\, U_{\pi/2}\boldsymbol{\psi}_i. \tag{5.6}$$

We now proceed to compare the numerical evaluation of Eq. 5.6 at different π-pulse application times, $T_R/2$. It suffices to consider the state vector immediately after the π-pulse, without the second period of free evolution (*i.e.* during which the refocusing would occur). To emulate the residual shot-to-shot variation of magnetic field,[4] we compute final wave functions ψ_f for a range of small perturbations of the Larmor frequency. These perturbations $\delta\omega_L \sim 2\pi \times 100\,\mathrm{Hz}$ are indeed small, and do not hamper the fidelity of a π *inversion* pulse (π-pulse applied to pure $m_F = -1$ state) or $\pi/2$-pulse for our Rabi frequencies. However, during the free evolution time of the refocusing pulse sequence, the *Larmor phase* in the stationary laboratory frame $\phi_L = (\omega_L + \delta\omega_L)T$ increases linearly with the time T. This is analogous to Ramsey interferometry: the second $\pi/2$-pulse of the Ramsey sequence converts the relative phase $(\omega_{\mathrm{rf}} - \omega_L)T$ to a relative population difference.

---

[4]Even with AC line synchronisation, small field fluctuations on the order of 0.1 mG from other sources still exist. We quantify this noise in Section 6.5.1.
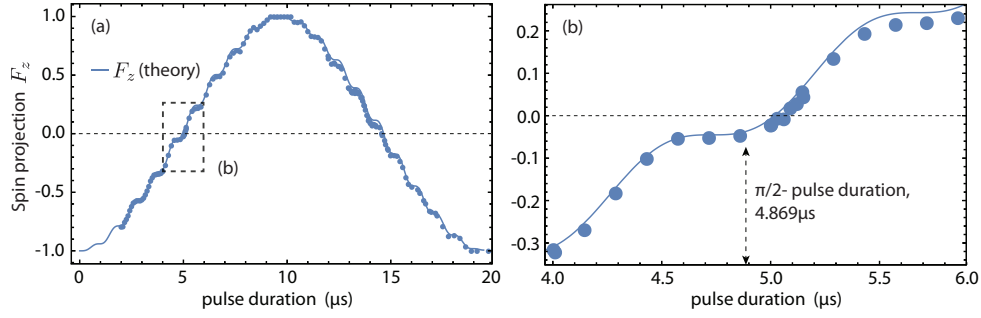
Figure 5.11: Time-domain Rabi oscillations beyond the rotating wave approximation (RWA), (a) over the course of an entire Rabi period and (b) specifically in the region of a $\pi/2$-pulse. When ω_L is no longer $\gg \Omega_{\mathrm{rf}}$, the effects of the counter-rotating term neglected in the RWA become apparent, leading to modulation of time-domain Rabi oscillations at $\omega_{\mathrm{rf}} + \omega_L$. The experimental parameters are $\omega_L = 2\pi \times 527.4\,\mathrm{kHz}$ and $\Omega_{\mathrm{rf}} = 2\pi \times 51.3\,\mathrm{kHz}$. The solid line corresponds to the solution for Equation 5.4 using the known Larmor frequency, with a correction imposed to account for saturation effects at high optical depths. Many complete Rabi periods are used to determine Ω_{rf}, which we use to define $T_{\pi/2} = \pi/2\Omega_{\mathrm{rf}} = 4.869\,\mathrm{\mu s}$; notably $F_z(T_{\pi/2}) \neq 0$ as a result of the high frequency modulation.

The small detunings $\delta\omega_L$ scramble the fringe phase for long Ramsey times, as mentioned in Section 4.5.

In the Bloch sphere representation (Figure 4.3), the Larmor phase after free evolution is the azimuthal phase of the Bloch vector. The Larmor frequency variations $\delta\omega_L$ originating from magnetic field fluctuations are sampled shot-to-shot in the case of our experiment. These variations are then amplified by the free evolution time T and give rise to Bloch vectors with substantially different azimuthal phases ϕ_L that are then acted on by a refocusing pulse. Since each shot of the experiment samples the instantaneous magnetic field strength the moment the π-pulse is applied, the spread of Larmor frequencies sampled in the following analysis corresponds to a series of experimental shots, each randomly sampling one of the Larmor perturbations. For typical magnetic noise of order $100\,\mathrm{\mu G}$, the spread of Larmor frequencies is $\sim 70\,\mathrm{Hz}$, the spread of Larmor phases at $T = 5\,\mathrm{ms}$ is on the order of π.

Figure 5.12 shows the simulated spin projection F_z and ρ_0 during a refocusing pulse, at evolution times $T_R/2$ of $5\,\mathrm{ms}$ and $15\,\mathrm{ms}$. Each coloured curve corresponds to a different Larmor phase, and the fast modulation of each is a consequence of the counter-rotating $\omega_{\mathrm{rf}} + \omega_L$ term included in the Hamiltonian. While other results, such as that in Figure 5.11 showed that a π- or $\pi/2$-pulse applied to the pure state $F_z = -1$ is exceptionally reproducible, the fidelity of a π-pulse applied to the $F_z = 0$ state after free evolution is inherently dependent on the Larmor phase. Since the Larmor phase changes

with free evolution time T, the outcome of the refocusing pulse depends on the time after the $\pi/2$-pulse, and since the Larmor phase sampled each run is essentially random, the outcome of a refocusing pulse observed shot-to-shot appears noisy.

We now explain why this is the case. The relative phase between the radiofrequency field and Larmor-precessing spin determine the axis about which the pulse rotates the spin – this is why changing the phase of the second $\pi/2$-pulse Ramsey interferometry results in fringes (Section 4.2). Within the RWA, the outcome of a resonant refocusing pulse of duration T_π applied to an equatorial spin with *any* azimuthal phase will always be $F_z = 0$. However, when the RWA is no longer valid, the non-negligible $\omega_{\mathrm{rf}} + \omega_L$ modulation of the spin trajectory apparent in Figure 5.11 renders the same resonant refocusing pulse intrinsically sensitive to the Larmor phase ϕ_L. *The phase of the modulation of the spin trajectory (the Larmor phase) therefore determines the outcome of a refocusing pulse.*

A spin-$1/2$ system exhibits much of the same behaviour for the spin projection F_z, and allows us to use the Bloch sphere directly to visualise the spin vector evolution; a similar depiction of the ρ_0 evolution is less obvious with the Bloch sphere. The evolution of the Bloch vector for an inversion pulse and refocusing pulses applied to equatorial spins with different azimuthal phases are shown in Figure 5.13. In the case of Figure 5.13, the Rabi frequency is scaled to unity, and the ratio $\omega_L/\Omega_{\mathrm{rf}} = 10$ corresponds to typical parameters used in refocusing pulse experiments. Depending on the initial azimuthal phase of the Bloch vector, a resonant refocusing pulse maps the vector to $F_z \neq 0$, *i.e* above or below the equator.

The free evolution time at which the refocusing pulse is applied corresponds to $\Omega_{\mathrm{rf}} T = 5000$, roughly $T = 16\,\mathrm{ms}$ for $\Omega_{\mathrm{rf}} = 2\pi \times 50\,\mathrm{kHz}$ as used in experiments. The 'rippling' of time-domain Rabi oscillations apparent in Figure 5.11 can be clearly seen; for higher Larmor frequencies, the helical pitch of the Bloch vector's trajectory is small and so the modulation of Rabi oscillations is faster, and of lower amplitude. Consequentially, any deviation above or below the equator (which maps to a non-zero F_z) is minimal, in contrast to that observed in Figure 5.13.

Were $\Omega_{\mathrm{rf}} \ll \omega_L$, the refocusing pulse would no longer be dependent on atomic phase and the effects described previously would be absent. However, the pulse fidelity would suffer from conventional magnetic detuning noise, since reducing Ω_{rf} reduces its power-broadening effect and fidelity would inevitably be variable shot-to-shot as well. Increasing ω_L whilst keeping Ω_{rf} high comes with the caveat of high quadratic Zeeman shifts, which suppress spinor collisions. In principle, strong off-resonant microwave dressing fields could be used to reduce q, making the scheme feasible at high Larmor
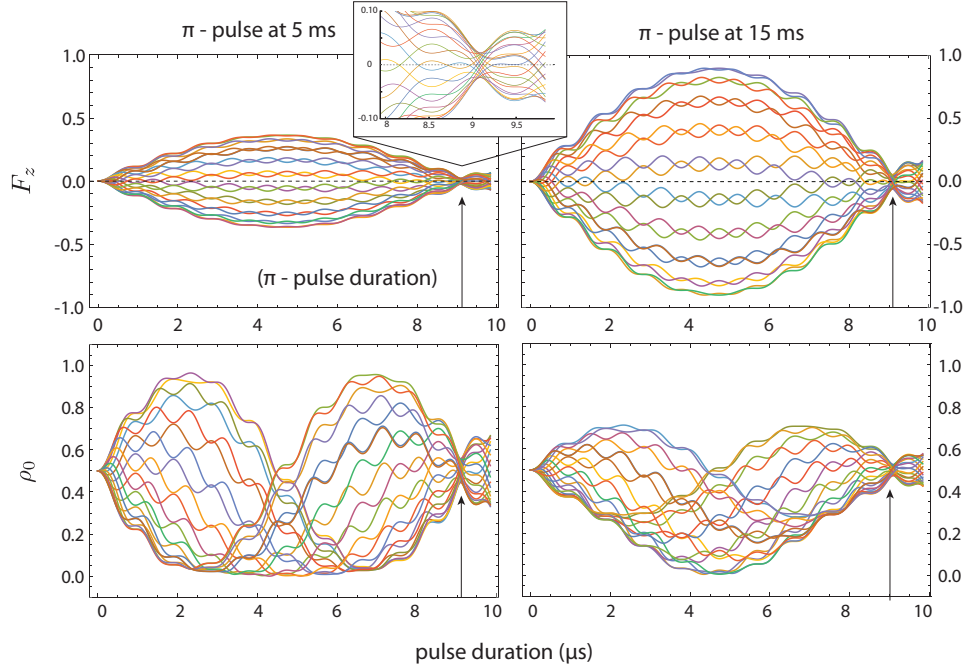
Figure 5.12: Larmor phase sensitivity of refocusing pulses on F_z (top) and ρ_0 (bottom) after $T_R/2 = 5\,\mathrm{ms}$ (left) and $15\,\mathrm{ms}$ free evolution times (right). The time axis is the duration of the pulse, which takes place after a $\pi/2$-pulse and free evolution for 5 or $15\,\mathrm{ms}$. Each coloured curve corresponds to a different Larmor phase $\delta\omega_L$, spread over $\pm 100\,\mathrm{Hz}$ in $10\,\mathrm{Hz}$ increments, which originates from magnetic noise not synchronous with the AC line. Arrows denote the π-pulse duration, $T_\pi = 2\pi/(2\Omega_{\mathrm{rf}})$. Unlike π-pulses applied to pure states ($F_z = -1$), refocusing pulses applied to the $F_z = 0$ superposition are sensitive to the Larmor phase: for a fixed π-pulse duration, sampling of individual Larmor phases shot-to-shot leads to an inherent variability in the outcome for both ρ_0 and F_z (inset). Each curve can then be regarded as an instantaneous magnetic field strength, which is then randomly sampled by each iteration of the experiment, even with line synchronisation. As the Larmor phase is dependent on free evolution time, the time at which the π-pulse is applied determines the fidelity.

frequencies.

Returning to numerical solutions of the spin-1 system, we can plot the variation of both F_z and ρ_0 after different evolution times by sampling a range of Larmor phases from a distribution of width $\sigma(\omega_L)T$ (with $\sigma(\omega_L) = \pm 2\pi \times 100\,\mathrm{Hz}$). We then compute the standard deviation of the resulting distribution of F_z and ρ_0 immediately after the π-pulse. Figure 5.14 shows the standard deviation of F_z and ρ_0 as a function of the time at which the π-pulse is applied.

The predictions of Figure 5.14 are confirmed experimentally by observing the shot-to-shot fluctuations in ρ_0 and F_z at different times. At $T = 6.6\,\mathrm{ms}$,
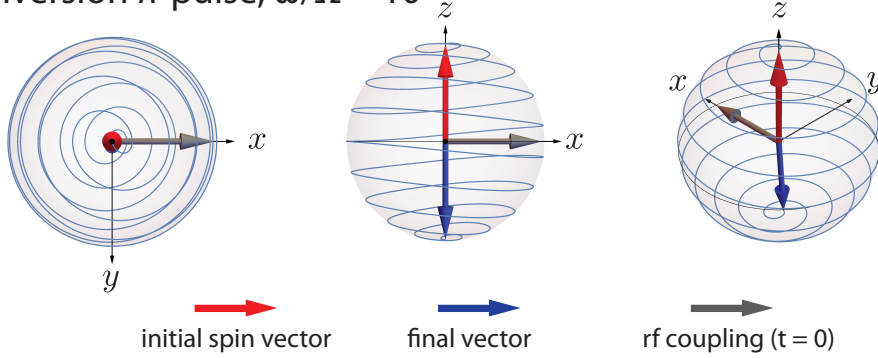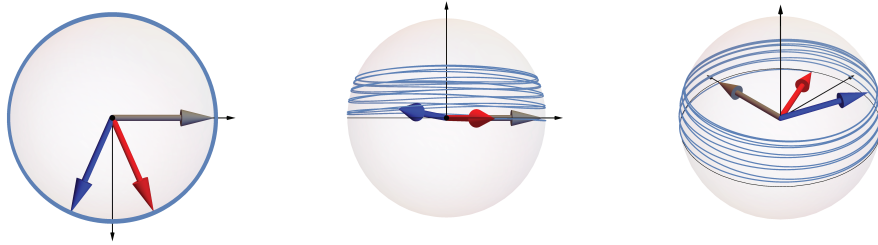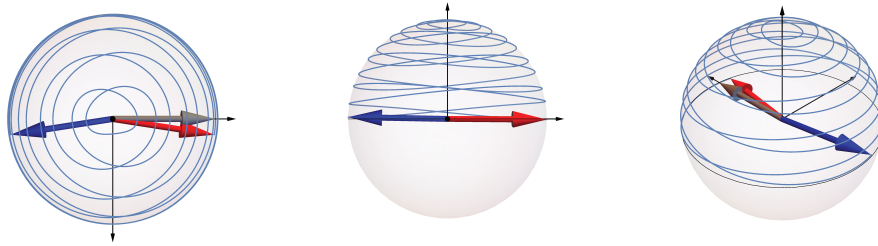
Figure 5.13: Inversion and refocusing pulses outside the validity of the RWA depicted on the Bloch sphere. In the RWA, the Bloch sphere rotates at the rf coupling frequency, and so resonant spin rotations execute lines of constant latitude about the sphere around the axis of the rf coupling field. Including the counter-rotating term in the coupling Hamiltonian (in the stationary frame) shows that during an inversion π-pulse (top) the Bloch vector executes a spherical spiral; the pitch of this spiral (which is vanishingly small for $\omega_L \gg \Omega_{\mathrm{rf}}$) is responsible for the 'rippling' of time-domain Rabi oscillations (Figure 5.11). For refocusing pulses (middle and lower frames), the azimuthal phase (Larmor phase) of the Bloch vector is different shot-to-shot due to small magnetic detuning perturbations amplified by long free-evolution times. Due to violation of the RWA, the atomic phase determines the fidelity of the refocusing pulses, resulting in $F_z \neq 0$ above or below the equator after a resonant refocusing pulse. The complicated trajectories shown in the middle and lower frames correspond to individual coloured curves in Figure 5.12.
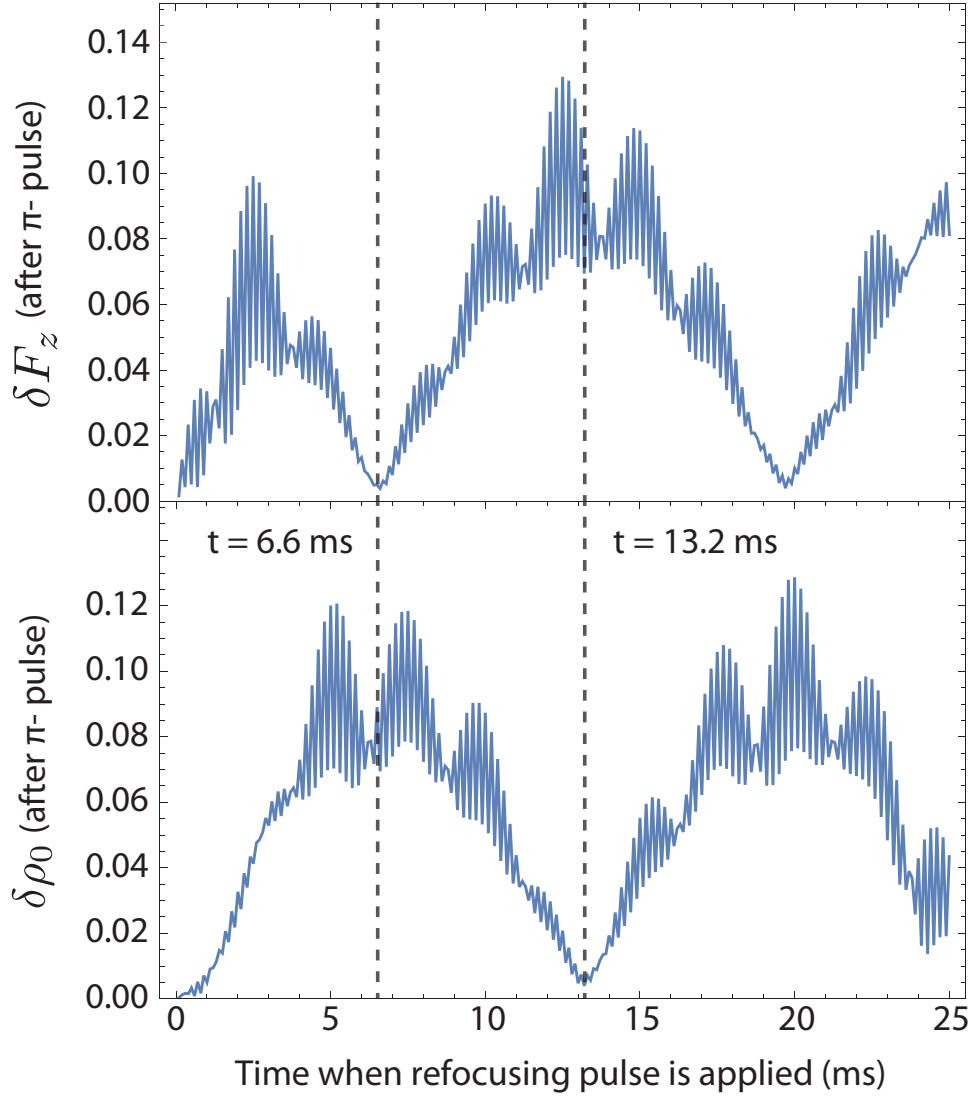
Figure 5.14: Simulated sensitivity of refocusing pulse fidelity to Larmor phase variations as a function of time at which the pulse is applied ($T_R/2$). Each point in this plot corresponds to a simulation of the refocusing pulse sequence, calculating the standard deviation of F_z and ρ_0 after a fixed pulse duration; *i.e.* the area shown in the inset of Figure 5.12, with $\omega_L = 2\pi \times 527.4\,\text{kHz}$, $\Omega_{\text{rf}} = 2\pi \times 51.3\,\text{kHz}$, $q = 2\pi \times 38\,\text{Hz}$ and 20 Larmor phases sampled over the range $(\pm 2\pi \times 100\,\text{Hz})T$ at each time. The uncertainty in F_z after a refocusing pulse is in quadrature with that of ρ_0, *i.e.* there are times at which a refocusing pulse is minimally perturbative to F_z and near-maximally perturbative to ρ_0, and vice versa.
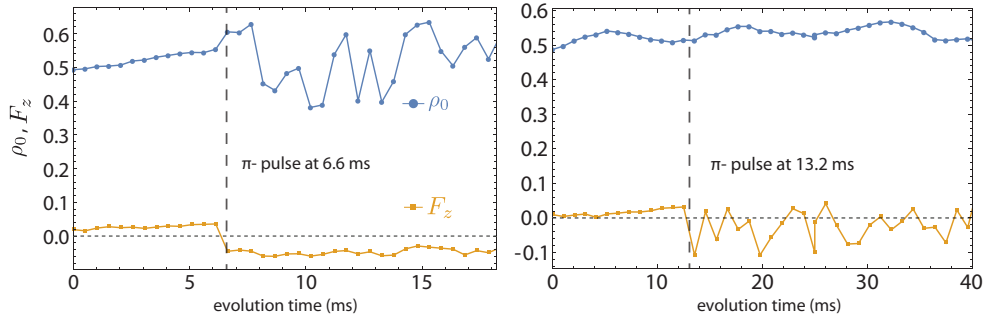
Figure 5.15: Onset of noise on ρ_0 (left) and F_z (right) from applying a π-pulse at $T_R/2 = 6.6\,\mathrm{ms}$ (maximum $\delta\rho_0$, minimum δF_z) and $T_R/2 = 13.2\,\mathrm{ms}$ (minimum $\delta\rho_0$, maximum δF_z). The observed shot-to-shot variations are concordant with the predictions of Figure 5.14.

variation in F_z is maximum and ρ_0 minimum, whereas for $T = 13.2\,\mathrm{ms}$ the case is reversed, with maximal F_z variations and minimum ρ_0 noise. The results are shown in Figure 5.15. We observe noise on one or the other quantity, in concordance with Figure 5.14. The variation of F_z is minimised at the points where the transverse spin length vanishes (due to the quadratic shift, $1/(4q) = 6.6\,\mathrm{ms}$), and variation of ρ_0 is minimsed at points in between. This appears to be an insurmountable difficulty for the success of the scheme: we can never minimise the fluctuations in F_z and ρ_0 simultaneously, particularly at short inter-pulse times.

In Figure 5.7, with a π-pulse applied at $T_R/2 = 25\,\mathrm{ms}$, we *observed* refocusing and *minimal* variation of ρ_0 and F_z. While the discussion in this Section explains why refocusing pulses applied at short evolution times exhibit noise, it *does not* account for the fact that refocusing pulses at longer times (25 ms and greater) do not: based on Figure 5.14, the variation of F_z should be near maximal at this point. The solution to this conundrum lies in effect we seek to mitigate – gradient dephasing. To best explain this, we return to Ramsey interferometry.

5.3.3 Gradient dephasing

Ramsey interferometry provides a direct method for quantifying the timescale on which gradient dephasing takes place. Gradients cause the contrast of Ramsey fringes to decay with evolution time: the Larmor frequency varies across the BEC, so that spins in different locations precess faster or slower. During the free-evolution stage of Ramsey interferometry, we can best view this as a group of vectors precessing at different rates around the equator of the Bloch sphere (Figure 4.4), with the contrast steadily decreasing until the

individual spin vectors are distributed isotropically around the Bloch sphere equator, at which point the overall contrast (vector sum) vanishes. We can express this mathematically by determining the time at which the spread of Larmor frequencies equates to 2π of relative phase,

$$\tau = \frac{2\pi}{2R_{TF,z'}\gamma B'} \tag{5.7}$$

with $R_{TF,z'}$ the largest (z') Thomas-Fermi radius of the BEC (38.8 µm), γ the gyromagnetic ratio ($2\pi \times 702$ kHz/G) and B' the field gradient. This dephasing time τ is not equal to the time it takes the contrast to vanish, due to the inhomogeneous density profile of the BEC, but provides a meaningful characteristic timescale for dephasing. We then consider Ramsey fringes obtained over 30 ms in the elongated trap at the same bias field $B_x = 727$ mG, as shown in Figure 5.16. The contrast of Ramsey fringes, which is essentially the overall transverse spin length, is amplitude modulated by the quadratic Zeeman shift. Including the effects of gradient dephasing, the Ramsey fringes take the form

$$F_z = \cos(\Delta t)\,\cos(qt)\,e^{-\left(\frac{t}{\tau}\right)^2}, \tag{5.8}$$

so that at $t = \tau$, the amplitude of Ramsey fringes has dropped to e^{-1} of the maximum value. In the nomenclature of NMR experiments this is roughly equivalent to the inhomogeneous broadening time, T_2^*, although this quantity in our case is dependent only on the background gradient strength as intrinsic dephasing mechanisms (dipolar relaxation) are weak. The data in Figure 5.16 has been overlaid with a gradient dephasing envelope, with $\tau = 12.5$ ms, from which we infer a field gradient of $B' \approx 14.7$ mG/cm. We note that after $T = 20$ ms, the Ramsey contrast has all but vanished, which corresponds to the typical times at which we have applied refocusing pulses without inducing noise on either ρ_0 or F_z and observed successful rephasing.

Recall that Figure 5.14 suggests that at typical refocusing pulse times of 25 ms we should observe shot-to-shot irreproducibility, when our experimental observations did not. Now we consider how gradient dephasing affects the sensitivity of refocusing pulses to Larmor phase. We can visualise this directly with the aid of the Bloch sphere and Figure 5.12. Consider the results of Figure 5.12, which are valid for a single atom subject to a small perturbation of the Larmor phase, and the resulting F_z sampled with each shot: each shot randomly samples a different coloured curve, and the outcome of each refocusing pulse varies.

In a uniform magnetic field, all spatial locations in the BEC experience the same overall Larmor perturbation due to magnetic field noise, and the overall spin (*i.e.* the vector sum of the constituent Bloch vectors) has the same azimuthal phase when the refocusing pulse is applied. This can be
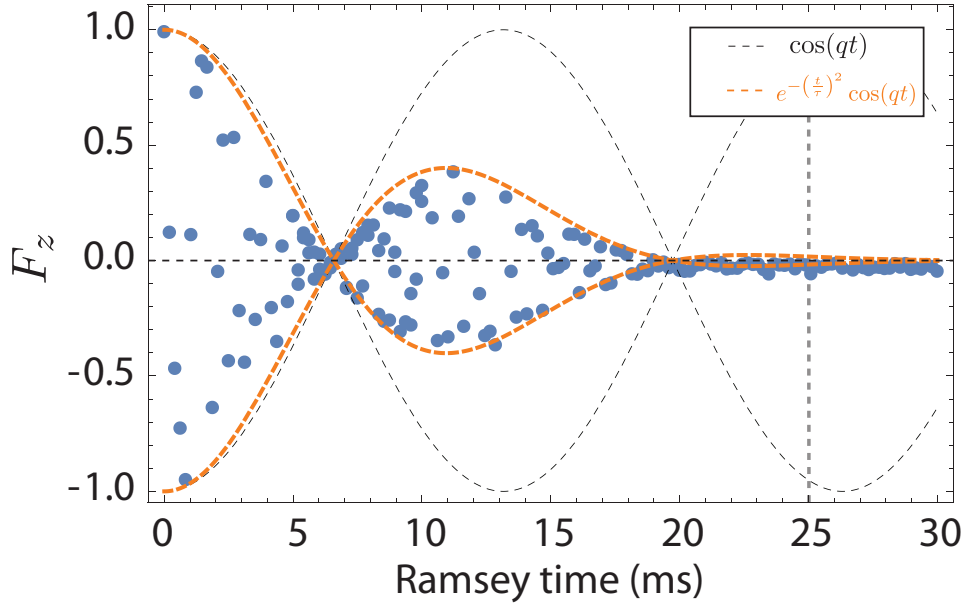
Figure 5.16: Time domain spin-1 Ramsey interference fringes. Shot-to-shot magnetic field fluctuations on the order of 0.1 mG quickly scramble the Larmor phase and hence fringe phase, but the quadratic Zeeman amplitude modulation and gradient-induced dephasing envelope remain obvious. The gradient dephasing envelope is approximated as Gaussian by Eq. 5.8, with $\tau = 12.5$ ms. After $T = 20$ ms, the fringe contrast is essentially zero.

represented by a single one of the coloured curves in Figure 5.12. After the refocusing pulse, the resultant spin (of the whole BEC) varies shot-to-shot. A BEC in a magnetic field gradient is composed of many spin vectors, precessing at different rates in different locations. Each spatial region of the condensate, and eventually every atom, is a Bloch vector with a different azimuthal phase: for long enough evolution time the azimuthal phases of the spin ensemble are isotropically distributed around the equator of the Bloch sphere. Now, the entire range of coloured curves in Figure 5.12 are sampled in a single shot by the BEC, and the resultant net variation out averaged out. This is why refocusing pulses work at longer evolution times.

In much the same way as the quadratic shift is responsible for the refocusing pulse noise on F_z being in quadrature with that of ρ_0, it is apparent that gradient dephasing is an effect common to each, and destroys the coherent effects responsible for shot-to-shot variation of refocusing pulses at short evolution times. The effect of spatial spin dephasing on spin-mixing oscillations also remains something of a mystery, as spin-mixing oscillations take place for some time after the Ramsey contrast has vanished (Figure 5.1), and somewhat unexpectedly, even when F_z exhibits strong variation (Figure

5.15, right).

The overriding result is that the optimum time to apply a refocusing pulse is dependent on the strength of the gradient. As we have seen (Figure 5.9), we have a gradient that varies run-to-run. It is clear therefore that without rectifying the drifting gradient, refocusing pulses are limited in effectiveness.

Sometime after the wrap-up of work on refocusing pulses, we became aware of work from the group of Takuya Hirano in Tokyo applying very similar techniques to $F = 2$ spinor BECs of ^{87}Rb [148]. This work focused on stabilising centre-of-mass motion of the magnetically sensitive components (a spin current) and preventing spin-mixing collisions at comparable magnetic field gradients of $\sim 15\,\text{mG/cm}$. The origin of a possible gradient-induced momentum manifesting as an effective quadratic Zeeman shift is described, although our results are largely at odds with the time varying quadratic shift they posit. Data from Ref. [148] shows refocusing pulses applied at $10\,\text{ms}$ intervals, with negligible perturbation of spin component populations. The authors also draw attention to the role of spin-mixing collisions on the fidelity of refocusing pulses, and consider them inherently detrimental, although our results suggest that this is not the case in general. A more thorough interpretation of these results, specifically for $F = 1$ with a detailed comparison with our own results is warranted.

5.4   The presence of vector light shifts

With the optimum π-pulse timing dependent on the magnitude of the gradient, a time varying gradient poses a serious problem. Figure 5.9 suggests that exactly such a time-dependent gradient exists. We therefore devoted considerable effort into tracking down its origin. It was not difficult to deduce that the time variation of the gradient was related to the experimental duty cycle: after $\sim 15\,\text{min}$ of continued operation, highly reproducible centroid separations were obtained. However, this is not a solution to the problem and constitutes an unsustainable way to run the experiment; especially given the tendency of the RFBlaster DDS devices to unpredictably crash requiring manual rebooting that interrupts contiguous shot execution.

We began a process of *exclusion testing*, where we would run the experiment for 10 min with one of the suspect devices deactivated: quadrupole driver, bias coils, and finally dipole trap AOMs, and then perform centroid separation experiments immediately after reactivating the device. This allowed the suspect device to remain in a 'cold state' during the time when the remainder of the experiment 'warmed up', and when reactivated for the centroid measurements would exhibit the trademark drift, allowing for easy

identification of the culprit.

The problem was localised to the dipole trap, specifically the thermal transients of the AOMs. This came as something of a surprise: photodiodes monitoring the dipole beam power showed no evidence of power drift, nor did the position of the condensate vary in trap, either of which would be red flags for the AOMs, which are well known to exhibit thermal drifts in diffraction efficiency and pointing direction. After some deduction we concluded that the thermal cycling of the AOMs was inducing a drifting polarisation anomaly in the dipole beams, which resulted in a time-dependent *vector light shift* (VLS).

The AC Stark shift imposed by the high power, off-resonant dipole beams forms a spin-state independent trapping potential provided the trapping light is linearly polarised; a VLS constitutes an effective Zeeman potential, or fictitious magnetic field [19, 149, 150]. The strength of the VLS is proportional to the intensity of the light, and since gravitational sag displaces the trap minimum (BEC position) from the intensity maximum of the light, the atoms additionally experience a gradient of the VLS. For a 0.1 mG VLS field at the peak intensity of the beam, typical gravitational sag of 10 μm results in the atoms sampling a 8 mG/cm gradient for a polarisation imperfection of a few degrees and our trapping beam parameters (see Section 7.1.1 of Chapter 7).

The full details about the nature of the VLS and our steps to eliminate it can be found in Chapter 7 of this thesis. Needless to say, optimising the π-pulse refocusing scheme with such an experimental systematic as a time-varying VLS gradient was essentially futile. We thus devoted our efforts to a series of experiments to characterise the magnetic gradient landscape our spinor BEC exists in, and developed a high-precision interferometric tensor gradiometer to this end, discussed in the following Chapter. We return to the theme of spinor dynamics in the presence of gradients in both these Chapters, with encouraging results.

5.5 Conclusions

In so many ways, our work in spin-mixing dynamics is largely unfinished. Questions remain unsolved: pressingly the origin of the observed drift of ρ_0 population at positive and negative values of q, and if this is a result of magnetic field gradients. Although we have speculated on the solutions to these questions in this thesis, more experimental data is needed to confirm any deduction. This reflects a general lack of detailed work in this area, and something that we feel requires thorough investigation.

Much time was spent pursuing the initial aims of this thesis, to control

spin-mixing collisions using dressing fields, the collated results of which could fill a standalone chapter of this thesis, at the expense of brevity. Indeed, it was after all the experimental difficulties had been identified and solved (Chapters 6 and 7) that the opportunity to pursue our coherent control ideas during the course of my candidature had all but expired. In Chapter 7, we had the opportunity to revisit spin mixing: we considered the effect of a vector-light shift gradient on an evolving spinor BEC, where the magnetic field gradient is minimised to below $6\,\mathrm{mG/cm}$ by applying a bias field in a particular direction (using the results described in Section 6.7.1). Under these conditions, we are able to observe coherent spin mixing to almost *one second* of evolution time, that tantalisingly retains the features of drifting $m_F = 0$ population and damping. The Chapters of this thesis that follow describe the means by which we characterise, measure and where possible eliminate the dephasing mechanisms we suffered during the investigations described in this Chapter.

Magnetic tensor gradiometry

<div style="text-align:right">6</div>

This Chapter describes our experiments using spinor BECs as ultra-small volume precision magnetic gradiometers. Our experiments to observe spin-mixing oscillations in the previous Chapter were largely thwarted by the presence of magnetic field gradients, we thus set out to perform the definitive characterisation of the magnetic landscape the BEC inhabits. As well as measuring field gradients for the purposes of eliminating them for future work with spin-mixing oscillations, the results of this Chapter further advance the use of spinor BECs as precision magnetic sensors.

The spatial extent of a magnetic sensor is an important characteristic that determines its sensitivity and applicability. Large sensors such as warm vapour magnetometers [30] are the most sensitive available, with sensitivities per unit bandwidth of order $fT\,Hz^{-1/2}$. However, the spatial interrogation region of order cm^3 is too large to precisely map out magnetic fields from microscopic biomagnetic or surface-science samples. Other sensors, which have considerably smaller sensor volumes, such as nitrogen-vacancy (NV) centre probes [34, 151] have comparatively modest sensitivities. Obtaining high field sensitivity in a microscopic sensing volume is an experimental challenge.

We have developed a magnetic gradiometer that measures the gradient of all the vector components of the magnetic field in all spatial directions. Formed from a pair of BECs, the spatial volume of the detector is only $2 \times 10^{-5}\,mm$ and it operates with a sensitivity of around $360\,pT\,Hz^{-1/2}$. With plausible improvements to the technique, sensitivities 500 times better could be obtained.

6.1 Introduction

The more information we can extract from a magnetic field measurement, the more we can know about a sample veiled by the diffraction limit of

light [152] or buried beneath the surface of the earth, waiting to explode [153]. Measuring the magnetic field of the brain provides high temporal resolution insights into the fundamental processes of cognition [154].

The most sensitive magnetometers today are atomic magnetometers, using *warm* atomic vapours. Atomic vapour magnetometers detect small rotations of optical polarisation due to magnetic fields in an spin-polarised atomic gas. There are several comprehensive reviews on the subject, the reader is directed to Refs. [20, 155] for a more detailed exposition. The most sensitive warm vapour magnetometers operate in the spin-exchange relaxation free (SERF) regime with field sensitivities of better than $1\,\mathrm{fT\,Hz^{-1/2}}$ [30, 156].

It is not just the strength of the magnetic field that is important; magnetic fields in nature change on length scales generally the size of the magnetic source. Broadly speaking, the most appropriate sensor for measuring such variation is also comparable to the size of the source. Measuring the variation of magnetic fields on the microscopic scale therefore demands equally microscopic sensors. In the case of atomic magnetometers, a challenge arises in obtaining high sensitivities within small sensor volumes: the atomic density must be increased to mitigate the decrease in atom number associated with lower volumes.

In warm vapour magnetometers, miniaturisation is not trivial – the vapour cell itself must be shrunk to the length scale of interest. Microfabricated vapour cells [31] attain $\mathrm{fT\,Hz^{-1/2}}$ sensitivities in sensing volumes of $1\,\mathrm{mm^3}$, but this is still too large to be used on the micro- or nanoscale. The smallest sensors to date are based on microfabricated SQUIDs [32] and NV-centres in diamond [33, 151, 155]. Although typically operating with sensitivities far below that of SERF warm vapour magnetometers, recent proposals suggest comparable sensitivities are possible [157].

What magnetic fields could we measure on the microscale? Consider, for example, the magnetic field of a single ^{87}Rb Bose-Einstein condensate. A 10^6 atom spin-polarised BEC is equivalent to a single magnetic dipole, with $|\mu| = 10^6(\mu_B/2)$. The magnetic field $50\,\mu\mathrm{m}$ from the condensate, along the quantisation axis, is $7\,\mathrm{pT}$. Although such fields are readily detectable by existing sensors, a sensor on the same scale as a BEC with high sensitivity is not quite so trivially obtained, much less implemented in vacuum. A BEC itself is one such sensor.

BECs and cold atoms have been heralded as precision magnetometers for much of their existence. The reason is high spatial resolution: cold atoms can be trapped on length scales of microns, and have long coherence times, meaning similarly long measurement times are possible to partially offset the loss in sensitivity from lower atom numbers. The increasing sophistication of optical trapping potentials allows for myriad spatial configurations of

ultracold and condensed atoms, allowing for novel microscale magnetic sensing configurations. Additionally, ultracold atoms offer calibration-free magnetic measurements in terms of well known fundamental constants. In warm atom vapour magnetometers, spin exchange relaxation collisions result in short coherence times; only by making the Larmor precession frequency much smaller than the collision rate (*i.e.* at low magnetic fields) can extremely high sensitivities be reached. Spin relaxation is largely absent in ultracold systems in their lowest hyperfine ground state (such as $F = 1$ ^{87}Rb), so there is no need for magnetic shielding to reach long interrogation times. Ultracold atomic magnetometers can be used in comparatively high magnetic fields, making them appropriate for real-world magnetic sensing.

BECs offer higher densities within a smaller sensing volume than ultracold thermal atoms. Although collisional shifts [158] are problematic for condensed clouds and have somewhat precluded their implementation as high-precision sensors, spinor BECs are unaffected by hyperfine collisional shifts [10] . Single-component BECs were first employed as magnetometers by observing how the density of a spatially elongated condensate varies due to the added potential of the magnetic field to be measured [9]. Spinor gases are a more direct magnetic field probe. Vengalattore et al. [10] made the most significant contribution to spinor BEC magnetometry to date by producing high-resolution spatial maps of the condensate magnetisation by imaging with a dispersive optical probe. The sensitivity of their technique ($8.3\,\mathrm{pT\,Hz^{-1/2}}$) was comparable to SQUID probes operating with comparable spatial resolution.

The primary means of extracting magnetic information from atomic samples – warm, ultracold and condensed – has been dispersive optical probes. In this scheme, the atomic polarisability rotates the linear polarisation of the light in the presence of the magnetic field through the Faraday effect. Faraday probing of thermal atoms in optical dipole traps has been used for magnetometry [11, 159–162], obtaining nT to $\mathrm{pT\,Hz^{-1/2}}$ sensitivity over spatial regions of 10-100 μm. Optical probing offers time-resolution within a single measurement,[1] but contends with photon shot noise and in some cases, perturbative effects on the trapped atoms from scalar and vector AC Stark shifts.

6.1.1 Tensor magnetometry

Measuring the variation of the total magnetic field *strength* in a microscopic region is often not enough. The magnetic field is a vector quantity, and so

---

[1]In contrast to shot-based projective measurements, such as rf spectroscopy and Ramsey interferometry.

a determination of the vector components of the magnetic field is a more
incisive measurement. The vector components of the magnetic field vary
separately in space as well. In total, there are nine magnetic field gradient
terms, forming the magnetic field gradient tensor

$$G_{ij} = \partial B_i / \partial x_j \tag{6.1}$$

for field components $B_i \in \{B_x, B_y, B_z\}$ along directions $x_j \in \{x, y, z\}$. Maxwell's
laws for magnetic fields in a vacuum, specifically Gauss's law for magnetic
fields[2] $\nabla \cdot \boldsymbol{B} = 0$ and Ampere's law $\nabla \times \boldsymbol{B} = 0$ restrict the number of in-
dependent gradient terms to five. In order to measure the components of
the gradient tensor, a *tensor* magnetometer must be sensitive to the variation
of all vector components across at least two spatial directions. Tensor gra-
dient magnetometers have to date been macroscopic devices – employing
SQUID [163–165] or fluxgate [166, 167] sensors – primarily applicable to
geophysics and ordinance detection [168]. Tensor gradiometers make use of
Maxwell's laws to infer the complete tensor from as few as five independent
gradient terms [163, 166, 167, 169, 170].

What are the benefits of a tensor measurement? The primary example
is the localisation of magnetic source distributions: the gradient tensor at a
single point in space determines the bearing, normalized source-strength,
and orientation of a single magnetic dipole source [171]. Complex real-world
magnetic source distributions are more readily inferred from arrays of tensor
measurements than from total magnetic field measurements [172, 173].

Knowledge of the full gradient tensor also confers numerous advantages
to experiments with spinor condensates. As we have discussed in Chapter 5,
the presence of magnetic field gradients is a serious problem for experiments
with spinor BECs, leading to component separation and suppression of spin-
mixing collisions. As we shall see later in this Chapter, knowing is half the
battle – once the gradient tensor has been determined, the field gradients
can be reduced by simply changing the orientation of the magnetic bias field
to a predicted minimising direction. Also, additional cancellation coils can
be tailored to cancel specific gradient terms, and applied inhomogeneous
magnetic fields can be precisely characterised by the atoms themselves.

The remainder of this Chapter is devoted to our realisation of an ultra-
small spatial volume, high sensitivity tensor gradiometer using spinor BECs.
In Section 6.2 we present basic magnetic relations that underlie the operation
of our gradiometer. Section 6.3 follows with a discussion of differential Ram-
sey interferometry (DRI), the scheme that allows us to measure the difference
in Larmor frequencies between nearby spins, an essential requirement for

---

[2]Also called the 'no magnetic monopoles' relation.

gradiometry. The apparatus layout and experimental procedure are detailed in Section 6.4 and the results of a measurement of the background field gradient tensor are presented in Section 6.6. A number of systematics limit our sensitivity; the most pressing of which is the presence of vector light shifts in the trapping beams which serve to limit the interferometric interrogation time. With realistic modifications detailed in Section 6.8, our technique could attain substantially higher sensitivities on the order of $600 \, \mathrm{fT} \, \mathrm{Hz}^{-1/2}$.

6.2 Magnetic field gradiometry

A scalar magnetometer measures the total magnetic field strength at a given point in space

$$B(\boldsymbol{r}) \equiv |\boldsymbol{B}(\boldsymbol{r})| = \sqrt{B_x(\boldsymbol{r})^2 + B_y(\boldsymbol{r})^2 + B_z(\boldsymbol{r})^2}. \tag{6.2}$$

A measurement of the magnetic field strength at two different points in space $\boldsymbol{r}_1, \boldsymbol{r}_2$ can then be used to approximate the magnetic field gradient for static magnetic fields across small enough baselines $\boldsymbol{r}_1 - \boldsymbol{r}_2$

$$B(\boldsymbol{r}_1) - B(\boldsymbol{r}_2) \approx (\nabla B(\boldsymbol{r})_{r=r_{12}}) \cdot (\boldsymbol{r}_1 - \boldsymbol{r}_2), \tag{6.3}$$

where the gradient is evaluated at the midpoint $\boldsymbol{r}_{12} = (\boldsymbol{r}_1 + \boldsymbol{r}_2)/2$ of the separation vector. The gradients of the total field strength take the form

$$\frac{\partial B}{\partial x_i} = \frac{B_x}{B} \frac{\partial B_x}{\partial x_i} + \frac{B_y}{B} \frac{\partial B_y}{\partial x_i} + \frac{B_z}{B} \frac{\partial B_z}{\partial x_i}. \tag{6.4}$$

where we have suppressed the dependence of all scalar fields on $\boldsymbol{r}$. We identify an important result, that underlies the vector sensitivity of this scheme: to achieve vector magnetic field sensitivity along a given direction, a bias field large compared to the two other orthogonal directions is applied. Specifically, to measure only gradients of the x-oriented component of the magnetic field, cancellation fields are applied along $\boldsymbol{y}$ and $\boldsymbol{z}$, leaving a magnetic field with $|B_x| \gg |B_y|, |B_z|$. Sensitivity to specific components of the gradient tensor is then achieved by changing the orientation of the separation vector between the two sensors. For example, if the two sensors are aligned along the x axis, separated by Δx, and a B_y bias is applied, the measured gradient is well approximated by

$$\underbrace{\frac{\Delta B}{\Delta x}}_{\text{Finite difference}} \approx \underbrace{\frac{\partial B}{\partial x}}_{|B_y| \gg |B_x|, |B_z|} \approx \overbrace{\frac{B_y}{B} \frac{\partial B_y}{\partial x}}^{|B_y| \approx |B|} \approx \mathrm{sign}(B_y) \frac{\partial B_y}{\partial x} \tag{6.5}$$

where the braces denote the separate approximations used and ΔB is the raw field difference measured by the two sensors. In order to measure the whole gradient tensor, three fields in orthogonal directions and two orthogonal separation directions are needed.[3] This configuration emulates commercially available tensor gradiometers based on SQUIDs [174]. As a SQUID is coupled to an inductive pickup loop in these designs (*i.e.* a coil normal to x only detects B_x), it is intrinsically vector sensitive and additional bias fields are not required.

In addition to providing vector sensitivity to a gradiometer, bias fields may be applied that pick out the smallest gradient of total field strength. Once the gradient tensor has been measured, linear algebra can determine the lowest possible gradient achievable without adding additional counter-gradients, we discuss this in more detail in Section 6.7.1. Applying only a single gradient then augments carefully chosen uniform bias fields to eliminate the gradient entirely.

6.3 Differential Ramsey interferometry

We now discuss the sensors themselves. Our tensor gradiometer uses a pair of spinor BECs to measure the difference in magnetic field strength across a dynamically reconfigurable baseline. We determine the difference in Larmor frequency using radiofrequency (rf) differential Ramsey interferometry (DRI), which simultaneously addresses the pair of BECs, each of which can be considered an independent, spin-1 Ramsey interferometer. The wavelength of the rf is much greater than the separation between the BECs and the Larmor frequency difference much smaller than the rf Rabi frequency; the BECs are thus simultaneously addressed by effectively the same rf pulse. The *difference* in the phase acquired by each interferometer $\Delta\phi$ is proportional to the magnetic field difference between the two condensates, and is thus immune to common-mode magnetic field noise. The pair of BECs can be aligned along two near-orthogonal axes, and bias magnetic fields may be applied along three orthogonal directions, allowing us to reconstruct the entire magnetic field gradient tensor.

The techniques used to prepare a pair of BECs separated along a particular axis will be discussed in Section 6.4. Recall that in Section 4.2 we introduced Ramsey interferometry for a spin-1 system, and that spin-1 Ramsey fringes in the rotating wave approximation take the form

$$F_z = \cos(qT)\,\cos(\Delta T), \tag{6.6}$$

---

[3]Or *vice-versa* – in general, it is easier to apply three orthogonal fields than to span all three separation directions in a typical BEC experiment.

with T the Ramsey interrogation time. For two spin-1 interferometers at positions x_1 and x_2 sampling a magnetic field difference $\Delta B = |B_1| - |B_2| \neq 0$, the Ramsey fringes from one interferometer will be faster than that of the other. Assuming that $q(B_1) \approx q(B_2) = q$, the fringes are written as $F_z^{(1)} = \cos(qT)\cos(\phi_1)$, $F_z^{(2)} = \cos(qT)\cos(\phi_2)$, with $\phi_{1,2} = \omega_{\rm rf} - \gamma B_{1,2}$, γ the gyromagnetic ratio and $\omega_{\rm rf}$ is the $\pi/2$-pulse frequency. Defining the relative phase

$$\Delta\phi = \phi_2 - \phi_1, \tag{6.7}$$

the magnetic field difference is related to the relative phase by

$$\Delta\phi = \gamma \Delta B\, T. \tag{6.8}$$

Combining Equations 6.5 and 6.8, DRI can be used to measure the components of the magnetic field gradient tensor, for example

$$\frac{\partial B_y}{\partial x} \approx \text{sign}(B_y)\frac{1}{\gamma\,\Delta x}\frac{\Delta\phi}{T}, \tag{6.9}$$

for a bias field in the y-direction and the two interferometers aligned along the x-axis, separated by Δx.

The relative phase can be extracted by comparing the Ramsey fringes from the two interferometers. The precision with which Ramsey interferometry can be used to infer $\Delta\phi$ improves linearly with interrogation time, and for a spinor BEC the Zeeman coherence time of order seconds does not limit the attainable precision [35]. We will discuss this more when we consider the sensitivity analysis of our system and prospective extensions in more detail (Section 6.8). As we shall see in Section 6.4 the scrambling of the absolute phase ϕ_i from each interferometer by common-mode field fluctuations necessitates a more sophisticated data analysis procedure than simply fitting sinusoids to the fringes from each interferometer.

The principle advantage DRI offers in simultaneous interrogation of the two interferometers is rejection of common-mode field fluctuations and variations in imaging efficiency. A related experiment measuring field gradients using simultaneous interrogation of two atomic fountains (where the interfered states have different momentum) was presented in Ref. [175]. Simultaneous interrogation renders the gradiometer insensitive to the bulk of AC magnetic field noise in the lab, assuming the sources are far away from the atoms compared to the separation baseline. Static field gradients can then be measured in a comparatively noisy magnetic environment. The gradiometer could attain AC field gradient sensitivity by the inclusion of spin-echo pulses [176].

As well as measuring components of the gradient tensor, the gradiometer can also be operated as a more conventional magnetometer, with the

gradiometer configuration eliminating common mode noise [177]. In this 'co-magnetometer' configuration, the spatial variation of magnetic field strength can be mapped by scanning the position of one BEC in the vicinity of a small magnetic source while keeping the other 'reference' BEC fixed. We return to this concept in Section 6.8, where we discuss the sensitivity limits of a prospective gradiometer.

6.4 Making two BECs

We produce two BECs by making a three-beam crossed optical dipole trap. Our standard dipole trap is modified by the inclusion of a second rf frequency into one of the AOMs that control the position and intensity of the dipole beams. The two radiofrequency signals are combined using a Minicircuits zf-sc-2-4-s+ splitter combiner and fed into the same rf amplifier. This configuration is mirrored on the other dipole beam, so that either beam may be independently split. The frequency difference between the two signals being combined is typically initiated at around 6 MHz – for splitting frequencies of 3 MHz or less, strong heating effects are observed in the trap.

The frequency of the AOM drive determines the pointing of the diffracted order, smoothly changing this frequency allows the trap position to be smoothly translated. To produce two BECs, we split a single dipole beam by about 100 μm, creating two crossed-beam dipole traps, and run our usual BEC experimental sequence (Section 3.9.4), resulting in cold atoms loaded into into two traps, which are then evaporated to BEC by reducing the power in both diffracted orders from the AOM. We then implement a transfer stage, where the separation between the two traps is smoothly increased over 2 s. A schematic illustration of the setup is shown in Figure 6.1.

Instead of using two rf sources to split a single trapping beam, one DDS may be programmed to quickly 'hop' between two (or more) frequencies, which has been used by the Kjaergaard group (Otago) for similar multi-well trapping potentials [178]. We initially experimented with this procedure and found that attaining a two-beam trap in this way was difficult, as it required a *single* trap to be smoothly deformed into two. In contrast to the simplicity of adding rf frequencies via a combiner, the 'hopping' method also requires very fast switching rates (~ 100 kHz) between frequencies to prevent heating of the condensate. This proved difficult to implement with a single RFBlaster DDS.

As described in Section 3.9, the two dipole beams define axes x' and z', and the two BECs are separated along one of these axes. For the purposes of gradiometry, the largest possible separation is desirable to maximise the
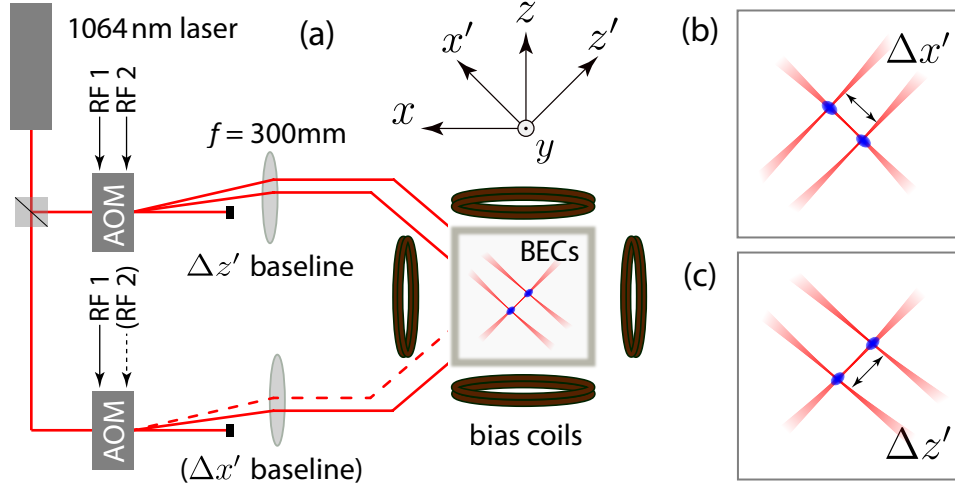
Figure 6.1: Diagram of the experimental setup (a). The gradiometer spans a spatial baseline given by $\Delta x'$ (b) or $\Delta z'$ (c) in the horizontal plane (gravity along $-\hat{\mathbf{y}}$) by translating the pair of condensates in opposite directions along one of the dipole beams, which are oriented at $\sim 45°$ to the magnetic bias field axes, x, z. This translation is achieved by feeding a second rf frequency into the AOMs that control the position and amplitude of the dipole beams. Note that in this Chapter, we have taken for the positive direction of the x'-axis the opposite to that in Figure 3.15 and that used in Chapter 7, where the propagation direction of the beam defines positive x'.

measured phase difference for a given gradient (Eq. 6.8). The maximum separation distance is limited largely by the diffraction bandwidth of the AOMs. Larger deviations from the centre frequency of 110 MHz are demanded for larger beam displacement, more rf power is required to maintain sufficient intensity to trap the atoms against gravity.

The baselines spanned by the gradiometer are along x' and z', but our magnetic axes are defined by the orientation of the magnetic bias coils, which are along x, y, and z. We must therefore measure gradients in the (x', y, z') coordinate system and transform these to the (x, y, z) frame. We calibrated the splitting of each dipole beam using a combination of side and top absorption images, through which we are able to determine the angle between the two beams and the baseline spanned for a given frequency difference of the two-tone AOM drive. Figure 6.2 shows absorption images of the BECs in the split configurations at the maximum achievable separations.

The (x', y, z') coordinate system is related to the (x, y, z) coordinate system by the following transformation:

$$\hat{\mathbf{x}}' = \sin\theta_{x'z}\,\hat{\mathbf{x}} + \cos\theta_{x'z}\,\hat{\mathbf{z}} \tag{6.10}$$

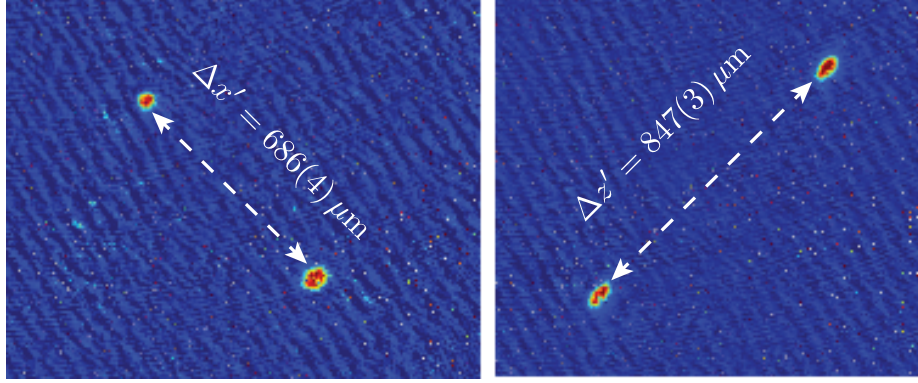$$\hat{\mathbf{z}}' = -\sin\theta_{z'z}\,\hat{\mathbf{x}} + \cos\theta_{z'z}\,\hat{\mathbf{z}} \tag{6.11}$$

Figure 6.2: False colour absorption images after 2 ms drop (from the top imaging camera, looking down on the x'-z' plane) of the two BECs split along the x' and z' directions.

with $\theta_{x',z} = 43.2(1)°$ the angle the x'-axis makes to the z-axis, and $\theta_{z',z} = 45.0(6)°$ the angle the z' axis makes to the z axis. The two approximately orthogonal baselines are not centred at the same point in space, largely due to asymmetric AOM diffraction bandwidths. Splitting the x'-dipole beam yields a separation baseline $\Delta z'$, and splitting the z' beam spans a baseline $\Delta x'$. The x' beam splitting was determined to be $\Delta z' = 28.3(1)\,\mu$m per MHz of frequency difference between the diffracted beams, with a maximum separation of $\Delta z'_{\text{max}} = 847(3)\,\mu$m. The z' beam splitting was found to be quite asymmetric, with each beam x'_0, x'_1 having a different separation calibration: $\Delta x'_0 = 19.8(1)\,\mu$m/MHz and $\Delta x'_1 = 22.9(1)\,\mu$m/MHz. The maximum separation along the x' axis is $\Delta x'_{\text{max}} = 686(4)\,\mu$m. More details on the transformation between primed and un-primed coordinates can be found in Section 6.6.3 of this Chapter.

As the frequency of the rf drive to the AOMs is changed, the diffraction efficiency also changes. The AOMs (Crystal Technologies 3110-197 units) are anti-reflection coated for 1064 nm light and have a centre frequency of 110 MHz, although we find optimum diffraction efficiency at 101.1 MHz. During the frequency ramp that extends the gradiometer baseline from the initial value of 100 μm to the maximum $\Delta x'$ or $\Delta z'$, we ramp up the rf power to compensate for the reduced diffraction efficiency. The transfer stage is optimised for equal atom number in each BEC, typically we have 5×10^4 to 1×10^5 atoms in each trap. The interferometry sequence is performed in freefall, as discussed in the following Sections, to prevent the contributions to the measured gradient from vector light shifts. Differences in the beam powers at the end of the transfer ramp would lead to different vector light shifts are thus not an issue, rendering in-depth characterisation of the final

optical powers unnecessary.

6.5 Differential interferometry

Once the pair of BECs has been translated to the maximum baseline separation, we initiate freefall Ramsey interferometry. To make the gradiometer sensitive to only spatial variations of *magnetic field*, we must extinguish all the trapping light, which eliminates the contribution of vector light shifts from the trapping beams, which plagued spinor dynamics experiments in Chapter 5. This also eliminates the systematics caused by time variation of the VLS shift. This simple solution is not without limitations. The maximum Ramsey interrogation time, and hence ultimate precision, is limited to a few milliseconds to prevent significant movement of the atoms through the interrogation region. After 3 ms of freefall interrogation time, the BEC has fallen 44 μm, about four Thomas-Fermi radii. This in turn increases the effective sensor volume: the sensing volume is now that which is swept out by a falling, expanding BEC. For the immediate application of measuring the background gradient tensor, where $|\partial B_i/\partial x_j|_{\min} \approx$ 1-2 mG/cm, we are not limited by either sensor volume or operating precision.

In order to measure a magnetic field difference, we perform DRI on the falling BECs, which begins 100 μs after the trapping light is switched off. The fringes from each interferometer take the form

$$F_z^{(1)} = \cos(qT)\cos([\Delta_1 + \delta\Delta]T + \varphi) \tag{6.12}$$
$$F_z^{(2)} = \cos(qT)\cos([\Delta_2 + \delta\Delta]T + \varphi) \tag{6.13}$$

where Δ_1 and Δ_2 are the detuning from resonance due to the magnetic field difference between the two BECs, φ is the phase of the second $\pi/2$-pulse and $\delta\Delta$ is a common-mode field fluctuation due to AC magnetic fields or shot-to-shot drift of the background magnetic field. Varying the pulse phase φ over 2π traces out phase-domain Ramsey fringes, and for short interrogation times where the product $\delta\Delta$ is small they are clearly sinusoidal, as shown in Figure 6.3. At longer interrogation times the absolute fringe phase is scrambled by the common-mode fluctuations $\delta\Delta$. The *relative* phase (Eq. 6.8) however, is preserved. Plotting the output from one interferometer against the other parametrically yields an ellipse – this has been known for some time by the gravity gradiometry community, who use exactly this method to eliminate common-mode baseline movement [179, 180]. As shown in Figure 6.3, there is negligible degradation of the plotted ellipse with interrogation time.
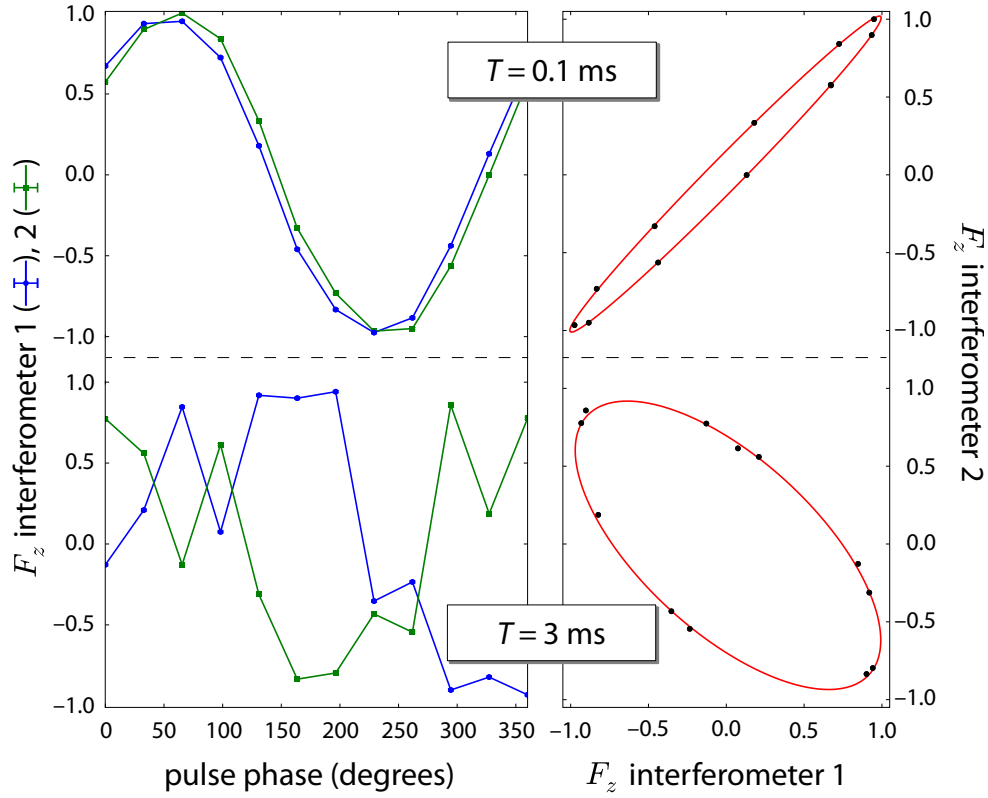
Figure 6.3: Common-mode rejection using ellipse fitting. Top: Phase domain Ramsey fringes at 100 μs (left) and 3 ms (right) Ramsey interrogation times. Common mode magnetic noise scrambles the absolute phase of the fringes from each interferometer, but a parametric plot of the outputs of one interferometer against another yields an ellipse (bottom). The relative phase, and hence magnetic field difference, between the two interferometers can the be determined by fitting an ellipse to the lower plots.

To determine the relative phase from such a plot, we fit an ellipse to the parametric dataset. The general form of a conic is

$$ax^2 + bxy + cy^2 + dx + ey + f = 0, \tag{6.14}$$

with an ellipse satisfying $b^2 - 4ac < 0$. The relative phase $\Delta\phi$ satisfies [179]

$$\cos\Delta\phi = \left(\frac{b}{2\sqrt{ac}}\right). \tag{6.15}$$

Note that the $\Delta\phi$ reported from an ellipse at some interrogation time T may be aliased, this must be accounted for when determining the actual magnetic field difference.

We use the widely cited method of Fitzgibbon et al. [181] to perform least-squares direct fitting of ellipses.[4] The uncertainty in the fitted relative phase

---

[4]Ellipse fitting is a very active field of research, least squares fitting of ellipses is non-trivial due to the constraint $b^2 - 4ac < 0$.

is not trivial. At short interrogation times, the entire ellipse is uniformly sampled in roughly 12 experimental shots (pulse phases φ), but at longer evolution times, this is generally not the case, with one or more axes of the ellipse sparsely sampled, this results in a poor estimate of $\Delta\phi$ that is not necessarily reflected in the reported uncertainty.

Although not a significant issue for $T = 3\,\text{ms}$ interrogation times, for the measurements in Chapter 7, where $T = 15\,\text{ms}$ it is significant. We initiated a collaboration with Zygmunt Szpak (in the group of Wojciech Chojnacki) at the University of Adelaide with the aim of using their fitting procedures [182] to report accurate uncertainties in the relative phase. Our own testing of this method (using simulated data drawn from randomly sampled distributions of phase noise, amplitude noise, etc.) indicated that the reported fit uncertainty in $\Delta\phi$ depends only on the number of points sampled on each ellipse, and the specified uncertainty of each point (*i.e.* the uncertainty in each measured F_z, typically 0.4%). The uncertainty in $\Delta\phi$ did *not* depend on the scatter of data (*i.e.* the width of the sample distributions) and as such is typically an underestimate of the error in the fit, although for uniformly sampled ellipses it is expected to be valid. More details on this can be found in Section 7.6.5 of Chapter 7.

6.5.1 Common-mode rejection

We can quantify the common-mode rejection of the gradiometer by observing the onset of phase noise with Ramsey time for each interferometer. For short Ramsey times, perturbations to the Larmor frequency from shot-to-shot magnetic field fluctuations – 'Larmor noise', $\sigma(\omega_L)$ – result in small deviations to the fringe phase, $\delta\phi = \sigma(\omega_L)T$ and in turn the measured spin projection F_z. The scatter of many repeated measurements of F_z at the mid-fringe point could thus be used to infer the magnitude of the Larmor noise. At longer Ramsey times though, the small Larmor noise is amplified, and can result in more than π of fringe phase shift ($\sigma(F_z) = 1$). The uncertainty in F_z should thus increase linearly with Ramsey time before saturating to $\sigma(F_z) = 1$. The true phase noise keeps increasing linearly with interrogation time, however, a consequence of circular statistics [183].

To determine the amplitude of the common-mode noise, we measure the standard deviation of F_z after a Ramsey sequence for one of the BECs from 15 repeated shots at two different phases (separated by 90°)

$$\bar{\sigma}(F_z) = \sqrt{\sigma(F_z(\phi))^2 + \sigma(F_z(\phi + 90°))^2} \tag{6.16}$$

and use this to determine the underlying 'mid-fringe' standard deviation.

Measuring in quadrature ensures that we retain sensitivity to phase noise even if one of the interferometers is at the fringe extrema, where for short Ramsey times δF_z is smaller. Figure 6.4 shows the results of this measurement as well as the fitted phase noise derived using numerical simulations of the quadrature noise sampling. The fitted function $\sigma(F_z) = f(\sigma(\omega_L)T)$ is an interpolation of phase noise *estimates* from 15 samples at each $\pi/2$-pulse phase (30 per Ramsey time) as a function of numerically simulated phase noise, $\sigma(\omega_L)T$. The Larmor noise $\sigma(\omega_L)$ is the only free parameter used to fit the numerical model to the data. After only 1 ms the phase noise has almost reached π rad, indicating the common-mode magnetic noise is significant, and that direct fitting of the phase domain fringes is impossible. The fitted Larmor frequency noise is equal to $\sigma(\omega_L) = 2\pi \times 191(13)$ Hz (or $\delta B = 271(18)$ µG), and at $T = 3$ ms corresponds to a phase uncertainty of ~ 3.6 rad.

The uncertainty in relative phase, which measures noise that is not common-mode in the two condensates is extracted from typical ellipse fits and determined primarily by the uncertainty in F_z. From repeated Stern-Gerlach absorption images of a condensate immediately after a $\pi/2$-pulse we are able to determine $\sigma(F_z) \approx 0.007$, which is ~ 3 times the standard quantum limit for 1×10^5 $F = 1$ atoms. Since the atomic shot noise of each interferometer is uncorrelated, the relative phase uncertainty is approximately $\sigma(\Delta\phi) \approx 2 \times 7$ mrad, and the common mode rejection ratio (CMRR) for our gradiometer operating at $T = 3$ ms is $20\log_{10}\frac{3.6}{2\times 7\,\text{mrad}} \sim 50$ dB.[5]

6.6 Measuring magnetic field gradients

Using Equation 6.5, it is enough to know the relative phase $\Delta\phi$ at a given Ramsey time and the separation baseline ($\Delta x'$ or $\Delta z'$) to determine the magnitude of the gradient. A magnetic field biases the gradiometer to be sensitive to only certain vector components of the field gradient: we use the components of the background magnetic field as the bias, cancelling the field in the two other directions with applied fields. We will discuss the implications of this more in Section 6.6.2.

Although the relative phase linearly increases with Ramsey time, aliasing of the measured phase must be taken into account. For this reason, we make a relative phase measurement at five different Ramsey times to determine the true phase. Figure 6.5 shows a measurement of a gradient dB_y/dz' at Ramsey times of 0.1 to 3 ms as well as the corresponding phase-unwrapping process,

---

[5]Since the phases considered in this calculation are effectively amplitudes of the magnetic field, the CMRR is given by $20\log_{10}\frac{B_{\text{c-m}}}{B_{\text{diff.}}}$, with $B_{\text{c-m}}$ and $B_{\text{diff.}}$ the common-mode and differential magnetic field noise.
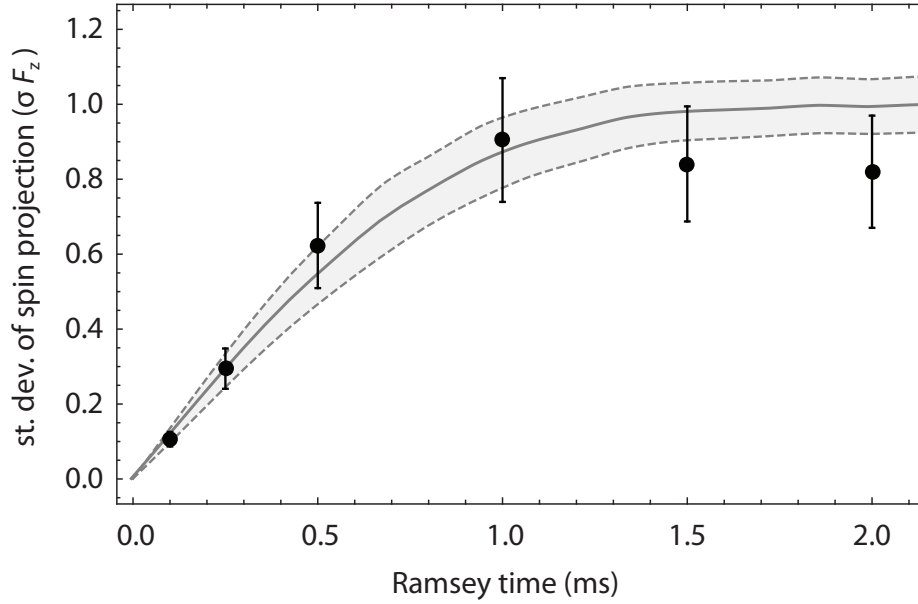
Figure 6.4: Measurement of phase noise onset in each interferometer with Ramsey time. Each data point corresponds to $\bar{\sigma}(F_z) = \sqrt{\sigma(F_z(\phi))^2 + \sigma(F_z(\phi + 90°))^2}$ from 15 repeated shots at two different phases separated by $90°$. Error bars are calculated from $\delta\bar{\sigma}(F_z) = \sigma(F_z)/\sqrt{30}$. Overlaid is a fitted interpolation function derived from numerical simulations of the quadrature phase sampling technique (dashed lines indicate one standard deviation from the mean of the fit line relating the spin projection fluctuations to the interrogation time via $\bar{\sigma}(F_z)) = f(\sigma(\omega_L)T)$. The discrepancy between the numerical model and the data at longer Ramsey times is due to quadratic Zeeman induced loss of contrast underestimating the maximum values of F_z in a measurement: all points have been approximately corrected based on the QZE contrast, which limits the maximum range of $\sigma(F_z)$ at longer Ramsey times. The fitted phase noise is $\sigma(\omega_L) = 2\pi \times 191(13)$ Hz.

which resolves the apparent similarity of the ellipses at $T = 0.1$ and $3\,\text{ms}$. The transformation between primed and un-primed coordinate systems is described in Section 6.6.3. The measured gradient is $-5.33(3)\,\text{mG/cm}$: we can determine the sign of the gradient by looking at phase domain fringes at short interrogation times: faster fringes correspond to higher amplitude of B, so one fringe leads the other (Figure 6.3, left).

We determine the magnetic field gradient from the slope of the linear fit in plots such as that in Figure 6.5 using Eq. 6.9. Doing so aids the measurement of the field gradient and assists the phase unwrapping, but also includes contributions from systematic errors due to the increasing drop time. We determine a nominal precision of $30\,\mu\text{G/cm}$ in the measured gradient from the uncertainty in the slope $\Delta B = \frac{1}{\gamma}\frac{d\Delta\phi}{dT}$ reported from such a fit. This
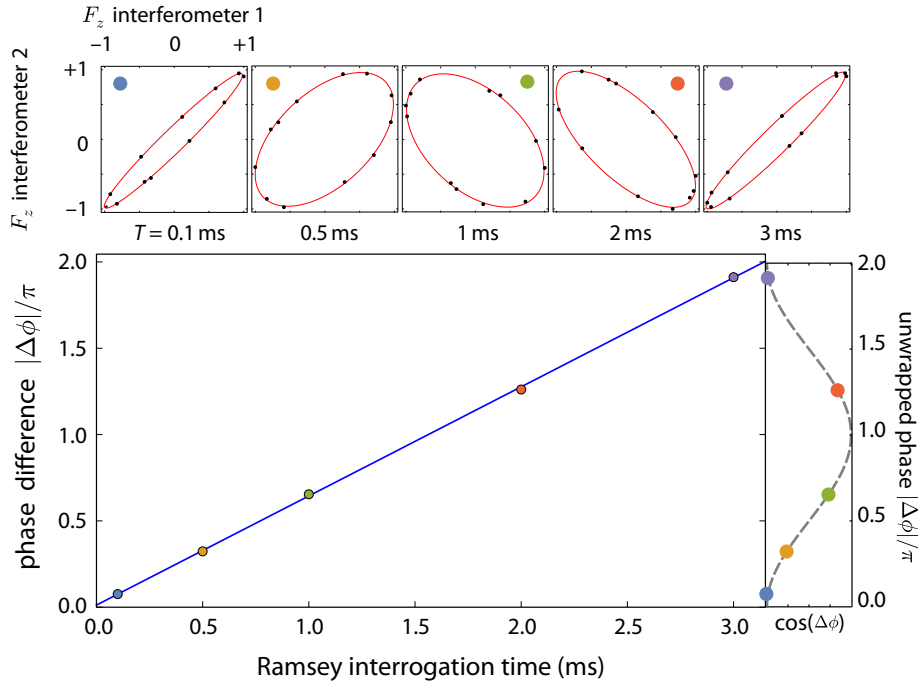
Figure 6.5: Measurement of magnetic field gradients. The relative phase $\Delta\phi$ is extracted from each ellipse, and plotted as a function of Ramsey time. Varying the Ramsey interrogation time accounts for aliasing – each ellipse is marked with a coloured pointer to indicate the phase-unwrapped value shown on the right of the plot.

method requires 60 iterations of the experiment to measure a single gradient. However, we stress that once the approximate un-aliased relative phase is known, substantially better precision is possible by fixing the interrogation time at the longest possible duration and measuring gradients at higher duty cycles. We discuss this more in Section 6.8.

6.6.1   Applied gradients: tensor sensitivity

The tensor sensitivity of our gradiometer hinges on two features. These are the ability to selectively measure field gradients by applying bias fields that render the gradiometer sensitive to gradients of one only field component and the ability to span baselines to discern gradients across a particular spatial direction. Due to the dipole beam axes being 45° rotated from the magnetic field axes, we are required to measure two gradients in the x', z' basis to measure a single x, y, z gradient. Ensuring that only single field component gradients are measured amounts to ensuring that only that one field component is large compared to the other two.

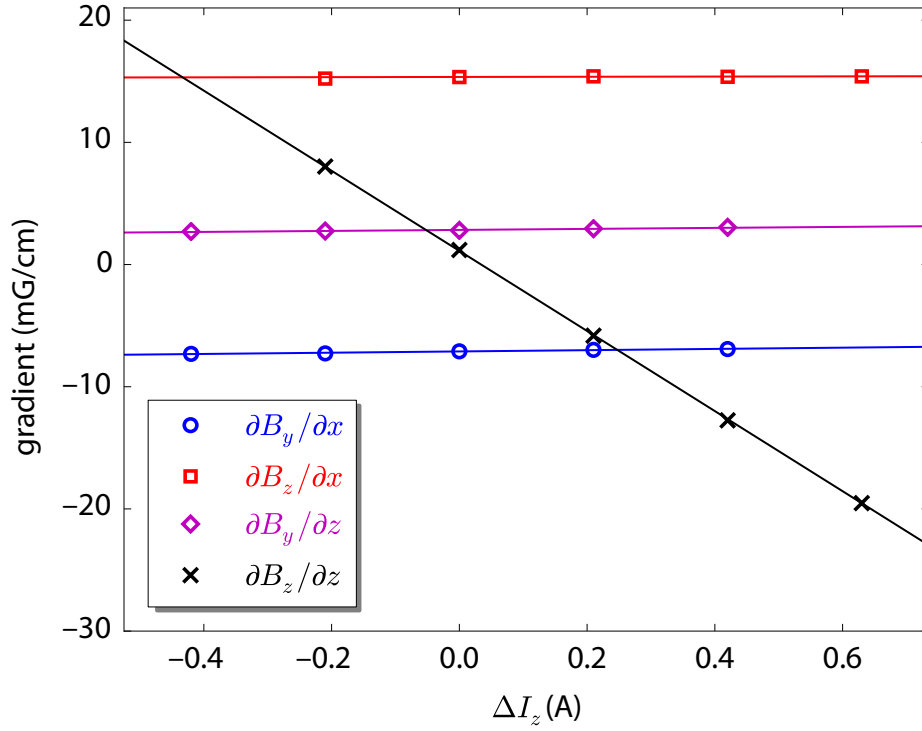Let us take a concrete example. The gradiometer is spanning a baseline

Figure 6.6: Measurement of an applied gradient, created by driving a current imbalance across the z-bias coils. We measure gradients with two different bias fields: a z-bias makes the gradiometer maximally sensitive to gradients $\partial B_z/\partial x_i$, including the applied gradient, and a y-bias that renders it insensitive to the applied gradient. The observed insensitivity of the $y-$biased measurements to the applied gradient quantifies the precision with which we can bias the gradiometer along that direction.

x_i and field-biased along y, rendering it sensitive to only $\partial B_y/\partial x_i$ gradients. A large $\partial B_z/\partial z$ gradient is applied by driving a current imbalance ΔI across the z-bias coils, no change should be detected in the measured field gradients with a y-bias. Driving a current difference across the z-coils creates primarily a $\partial B_z/\partial z$ gradient, as well as two other diagonal terms that satisfy Gauss's law $\nabla \cdot \boldsymbol{B} = 0$: all of the applied gradients can be eliminated from the measurement by y-bias choice (eliminating $\partial B_x/\partial x$, $\partial B_z/\partial z$) and baseline configuration (eliminating $\partial B_y/\partial y$). Figure 6.6 shows the result of this measurement, with a $y-$bias and a $z-$bias applied to the gradiometer in the presence of an applied gradient measured to be $\partial B_z/\partial z = 32.8(1) \times \Delta I$ mG/cm. The insensitivity of the other measured gradients ($< 0.51(6)$ mG/cm) indicates that biasing the gradiometer with magnetic fields renders tensor sensitivity to gradients, as the coils are reasonably well aligned to their respective axes.

6.6.2 Coil-induced gradients

The bias coils themselves produce magnetic field gradients if the condensates are not positioned about the geometric centre of the coils. To measure the true background gradient tensor, these coil-induced gradients must be considered. Numerical magnetostatic calculations of the of the field profile from each coil pair as a function of current and displacement from geometric centre were performed by R. P. Anderson using the *Mathematica* package *Radia*. The gradients induced by each coil pair can be conveniently expressed as individual tensors and depend on coil current $I_{x,y,z}$ and displacement $(\delta x, \delta y, \delta z)$ from the geometric centre of the coils:

$$G_{\mathrm{x-coils}} = \begin{pmatrix} 1.88\,\delta x & -1.74\,\delta y & -0.14\,\delta z \\ -1.74\,\delta y & -1.74\,\delta x & 0 \\ -0.14\,\delta z & 0 & -0.14\,\delta x \end{pmatrix} I_x \qquad (6.17a)$$

$$G_{\mathrm{y-coils}} = \begin{pmatrix} -1.11\,\delta y & -1.11\,\delta x & 0 \\ -1.11\,\delta x & 1.07\,\delta y & 0.04\,\delta z \\ 0 & 0.04\,\delta z & 0.04\,\delta y \end{pmatrix} I_y \qquad (6.17b)$$

$$G_{\mathrm{z-coils}} = \begin{pmatrix} -0.77\,\delta z & 0 & -0.77\,\delta x \\ 0 & -0.63\,\delta z & -0.63\,\delta y \\ -0.77\,\delta x & -0.63\,\delta y & 1.40\,\delta z \end{pmatrix} I_z\,, \qquad (6.17c)$$

where currents are in A, gradients are in mG/cm and displacements are in mm. In general, we can minimise the contribution of coil-induced gradients by using the background magnetic field to bias the gradiometer. Bias fields are applied that approximately cancel the field in the two other directions, leaving only a single vector component of the background magnetic field– the current is zero in the bias coils in the direction used to bias the gradiometer (B_y, for example). This makes $G_{\mathrm{y\text{-}coils}} = 0$, and eliminates these current-induced gradients. Alternatively, we could apply a sufficiently large bias field such that $|B_y| \gg |B_x|, |B_z|$ but to sufficiently overcome the orthogonal background fields, $I_y > 6\,\mathrm{A}$ would be needed, making coil-induced gradients significant. All components of the background field are $> 300\,\mathrm{mG}$, and we are able to null the background fields to lower than $\sim 5\,\mathrm{mG}$ in any given direction.

For displacements from geometric centre of $< 1\,\mathrm{mm}$, all coil induced gradients (Eq. 6.17) are $< 1\,\mathrm{mG/cm}$. When we measure components of the gradient tensor, we measure both the background gradient terms as well as the coil-induced gradients, *i.e.* $G_{\mathrm{tot}} = G_{\mathrm{bg}} + G_{\mathrm{coils}}$. Specifically, for the case of a measurement of $\partial B_y / \partial x$, we measure

$$\left(\frac{\partial B_y}{\partial x}\right)_{\mathrm{meas}} = \left(\frac{\partial B_y}{\partial x}\right)_{\mathrm{bg}} + \left(\frac{\partial B_y}{\partial x}\right)_{\mathrm{x-coils}} + \left(\frac{\partial B_y}{\partial x}\right)_{\mathrm{z-coils}}. \qquad (6.18)$$
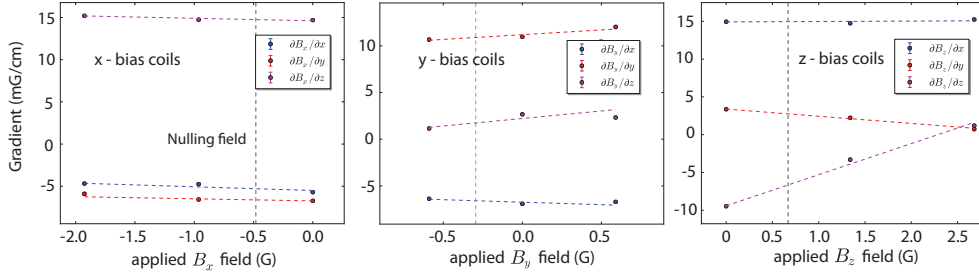
Figure 6.7: We can measure the coil-induced magnetic field gradients by applying different bias fields, corresponding to different coil currents. The variation with applied field is a measure of the coil induced magnetic field gradient, which is principally due to misalignment of the gradiometer interrogation region from the geometric centre of the coil pairs. Dashed lines denote the applied field which cancels the background field. The largest of the coil induced gradients is $\partial B_z/\partial z$. As is seen from Equation 6.17 this suggests that the primary misalignment from geometric centre is along z. In this Figure, the measurements for gradient terms $\partial B_i/\partial y$ were made using the technique described in Section 6.6.4, and not by inference from the symmetry of the gradient tensor.

As can be seen from the tensors in Eq. 6.17, one of the coil induced gradient terms always vanishes for non-diagonal gradients: the z-coils do not produce a $\partial B_y/\partial x$ in this case.

We quantified the effects of coil induced gradients by measuring gradients with bias fields corresponding to the background field $|B_i|\,\hat{x}_i$ ($I = 0$), and applied fields $-|B_i|\,\hat{x}_i$ and $-1.5|B_i|\,\hat{x}_i$ ($I = 2\,I_N$ and $I = 3I_N$ respectively, I_N the current required to null the background field). The measurements are shown in Figure 6.7, from which we see that the z-coils exhibit a significant current dependent gradient, principally $\partial B_z/\partial z$; primarily due to the larger currents needed (coil calibration is $0.38\,\text{G/A}$). Our dominant misalignment from geometric centre may also be along the z-axis, which is plausible as the z-coils are the smallest coils separated by the greatest distance. The contribution from coil-induced gradients is small enough not to hamper a reasonably accurate determination of the background tensor, but may play a role in explaining several minor discrepancies, discussed in the next Section.

6.6.3 Measurement of the field tensor

We measure derivatives of magnetic fields (B_x, B_y, B_z) in the (x', y, z') basis. In order to express the gradient tensor G_{ij} in the (x, y, z) basis, we must transform the measured field gradients from the form $\partial B_i/\partial x'_j$ to $\partial B_i/\partial x_j$, where $x_i \in \{x, y, z\}$ and $x'_j \in \{x', y, z'\}$. Referring to the coordinate transformation (Eq. 6.11), we can construct the transformation matrix $S_{ij} = \partial x_i/\partial x'_j$, so that the transformation of a vector from one basis to the other may be written

as $x_i = S_{ij} x'_j$ and $x'_i = A_{ij} x_j$, with $A_{ij} = (S_{ij})^{-1}$. Derivatives transform as follows

$$\frac{\partial B_i}{\partial x_j} = \frac{\partial x'_k}{\partial x_j} \frac{\partial B_i}{\partial x'_k} \tag{6.19}$$

$$= (A^T)_{jk} \frac{\partial B_i}{\partial x'_k} \tag{6.20}$$

$$= T_{jk} \frac{\partial B_i}{\partial x'_k}, \tag{6.21}$$

with the transformation matrix given by

$$T = \begin{pmatrix} \cos\theta_{z'z} \csc\theta_{x'z'} & 0 & -\cos\theta_{x'z} \csc\theta_{x'z'} \\ 0 & 1 & 0 \\ \sin\theta_{z'z} \csc\theta_{x'z'} & 0 & \sin\theta_{x'z} \csc\theta_{x'z'} \end{pmatrix} \tag{6.22}$$

with $\theta_{x'z'} = \theta_{x',z} + \theta_{z',z} \sim 90°$.

Maxwell's laws for magnetic fields in a vacuum are used to complete the measurement of the gradient tensor. Gauss's law for magnetic fields ($\nabla \cdot \boldsymbol{B} = 0$) and Ampere's law ($\nabla \times \boldsymbol{B} = 0$) require the gradient tensor to be traceless and symmetric respectively. We can infer the three gradient terms we cannot measure with x'-z' baselines, namely $\partial B_i / \partial y$, from measurements with a y-bias field. The measured background gradient tensor is

$$G = \begin{pmatrix} -5.71(7) & -6.92(4) & 14.70(7) \\ -6.92(4) & 15.18(8) & 2.66(4) \\ 14.95(3) & 2.66(4) & -9.47(3) \end{pmatrix} \mathrm{mG\,cm^{-1}}, \tag{6.23}$$

where the inferred gradients have been outlined in the dashed region. Our measurements additionally give us a direct measure of the systematic uncertainty inherent to the technique by independently quantifying two of the off-diagonal matrix elements that should be equal by Maxwell's laws: $\partial B_x / \partial z$ and $\partial B_z / \partial x$, which disagree by $0.25\,\mathrm{mG/cm}$, which should be compared to the statistical uncertainty of $\sim 0.08\,\mathrm{mG/cm}$. As discussed in Section 6.6.2, the possible causes of such a discrepancy may include coil induced gradients, which are of a comparable magnitude. As will be discussed in Section 6.6.4, we performed a consistency check for the inferred values by directly measuring gradients along a y-oriented baseline, with mixed results.

6.6.4 A consistency check

We also measured gradients using a baseline primarily along gravity. Using single-axis AOM deflection limits the available gradiometer baselines to the x-z plane in our experiments. To spatially separate two BECs along gravity

(the y-direction), we simply create two BECs with a small separation along x' or z', drop one, and release the second once the first has fallen a suitable distance. The complete scheme is shown in Figure 6.8. Typically, we use the z' axis for the initial horizontal baseline; the two BECs are separated by $\sim 100\,\mu\text{m}$. We then drop the positive z' BEC first, denoted as BEC_1, by extinguishing one of the dipole beams, the single remaining beam is insufficient to overcome gravity and so BEC_1 falls. We delay the drop of BEC_2 by $T_{\text{freefall}} = 10\,\text{ms}$ to allow BEC_1 enough time to reach a distance comparable with the other baselines ($> 500\,\mu\text{m}$). Additionally, the residual trapping force from the single remaining dipole beam distorts BEC_1, spatially extending the cloud along the dipole beam. This limits the maximum drop delay time, and hence baseline distance, due to the need to obtain reasonable Stern-Gerlach separation for absorption imaging. We define the separation axis between the two BECs as $\hat{\mathbf{y}}'$, with the position along this axis given by $\alpha \hat{\mathbf{z}}' + \beta(t)\hat{\mathbf{y}}$, with $\alpha \approx 70\mu\text{m}$, $\beta \approx 536\,\mu\text{m}$ at $t = 10\,\text{ms}$. Once the two BECs have reached a separation $\Delta y' = 545\,\mu\text{m}$, we drop BEC_2 and perform freefall DRI.

There are further complications that must be considered. The sensor volume increases, and *asymmetric* displacement from bias coil geometric centre (in contrast to x' and z' separations) results in a greater contribution from coil-induced gradients. The falling BEC_1 is by no means stationary during the Ramsey interrogation, either. In comparison to the other freefall experiments, it falls *relative* to BEC_2; over the course of a $3\,\text{ms}$ Ramsey interrogation, BEC_1 will fall a further $300\,\mu\text{m}$. This means that as the Ramsey time increases, the vertical baseline does as well. We accommodate this difficulty by altering the timing of the experiment so that BEC_2 is dropped earlier or later depending on the Ramsey interrogation time, which ensures that the average y separation is constant. Specifically, we drop BEC_2 $t = 100\,\mu\text{s} + \frac{1}{2}\left(T - T_{\text{max}}\right)$ before the first $\pi/2$-pulse, with $T_{\text{max}} = 3\,\text{ms}$ the maximum Ramsey time. The vertical separations for two different Ramsey times are shown in Figure 6.9: the average separation vector remains constant.

We measured gradients using all three bias fields and spatial baselines along y' in order to compute the gradient tensor terms $\partial B_i / \partial y$. The coordinate transform matrix becomes rather more complicated, the y' unit vector is written as

$$\hat{\mathbf{y}}' = -\tan\theta_{z'z}\frac{\Delta z}{\Delta y'}\,\hat{\mathbf{x}} + \frac{\Delta y}{\Delta y'}\,\hat{\mathbf{y}} + \frac{\Delta z}{\Delta y'}\,\hat{\mathbf{z}}, \qquad (6.24)$$

with $\Delta y'^2 = \Delta y^2 + (1 + \tan^2\theta_{z'z})\Delta z^2$. The transformation matrix, in analogy with Eq. 6.22, for transforming gradients from the x', y', z' basis to the x, y, z
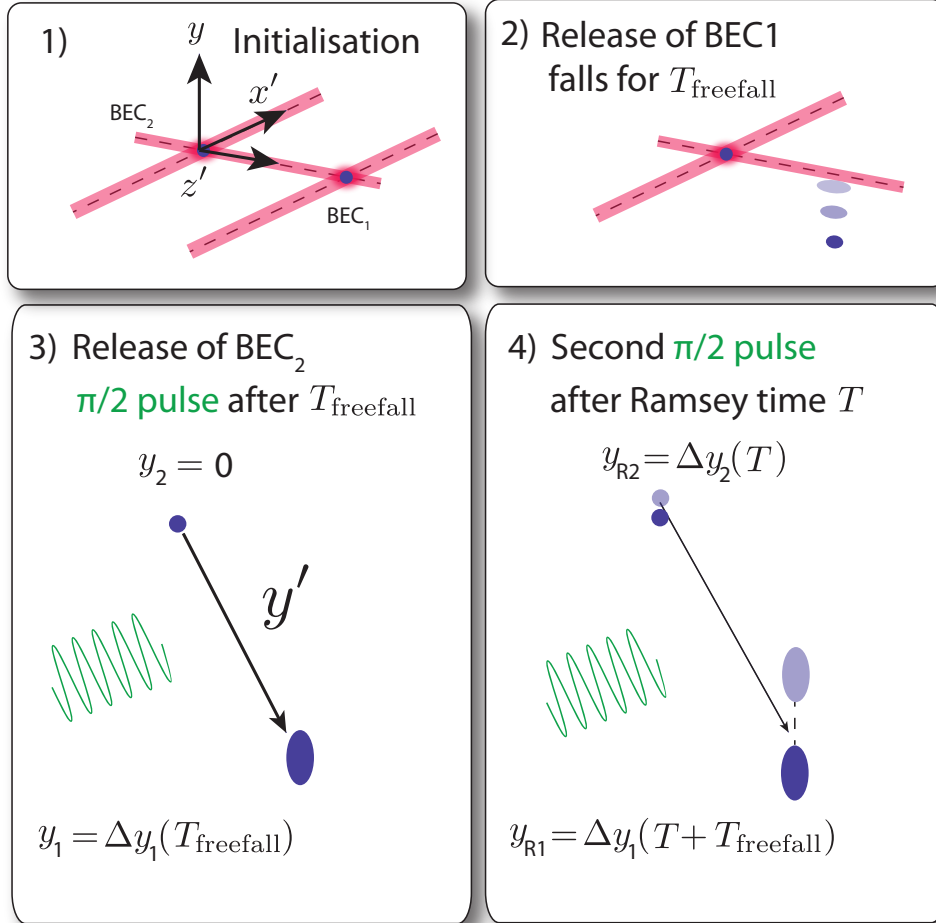
Figure 6.8: Measuring gradients along a vertical baseline: 1) Two BECs are initialised along a z' baseline of $\sim 70\,\mu$m. 2) We drop one of them by extinguishing the split beam that confines it, it then expands slightly along the length of the beam before falling out. 3) The Ramsey sequence is initiated after dropping BEC_2 for T_{freefall}. 4) At the conclusion of the Ramsey pulse sequence, the two BECs have fallen a further distance. Figure 6.9 describes how we account for this effect to ensure the average y baseline remains constant.
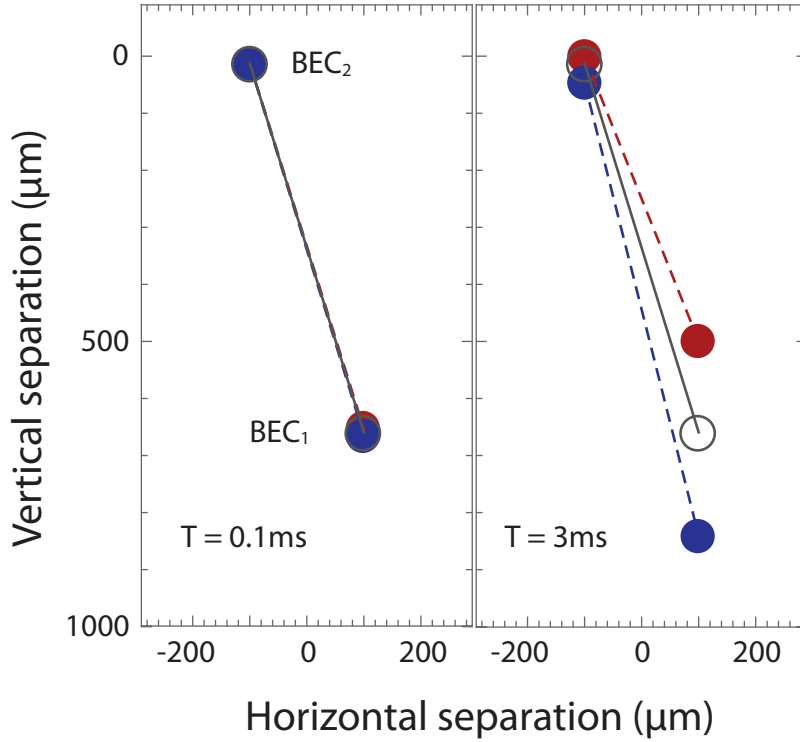
Figure 6.9: The y'-baseline for two different Ramsey times, with BEC_2 dropped as a function of Ramsey time before the first $\pi/2$-pulse. Red and blue colours indicate positions of the two falling BECs at the beginning and end of the Ramsey sequence, respectively. The average y' baseline (grey) is constant for the two different Ramsey times.

basis now takes the form

$$
T = \begin{pmatrix}
\cos\theta_{z'z}\csc\theta_{x'z'} & 0 & -\cos\theta_{x'z}\csc\theta_{x'z'} \\
0 & \Delta y'/\Delta y & -\Delta z\sec\theta_{z'z}/\Delta y \\
\sin\theta_{z'z}\csc\theta_{x'z'} & 0 & \sin\theta_{x'z}\csc\theta_{x'z'}
\end{pmatrix} \tag{6.25}
$$

and the gradient tensor is

$$
G = \begin{pmatrix}
-5.71(7) & -6.68(9) & 14.70(7) \\
-6.92(4) & 10.97(5) & 2.66(4) \\
14.95(3) & 3.37(2) & -9.47(3)
\end{pmatrix} \text{mG/cm.} \tag{6.26}
$$

We immediately notice that $\text{Tr}(G) = -4.21(9)\,\text{mG/cm} \neq 0$, and that in contrast to the symmetric terms directly measured using in-plane gradiometry $\left(\frac{\partial B_x}{\partial z}, \frac{\partial B_z}{\partial x}\right)$, the gradients $\frac{\partial B_y}{\partial z}$ and $\frac{\partial B_z}{\partial y}$ differ by almost $1\,\text{mG/cm}$. In order to resolve this, we performed a number of variations of the experiment, specifically focused on measuring $\partial B_y/\partial y$. These included:

- Measuring the diagonal gradient $\partial B_y / \partial y'$ across a smaller vertical baseline: $\partial B_y / \partial y = 11.51(9)\,\text{mG/cm}$

- Using an x' initial splitting rather than z': $\partial B_y / \partial y = 9.78(2)\,\text{mG/cm}$

- Dropping BEC_2 first (*i.e.* on the $-z'$ side) and comparing with dropping BEC_1 first; this changes the direction the y' axis points: $\partial B_y / \partial y = 10.99(9)\,\text{mG/cm}$

All of these measurements are remarkably consistent given the different configurations, and cannot account for the observed $\text{Tr}(G) \neq 0$. We are thus lead to conclude that the technique is sound, and that the large $\sim 500\,\mu\text{m}$ displacement from geometric centre of the coils may result in sampling significant coil-induced gradients that are largely absent for the x'-z' in-plane measurements. We could modify the in-plane measurements so that they interrogate the same spatial region as the y' measurements do, *i.e.* by dropping both BECs at the maximum x' or z' splitting so that the x'-z' plane is now at the centre point of the $\Delta y'$ baseline. However, the comparative simplicity of using in-plane measurements to infer y-gradients is certainly preferable to any of the complicated schemes used to span all spatial axes. Indeed, we have no reason to believe that even in the presence of such systematics as small coil induced gradients the inference is at all invalid. We draw attention to the observed near equality of the off-diagonal terms to within $0.25\,\text{mG/cm}$ in Eq. 6.23, which provides a reasonable estimate of the magnitude of systematic error encountered in in-plane measurements.

6.7 Applications of the gradient tensor

Here, we make note of a few specific applications that motivate measuring the gradient tensor. Magnetometers in various forms have long held interest as devices for remote sensing: detecting buried ore bodies and unexploded ordinance are a few of the popular applications. As such, magnetometers can be operated in a number of configurations, which in general correspond to measurements of the scalar, vector and tensor characteristics of the magnetic field. A scalar magnetometer is used to produce maps of $|B|$, the 'total magnetic field intensity' (TMI), whereas a vector magnetometer is able to resolve vector components of the magnetic field. Both scalar and vector magnetometers can be operated in gradiometer configurations, which are more informative and robust to environmental noise.

Measuring the full gradient tensor has several advantages. In geophysical applications, where the sensor is large and often towed behind aircraft, tensor gradiometers are immune to orientation errors, unlike vector component

magnetometers [171,173]. The real power of tensor gradiometry lies in the ability to localise and characterise magnetic dipole sources. The eigenvectors of the gradient tensor are mutually orthogonal and correspond to a coordinate system where all off-diagonal gradient terms vanish. The eigenvector corresponding to the largest magnitude eigenvalue points at the dipole moment of the magnetic source [171,173,184]. When we perform this analysis on our own measured tensor, we find that the corresponding eigenvector points within $20°$ of the large unshielded magnets of the Perkin-Elmer ion pump attached to the oven end of the apparatus. Since our gradients originate from multiple sources (there are two other ion pumps), the discrepancy is not really surprising; for multiple sources the principal eigenvector points at the equivalent point source.

6.7.1 Minimising gradients with bias direction alone

As a consequence of applying different bias fields to measure certain gradients, we discovered that the net gradient can be reduced substantially by changing the direction of the bias field. A scalar figure of merit that quantifies the magnitude of the field inhomogeneity is the norm of ∇B, termed the 'total gradient'. From Eq. 6.4, the total gradient is given by

$$|\nabla B| = \frac{|G \cdot B|}{B}. \tag{6.27}$$

The BEC itself is spatially anisotropic, meaning that gradients along certain directions (such as the longest elongation direction) will be more important than others in determining the dephasing time of the spinor BEC. To encapsulate the anisotropic spatial extent of the BEC, we can weight the components of ∇B according to the measured Thomas-Fermi radii (Table 3.1) before taking the norm using

$$|\nabla \tilde{B}| = \left| \left(\mathcal{R}_y \left(\frac{\pi}{4} \right) \frac{G \cdot B}{B} \right) * \frac{R_{TF}}{|R_{TF}|} \right|, \tag{6.28}$$

where $R_{\text{TF}} = (R_{x'}, R_y, R_{z'})$ are the Thomas-Fermi radii (Eq. 3.16) and $*$ denotes element-wise multiplication.[6] The Thomas-Fermi weighted gradient gives an indication of the coherence time of the BEC in the presence of certain gradients. For instance, a large $\partial B / \partial x'$ will not cause fast dephasing for a 1D-BEC elongated along the z'-axis, due to the small spatial extent in the x' direction (Equation 5.7).

---

[6]Note that we have rotated (about y) the product $G \cdot B$ into the x', z' basis using the rotation matrix $\mathcal{R}_y$ before weighting by the Thomas-Fermi radii, as this basis defines the axes of symmetry of the condensate.
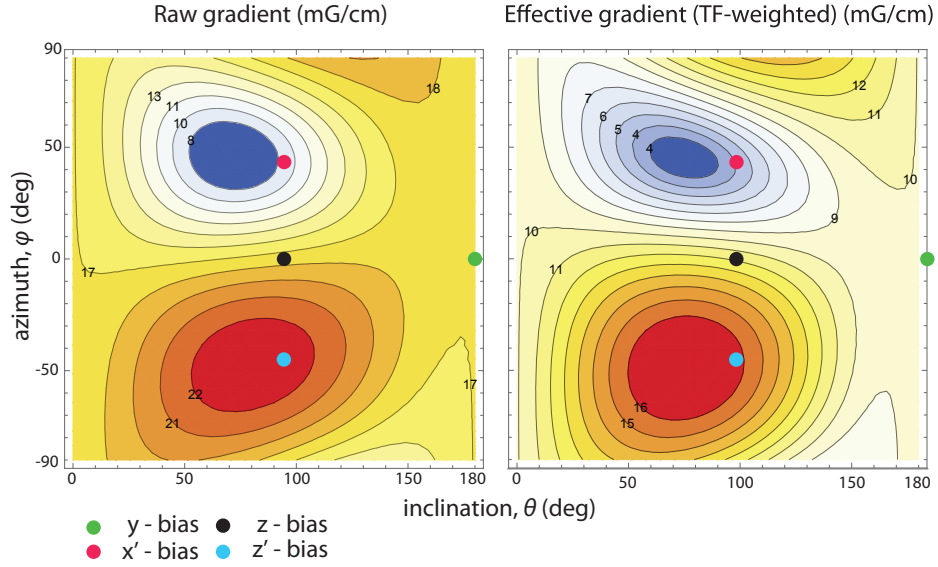
Figure 6.10: Left: Magnitude of total gradient $|\nabla B|$ (Eq. 6.27), and right: weighted effective gradient $|\nabla \tilde{B}|$ (Eq. 6.28) as a function of bias field direction for our measured gradient tensor in Eq. 6.23. The coloured markers denote bias fields routinely used, note that the typical x'-bias gradient is on the order of $8\,\text{mG/cm}$ whereas a z'-bias gradient is over $23\,\text{mG/cm}$. The effective (Thomas-Fermi weighted) gradient can be reduced to as little as $2\,\text{mG/cm}$.

We evaluate the total gradient (Eq. 6.27) and Thomas-Fermi weighted effective gradient (Eq. 6.28) as a function of the bias field direction by expressing the magnetic bias field vector in spherical coordinates:

$$\boldsymbol{B} = B\left(\cos\phi\sin\theta, \cos\theta, \sin\theta\sin\phi\right), \tag{6.29}$$

with θ the inclination angle ($\theta = 0, \boldsymbol{B}/B = +\hat{\mathbf{y}}$) and ϕ the azimuthal angle, measured from the z axis ($\theta = -45°, \boldsymbol{B}/B = \hat{\mathbf{z}}'$). The two metrics are plotted as a function of θ and ϕ in Figure 6.10 for the measured magnetic field gradient tensor, Eq. 6.23.

Minimising the gradients with a particular choice of bias direction alone allows us to substantially reduce the background gradient from that typically present at the usual orthogonal lab-frame bias fields B_x, B_y and B_z. Simply aligning the field along x' reduces the total gradient to $< 8\,\text{mG/cm}$; later in Chapter 7 we show how this allows observation of coherent spin-mixing for $> 500\,\text{ms}$, in contrast to most other experiments that had $|\nabla B| \sim 15$-$20\,\text{mG/cm}$. A future work will consider the prospect of *applying* gradients with the bias coils (*i.e.* imbalanced currents in each coil) and changing the direction of $\boldsymbol{B}$ to actually obtain $|\nabla B| \approx 0$. Initial theoretical efforts in this direction have shown that it is possible with two *diagonal* gradients from the

y and *z*-coil pairs applied and a bias field with varying direction to make $|\nabla B| = 0$. To make $G = 0$ would require two coils aligned along the principal eigenaxis of *G*, an elegant but technically awkward approach.

6.8 Operational and prospective sensitivity

We now discuss the operational sensitivity of our gradiometer, as well as the potential sensitivity that could be achieved with modest improvements. The fundamental sensitivity of an atomic magnetometer amounts to how well the Larmor precession frequency can be measured. The standard quantum limited precision of a measurement of the atomic spin is given by

$$(\delta F_z)_{\text{SQL}} = \sqrt{\frac{F}{2N}} \qquad (6.30)$$

for *N* spin-*F* atoms [155]. The corresponding Larmor phase uncertainty for $F = 1$ atoms is

$$(\delta\phi)_{\text{SQL}} = \sqrt{\frac{1}{2N}}. \qquad (6.31)$$

The quantum limited magnetic field sensitivity of a sample with *N* spin-1 atoms used in a projective measurement of atomic phase of duration *T* is thus

$$\delta B_{\text{SQL}} = \frac{1}{\gamma} \frac{\delta\phi_{\text{SQL}}}{\sqrt{N}T}. \qquad (6.32)$$

We can improve the precision by making repeated measurements; *M* repeated measurements increases the signal-to-noise by a factor $1/\sqrt{M}$. In our specific case, multiple experimental shots are required to impute a differential phase from the elliptical data reduction, a single-shot phase sensitivity is ill-defined. However, the differential phase uncertainty from fitting an ellipse with *M* points[7] scales with $1/\sqrt{M}$, and thus the quantum limited field sensitivity is

$$\delta B_{\text{SQL}} \sim 1/\left(\gamma\sqrt{NTDT_{\text{int}}}\right) \qquad (6.33)$$

for $N/2$ atoms per condensate, a total integration time $T_{\text{int}} = MT_{\text{shot}}$, a duty cycle $D = T/T_{\text{shot}}$, and a single-shot duration of T_{shot}, the time for the total duration of an experiment cycle, in our case the $\sim 25\,\text{s}$ it takes to make BEC. The units of the *sensitivity per unit bandwidth* $\delta B\sqrt{T_{\text{int}}}$ are tesla per root Hertz, $\text{T}\,\text{Hz}^{-1/2}$, and have a simple physical interpretation that belies the unusual (for a non-specialist) units: if one makes a single measurement of a *B* per second, and the sensitivity is δB_0, taking *M* measurements per second improves this to $\delta B_0/\sqrt{M}$, this metric incorporates the highly important

---

[7]As determined from Monte-Carlo simulations of the data analysis procedure. See also the discussion in Section 7.6.5.

nature of measurement dead time and cycle duration, which can be quite large for a BEC experiment.

Since our measurement does not use a dispersive optical probe and has negligible spin relaxation, it is atomic spin-projection noise that limits the fundamental sensitivity of the technique. From repeated absorption images we determine that the uncertainty in our measurements of spin projection δF_z are ~ 3 times that of the standard quantum limit for $N = 10^5$ atom condensates, and this ultimately determines the relative phase uncertainty extracted from ellipse fits, following on the discussion in Section 6.3.

Our scheme could also function as a precision *co-magnetometer*, where one BEC placed in the vicinity of a small magnetic source acts as a sensor and the other condensate a reference interferometer, providing substantial common-mode rejection of ambient noise. This is the mode of operation we have in mind when specifying field sensitivities per unit bandwidth or per unit spatiotemporal bandwidth below, the standard metrics for magnetometry. No common, analogous metric is in use for *gradient* sensitivity per spatiotemporal bandwidth. The corresponding field sensitivity per unit bandwidth of this gradiometer – or co-magnetometer – is $\delta B \sqrt{T_{\text{int}}} = 360\,\text{pT}\,\text{Hz}^{-1/2}$ for $T_{\text{shot}} = 25\,\text{s}$ and $T = 3\,\text{ms}$.

6.8.1   A prospective gradiometer using trapped condensates

A gradiometer formed from a pair of *trapped* condensates offers a smaller sensor volume and a significant improvement in sensitivity compared to one formed from condensates in freefall. For an in-trap magnetometer to be possible, it is necessary to suppress the vector light shift (VLS) induced by residual elliptical polarisation of the trapping beams. The vector light shift may be reduced several orders of magnitude by ensuring the dipole trapping beams are linearly polarised with conventional *ex vacuo* polarimetry, but accounting for the birefringent glass vacuum cell requires an atomic measurement [19, 185]. Chapter 7 of this thesis describes exactly such a measurement to eliminate vector light shifts, and discusses in more detail the spurious contribution of the VLS to gradiometry experiments.

The sensitivity of each interferometer scales with the evolution time and the atomic density. A Zeeman coherence time approaching one second was observed in a spin-1 ^{87}Rb condensate [35], limited by losses due to density-dependent collisions. For a ^{87}Rb BEC with a peak number density of $10^{14}\,\text{atoms}\,\text{cm}^{-3}$ (corresponding to $N = 10^6$ atoms for the current trap), an interrogation time of $T = 200\,\text{ms}$ is foreseeable, an order of magnitude lower than the three-body limited lifetime. Recently, experiments with a dispersive Faraday probe beam by M. Jasperse have shown coherence times on the order

of several seconds [95]. The prospective in-trap magnetometer could thus achieve a differential field sensitivity per unit bandwidth of $600\,\mathrm{fT\,Hz^{-1/2}}$ at the standard quantum limit, even with the same trap and single-shot duration used here (corresponding to a non-unity duty cycle of $D = 0.008$). The duty cycle of BEC experiments is often lamented, but is improving and newer generations of experiments offer higher duty cycles: all-optical experiments can reduce cycle times to a few seconds [118].

6.8.2 Spatial resolution

Spatial resolution is conventionally quantified by the sensing volume. For the freefall measurements described herein, the sensor volume $V = 2 \times 10^{-5}\,\mathrm{mm^3}$ is that swept out by a falling, expanding condensate during the Ramsey interrogation. Using the metric of *spatiotemporal* sensitivity, which incorporates the sensing volume V, our demonstrated measurement has $\delta B_{\text{s-t}} = \delta B \sqrt{T_{\text{int}}} \sqrt{V} = 51\,\mathrm{fT\,cm^{3/2}\,Hz^{-1/2}}$. The above *prospective* in-trap magnetometer has $V = (20\,\mathrm{\mu m})^3$, corresponding to $\delta B \sqrt{T_{\text{int}}} \sqrt{V} = 0.05\,\mathrm{fT\,cm^{3/2}\,Hz^{-1/2}}$. This large prospective improvement is not unprecedented in microscale magnetometry where first demonstrations are far from fundamental limits; warm vapour magnetometry in microfabricated cells ($V \sim \mathrm{mm^3}$) was first demonstrated at a sensitivity of $5500\,\mathrm{fT\,cm^{3/2}\,Hz^{-1/2}}$ [186], rapidly developed to $5.0\,\mathrm{fT\,cm^{-3/2}\,Hz^{-1/2}}$ [187], and more recently $0.16\,\mathrm{fT\,cm^{3/2}\,Hz^{-1/2}}$ [31].

Another comparative metric for precision magnetometers is the magnetic field energy resolution per unit bandwidth $\epsilon = (\delta B_{\text{s-t}})^2/2\mu_0$ [31,156,188] with $\delta B_{\text{s-t}}$ the spatiotemporal sensitivity per unit bandwidth. The units of ϵ are J s; ϵ is thus often quoted in multiples of $\hbar$. The value of ϵ for the freefall measurements described here is about $10^7\,\hbar$, but with the improvements described can be substantially reduced. Compared to vapour cell magnetometers, the low duty cycle and number of spins in BEC based measurements limits their sensitivity per unit bandwidth, but the far smaller volume results in a comparable ϵ, with $\epsilon \sim 50\text{–}100\hbar$ for vapour magnetometers [31,156], and $\epsilon \sim 10\hbar$ for the prospective magnetometer described above, using the in-trap volume of a spherical BEC with Thomas Fermi radii $\sim 10\,\mathrm{\mu m}$.

A note on prospective improvements

The improvements we have detailed here to realise a high precision co-magnetometer are eminently feasible and highlight the power of the DRI technique in achieving high spatiotemporal sensitivities. After the preparation of the manuscript detailing our gradiometry experiments, we became aware of a related work performed in the group of M. Oberthaler, their experiment uses spin-squeezing enhanced Ramsey interferometry to measure

magnetic field gradients across a 1D array of 30 BECs in a lattice [189]. Many of the techniques used are familiar, including common-mode rejection from simultaneous addressing of multiple quantum gases. Here, a few of the key differences between our work and that of Ref. [189] are outlined to highlight the relative advantages and disadvantages of each technique.

The use of squeezing in Ref. [189] is an interesting and encouraging result, though its metrological advantages are not yet demonstrated. The squeezing-enhanced sensitivity is only present for Ramsey interrogation times of $< 350\,\mu s$ due to field fluctuations, which results in a magnetometric sensitivity of $1.9\,nT\,Hz^{-1/2}$. The scheme is claimed to be 'scalable': the squeezing demonstrated in ensembles of 1×10^4 atoms can be realised with $N \sim 10^7$ atoms, and interrogation times of up to $250\,ms$ are possible, giving sensitivities approaching $1pT\,Hz^{-1/2}$.

The sensitivities, both demonstrated and prospective, attained by Ref. [189] are less than that of our own work and come with the formidable experimental complexity required for squeezing. Such a squeezing scheme only operates at magnetic Feshbach resonances (a $9.1\,G$ bias field is required) making actual metrological application very limited. The proof-of-principle demonstration of sub-shot noise magnetometry is interesting and a considerable achievement, but we highlight the fact that the improvements offered by squeezing still result in unremarkable metrological sensitivity that can be bettered by more conventional means.

6.9 Conclusions

In this Chapter, we have demonstrated magnetic tensor gradiometry using differential Ramsey interferometry of spatially separated BECs in freefall. The gradiometer senses vector components of the magnetic field, rejecting gradient components orthogonal to the biasing direction. The gradiometer is immune to common-mode magnetic noise orders of magnitude larger than the field difference, and operates without field cancellation or screening. The gradiometer could also be used as a high-precision co-magnetometer with substantial common-mode rejection, allowing for microscale magnetic sensing *in vacuo*. In the penultimate Chapter of this thesis, we describe the adaptation of differential Ramsey interferometry to measuring and eliminating the vector light shift from the optical dipole trap. The successful outcomes of these experiments make the prospective improvements detailed herein one step closer to being experimentally feasible. In the final Chapter of this thesis, we discuss future work and prospective applications of the techniques described here.

7

Vector light shifts in optical dipole traps

This Chapter describes our experiments measuring and cancelling the vector light shift (VLS) from the optical dipole trap. As discussed in Section 5.4, we first became aware of the VLS when experiments suggested the presence of irreproducibility in the $m_F = \pm 1$ state component separation. This was consistent with a time-varying magnetic field gradient: a gradient that changes magnitude (and even sign) on a shot-to-shot basis is a daunting prospect when considering gradient cancellation strategies. A time-varying gradient also poses a severe difficulty for spin-echo rephasing experiments. As discussed in Chapter 5, the time variation of this 'magnetic' field gradient was duty-cycle dependent and localised to the dipole trap. We deduced that it was an *effective* magnetic field gradient caused by a non-zero VLS due to residual circular polarisation of the dipole trapping beams. The transient nature of the effect was due to thermal effects in the dipole trap AOMs causing shot-to-shot drifts of the polarisation and intensity of the dipole beams at the atoms.

The presence of the VLS was also a complication for magnetic tensor gradiometry, discussed in the previous Chapter, which necessitated performing interferometry on atoms in freefall. The many interesting extensions to differential interferometry – in particular, the realisation of high magnetometric sensitivities – are thwarted by the VLS, which contributes a spurious, transient magnetic field to a measurement. The dipole trap VLS posed a serious impediment to the observation of spinor physics as well as magnetic gradiometry, its elimination is thus of paramount importance.

We resolved to characterise the VLS and develop a measurement scheme based on differential Ramsey interferometry to achieve high-precision cancellation. This Chapter is sectioned as follows: Section 7.1 provides an introduction to the VLS, including details of previous work. Section 7.2 provides a theoretical treatment, and Section 7.4 covers the experimental setup, which is similar to that described in Chapter 6. We then proceed to describe two specific measurements: Section 7.5 explains the first set of
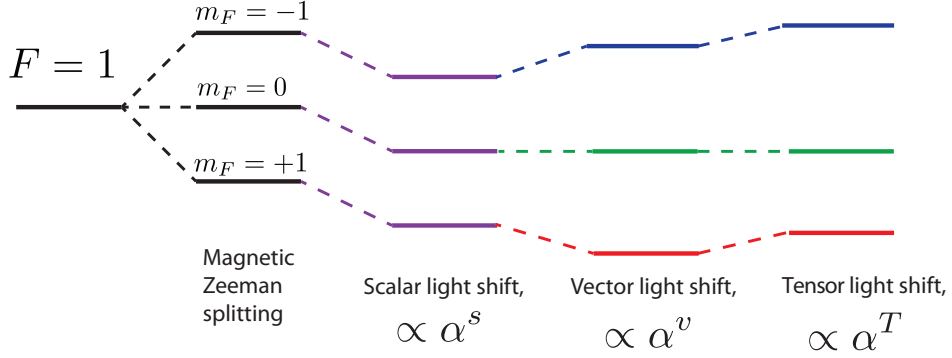
Figure 7.1: Spin-1 energy level shifts due the scalar, vector and tensor light shifts. The scalar light shift, responsible for optical trapping of atoms, shifts all m_F states equally. The VLS can be seen to act as an effective linear Zeeman shift, whereas the tensor shift has the characteristics of a quadratic Zeeman shift.

measurements we performed to characterise the VLS, using a pair of BECs to sample different positions of the radial intensity profile of a single trapping beam, and thus determining the differential VLS across the beam. These measurements are then used to determine the effective VLS-induced field gradient. In Section 7.5.2, we return to our previous studies of spin-mixing dynamics and examine how collisional dynamics are affected by the application (and elimination) of the VLS gradient. Finally, in Section 7.6 we present a high precision measurement technique using differential interferometry over interrogation times of 15 ms. This allows us to ultimately cancel the VLS to within 0.04 %. An extension to the techniques demonstrated in this thesis is presented in Section 7.7.

7.1   Introduction

Light shifts were first outlined (experimentally and theoretically) by Barrat and Cohen-Tannoudji in 1961 [190], and an in-depth theoretical description was performed by Happer and Mathur in 1967 [191]. The early article [192] and review [193] concern calculation of atomic polarisabilities, and Ref. [194] conducted the first experimental study of inhomogeneous broadening due to light shifts. Figure 7.1 illustrates the energy level shifts due to the scalar, vector and tensor light shifts; the vector and tensor shift can be regarded as effective linear and quadratic Zeeman shifts respectively. The magnitude and sign of each energy shift is proportional to intensity and varies with frequency ω of the light as described by the atomic polarisability $\alpha_\eta^k(\omega)$, with $k = s, v, T$ denoting the scalar vector and tensor shift of the atomic energy

level η. For a given atomic state, all three polarisabilities vary considerably with frequency, as we shall see in the following Section where we calculate scalar and vector polarisabilities at several important wavelengths. In the far-detuned limit relevant to most dipole traps, it is generally the case that $\alpha_\eta^S \gg \alpha_\eta^v \gg \alpha_\eta^T$. Aside from the integral presence of the scalar light shift in optical trapping, the vector shift (and to a lesser extent the tensor shift) are frequently encountered in atomic physics experiments, particularly in the same way we discovered them: as small, unwanted frequency shifts.

In the brief and by no means exhaustive review of the literature that follows, we make note of a few key themes that motivate precise measurement of vector light shifts. Our specific motivations for eliminating the vector light shift mirror the general aims of many related works concerning light shifts. The principal benefits trapped neutral atoms bring to metrology and quantum information processing is absolute, calibration free measurements in terms of precisely known physical constants and long coherence times by virtue of isolation from their environment. The ability to manipulate atomic ensembles with lasers or magnetic fields with high precision has applications in spatial addressing of individual qubits as well as microscopic-volume magnetometry. However, the spin-dependence of the optical trapping potential is an issue. While we have concerned ourselves with effective Zeeman shifts *within* a given hyperfine level, different hyperfine levels experience slightly different energy shifts – the *differential* light shift (DLS). The DLS typically refers to the different *scalar* shift of hyperfine levels *F*, due to the different detuning of the light from a transition to an excited state. This results in a intensity-dependent frequency shift of a hyperfine transition, a well-known example is that of the alkali metal clock transitions, which are used abundantly in metrology for reasons of magnetic field insensitivity.

Unlike optical transitions, where 'magic' or 'tune-out' wavelengths exist that allow differential AC Stark shifts to be eliminated – allowing the realisation of ultra-high precision optical clocks [195, 196] – there are no such magic wavelengths for hyperfine differential shifts. Ideally a hyperfine transition for precision metrology would be DLS free and magnetic field insensitive [197]. The effects of differential light shifts on sensitive metrology experiments, including the SI definition of the second, are important – these small frequency shifts can hamper the accurate determination of time. Much experimental [197, 198] and theoretical [199–202] work has been devoted to characterisation and elimination of the clock state DLS, including methods using the vector light shift to cancel the DLS [197, 203–205]. In addition to an overall frequency shift, the light shift also varies spatially – due to inherently inhomogeneous optical intensities of trapping potentials and lattices – meaning that the DLS is also spatially dependent. It therefore has the effect of an

inhomogeneous dephasing mechanism, limiting the coherent evolution time of a quantum system [144, 198, 206], in much the same way as was discussed in Chapter 5 of this thesis. Indeed, the first demonstration of spin-echo rephasing of Ramsey fringes in ultracold gases was concerned with reducing DLS-induced inhomogeneous dephasing [143].

Optical magnetometry experiments – including vapour-cell magnetometers – also contend with the VLS and DLS, where pump-probe intensity fluctuations give rise to noise of the effective VLS magnetic field [20] and thus limit attainable sensitivities unless steps are taken to eliminate the shift [207]. The VLS (and tensor shift) were identified early on as sources of systematic uncertainty in prospective atomic measurements of the permanent electron electric dipole moment [19]. The first spinor-condensate magnetometer actually measured a VLS, and cannot in practice distinguish real magnetic fields from trapping beam-induced vector light shifts [10].

Since we use the Zeeman states of a single hyperfine level in our interferometry experiments, we have no sensitivity to the differential light shift and instead suffer from the vector light shift of the Zeeman states within a single hyperfine level. The manifestation of this shift in our experiments is essentially identical to those described above, however. The inhomogeneity of the VLS (accentuated by gravitational sag of the trap minimum from the intensity maximum) leads to dephasing that is manifest in both Ramsey interferometry and spin-mixing experiments, and the overall intensity-dependent VLS contaminates magnetic gradiometry experiments.

There is somewhat less published experimental work in considering the effects of VLS-shifted Zeeman states on metrology, primarily due to the more widespread use of approximately field-insensitive microwave transitions for precision measurements where the DLS is significant. A recent experiment concerning measurement of the VLS of ground-state Zeeman sublevels of ^{133}Cs was performed by Zhu et al. in Ref. [208]. This experiment essentially uses the VLS as a metric of absolute polarisation purity in an optical lattice, using a variation of the Hanle effect to measure inhomogeneous dephasing times. This is used as a diagnostic for improving the polarisation linearity, obtaining a minimum VLS corresponding to a polarisation imperfection $\sim 10^{-5}$ rad.

A VLS can also serve beneficial purposes. Early work highlighted the usefulness of spin-dependent optical trapping potentials [209], and using the VLS for state control in optical lattices [210–212]. More recent advances have used the VLS in the demonstration of synthetic gauge fields and spin-orbit coupling in BECs [5, 213]. Here, the synthetic field is generated by momentum-dependent Raman coupling originating from the effective vector potential provided by the VLS, rather than the 'fictitious' field. One can indeed use the

VLS for constructive purposes in a spinor condensate, as a means of tailoring a spatially-varying Zeeman potential: a laser tuned to the magic wavelength ($\sim 790\,\mathrm{nm}$ for ^{87}Rb, eliminating the scalar shift) can be focused onto particular regions of the condensate. This has application in the continuing study of magnetic excitations (magnons) in spinor condensates [214].

7.1.1 Magnetometry and the vector light shift

Before delving into an in-depth theoretical treatment of atomic vector polarisabilities, we provide a simple explanation for how the VLS contributes to a spurious magnetometry signal as well as a spatially inhomogeneous energy shift in our experiments. For the former, we consider our magnetic tensor gradiometer described in the previous Chapter, with the interferometry sequence performed in-trap, so that the atoms are sensitive to the VLS. We consider the case of one dipole beam split into two (beams 1 and 2) forming two traps along a third crossing beam. The atoms sample the differential VLS due to a polarisation imperfection common to beams 1 and 2 (we ignore the VLS from the third crossing beam). As shown in Figure 7.2, the VLS contributes to a gradiometric measurement when the light is not linearly polarised and an intensity difference exists between the two traps.

The polarisation of laser light used for optical trapping can be made highly linear using polarisers, etc. but passage through birefringent materials such as the glass vacuum cell result in elliptically polarised light at the atoms. We assume that the spatial dependence of the birefringence responsible for the polarisation imperfection is weak, and can be neglected for small beam separations of several hundred microns between beams 1 and 2. We therefore consider an approximately common polarisation imperfection of beam 1 and 2 sampled by each trapped cloud, resulting in a VLS difference proportional to the intensity difference of the beams. In this example, we have assumed that beams 1 and 2 initially have a common polarisation, in reality, each beam originates from the diffracted order of an AOM, and this assumption is not immediately applicable. However, a polarising element (such as a Glan-laser prism) placed immediately after the AOM ensures the beams are linearly polarised, and have negligible *relative* polarisation imperfection. This does not rectify the problem of passage through birefringent optics *after* the polariser however. The two linearly polarised dipole beams propagate through the same birefringent optics, so we can assume the two beams in general share a common polarisation imperfection.

In principle, we could eliminate the differential VLS from a gradiometry experiment by means of common-mode rejection with balanced intensities with a common polarisation imperfection, as in Figure 7.2(c). Equal beam
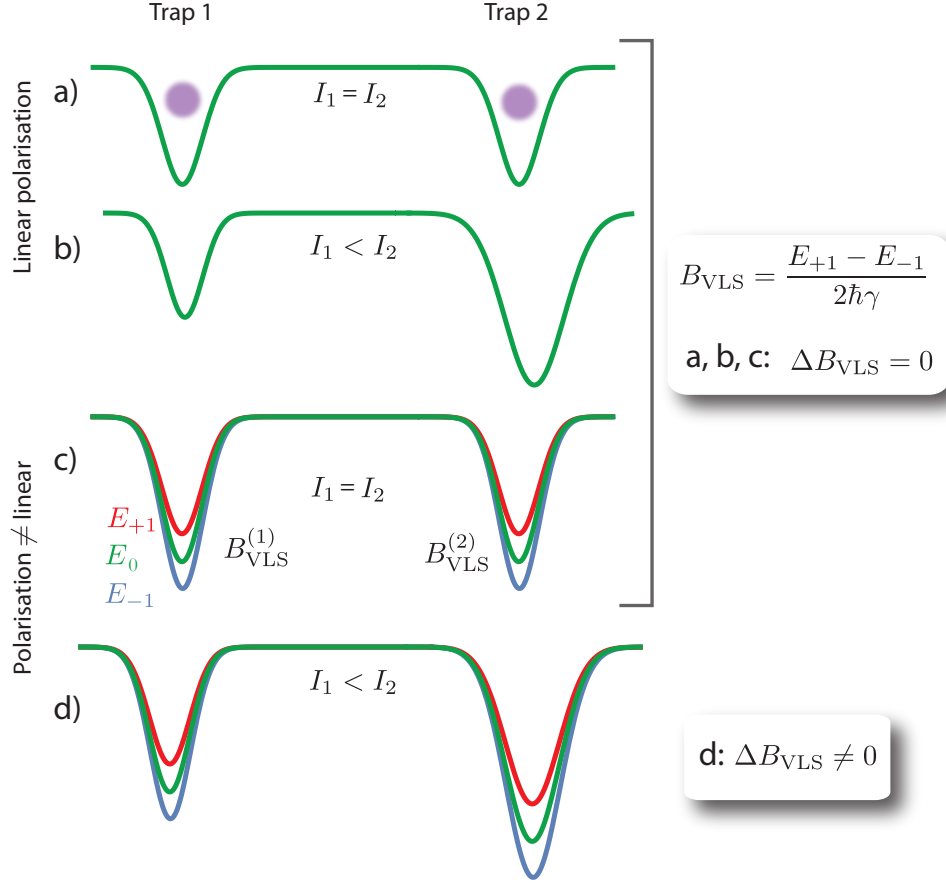
Figure 7.2: Contribution of the VLS to magnetic gradiometry experiments, which detect the differential magnetic field as well as the differential VLS field $\Delta B_{\mathrm{VLS}} = B_{\mathrm{VLS}}^{(2)} - B_{\mathrm{VLS}}^{(1)}$. In (a), we have the ideal case of equal intensities in each trap and linear polarisation: $B_{\mathrm{VLS}}^{(1)} = B_{\mathrm{VLS}}^{(2)} = 0$. (b) When an intensity difference exists but there is no polarisation imperfection, the VLS-induced field difference is still zero. We assume that the two dipole beams exhibit a common polarisation imperfection from passage through birefringent optics, allowing us to consider the VLS as arising from an overall polarisation common to each trap. (c) For equal intensities, common-mode rejection eliminates the VLS field contribution. In most experimental cases however a VLS difference exists between the two traps due to the cloud sampling a *local* intensity difference (originating from misalignment or power differences) and non-linear polarisation (d), this contributes a spurious magnetic field to a gradient measurement.

intensities are still challenging to achieve *in vacuo*, however, and susceptible to relative intensity fluctuations and drift. This is also complicated by the fact that the atomic cloud samples a local intensity, due to gravitational sag and beam misalignments: equal beam *power* does not correspond to equal intensity. Only by making the trapping light exceptionally linearly polarised *in vacuo* can we truly eliminate the VLS.

The VLS results in inhomogeneous frequency shifts which lead to dephasing. The origin of this effect in our case is due to gravitational sag of the trapping potential: the atoms *do not* sample the peak intensity of the trapping beam where the VLS is roughly spatially constant; instead gravity pulls the trap minimum below the peak intensity. The BEC, located at the trap minimum, therefore samples a spatially varying vector shift, as shown in Figure 7.3. The vector polarisability at 1064 nm in ^{87}Rb is 0.336 Hz W^{-1} cm^2; for an intensity of $I \sim 10^4$ W cm^{-2} this gives an overall effective field of 0.3 mG for a polarisation imperfection of $2°$. For in-trap magnetometry, a 1 % intensity imbalance thus results in a 3 μG field difference, which is readily detectable by the prospective magnetometer described in Chapter 6. The tight focusing of the trapping beams results in an effective gradient of 24 mG/cm for the example shown in Figure 7.3. This gradient is greater than any of the background gradients measured in the previous Chapter, and as such is a considerable problem for both spinor dynamics and gradiometry experiments, where it limits the maximum interrogation time and hence the ultimate achievable precision.

7.2   Theory of atomic polarisabilities

Here we consider the theoretical description of atomic polarisabilities with the aim of computing numerical values of the proportionality constant $\alpha_V(\omega)$. First principles calculation of atomic polarisabilities is non-trivial, requiring demanding simulations of atomic structure [215]. However, dynamic (frequency dependent) polarisabilities can be determined from numerically calculated dipole matrix elements and angular momentum algebra. We follow the formulation given in Ref. [216] to derive the theoretical formulation of scalar and vector light shifts, while also drawing from Refs. [202, 217]. We will ignore tensor shifts where appropriate for specific results concerning the VLS and scalar shifts.

The VLS is just part of the AC Stark shift that arises from the interaction between an atom and electric field of the form

$$\mathbf{E}(t) = \frac{1}{2} E_0 \, \hat{\epsilon} \, e^{-i\omega t} + \text{c.c.} \, , \tag{7.1}$$
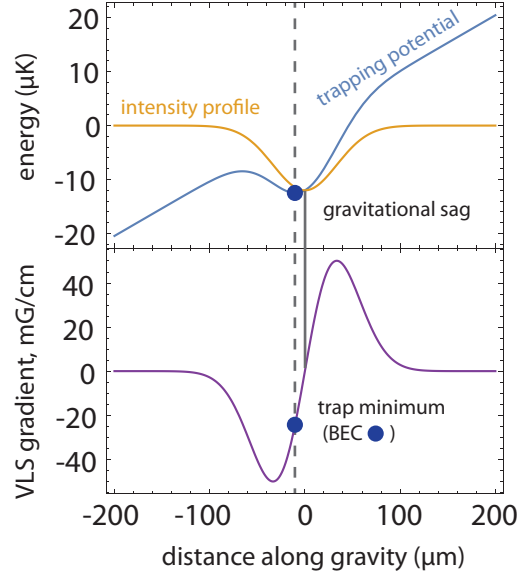
Figure 7.3: Gravitational sag of the dipole potential results in a trapped BEC sampling a spatially varying VLS. Top panel shows the potential from only the intensity profile of the laser along the direction of gravity. The blue trace is the dipole potential shifted by gravity. The BEC (positioned at the trap minimum) therefore samples an accentuated intensity gradient which results in a gradient of the VLS (lower panel). The magnitude of the effective field is 0.3 mG for a 0.6% polarisation imperfection, but the gradient is 24 mG/cm due to the tightly focused beams. We have used the scalar ($31.73 \, \mathrm{Hz \, W^{-1} \, cm^2}$) and vector ($0.336 \, \mathrm{Hz \, W^{-1} \, cm^2}$) light shifts at 1064 nm, calculated in the following Section. The beam power (550 mW) and waist (67 μm) are approximately what is used in experiments, for these parameters gravity shifts the trap minimum by 10 μm from the peak intensity point.

with $\hat{\epsilon}$ the polarisation vector and E_0 the complex electric field amplitude. The light-atom interaction in the dipole approximation is written as

$$U^E = -\mathbf{E} \cdot \mathbf{d}, \tag{7.2}$$

with the $\mathbf{d}$ the atomic electric dipole operator, $\mathbf{d} = \alpha \mathbf{E}$; the constant of proportionality α is the *atomic polarisability*. The AC Stark shift is the time-average of the induced dipole moment:

$$\delta E_a = -\frac{1}{2} \langle \mathbf{E}(t) \cdot \mathbf{d}(t) \rangle, \tag{7.3}$$

where the factor of $\frac{1}{2}$ accounts for the fact that this is an induced dipole moment [115, 216]. For atomic eigenstates with energies $\hbar\omega_a$ and $\hbar\omega_b$, which we call the ground state and excited state respectively, the second-order energy shift of the ground state ($E_1 = 0$ if the states are eigenfunctions of the

parity operator [215]) is given by [216]:

$$\delta E_a = -\frac{|E_0|^2}{4\hbar} \sum_b \text{Re} \left(\frac{|\langle b|\hat{\bm{\epsilon}} \cdot \mathbf{d}|a\rangle|^2}{\omega_b - \omega_a - \omega - i\Gamma_{ba}/2} + \frac{|\langle a|\hat{\bm{\epsilon}} \cdot \mathbf{d}|b\rangle|^2}{\omega_b - \omega_a + \omega - i\Gamma_{ba}/2} \right)$$

(7.4)

where the sum runs over multiple excited states b and Γ_{ba} is the natural linewidth of the $a \leftrightarrow b$ transition. When considering only fine-structure coupling (*i.e.* when detuning from resonance is much larger than fine structure splitting), the Stark shift operator can be expressed as [202, 216]

$$U^E = -\frac{|E_0|^2}{4} \left[\alpha_{nJ}^s - i\alpha_{nJ}^v \frac{\hat{\bm{\epsilon}}^* \times \hat{\bm{\epsilon}} \cdot \mathbf{J}}{2J} + \alpha_{nJ}^T \frac{3\left((\hat{\bm{\epsilon}}^* \cdot \mathbf{J})(\hat{\bm{\epsilon}} \cdot \mathbf{J}) + (\hat{\bm{\epsilon}} \cdot \mathbf{J})(\hat{\bm{\epsilon}}^* \cdot \mathbf{J}) \right) - 2\mathbf{J}^2}{2J(2J+1)} \right],$$

(7.5)

with the scalar, vector and tensor reduced polarisabilities given by

$$\alpha_{nJ}^s = \frac{1}{\sqrt{3(2J+1)}} \alpha_{nJ}^{(0)}$$

(7.6a)

$$\alpha_{nJ}^v = -\sqrt{\frac{2J}{(J+1)(2J+1)}} \alpha_{nJ}^{(1)}$$

(7.6b)

$$\alpha_{nJ}^T = -\sqrt{\frac{2J(2J-1)}{3(J+1)(2J+1)(2J+3)}} \alpha_{nJ}^{(2)}$$

(7.6c)

with

$$\alpha_{nJ}^{(K)} = (-1)^{K+J+1} \sqrt{2K+1} \sum_{n',J'} (-1)^{J'} \begin{Bmatrix} 1 & K & 1 \\ J & J' & J \end{Bmatrix} |(n'J'\,\|\mathbf{d}\|\,nJ)|^2$$

$$\times \frac{1}{\hbar} \text{Re} \left(\frac{1}{\omega_{n'J'nJ} - \omega - i\Gamma_{n'J'nJ}/2} + \frac{(-1)^K}{\omega_{n'J'nJ} + \omega - i\Gamma_{n'J'nJ}/2} \right), \quad (7.7)$$

where $\omega_{n'J'nJ} = \omega_{n'J'} - \omega_{nJ}$, $\Gamma_{n'J'nJ}$ is the natural linewidth of the $nJ \to n'J'$ transition, $\{j_n...\}$ denotes the Wigner 6j-symbol and $|(n'J'\,\|\mathbf{d}\|\,nJ)|^2$ the reduced dipole matrix element. In order to accurately calculate atomic polarisabilities, one must know the dipole matrix elements, which are measured experimentally or numerically calculated [79, 215, 217].

We note here a few results from Eqs. 7.6 specific to our case ($J = 1/2, K = 0, 1$; we omit the result for the tensor shift[1]), namely

$$\alpha_{nJ}^s = \frac{1}{\sqrt{6}} \alpha_{nJ}^{(0)}$$

(7.8a)

$$\alpha_{nJ}^v = -\frac{1}{\sqrt{3}} \alpha_{nJ}^{(1)}.$$

(7.8b)

---

[1]The tensor shift is typically $< 1\,\text{Hz}$, at dipole-trapping wavelengths and is of little consequence for our experiments.

Also, by permuting columns of the 6*j*-symbol in Eq. 7.7, it is clear it is equivalent to Eqs. 7–9 of Ref. [217]. We note this specifically since the theoretically calculated polarisabilities in Ref. [217] are in fact $\alpha_{nJ}^{s,v,T}$ and do not include hyperfine terms. The fine-structure polarisabilities can be related to hyperfine polarisabilities as we shall see below.

The previous discussion focused on polarisabilities and energy shifts of atoms in a single fine structure state $|nJ\rangle$. However, to determine the light shifts of atoms in hyperfine states, we must include the effects of the hyperfine interaction and recast Eqs. 7.5, 7.7, and 7.6 in the $|nJF\rangle$ basis. The Stark shift operator becomes

$$U^E = -\frac{|E_0|^2}{4}\left[\alpha_{nJF}^s - i\alpha_{nJF}^v\frac{\hat{\boldsymbol{\epsilon}}^* \times \hat{\boldsymbol{\epsilon}} \cdot \mathbf{F}}{2F}\right. \tag{7.9}$$

$$\left. +\alpha_{nJF}^T\frac{3\left((\hat{\boldsymbol{\epsilon}}^* \cdot \mathbf{F})(\hat{\boldsymbol{\epsilon}} \cdot \mathbf{F}) + (\hat{\boldsymbol{\epsilon}} \cdot \mathbf{F})(\hat{\boldsymbol{\epsilon}}^* \cdot \mathbf{F})\right) - 2\mathbf{F}^2}{2F(2F+1)}\right], \tag{7.10}$$

and the corresponding polarisabilities become

$$\alpha_{nJF}^s = \frac{1}{\sqrt{3(2J+1)}}\alpha_{nJ}^{(0)} \tag{7.11a}$$

$$\alpha_{nJF}^v = (-1)^{J+I+F}\sqrt{\frac{2F(2F+1)}{F+1}}\begin{Bmatrix} F & 1 & F \\ J & I & J \end{Bmatrix}\alpha_{nJ}^{(1)} \tag{7.11b}$$

$$\alpha_{nJF}^T = (-1)^{J+I+F}\sqrt{\frac{2F(2F-1)(2F+1)}{3(F+1)(2F+3)}}\begin{Bmatrix} F & 2 & F \\ J & I & J \end{Bmatrix}\alpha_{nJ}^{(2)}. \tag{7.11c}$$

Picking the specific cases of interest, $F = 1, I = 3/2$, we see that

$$\alpha_{nJF}^s = \alpha_{nJ}^s \tag{7.12a}$$

$$\alpha_{nJF}^v = \frac{1}{2\sqrt{3}}\alpha_{nJ}^{(1)} = -\frac{1}{2}\alpha_{nJ}^v. \tag{7.12b}$$

As is clear from Eq. 7.12a, the scalar shift is independent of *F*, whereas in reality there is a small *F* dependence (the differential scalar shift mentioned earlier). According to Ref. [216], when the hyperfine interaction is much larger than the Stark shift the atomic polarisabilities are given by

$$\alpha_F^s = \frac{1}{\sqrt{3(2F+1)}}\alpha_F^{(0)} \tag{7.13a}$$

$$\alpha_F^v = -\sqrt{\frac{2F}{(F+1)(2F+1)}}\alpha_F^{(1)} \tag{7.13b}$$

where

$$\alpha_F^{(K)} = (-1)^{K+F+1}(2F+1)\sqrt{2K+1}\sum_{n',J'}\left|\left(n'J'\left\|\mathbf{d}\right\|nJ\right)\right|^2$$

$$\times\sum_{F'}(-1)^{F'}(2F'+1)\left\{\begin{array}{ccc}1 & K & 1\\ F & F' & F\end{array}\right\}\left\{\begin{array}{ccc}F & 1 & F'\\ J' & I & J\end{array}\right\}^2$$

$$\times\frac{1}{\hbar}\mathrm{Re}\left(\frac{1}{\omega_{n'J'F'nJF}-\omega-i\Gamma_{n'J'F'nJF}/2}+\frac{(-1)^K}{\omega_{n'J'F'nJF}+\omega-i\Gamma_{n'J'F'nJF}/2}\right).$$
(7.14)

The atomic energy level shift due to the scalar and vector light shifts is given by

$$U^E = -\frac{|E_0|^2}{4}\left[\alpha_F^s - i\alpha_F^v\frac{\hat{\boldsymbol{\epsilon}}^*\times\hat{\boldsymbol{\epsilon}}\cdot\mathbf{F}}{2F}\right].$$
(7.15)

It is this final form that we numerically calculate polarisabilities and light shifts for $F=1$ ^{87}Rb with, using only the contribution from the D$_1$ (^{2}5S$_{1/2}\to^2$ 5P$_{1/2}$) and D$_2$ (^{2}5S$_{1/2}\to^2$ 5P$_{3/2}$) lines. Note that the dipole matrix elements used for ^{87}Rb in Table 7 of Ref. [79] must be multiplied by $\sqrt{2J+1}$ to use the same convention as those in other works [215–217]. Reduced dipole matrix elements are *generally* represented by the notation $(J\left\|\mathbf{d}\right\|J)=\sqrt{2J+1}\left\langle n'J'\left\|\mathbf{d}\right\|nJ\right\rangle$, although adherence to this convention is inconsistent. The result is that reduced dipole matrix elements in one source may differ from another by a factor $\sqrt{2J+1}$.[2]

In this work, and in the works cited, the dipole matrix elements used are the *reduced* dipole matrix elements and we use the round-bracket convention to be explicit: the dipole matrix elements in Ref. [79] (Steck) are multiplied by $\sqrt{2J+1}=\sqrt{2}$ to conform to the convention used in Ref. [216].

We now explicitly define the polarisation vector $\hat{\boldsymbol{\epsilon}}$, in the spherical basis in terms of the unit vectors $\hat{\boldsymbol{\epsilon}}_R=(\hat{\mathbf{x}}-i\hat{\mathbf{y}})/\sqrt{2}$, $\hat{\boldsymbol{\epsilon}}_L=(\hat{\mathbf{x}}+i\hat{\mathbf{y}})/\sqrt{2}$ it is given by

$$\hat{\boldsymbol{\epsilon}} = \sin\left(\theta+\frac{\pi}{4}\right)\hat{\boldsymbol{\epsilon}}_L + e^{2i\phi}\cos\left(\theta+\frac{\pi}{4}\right)\hat{\boldsymbol{\epsilon}}_R.$$
(7.16)

The angle θ characterises degree of circular polarisation and ϕ the orientation of the polarisation axis, for $\theta=0$ the light is linearly polarised. The previous analysis has assumed that the quantisation axis, determined by

---

[2]The resolution of this ambiguity lies in the definition of the natural linewidth:

$$\Gamma_{J\to J'} = \underbrace{\frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}\frac{2J+1}{2J'+1}|\langle J'\left\|\mathbf{d}\right\|J\rangle|^2}_{\text{Steck, Ref. [79]}} = \overbrace{\frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}\frac{1}{2J'+1}|(J'\left\|\mathbf{d}\right\|J)|^2}^{\text{Kien, Ref. [216]}},$$ although frustratingly

Ref. [216] does not use the round-bracket notation.

an external magnetic field, is parallel to the wavevector of the laser, which is also the direction of the VLS 'fictitious' field. When this is not the case, the VLS shift is given by the projection of the effective VLS magnetic field onto the quantisation axis [201]. For a bias field making an angle φ to the wavevector of the laser beam the cross product in Eq. 7.15 thus reduces to $-i\hat{\epsilon}^* \times \hat{\epsilon} \cdot \mathbf{F} = \sin 2\theta \cos \varphi\, m_F$. Varying φ is equivalently a method of minimising the VLS, but of limited utility to us as it restricts the vector sensitivity of our magnetometry experiments to only one direction where the VLS is cancelled.

We make the following substitution, expressing the squared electric field amplitude as an optical intensity $|E_0|^2 = 2I_0/c\,\epsilon_0$:

$$U^E/h = -I_0 \left(\alpha_S + \alpha_V \sin 2\theta \, \cos \varphi\, m_F \right), \tag{7.17}$$

where we have defined the scaled polarisabilities as a frequency shift per unit intensity

$$\alpha_S = \frac{\alpha_F^s}{2\epsilon_0 c} \tag{7.18a}$$

$$\alpha_V = \frac{\alpha_F^v}{4\epsilon_0 c} \tag{7.18b}$$

for the specific case of $F = 1$. The scaled polarisabilities are plotted as a function of wavelength in Figure 7.4. We calculate $\alpha_V = -0.336\,\mathrm{Hz\,W^{-1}\,cm^2}$ at $\lambda = 1064\,\mathrm{nm}$, thus the energy shift for the vector light shift is

$$E_{\mathrm{VLS}}/h = -0.336\,\mathrm{Hz\,W^{-1}\,cm^2} \times I_0 \sin 2\theta \, \cos \varphi\, m_F. \tag{7.19}$$

7.2.1 Polarisabilities in different unit systems

Obtaining previously calculated theoretical values of the vector polarisability proved a challenge. To our knowledge, there has been no published experimental or theoretical work on the vector polarisability of $F = 1$ ^{87}Rb at 1064 nm. In adapting Ref. [216] to calculate vector shifts for ^{87}Rb, we can compare our calculated values for *other* polarisabilities against those previously presented in the literature; namely Ref. [217]. This then requires a unit conversion from atomic units (a.u.) into SI units, $\mathrm{Hz\,W^{-1}\,cm^2}$. Additionally, we must make use of several relations derived in the preceding Section to convert polarisabilities between unit systems and angular momentum bases, specifically Eqs. 7.12, since Ref. [217] calculates $\alpha_{nJ}^{s,v}$ rather than $\alpha_F^{s,v}$.

To begin, we note that the form we have used (Eq. 7.14) actually accounts for the differential (F-dependent) shift. If we assume this is a small correction,
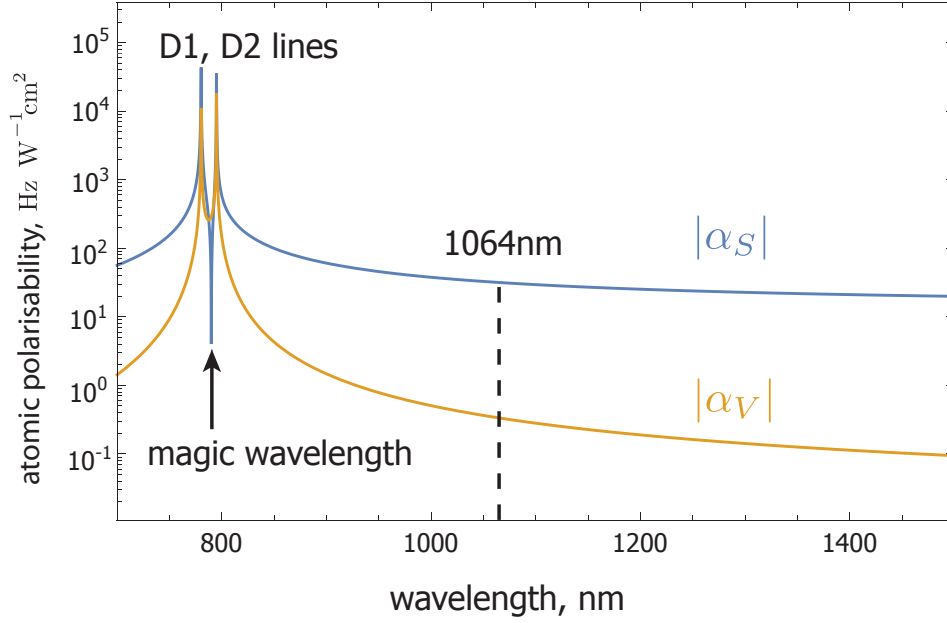
Figure 7.4: Scalar (α_S) and vector (α_V) polarisabilities, calculated using Eqs. 7.18 and 7.14. The sharp dip between the D1 and D2 lines (794 nm and 780 nm) corresponds to the magic wavelength ($\lambda = 790$ nm) where the scalar light shift vanishes. At $\lambda = 1064$ nm, the scalar and vector polarisabilities are 31.73 Hz W^{-1} cm^2 and 0.336 Hz W^{-1} cm^2 respectively.

we still identify Eqs. 7.12 as holding, so that $\alpha_{nJ}^s \approx \alpha_F^s$ and vector polarisabilities α_{nJ}^v (calculated in Ref. [217]) are related to the hyperfine polarisabilities calculated in Eq. 7.14 by

$$\alpha_F^v \approx -\frac{1}{2}\alpha_{nJ}^v. \tag{7.20}$$

Using the conversion factors in Ref. [215], we can convert a calculated $\alpha_{nJ}^{s,v}$ in a.u. into an $\alpha_{S,V}$ in Hz W^{-1} cm^2 using

$$\alpha_S = \frac{0.2488}{10^3}\frac{\alpha_{nJ}^s}{2\epsilon_0 c} \qquad \alpha_v = \frac{0.2488}{10^3}\frac{\frac{1}{2}\alpha_{nJ}^v}{4\epsilon_0 c}. \tag{7.21}$$

We can now compile Table 7.1, which shows good agreement between our calculated values and those of Ref. [217].

7.3 Controlling polarisation

Polarisation is our principal knob for tuning and ultimately eliminating the VLS. The VLS exists in our experiments because the otherwise linearly polarised dipole trapping light acquires a small polarisation imperfection from reflecting off and passing through a number of optical components, such

| λ | α_S | | α_V | |
(nm)	Calc.	Ref. [217]	Calc.	Ref. [217]
770	431.44	430.22(5)	52.29	52.29(1)
1064	31.73	32.21(2)†	0.336	–

Table 7.1: Scalar and vector polarisabilities, in units of $\mathrm{Hz\,W^{-1}\,cm^2}$ for ^{87}Rb computed using Eqs. 7.18 and 7.14. Good agreement is found between the calculated values and the polarisabilities calculated by Ref. [217] using the conversion in Eq. 7.21. †An earlier experiment [218] measured the value of α_S at $\lambda = 1064$ nm to be $36(3)\,\mathrm{Hz\,W^{-1}\,cm^2}$.

as the dichroic mirrors and most importantly of all, the glass cell. No matter what effort we go to to make the light purely linearly polarised *before* the cell (*i.e.* where it is easy to do so), the light at the atoms will have some slight ellipticity due to the birefringence of the vacuum cell. We therefore require an *in vacuo* polarimetry technique to measure the VLS and diagnose the linearity of the light, which is adjusted using an *ex vacuo* quarter-wave plate. A previous work used the differential vector shift imparted on a cloud of ^{133}Cs in a dipole trap with $\lambda = 866$ nm (where the vector shift is substantial) to essentially measure the birefringence of vacuum windows [185].

In this Section, we will address two issues before plunging into the experimental details. The first concerns the identification of the angle θ in Eqs. 7.16 and 7.19 with that of a quarter-wave plate rotated by the same angle, and the second the is the quarter-wave plate's ability to compensate the birefringence of down-stream optical components.

To describe polarised light, we use the Jones calculus. Polarisation states are represented by Jones vectors $\tilde{E}$, and optical components such as wave-plates and polarisers by Jones matrices. Jones matrices act on Jones vectors, transforming the polarisation state of the light. Polarised light propagating along z with orthogonal electric field components $E_x(t)$ and $E_y(t)$ is described by the Jones vector [219]

$$\tilde{E} = \begin{pmatrix} E_x(t) \\ E_y(t) \end{pmatrix} = e^{i\phi} \begin{pmatrix} E_x^{(0)} e^{i\phi_x} \\ E_y^{(0)} e^{i\phi_y} \end{pmatrix} \tag{7.22}$$

with $\phi_{x,y}$ the phase of each electric field component in addition to the overall phase $\phi = kz - \omega t$. In practice, we neglect the overall phase ϕ, normalise the electric field amplitude (and hence each Jones vector), and define the relative phase between components $\delta\phi = \phi_y - \phi_x$. Elliptically polarised light with an arbitrary orientation can then be written as

$$\tilde{E}(\delta\phi, \theta) = \begin{pmatrix} \cos\theta \\ e^{i\delta\phi} \sin\theta \end{pmatrix} \tag{7.23}$$

where the angle θ is the angle the Jones vector makes to the x axis. We can thus define basis vectors corresponding to vertical and horizontal linear and right and left circularly polarised light:

$$\text{Linear:} \quad \hat{\boldsymbol{\epsilon}}_V = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \hat{\boldsymbol{\epsilon}}_H = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \tag{7.24}$$

$$\text{Circular} \quad \hat{\boldsymbol{\epsilon}}_R = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix}, \quad \hat{\boldsymbol{\epsilon}}_L = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}. \tag{7.25}$$

The action of a birefringent optical component on polarised light can be described by the action of a Jones matrix on a Jones vector: the Jones matrix for a quarter-wave plate $\mathcal{A}_{\lambda/4}$ whose fast axis is an angle θ measured clockwise from vertical is given by

$$\mathcal{A}_{\lambda/4}(\theta) = \begin{pmatrix} \cos^2\theta + i\sin^2\theta & (1-i)\cos\theta\sin\theta \\ (1-i)\cos\theta\sin\theta & i\cos^2\theta + \sin^2\theta \end{pmatrix}. \tag{7.26}$$

We can then compute the projection of the product $\tilde{\boldsymbol{E}}_{\text{out}} = \mathcal{A}_{\lambda/4}(\theta)\tilde{\boldsymbol{E}}_{\text{lin}}$ onto the circular basis states, with the quarter-wave plate fast axis an angle θ from vertical and the Jones vector $\tilde{\boldsymbol{E}}_{\text{lin}}$ a linearly polarised input state, $\hat{\boldsymbol{\epsilon}}_H$ or $\hat{\boldsymbol{\epsilon}}_V$. The resulting light expressed in the circular basis has the form $\tilde{\boldsymbol{E}}_{\text{out}} = \alpha(\theta)\hat{\boldsymbol{\epsilon}}_L + \beta(\theta)\hat{\boldsymbol{\epsilon}}_R$, with the *circularity* given by $\mathcal{C}(\theta) = |\alpha(\theta)|^2 - |\beta(\theta)|^2$, the normalised difference of left and right circular polarisation projections.[3] The circularity essentially tells us the relative fraction of right and left circularly polarised light, the normalised length of the effective VLS field along the direction of the wavevector, and using Jones calculus it can be shown that

$$-i\tilde{\boldsymbol{E}}^* \times \tilde{\boldsymbol{E}} \cdot \hat{\boldsymbol{z}} = \mathcal{C} = |\alpha|^2 - |\beta|^2 \tag{7.27}$$

for an arbitrary polarised state $\tilde{\boldsymbol{E}} = \alpha(\theta)\hat{\boldsymbol{\epsilon}}_L + \beta(\theta)\hat{\boldsymbol{\epsilon}}_R$.

Computing $\mathcal{C}$ for the polarisation vector in Equation 7.16 and for a quarter-wave plate acting on a linearly polarised state reveal both to have exactly the same functional form, $\sin 2\theta$. This result is general; an arbitrary elliptically polarised state can be written in the linear basis as $\tilde{\boldsymbol{E}}_{\text{arb.}} = \frac{1}{\sqrt{2}}(\hat{\boldsymbol{\epsilon}}_V + e^{2i\theta}\hat{\boldsymbol{\epsilon}}_H)$, and is equivalent up to a rotation and overall phase factor to $\hat{\boldsymbol{\epsilon}}$ in Eq. 7.16

---

[3]We use *circularity* for the single number that describes the polarisation state of light in the circular basis, which is equivalently the Stokes parameter S_3 [219]. The *ellipticity* as such a number is not a widespread convention to our knowledge; the ellipticity is occasionally referenced as the ratio of the semi-major to semi-minor axes (b/a) of the polarisation ellipse (related to the eccentricity, $e = \sqrt{1 - (b/a)^2}$); or simply as the relative phase between electric field components. The ellipticity angle $\chi = \tan^{-1}\left(\frac{b}{a}\right)$ is also used. The circularity as we have defined it naturally incorporates the handedness of the light and is a natural extension of the Jones formalism used here.

and the state $\tilde{E}_{\text{out}}$. Computing the circularity and the cross product for the state $\tilde{E}_{\text{arb.}}$ also yields $\sin 2\theta$.

We are thus able to identify the mixing angle in Eq. 7.16 with the rotation angle of a quarter-wave plate, and deduce that for a $180°$ rotation of the quarter-wave plate the VLS varies between $-\alpha_V I_0 < E_{\text{VLS}}/h < \alpha_V I_0$. This also gives an indication of the experimental requirements needed to precisely eliminate the polarisation imperfection: the degree to which the VLS can be cancelled then requires a waveplate rotation mount capable of very small, reproducible rotations.

Before we conclude, we note how Jones calculus can be applied to confirm a quarter-wave plate can be used to 'undo' the elliptical polarisation imperfection induced on light from an arbitrary phase shifter, such as the combined effect of the post-waveplate optics including the glass cell. The arbitrary phase shifter can be represented by the matrix $\mathcal{A}_{\phi_k}$. The light in the vacuum cell seen by the atoms is given by $\tilde{E}_{\text{atoms}} = \mathcal{A}_{\phi_k} \mathcal{A}_{\lambda/4} \tilde{E}_{\text{lin}}$, with $\mathcal{A}_{\phi_k}$ taking the form of a Jones matrix that imparts a relative phase shift of some inconvenient amount ϕ_k, with fast axis oriented at some angle θ_k;

$$\mathcal{A}_{\phi_k} = \begin{pmatrix} \cos^2 \theta_k + e^{i\phi_k} \sin^2 \theta_k & (1 - e^{i\phi_k}) \cos \theta_k \sin \theta_k \\ (1 - e^{i\phi_k}) \cos \theta_k \sin \theta_k & e^{i\phi_k} \cos^2 \theta_k + \sin^2 \theta_k \end{pmatrix}. \tag{7.28}$$

Eliminating the VLS therefore amounts to finding a quarter-wave plate rotation angle θ_{null} such that for $\tilde{E}_{\text{atoms}}$, $\mathcal{C}(\theta_{\text{null}}) = 0$. The circularity of $\tilde{E}_{\text{atoms}}$ is given by

$$\mathcal{C} = \cos \phi_k \sin 2\theta - \cos 2\theta \sin 2(\theta - \theta_k) \sin \phi_k. \tag{7.29}$$

Analytically proving that there always exists some quarter-wave plate angle θ that satisfies $\mathcal{C}(\theta) = 0$ is not trivial. However, after visual inspection of the effect of each matrix[4] it is apparent that any circularity generated by an arbitrary phase shifter at some fixed angle can equivalently be generated by a quarter-wave plate at some rotation angle, which produces $-1 \leq \mathcal{C} \leq +1$ for $\theta \in [0, \pi]$. The unknown phase shifter can always be cancelled by a quarter-wave plate rotation. Colloquially the mechanism is called 'pre-fudge[5] – un-fudge', and is depicted in Figure 7.5.

7.4   Experimental methods

We now describe the experimental protocol used to measure vector light shifts. We adapt the differential Ramsey interferometry technique described

---

[4]One can examine polarisation ellipses or plots of $\mathcal{C}$.

[5]This is perhaps the most phonologically similar and acceptable version of an adverb that can take another form, particularly when the phase shift to be cancelled is especially insidious.
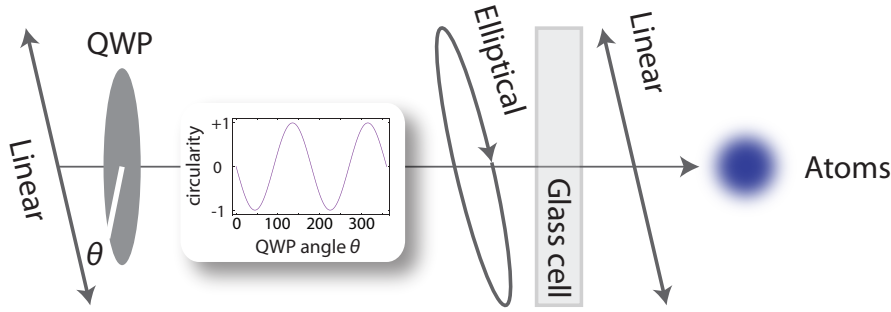
Figure 7.5: A quarter-wave plate (QWP) is used to correct for the bire-fringence due to optical components and the glass cell; allowing for the preparation of high-quality linear polarisation *in vacuo*. The initial state need not be linear; the method can be equivalently used to correct residual elliptical polarisation in the incident beam.

in Chapter 6 to measure the VLS by performing the Ramsey interrogation whilst the atoms are still exposed to trapping light. To detect a VLS as a differential interferometric signal, the BECs in the two dipole traps must experience either a different polarisation in each beam, or a different intensity with a common polarisation, as shown in Figure 7.2. It is also useful to be able to measure the VLS from each dipole beam separately.

We devised two methods for measuring the VLS using differential in-terferometry, referred to as 'delayed drop' and 'in-trap' for simplicity. The optical experimental configurations for measuring the VLS and schematic illustrations of the two measurement schemes is shown in Figure 7.6.

7.4.1 Delayed drop technique

In the 'delayed drop' technique, one dipole trapping beam is split into two, *A* and *B*, positioning two BECs axially along the crossing beam (beam *C*), as shown in Figure 7.6 (b). It is the VLS from the crossing beam that we will measure as a differential signal:

1. We extinguish one of the split beams (beam *A*) first: since the crossing beam alone is insufficient to support the BEC against gravity, BEC_A falls.

2. After a short (up to 5 ms) delay, we extinguish the second split beam (beam *B*), so that the two clouds are now *vertically* separated along the radial direction of beam *C* by the distance fallen by BEC_A, typically 10-20 µm.

3. The two clouds now sample the Gaussian intensity profile of beam *C*, and when beam *C* has some circularity, the two BECs thus sample the differential VLS across the radial extent of the beam.

4. Immediately after extinguishing beam *B* we begin the Ramsey sequence, while beam *C* is still on. During the Ramsey interrogation time, the two clouds continue to fall through the beam, but for short enough times ($\sim 250\,\mu$s) this can be neglected. The second $\pi/2$-pulse is then applied, closing the interferometric sequence.

The measured relative phase is thus proportional to the absolute VLS shift of beam *C*. Neglecting the axial dependence of intensity and any background magnetic field differences, the intensity difference the BECs sample by virtue of falling distances y_A and y_B results in a differential phase that takes the form

$$\frac{\Delta\phi}{2\pi T} = \alpha_V \left[I_C(y_A) - I_C(y_B) \right] \sin 2\theta \cos \varphi, \tag{7.30}$$

with y_A, y_B the y position of BEC_A and BEC_B within beam *C*, θ the polarisation imperfection of beam *C* (equivalently the rotation angle of the quarter-wave plate in beam *C*) and φ the angle the magnetic bias field makes to the beam *C* wavevector. This configuration can be implemented on either dipole beam, allowing for independent measurement of the VLS from each beam. It is also a spatially resolved measurement that can determine the VLS gradient if the vertical distance between the BECs is known.

There are a number of undesirable caveats associated with the delayed drop technique. For example, it is limited to short Ramsey interrogation times of $\sim 250\,\mu$s to prevent motional blurring of the measured VLS, in much the same way as the freefall gradiometry experiments were in Chapter 6. Although BEC_A feels some residual trapping force due to beam *C*, and as such falls less distance than if it were in freefall, there remains the question of how to identify a metric for a *nulled* VLS. Any difference in the Zeeman shift between the two BECs (with beam *C* on, beams *A* and *B* off) is due to the VLS of beam *C* alone and a background magnetic field gradient. When the relative phase between them is equal to the relative phase measured with the same spatial separations and no light on, the VLS is zero. The background magnetic field difference between the two BECs can be measured using freefall gradiometry.

Consider the first three stages of the experimental procedure outlined above. If we replicate the same spatial separation configuration for a VLS measurement with these steps, but extinguish the crossing beam in the final step, we perform a measurement of the background relative phase, without the VLS. The spatial separation between the two clouds for a measurement

performed with beam C on is essentially the same for a measurement performed with beam C off for sufficiently short Ramsey times ($\sim 100\,\mu s$). The background gradient, $\Delta\phi_{BG}$ can then be determined directly, and when $\Delta\phi_{\text{meas}} \equiv \Delta\phi_{BG} + \Delta\phi_{VLS} = \Delta\phi_{BG}$ (*i.e.* at the *nulling* angle θ_N), there is no VLS. The Ramsey time must still be kept short so that $\Delta\phi_{BG}$ remains a valid measurement of the background gradient and a suitable metric for defining $\Delta\phi_{VLS} = 0$.

7.4.2   In-trap measurement

An alternative measurement to delayed-drop is to keep the two BECs in trap and create an intensity difference between beams A and B, with both beams having common polarisation defined by a quarter-wave plate (*i.e.* both beams pass through it, and travel the same optical path). Aligning the magnetic bias field parallel to the split beams ($\cos\varphi = 1$ in Eq. 7.19) ensures maximum sensitivity to the VLS difference between beams A and B, whilst minimising any contribution to the differential signal from beam C. The interferometry sequence is then performed in-trap, the scheme is depicted in Figure 7.6(c) and proceeds as follows:

1. An intensity difference is applied between the two beams, so that the measured phase depends on the polarisation imperfection and intensity difference as $\Delta\phi \propto \Delta I \sin 2\theta' \approx 2\Delta I \theta'$ for a small imperfection θ'.

2. For a fixed common-mode polarisation, the VLS is revealed in the differential intensity dependence of the relative phase between the two BECs. The relative phase is related linearly to the intensity difference; the gradient of this line depends on the polarisation imperfection of the light.

3. When the quarter-wave plate is turned to some angle θ_N, the 'nulling angle' where $\Delta\phi$ is independent of ΔI, the light is linear and the VLS is cancelled.

4. The experimental aim is therefore to reduce the gradient $\partial(\Delta\phi)/\partial(\Delta I)$ to zero, which amounts to finding the nulling angle θ_N.

For these measurements, the relative phase is given by:

$$
\begin{aligned}
\frac{\Delta\phi}{2\pi T} &= \alpha_V \left[I_A(r) - I_B(r) \right] \sin 2\theta' \\
&= 2\alpha_V \Delta I \theta',
\end{aligned}
\tag{7.31}
$$

where we have assumed that the bias field is in the direction of beams A and B, so that $\cos\varphi = 1$. In the intervening step we have also suppressed

the dependence of the intensity in each beam on the position of the BEC; $\Delta I = I_A(r) - I_B(r)$. This assumes each beam is identical and each BEC is in exactly the same position within each beam, with the intensity difference driven by imbalancing the optical power. This assumption is invalid in practice for experiments, as misalignments and different beam waists mean each BEC samples a different local intensity. More will be said on this in Section 7.6, Eq. 7.31 is still valid for an intensity difference due to misalignments. The gradient of a $\Delta\phi$ vs. ΔI plot for a given polarisation imperfection θ' is thus

$$
\begin{aligned}
\frac{\partial(\Delta\phi)}{\partial(\Delta I)} &= (2\pi T)\, 2\alpha_V \theta' \\
&= (2\pi T) 2\alpha_V (\theta - \theta_N).
\end{aligned} \tag{7.32}
$$

with $\theta' = \theta - \theta_N$ now represented as the difference between the quarter-wave plate rotation angle θ and the nulling angle θ_N. This gradient $\partial(\Delta\phi)/\partial(\Delta I)$ scales linearly with interrogation time, typical background gradient dephasing times of $\sim 15\text{-}30$ ms allow for substantially improved precision and thus more comprehensive nulling of the VLS than delayed drop measurements. The contribution from background magnetic fields is eliminated and phase aliasing does not need to be accounted for sufficiently small polarisations. In this sense, the delayed drop measurements serve as a useful precursor measurement that roughly obtains θ_N, an estimate that can be substantially improved on with the more sensitive in-trap method.

In contrast to the delayed drop measurements, the scheme discussed in this Section is a measurement of the relative polarisation between two beams. Both beams travel through similar optical paths and can be made to originate in a well-characterised common polarisation state (from a Glan-laser prism immediately after the AOM, as discussed in Section 7.1.1). We can approximate the absolute polarisation of a *single* beam passing through essentially the same optical path by the differential polarisation obtained by in-trap differential VLS measurements. We discuss the validity of this approximation at length in Section 7.6.5.

7.5   Characterisation of the VLS: delayed drop

The delayed drop technique allows us to measure the VLS from each beam independently with a dynamic range spanning the maximum and minimum VLS. Figure 7.7 shows the results of the measurement for both dipole beams, for different magnetic bias field directions. Each data point corresponds to a relative phase extracted from an ellipse containing 12 points, taken every 10-20° of quarter-wave plate rotation. The angle of the waveplate is not the
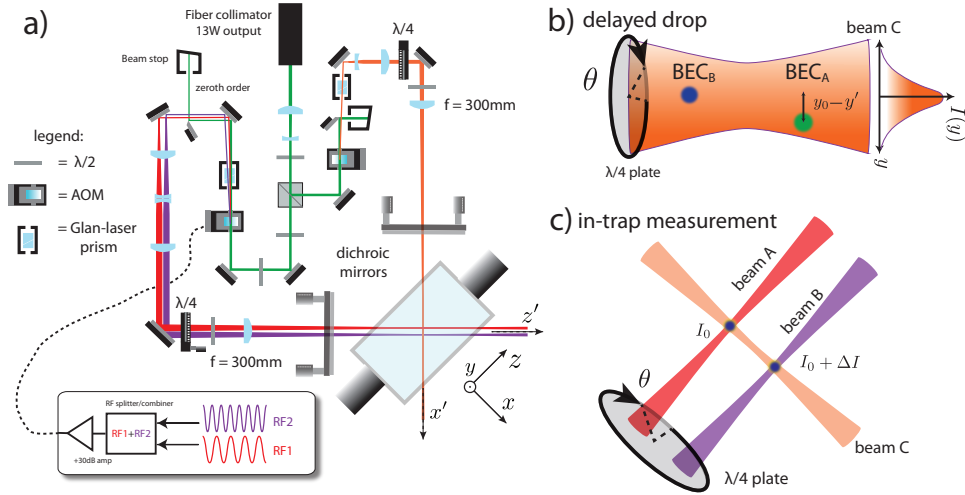
Figure 7.6: Schematic of the experimental apparatus (a) and two configurations to measure vector light shifts (b-c). The position and amplitude of each dipole beam is controlled by an AOM, introducing a second rf frequency to the AOM allows independent control of two dipole beams. Splitting either the x' or z' oriented dipole beams allows us to measure the difference in vector light shift at different locations of the same beam (b): the two BECs are positioned along the axial extent of the crossing beam C by splitting one beam into two; beams A and B. We then turn the beams off at different times to allow each BEC to fall a different distance, and sample a different intensity. Rotating the quarter-wave plate (QWP) in the beam then varies the VLS sampled by each BEC. Higher precision measurement is achieved by keeping the BECs in trap (c): we vary the power difference between beams A and B for a common polarisation of the two split beams, which then measures the intensity dependence of the VLS at a given quarter-wave plate angle.

orientation of the fast axis, and corresponds to its arbitrary alignment in the rotation mount; this manifests as some overall offset to the quarter-wave plate rotation angle in the Jones calculus discussed previously.

It should be noted that for these data, several differences exist between this setup and that described in Figure 7.6. The two Glan-laser polarising prisms were not present,[6] and the rotation mounts used were low-precision Thorlabs CRM1 stages with 2° increments. The quarter-wave plates used in all VLS experiments were zero-order (WPF-4225-1064) from Union Optics, the half-wave plate after each quarter-wave plate is also a zero-order type (WPF-2225-1064).

We measured the relative phase versus waveplate angle for three different

---

[6]The Glan-laser prisms were added incrementally after the completion of delayed-drop measurements, as we probed ever more precisely the VLS with in-trap measurements it became clear that the Glan-laser prisms were required for suppression of residual thermal transients.
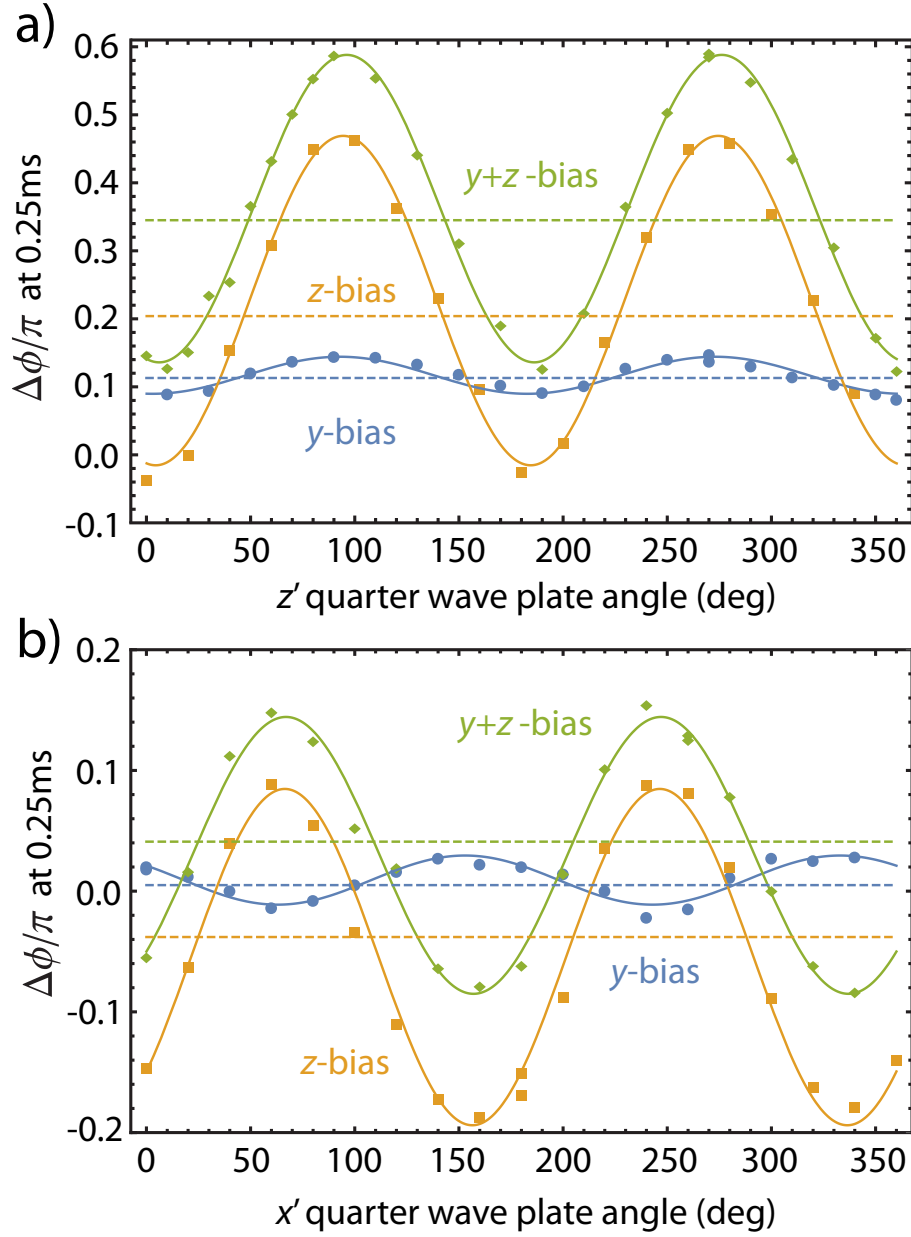
Figure 7.7: Measurement of vector light shifts from individual dipole beams: a) z'-oriented dipole beam and b) x'-beam, for three different bias field directions. Each BEC samples a different region of the radial intensity profile of the dipole beam, and thus a differential VLS. Rotating the quarter-wave plate in the dipole beam then changes the differential VLS. The applied bias fields are $B_z = (0, 0, 670\,\text{mG})$, $B_y = (\delta B, 293\,\text{mG}, 0)$ and $B_y + B_z$, where δB is a stray transverse component. The dashed lines represent the background magnetic field gradient contribution to $\Delta\phi$ from each field direction. Overlaid is a phenomenological fit from which we determine the amplitude of the VLS gradient (Eq. 7.30). Error bars on each point ($\sim 0.003\,\pi$) are smaller than the data points.

bias field directions: $B_z = (0, 0, 670\,\text{mG})$, $B_y = (\delta B, 293\,\text{mG}, 0)$ aligned along the z and y axes respectively (δB is a stray field component discussed below) as well the combined field $|B_z + B_y| = 731\,\text{mG}$. The Ramsey interrogation time is set at $250\,\mu s$. Each VLS measurement has a corresponding background gradient measurement determined by performing the delayed drop Ramsey interferometry sequence with the crossing beam (beam C) extinguished, as described in Section 7.4.1. Immediately it is observed that VLS has the expected $180°$ periodicity with waveplate angle.

However, when we consider the circularity for the light the atoms experience, represented by the Jones vector

$$\tilde{E}_{\text{atoms}} = \mathcal{A}_{\phi_k}(\theta_k)\,\mathcal{A}_{\lambda/4}(\theta)\,\tilde{E}_{\text{lin}}, \tag{7.33}$$

the functional form is more complicated (Eq. 7.29). Consider the matrix $\mathcal{A}_{\phi_k}(\theta_k)$, the phase retarder defined in Eq. 7.28 as representing the glass vacuum cell and any birefringent optics between the quarter-wave plate and the atoms. The Jones vector $\tilde{E}_{\text{lin}}$ is assumed to be linearly polarised. Based on Eq. 7.29, the general form of the circularity (and hence VLS) is more complicated than simply $\sin 2\theta$. Indeed, it is only the case for small ϕ_k that the light transmitted through the cell has a circularity that varies as $\sin 2\theta$: the small angle approximation for small ϕ_k in Eq. 7.29 yields:

$$\mathcal{C} \approx \sin 2\theta - \phi_k \cos(2\theta)\sin 2(\theta - \theta_k). \tag{7.34}$$

We can therefore conclude, from the data in Figure 7.7, that the cell is indeed only weakly birefringent. This is commensurate with what has been reported in other works [220]. The action of the quarter-wave plate in eliminating the VLS is twofold. The input state, which has been subjected to reflection off various optical elements after the AOM is slightly elliptically polarised. The quarter-wave plate corrects this, but does not make the output state linearly polarised, rather it prepares a state that when subjected to the slight birefringence of the subsequent optics and cell, becomes linearly polarised at the atoms.

We observe that for a y-bias field, which should eliminate the VLS by virtue of being perpendicular to the propagation direction of the laser ($\cos \varphi = 0$ in Eq. 7.19), there is a measurable VLS. This all but confirms our suspicions noted in Chapter 6 that a stray field component perpendicular to B_y exists. The magnitude of this stray field was later found to be of order $70\,\text{mG}$. The amplitude of the VLS in the x'-oriented dipole beam (waist $\sim 100\mu m$, z'-beam $w_0 \sim 70\,\mu m$) is less than that of the higher-intensity z' beam, as expected. We also note that the phase of the B_y data in the z' and x' plots: it is in phase and $180°$ out of phase respectively with the primarily
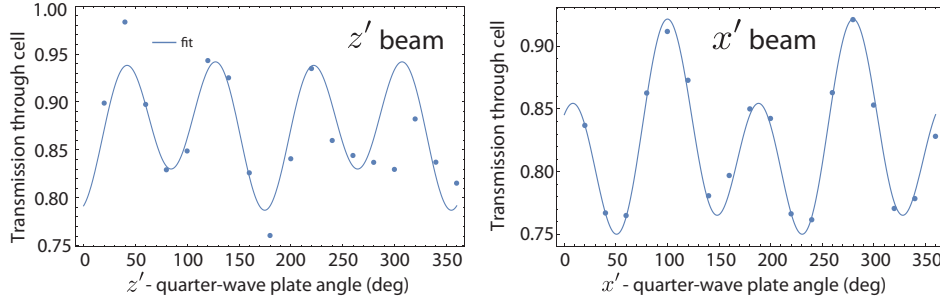
Figure 7.8: Fresnel losses through cell as a function of quarter-wave plate angle for both the x' and z' dipole beams. Overlaid is a fit that models the transmission of light through the birefringent optical path (cell, optics, etc.) and reproduces the observed variation of transmitted power with waveplate angle well. The larger residuals in the z' dataset is a result of limited optical access to the beam, making it considerably more difficult to measure power reproducibly with an optical power meter.

z-biased data. This can be understood by considering the projection of the z-bias field on the wavevector $\hat{\mathbf{k}}$ of each beam, which determines the sign of the VLS. The sign of $\hat{\mathbf{k}}_{x'} \cdot \hat{\mathbf{y}}$ is the same for both bias configurations (which suggests the stray component for a y-bias field is most likely due to imperfect nulling of the x-bias field), whereas for the z-biased measurements $\hat{\mathbf{k}}_{z'} \cdot \hat{\mathbf{z}}$ and $\hat{\mathbf{k}}_{x'} \cdot \hat{\mathbf{z}}$ are opposite sign. Since the y-bias field used had only $300\,\mathrm{mG}$ amplitude, any stray component perpendicular to it of only small amplitude is enough to tilt the field direction so that we measure a VLS shift.

A small confounding effect is that the variation of the circularity of the light with waveplate angle leads to a varying transmission of light through the cell. This results in a changing intensity of the trapping light. The reflection of light off the vacuum cell is quite severe: it is *not* anti-reflection coated. The intensity variation is determined by measuring powers before and after the cell and using the Fresnel equations to determine the fraction of light transmitted through the cell. Shown in Figure 7.8 is the variation of the transmitted intensity fraction with waveplate angle. We can use the Jones matrix formulation of polarised light – specifically that described in Eq. 7.33 – in conjunction with the Fresnel equations to determine the functional form of the transmitted intensity.

We can then fit to the transmission data in Figure 7.8 a model with free parameters for the orientation of the linear polarisation incident upon the quarter-wave plate and the phase ϕ_k and orientation θ_k of the unknown composite phase retarder (this does not include the cell), as well as an offset angle to account for the orientation of the waveplate axis. This is overlaid with the transmission data in Figure 7.8, showing agreement for both beams. The

optical power inside the cell can then be determined using the fitted values and considering only passage through one of the cell walls. This calculation will be important later when we consider the intensity calculation for in-trap measurements of the VLS. The overall intensity variation is of order 10-20%, which maps to a commensurate variation of the VLS amplitude in the delayed drop measurements across the full rotation range of the waveplate.

There are several complications to this analysis that are worth mentioning. The analysis does not take into account the birefringence of the cell, with polarisation state of light remaining unchanged during passage through the cell. This is concordant with the observations of Figure 7.7 and Eq. 7.34, and is not expected to adversely affect the inferred transmittance. The fitted parameters extracted for the phase shift ϕ_k and angles θ_i, θ_k do not correspond to uniquely defined permutations of optical elements, nor do they serve to definitively inform the circularity inside the cell independent of VLS measurements.

Without adjusting the rotation axis of the approximately linear polarisation state incident on the quarter-wave plate it is not possible to achieve fully right or left circularly polarised light at the atoms: this reduces the overall amplitude of the VLS as a function of quarter-wave plate angle. An elliptically polarised state input to a quarter-wave plate reduces the maximum range of the output circularity. Rotating this input state (with a half-wave plate, for example) can rectify this. In our experiment, however, the half-wave plate is located after the quarter-wave plate:[7] our aim is to eliminate the VLS, not maximise it. Both configurations can be used to eliminate the VLS (only a quarter-wave plate is essential), but only the latter can achieve maximum circularity. This will be discussed more in Section 7.6.5.

We extract the quarter-wave plate angle that minimises the VLS by fitting sinusoids to the data and determining the point at which the differential phase from the delayed drop VLS measurement is equal to that measured from background gradients alone: $\Delta\phi \approx \Delta\phi_{\text{BG}}$. For the z-biased data, the nulling angles are measured to be $46.6(6)°$ and $24.9(1.1)°$ for the z' and x' beams respectively. At best, we are able to match $\Delta\phi_{\text{meas}} = \Delta\phi_{\text{BG}}$ within fitted ellipse uncertainties, so for an absolute upper bound to VLS suppression at best we could achieve around $0.4\,\%$ of the maximum value, assuming $I_0 = 8 \times 10^3\,\text{W cm}^{-2}$, $\alpha_V = 0.336\,\text{Hz W}^{-1}\,\text{cm}^2$ and $\sigma(\Delta\phi) \approx 0.003\,\pi$. In practice the suppression is limited to $2\,\%$, the typical fitted uncertainty divided by the maximum imparted VLS amplitude. We use the higher precision in-trap measurements discussed later in this Chapter to rigorously eliminate the VLS.

---

[7]For space reasons and in the interests of preparing a state that is reflected maximally off the dichroic mirrors

The VLS-induced gradient is determined by measuring the differential distance $\delta y = y_A - y_B$ between the BECs along and converting the relative phase to a frequency difference: the radial intensity gradient along y is assumed to be responsible for the VLS gradient, the axial separation along the beam is not relevant.[8] Due to the residual trapping force of beam C retarding the fall of BEC$_A$, the intensity variation with waveplate angle means that this distance varies slightly with waveplate angle as well, although our observations suggest this effect to be small. Around the vicinity of the nulling angle, the differential distances fallen in the z' and x' beams respectively were determined to be 41.7 μm and 35.7 μm respectively. The VLS gradient induced for small angle changes around the nulling angle can then be determined as $8.2(2)\,\text{mG}\,\text{cm}^{-1}\,\text{deg}^{-1}$ for the z' beam and $5.5(2)\,\text{mG}\,\text{cm}^{-1}\,\text{deg}^{-1}$ for the x' beam. This calculation assumes that the radial intensity of the dipole beam varies linearly between the two BECs, which is by no means guaranteed. The reported gradients are almost certainly underestimates, but in any case are severe. The polarisation imperfections induced by thermal transients of the AOM need only be as small as 1-2° to create gradients comparable to the background gradients measured in Chapter 6.

7.5.1   Interferometer contrast

We can observe the contrast of Ramsey fringes as the quarter-wave plate is rotated around the nulling angle as further confirmation that our technique eliminates the VLS. The fringe contrast is affected by inhomogeneous dephasing brought about by spatial variation of the VLS shift as well as background magnetic gradients. Varying the waveplate angle about the nulling angle at fixed Ramsey time allows the effect of the VLS gradient to be observed. The results for this measurement are shown in Figure 7.9. Each contrast measurement is composed of 20 measurements of a phase domain Ramsey fringe[9] at an interrogation time of 15 ms, with a B_z bias field. The contrast is clearly seen to peak around the nulling angle, indicating suppression of inhomogeneous dephasing when the VLS is nulled. The offset of the peak contrast point from the nulling angle is due to the fact that the VLS gradient may partially cancel the background magnetic field gradient contribution, leading to a higher contrast.

---

[8]Axial intensity variation along the single beam is weak, with the Rayleigh range on the order of 1 cm.

[9]The absolute fringe phase sampled with each measurement at $T = 15\,\text{ms}$ is random; the contrast is the maximum fringe amplitude from 20 repeated measurements at fixed Ramsey time.
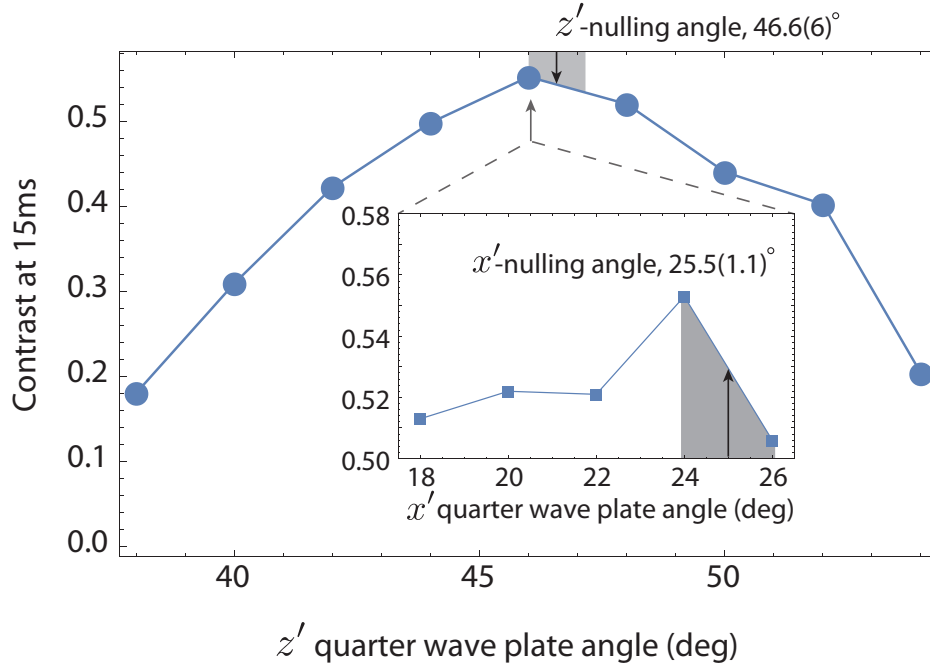
Figure 7.9: Interferometric contrast improved by VLS cancellation. The vector light shift is spatially inhomogeneous due to gravitational sag, resulting in spatial dephasing that reduces the contrast of Ramsey fringes. Varying the quarter-wave plate that tunes the VLS amplitude shows the fringe contrast peak near the nulling angle for z' dipole beam (main figure) and the x' beam (inset). The small shift from the measured nulling angle may be the result of partial cancellation of a component of the background gradient, resulting in a non-zero VLS but higher contrast. Shaded regions represent uncertainty in the nulling angle, calculated in previous Section.

7.5.2 Spin-mixing dynamics in a VLS gradient

We drew special attention to the effects of gradients on coherent spin-mixing dynamics in Chapter 5, in particular the nefarious time-varying nature of VLS gradients originating from thermal AOM-induced polarisation transients. The delayed drop measurements described previously did not appear to suffer from these effects, which would manifest as shot-to-shot horizontal scatter of points. We did not observe this, which may be due to the lower rf powers dissipated in the AOMs, by virtue of attenuation in the splitter-combiner used in these experiments. We did not use splitter combiners during spin-mixing experiments, and our measurements of irreproducible drift behaviour suggested the effect was strongest when the *initial* rf power was highest. Initialising the dipole trap AOMs to have maximum rf amplitude at the magnetic trap switch-on, rather than the final power at the end of optical evaporation lead to the most severe irreproducibility. The dipole

beams stay at their maximum amplitude for the majority of the time they are on (6 s every run), this led to a gradual heating of the AOMs over the course of ~ 10 experimental cycles. The polarisation state exiting the AOM became more elliptical with each run as the temperature increased. The splitter-combiners attenuate the input power by approximately 10 dB, hence reducing the overall power in the AOMs and thus the thermal polarisation transient.

We can now show the effect of a VLS induced gradient on spin-mixing dynamics directly. We used the results of Section 6.7.1 to find a magnetic field direction that minimises the background magnetic gradient so that the contribution from the VLS is easily seen. From Figure 6.10, we see that the effective gradient is only ~ 4 mG/cm for a x'-oriented bias field, which additionally allows for maximum sensitivity to the x'-oriented beam VLS. With this reduced gradient, we are able to see spin-mixing dynamics in the trap used for the VLS measurements – the semi-isotropic trap – despite the small value of $c = -2\pi \times 3.2$ Hz. Shown in Figure 7.10 are the spin-mixing dynamics and spatial separations with approximately nulled VLS at a bias field of 372 mG ($q = 2\pi \times 10$ Hz), and a measurement with a large VLS shift (and hence a VLS induced gradient of $132(6)$ mG/cm). These observations are consistent with our findings in Chapter 5: spin-mixing dynamics are preserved when spatial separation between the spin components is minimised. With a large VLS gradient applied, strong component separation along the y-axis is observed, that has consequentially deleterious effects on spin-mixing dynamics. Not shown is the component separation along the axis perpendicular to y, which is prevalent in both low-VLS and high-VLS cases, a result of the residual background gradients. This possibly accounts for the observed drift of $m_F = 0$ population present in the spin-mixing data consistent with that seen in Chapter 5.

Further investigation into spin-mixing dynamics in various, well-calibrated gradients would be particularly interesting to consider. The application of a range of gradients of varying magnitude could potentially reveal the mechanism that drives the evolving spinor to the ground state: note the slope towards *lower* ρ_0 after the initial peak near 25 ms, which differs from that observed in previous experiments. Also, we observe the spin projection (magnetisation) F_z appears to no longer be conserved. This has been the study of a previous theoretical work that explores the consequences of non-conserved magnetisation on spin-mixing dynamics [221], and may be interesting to pursue.

As an aside, we mention here the vastly longer spin-mixing times available with a nulled VLS and the x' bias field. Figure 7.11 shows population oscillations persisting to almost one second. Under these circumstances, the
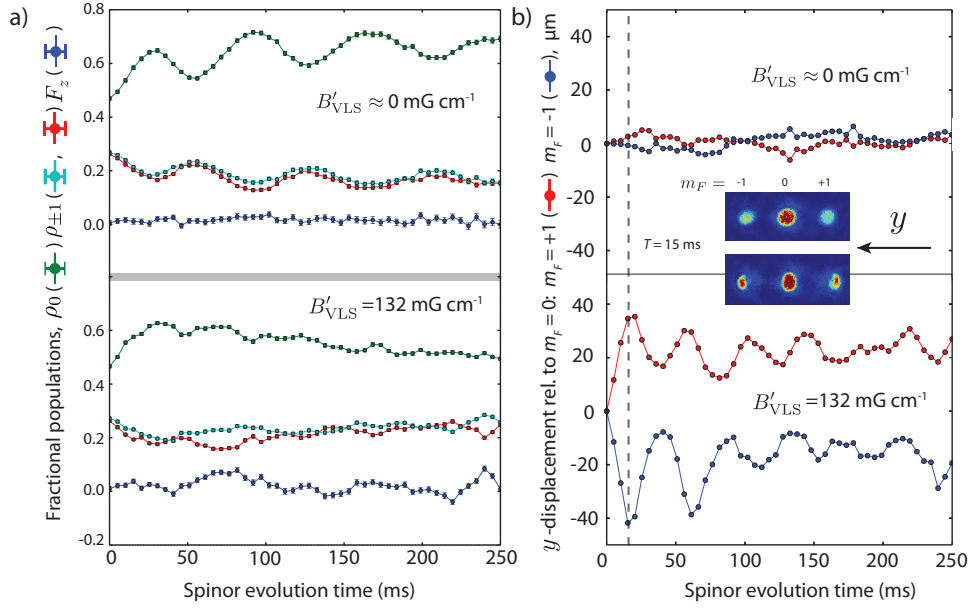
Figure 7.10: Spin-mixing dynamics of a spinor BEC with the VLS approximately nulled (top) and when set to $132(6)\,\mathrm{mG\,cm^{-1}}$: fractional population evolution (left) and centroid displacement, relative to the $m_F = 0$ centroid position (right). The VLS induces an effective magnetic field gradient along the gravitational direction (y), driving strong component separation. In this regime, the spatial overlap of the three Zeeman sublevels is reduced, suppressing spin-mixing from the $m_F = \pm 1$ states into the $m_F = 0$ state.

effects of gradients are sufficiently weak to render this system eminently suitable for realising weak spinor physics, allowing long-timescale dynamics to be observed. This result highlights the power of performing a rigorous characterisation of the magnetic landscape of a quantum gas.

7.6 High precision cancellation: In-trap VLS measurements

The sensitivity of differential Ramsey interferometry to a magnetically or optically induced Zeeman shift increases linearly with interrogation time. The maximum interrogation time that can be reached in delayed drop is limited by the effect of motional blurring. Moreover, background magnetic field gradients contribute to the measured differential phase; these must be subtracted to infer the VLS suppression. The in-trap measurements discussed in Section 7.4.2 allow much longer interrogation times, limited by decay of contrast (dephasing time). Measuring differential phase as a function of differential intensity naturally eliminates any contribution from background
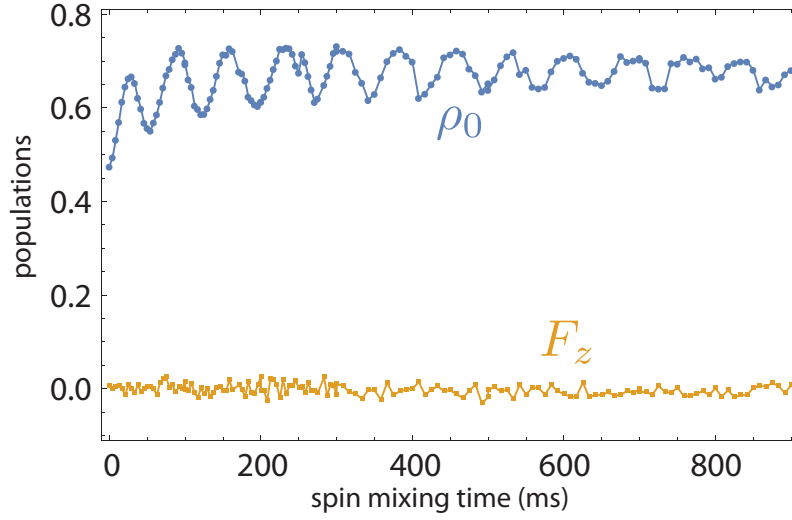
Figure 7.11: Spin-mixing oscillations with an x' bias field and nulled VLS. The effective background gradient is around $4\,\mathrm{mG/cm}$, resulting in spin-mixing oscillations preserved up to almost $1\,\mathrm{s}$.

magnetic gradients. The task of eliminating the VLS is reduced to finding a quarter-wave plate angle that minimises the dependence of the relative phase on the intensity difference between the two traps.

We assume that the quarter-wave plate angle defines the common polarisation of each of beam A and B. Unlike the delayed drop method, this is fundamentally a relative polarisation measurement: both beam A and B to a good approximation samples the same optical path dependent birefringence as the other. There will inevitably be a small polarisation difference between the two beams due to spatially varying birefringence of the cell, or different angles of incidence onto birefringent optical components between the quarter-wave plate and cell. Eliminating the VLS with the in-trap technique may not map directly to true linear polarisation in the case of a single un-split beam, but it *does* eliminate optical Zeeman shifts in a magnetic gradiometry experiment. Later in Section 7.6.5 we make comments on how the magnitude of differential polarisation may be accounted for.

7.6.1   Procedure for in-trap VLS measurements

The experimental procedure for the in-trap technique is as follows. For each quarter-wave plate angle, we incrementally vary the intensity in one of the split beams (beam A) by changing the corresponding rf power in channel A, and for each optical power imbalance $\Delta P = P_A - P_B \propto \Delta I$ we measure a relative phase $\Delta \phi$ by performing in-trap Ramsey interferometry and fitting

to an ellipse. Based on Eq. 7.32, $\Delta\phi$ should vary linearly with intensity difference and (in the small angle regime, near the nulling angle) linearly with quarter-wave plate angle θ. For each waveplate angle we thus extract the gradient of $\Delta\phi$ vs ΔP, which is proportional to the difference in VLS between the beams. When this gradient is zero, the beams are linearly polarised and the VLS is zero.

Assuming that the quarter-wave plate indeed defines the overall polarisation of beams A and B, when the intensities of the beams are equal, the measured phase will be the same for all waveplate angles. Each BEC samples the *local* intensity of the trapping light at its respective trap minimum, posing a complication: the local intensity is dependent on the optical power, beam waist and beam alignment (*i.e.* where the crossing beam intersects). The local intensity is different for both traps: equal optical powers and beam waists does *not* correspond to equal intensity. There is no simple, independent way to measure the intensity sampled by each BEC. A full numerical calculation of the trapping potential from measured optical powers and beam waists is non-trivial, and there is also the problem of astigmatism induced by the glass cell that is difficult to account for with *ex vacuo* measurements and calculation. However, any local intensity mismatch between the two traps can be accommodated by an optical power difference ΔP between the beams, so that $I_A(P_A, w_A, \mathbf{r}_A) = I_B(P_B, w_B, \mathbf{r}_B)$, with I_A, I_B the intensity profiles of beams A and B, with waists w_A, w_B and $\mathbf{r}_{A,B}$ defines the position of each BEC within each beam. The optical power difference ΔP_0 at which $I_A = I_B$ can be determined by measuring $\Delta\phi$ vs ΔP for different quarter-wave plate angles: the point at which all linear curves $\Delta\phi(\Delta P)$ intersect defines ΔP_0. Measuring and calculating the local intensity is discussed more in Section 7.6.3.

Optical setup differences

Two changes were made to the optical setup for in-trap VLS measurements, both of which reflect the higher sensitivity of the technique. We replaced the quarter-wave plate rotation mount with a high-precision Thorlabs CRM1-P rotation mount, which has 2.5 arcmin resolution on a micrometer drive. This change inevitably means the absolute angle determined in the previous Section had to be re-measured, owing to the arbitrary alignemnt of the waveplate fast axis within the rotation mount. The new nulling angle for the z' dipole beam was determined using delayed drop measurements to be in the vicinity of $\sim 340°$. For all in-trap measurements, we measure the differential VLS from only the z'-oriented dipole beam, applying a bias field parallel to this direction for these measurements for maximum sensitivity

but also for maximum *insensitivity* to the approximately nulled[10] VLS of the orthogonal x'-beam.

Initial experiments used Ramsey times of 3 ms and showed very strong fluctuations of $\Delta\phi$ consistent with drift as the experiment 'warmed up'. This drift was comparable to the change in $\Delta\phi$ imparted with changing the power imbalance, and thus represented a significant barrier to progress. We immediately identified this with the same drifting VLS observed in spin-mixing experiments and concluded it was due to thermal *polarisation* transients from the AOM, and not power drifts[11] as was previously thought. This drift was not observed in the delayed drop measurements, presumably due to the reduced sensitivity and coarser sampling of waveplate rotation angles: the 3 ms in-trap technique is $12\times$ more sensitive by virtue of interrogation time. We estimate that the AOM-induced polarisation drift is on the order of no more than several degrees: too small to affect delayed-drop measurements but large enough to be problematic for 3 ms in trap measurements, where the waveplate is rotated by only $5°$. In Section 7.5, we saw how polarisation imperfections on the order of several degrees resulted in effective gradients of $\sim 10\,\mathrm{mG/cm}$. This is commensurate with the observed drift that thwarted reproducible spin-mixing experiments in Chapter 5, providing a satisfying (if still disappointing) resolution to this problem.

Two Glan-laser (GL) polarisers (Thorlabs GL-5 and GL-10 respectively) were placed immediately after the AOMs on the z' and x' dipole beams. The GL polariser removes any polarisation drifts, differential or otherwise, induced by the AOM, resulting in an exceptionally well defined linear polarisation state for both beams A and B. The polarisation deviations induced by the AOMs were therefore converted to common-mode power variations on the same order ($< 2\%$) which are far less problematic.

7.6.2 Results of in-trap VLS measurements

Our initial experiments used Ramsey interrogation times of 3 ms and quarter-wave plate rotation ranges of $5°$. We used equal rf drive powers for each beam, which corresponds to $\sim 550\,\mathrm{mW}$ in beam A, and $730\,\mathrm{mW}$ in beam B, a 33% optical power imbalance due to asymmetry of the AOM diffraction bandwidth. We vary ΔP_{rf} between -0.04 and 0.04, and for each value of ΔP_{rf} we extract the relative phase from an ellipse formed from 12 varied Ramsey $\pi/2$-pulse phases. The results are shown in Figure 7.12. In order to determine the rf drive power imbalance required for balanced intensities, we

---

[10]Approximately nulled using the delayed drop technique.

[11]Previously we assumed that the AOM was minimally perturbative to the polarisation, and the observed drift was due to to power drifts at a fixed, non-zero polarisation.

measure the intersect of several $\Delta\phi$ vs. ΔP_{rf} curves (top panel, Figure 7.12), where ΔP_{rf} is the rf drive power imbalance (in arbitrary units). The ΔP_{rf} axis in the Figure is scaled so the intersect point (approximately $\Delta P_{rf} = 0.0052$ in absolute units) is zero. Given that this corresponds to an optical power imbalance of around 30%, the degree to which beam alignment determines the local sampled intensity is clearly evident.

The suppression of the VLS is determined by computing the ratio of the smallest achieved $\frac{\partial(\Delta\phi)}{\partial(\Delta P_{rf})}$ to the maximum VLS. We define the gradient of each $\Delta\phi$ vs ΔP_{rf} plot as

$$
\begin{aligned}
M(\theta) = \frac{\partial(\Delta\phi)}{\partial(\Delta P_{rf})} &= (2\pi T) \times \alpha_V k \sin 2(\theta - \theta_N) \\
&\approx (2\pi T) \times 2\alpha_V k(\theta - \theta_N) \quad\quad (7.35)
\end{aligned}
$$

where k is the constant of proportionality that converts power imbalances to intensity differences; $\Delta I = k\Delta P_{rf}$ (discussed in the next Section). The slope of the M vs. θ plot is $S = 2\pi T \times 2\alpha_V k$. Denoting the smallest achieved gradient as $M_{min}(\theta)$, the suppression ratio is

$$
\begin{aligned}
M_{min}(\theta)/M_{max} &= \frac{(2\pi T) \times k\alpha_V 2(\theta_{min} - \theta_N)}{(2\pi T) \times k\alpha_V} \\
&= 2M_{min}/S. \quad\quad (7.36)
\end{aligned}
$$

The suppression ratio is thus the ratio of the smallest achieved M to half the slope S. Using this metric, we determine that the VLS is suppressed by a factor of 230, or to almost 0.4 %.

Higher precision measurement and nulling is obtained by increasing the Ramsey interrogation time. For the z' bias field, the interferometric contrast peaks around $T = 15$ ms: quadratic Zeeman contrast nulls at 7 and 22 ms and background gradient dephasing after 20 ms makes 15 ms a reasonable trade-off between long evolution time and contrast. These measurements are therefore 5× more sensitive than the previous data taken at 3 ms Ramsey time.

During initial experiments using 15 ms interrogation times we observed that the uncertainty on fitted gradients close to the nulling angle increased. We attributed this to polarisation fluctuations in the x' (crossing) dipole beam, which we assumed we were insensitive to by virtue of bias field orientation. A GL polarising prism was added to the x'-dipole beam setup, and the VLS nulled again using the delayed-drop technique: the addition of the GL prism resulted in an overall drop in power of the x'-beam that was compensated for with an rf drive power increase to the AOM. The optical power change was not precisely calibrated and caused a shift in the measured VLS nulling angles for both beams by about 1°, presumably due to a different local intensity
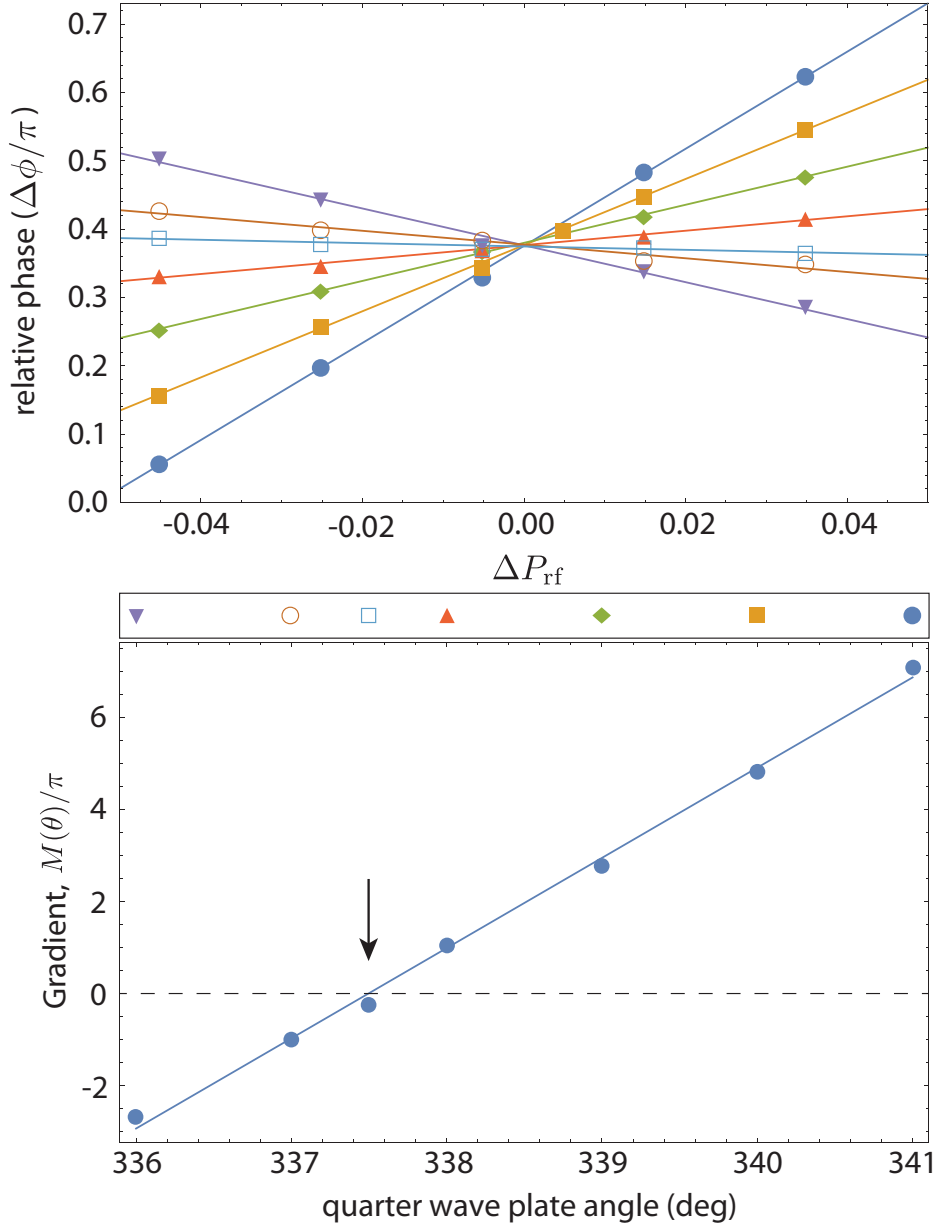
Figure 7.12: In-trap measurement and nulling of the VLS at 3 ms Ramsey interrogation time. The top panel shows the variation of $\Delta\phi$ with power imbalance ΔP_{rf} for different quarter-wave plate angles (symbols map to angles in lower frame). In the lower frame, the gradient $M(\theta)/\pi$ (Equation 7.35) of each of the $\Delta\phi$ vs.ΔP_{rf} datasets is plotted against waveplate angle. When $\Delta\phi$ is independent of intensity, the gradient is zero and the VLS is nulled. This occurs around an angle of $\theta_N = 337.5°$ (arrow). At this angle, the VLS is suppressed by a factor of almost 230. Error bars on each plot are smaller than data points.

being sampled in the split beams, which is dependent on the alignment of the crossing beam, or removal of residual x' VLS. We will forgo presentation of the final 15 ms dataset until after we have discussed the calculation of the intensity in the following Section.

7.6.3   Calculation of the intensity

In order to determine the values of vector shifts in more useful units (namely $\mathrm{Hz\,W^{-1}\,cm^2}$) we need to know the optical intensity difference sampled by the BECs, which is no trivial undertaking. Optical power and beam waists can be determined very accurately *ex vacuo* using standard optical test equipment (power meters and beam profilers). However, is not simple to estimate the overall trapping potential, including the crossing beam and gravitational sag, which ultimately determines the trap minimum and hence the position of the BEC in each of the traps. As we have indirectly seen, the relative alignment of the crossing beam with beams A and B dramatically affects the local intensity sampled by each cloud.

We considered a number of strategies to infer the intensity sampled by each BEC. To begin, we measure the waists of beams A and B, and optical power as a function of ΔP_{rf}. Shown in Figure 7.13 are the results of these beam profiling measurements. Waists are measured *ex vacuo* by projecting both dipole beams upwards with a fold mirror placed in the path of the beams prior to the cell and translating a Thorlabs beam profiler (BC106-VIS) along both beams. The waists are calculated from fitted $1/e^2$ Gaussian widths. We see that the beams are both astigmatic and divergent, almost certainly a result of aberrations from the 300 mm dipole beam focusing lens, which may be severe given the beam spots are displaced from the optical axis of the lens. We determine the beam width 'at the atoms' by measuring the distance between the beam centroids on the CCD profiler. When this distance is equal to the distance between the two BECs (342 µm, measured from absorption images, as described in Section 6.4), the beam widths at that point are approximately equal to the beam widths at the atoms.

We see immediately that the trapping beam widths at the atoms $\bar{w}_A$, $\bar{w}_B = 66.25, 69.75$ µm are different (bar denotes mean beam width). This, in addition to the beam alignment difficulty contributes to the local intensity each BEC samples. Shown in Figure 7.14 is the optical power in each beam as a function of ΔP_{rf}. Due to the addition of the GL polariser in the x' beam, the equal-intensity intersect point had to be remeasured and was found to be $\Delta P_{\mathrm{rf}} \approx 0.027$. This corresponds to an optical power imbalance of about 15 %, and is not accounted for by the measured difference in beam widths as described below.
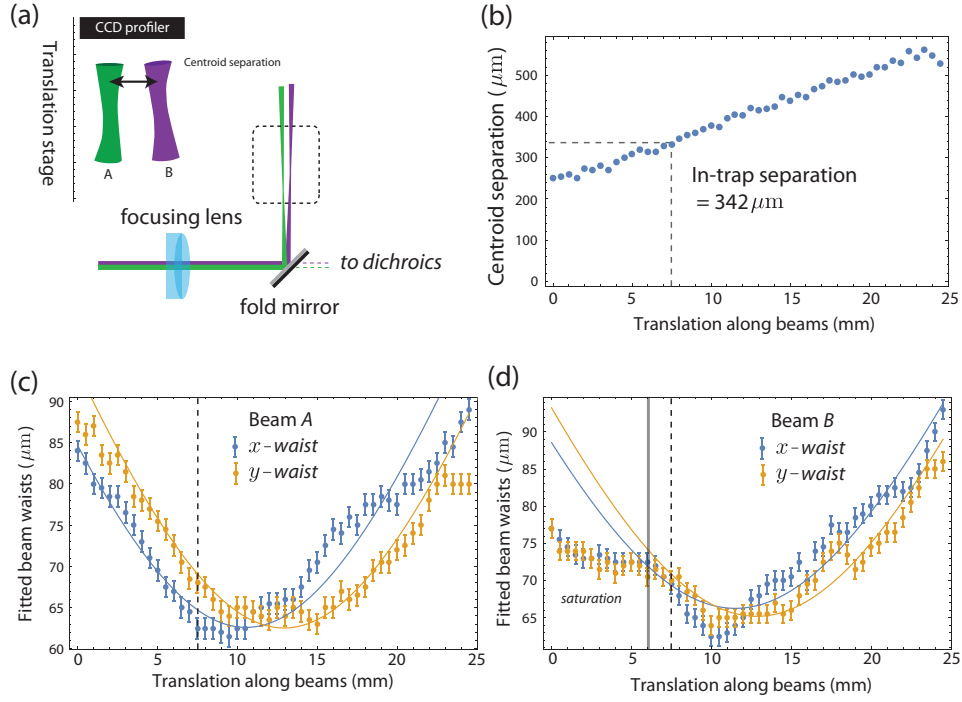
Figure 7.13: Determination of dipole beam waists at the positions of the atoms, (a) schematic of configuration to measure beam waists along the propagation directions of the two split dipole beams. The CCD beam profiler is used to measure the beam waists (c, d) as well as the separation between the fitted centroids (b). When the measured separation is equal to the inter-trap distance measured using absorption images *in vacuo*, the waists of beam *A* (c) and *B* (d) can be determined, at the position of the atoms. The measurements to the left of the solid line in (d) exhibit saturation effects in the beam profiler due to higher optical powers (which were subsequently reduced).

We estimate the trapping light intensity seen by the atoms by defining effective beam intensities $I_{A,B} = 2P_{A,B}/\pi w_{A,B}^2$, and assuming that at the balanced intensity point $I_A = I_B$ so that $w_A/w_B = \sqrt{P_A/P_B}$. We then substitute $w_A = 66.25\,\mu\text{m}$, the measured beam waist of beam A and infer $w_B = 71.2\,\mu\text{m}$. We therefore are able convert dimensionless DDS control units ΔP_{rf} into a differential intensity, ΔI. This is a simplistic assumption, which is expected to overestimate the intensity since we assume that BEC_A is subjected to the peak intensity of the dipole beam. In a single beam trap such an assumption is manifestly inadequate, but in a cross beam trap the additional trapping force of the second dipole beam *can* indeed result in a trap minimum that samples the peak intensity of one beam but not the other, in which case our assumption has a degree of validity.
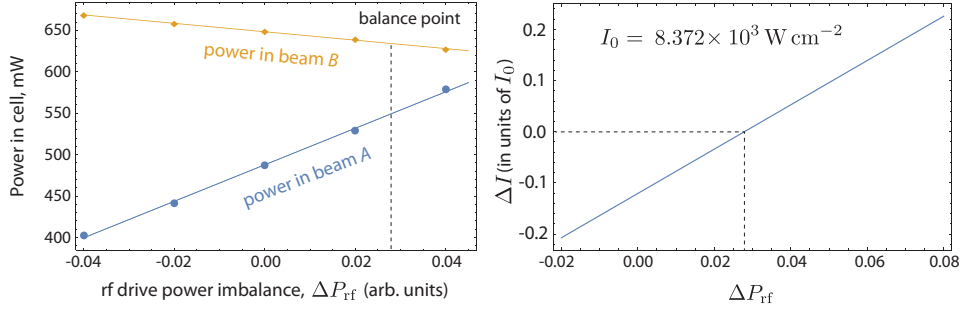
Figure 7.14: Calibration of rf drive power imbalance ΔP_{rf} to intensity imbalance. Left: optical power in the cell, calculated from measuring beam powers and accounting for the transmission through the cell ($T \approx 0.965$), as a function of rf drive power imbalance. Assuming the intersect point of different $\Delta\phi$ vs. ΔP_{rf} curves determines the power imbalance required for equal intensities, we can define effective beam intensities sampled by each BEC. For this data, the balanced intensity point is at 0.027, the ratio $w_A/w_B = \sqrt{P_A/P_B}$ allows us to determine the ratio of beam waists required for equal intensities based on our measurement of w_A. The differential intensity $\Delta I = I_A - I_B$ is then plotted as a function of ΔP_{rf} (right).

Absorption images of thermal atoms allowed to expand along the length of the dipole beams can be used to trace the position of the dipole beams, and hence characterise misalignment. The optical power in a single beam at the drive powers used for typical experiments is insufficient to support the atoms against gravity, so higher optical powers (and hence higher rf drive powers) are used. Thermal effects in AOMs cause the position of the dipole beams to change reproducibly as the rf power is changed (during optical evaporation, for instance). It is known that the position of the dipole beams change slightly as the rf power is is changed: the beam positions inferred from atoms tracing out the beam path at higher optical powers are therefore only approximate, and at lower powers the effect of gravitational sag is stronger. In any case we see that the AOM splitting axis is slightly tilted with respect to the crossing dipole beam C (possibly by as much as 8°), with the result that the crossing beam 'holds' the BECs at different positions relative to the peak intensity point of each of the A, B dipole beams. In conjunction with the different beam waists this effective misalignment accounts for the intensity difference between the two beams for equal drive powers.

A better estimate of the intensity could be made with the aid of a simple model of the trapping potential and our measured trap frequencies. The trap potential is given by

$$U(\mathbf{r})/h = -I(\mathbf{r})\left(\alpha_S + \alpha_V \sin 2\theta \, m_F\right), \tag{7.37}$$

with α_S at 1064 nm given in Table 7.1 and $I(\mathbf{r})$ the trapping laser profile

(Eq. 3.14). This model can be used to determine the trap depth, trap minimum (x_0, y_0, z_0) (*i.e.* position of the BEC in the beam), and in the harmonic approximation, trap frequencies from $\omega_x = (\frac{1}{m}\frac{\partial^2}{\partial x^2}U|_{x=x_0})^{1/2}$. A model with free parameters for beam alignment can be used to determine the configuration that most accurately matches the trap frequencies we measure using parametric resonance and sloshing, described in Chapter 3. We attempted this and found that simply calculating the trap frequencies in the harmonic approximation with our measured trap parameters typically gave trap frequencies ~ 2 times what we measure. The source of this error was never definitively identified. While phenomenological factors can be introduced to give trap frequencies (and intensities) consistent with what was measured, the discrepancy speaks to the simplicity of the model and the rudimentary way trap frequencies were calculated. The dipole potential, complete with misalignments and gravitational sag, is also anharmonic. A more detailed simulation that directly probes the oscillatory behaviour of a perturbed BEC may result in a better estimate of trap frequencies and hence an improved intensity estimate.

We note here that the problem of intensity overlap encountered here is one that frequently vexes experiments that seek to determine optically-induced energy shifts [197, 215]. For our purposes it is a systematic that frustrates efforts to quote values of light shifts in terms of physical units. This does not hamper our elimination of the VLS: when the measured $\Delta\phi$ is independent of ΔP_{rf} at some waveplate angle, the VLS is zero in any case. More will be said about the estimation of intensity in Section 7.6.5.

7.6.4   Elimination of the VLS and determination of α_V

With our rf-drive power imbalance ΔP_{rf} now able to be converted to an approximate intensity imbalance ΔI, we are able to determine the amplitude of the VLS in units of $\text{Hz}\,\text{W}^{-1}\,\text{cm}^2$. Figure 7.15 shows the results of measurements at a Ramsey time of 15 ms. Uncertainties on the measured gradients are derived from fitted error in each gradient. We see a number of differences between the data and that taken at 3 ms, the most obvious of which is the inverted sign of the gradient vs. θ plots, a result of aliasing.[12] The balanced intensity point is less clearly defined also, the scatter of intersection points is discussed later in this Chapter. The precision with which the technique senses small magnetic fields is also apparent. For two measurements in the top panel of Figure 7.15 separated by $0.02\,\pi$ (a frequency difference of 0.67 Hz),

---

[12]We assume that the sign of the vector polarisability is negative, and could confirm this by observing short Ramsey time fringes directly, and gradually increasing the Ramsey time to account for aliasing, as in Figure 6.5. We would necessarily also need to know the helicity of the light.

the corresponding effective magnetic field difference clearly resolved in this measurement is approximately 94 pT.

Each point in the lower plot of Figure 7.15 is an effective vector polarisability, $\alpha'(\theta) \approx 2\alpha_V\theta$, which follows directly from Eq. 7.36 once relative phases are converted to frequency differences and ΔP_{rf} converted to ΔI, $k = 4.03 \times 10^8\,\mathrm{W\,cm^{-2}}(\Delta P_{rf})^{-1}$. The minimum achieved effective VLS is $\alpha'_{min} = -1.5(5) \times 10^{-4}\,\mathrm{Hz\,W^{-1}\,cm^2}$. The slope of the line in the lower panel of Figure 7.15 is a measurement of the vector polarisability of ^{87}Rb at $\lambda = 1064\,\mathrm{nm}$, which we determine to be $\alpha_V(\mathrm{meas.}) = -0.331(18)\,\mathrm{Hz\,W^{-1}\,cm^2}$ in remarkable agreement with the theoretical value calculated in Section 7.2, $\alpha_V = -0.336\,\mathrm{Hz\,W^{-1}\,cm^2}$. The suppression of the VLS from Eq. 7.36 is $\alpha'_{min}/\alpha_V(\mathrm{meas.}) = 4(1) \times 10^{-4}$, independent of the uncertainty in optical intensity. Assuming the polarisation purity determined here is an equivalent metric of the polarisation of a single, un-split dipole beam with intensity $I = 8.37 \times 10^3\,\mathrm{W\,cm^{-2}}$ (the beam intensity at the intersect point calculated previously), the overall VLS is only 1.2(4) Hz, corresponding to an effective magnetic field of 1.7(6) μG. The nulling angle that eliminates the VLS is determined to be 337.115(3)°.

7.6.5 Experimental systematics and potential improvements

Intensity calibration

We make note here of several important systematics. We have chosen to express light shifts and polarisabilities in units of $\mathrm{Hz\,W^{-1}\,cm^2}$, which introduces a large systematic error due to the assumptions associated with the local intensity experienced by each BEC. It is difficult to determine the local intensity of each beam (beam A and B, the split dipole beams), and the relative alignment of the crossing beam could create a trap minimum very near the peak intensity point. In any case we assume that the method used here overestimates the intensity, which may be less than the peak intensity used here.

Birefringence of the cell

A competing systematic is associated with the maximum slope of the plot used to infer α_V (Figure 7.15), which is representative of degree to which one can achieve purely right (or left) circularly polarised light with only quarter-wave plate adjustments. The VLS we measure is essentially proportional to the circularity of an initially linearly polarised state, oriented at an angle θ_i from the vertical, passing through a quarter-wave plate and a ϕ_k-phase retarder, representing post-waveplate birefringent optics. The functional
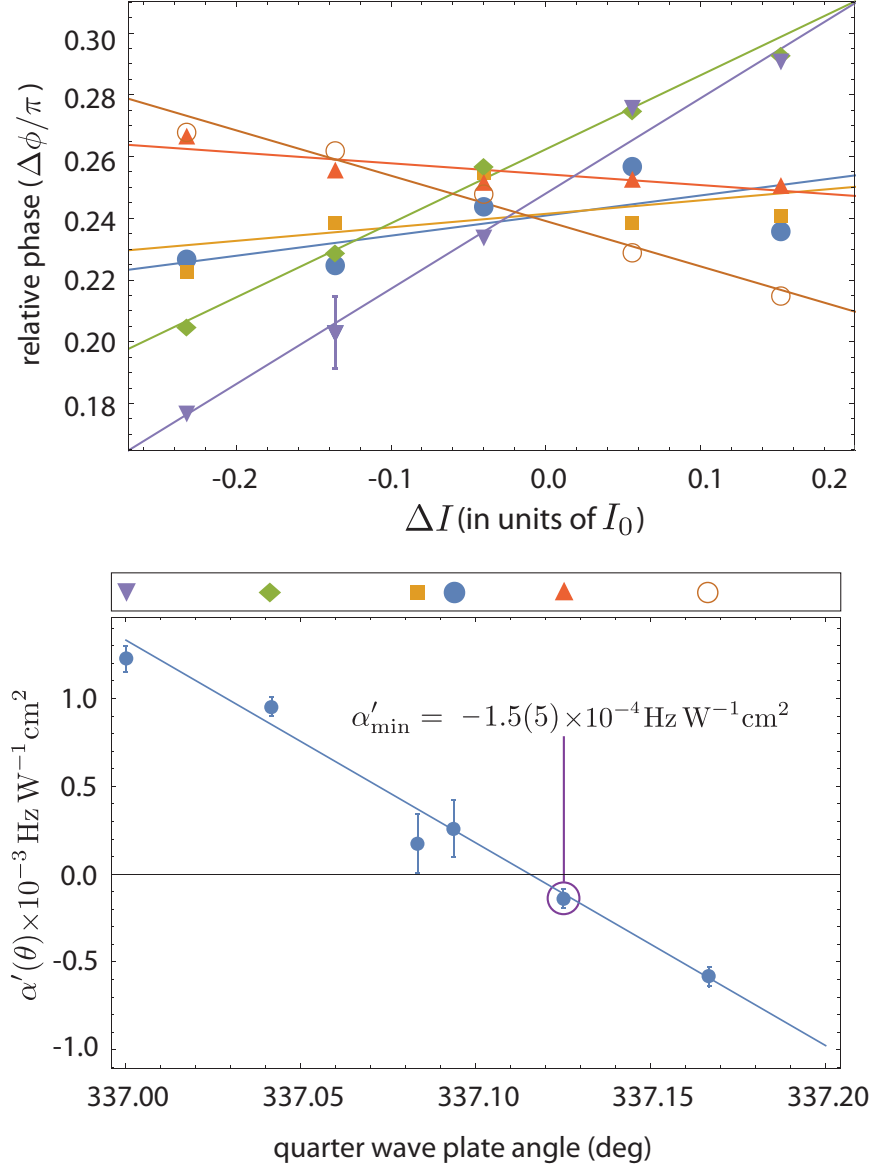
Figure 7.15: High precision measurement and cancellation of the VLS. The upper plot shows the variation of phase with intensity difference, for different quarter-wave plate rotations (this time over a range of 10 arcmin). The single error bar is representative of typical uncertainty reported from ellipse fits, $\delta \approx 0.011\,\pi$. The gradient of each line M is an effective polarisability $\alpha'(\theta)$, and is plotted vs. waveplate angle in the lower plot, error bars represent the uncertainty in each fitted gradient. The slope of this curve is equal to $2\alpha_V$. We can determine the minimum achieved VLS (indicated) and find the suppression ratio (Eq. 7.36) $\alpha'_{\min}/\alpha_V(\text{meas.}) = 4(1) \times 10^{-4}$.

form of the circularity is in general complicated (Equation 7.29), but for small phase shifts ϕ_k can be represented by Equation 7.34. The delayed drop VLS measurements show that the functional form of the circularity is very close to $\sin 2\theta$, indicating that ϕ_k is indeed small. The circularity when the polarisation state incident on the quarter-wave plate is rotated by an angle θ_i is

$$\mathcal{C} = \sin 2(\theta - \theta_i) - \cos 2(\theta - \theta_i) \sin 2(\theta - \theta_k)\phi_k, \tag{7.38}$$

with θ_k the orientation of the fast axis of the ϕ_k-phase retarder. The input linear state rotation angle is just an offset to the quarter-wave plate angle, substituting $\theta \rightarrow \theta + \theta_i$ yields

$$
\begin{aligned}
\mathcal{C} &= \sin 2\theta - \cos 2\theta \sin 2(\theta - \theta_k + \theta_i)\phi_k \\
&= \sin 2\theta - \cos 2\theta \sin 2(\theta - \delta\theta)\phi_k
\end{aligned}
\tag{7.39}
$$

where we have defined the angular difference between the fast axis of the ϕ_k-retarder and θ_i as $\delta\theta$. Assuming θ to be near the nulling angle, we expand Eq. 7.39 in a power series, retaining terms linear in θ:

$$\mathcal{C} = 2\theta - 2\theta\phi_k \cos 2\delta\theta + \sin 2(\delta\theta)\phi_k + \mathcal{O}(\theta^2). \tag{7.40}$$

The in-trap VLS measurements infer the slope S of a plot of $\partial E_{\mathrm{VLS}}/\partial \Delta I$ against θ to be $2h\alpha_V$. Since this is proportional to $\partial \mathcal{C}/\partial\theta \approx 2$, the correction to S due to a small birefringent phase shift is

$$\frac{\partial \mathcal{C}}{\partial \theta} = 2 - 2\phi_k \cos 2\delta\theta. \tag{7.41}$$

For $\delta\theta = \pi/4$ or $\phi_k = 0$, the maximum slope is achieved. This effect can reduce or increase the inferred α_V, depending on the quarter-wave plate orientation.[13] We believe this effect to be small, however.

Based on the possible systematic errors associated with intensity our measurement is a *lower* bound on the value of α_V. The combined uncertainty in the estimated intensity is almost 25 %, highlighting the difficulty in obtaining accurate intensity measurements. We note that in the face of these systematics it is quite remarkable, if not coincidental, to get a value for α_V close to what is expected from theory. This either suggests the systematics are small or that some other factor, such as the reduced/increased slope of S from Eq. 7.41 results in measuring a larger α_V, partially cancelling the intensity systematics mentioned previously. The alignment of the magnetic bias field along z' also reduces the measured gradient. The necessity of further work in rigorously quantifying these systematics is not a pressing matter for our application. It

---

[13]The circularity begins to resemble an asymmetric sawtooth as opposed to $\sin 2\theta$, so the slope at the waveplate angles where $\mathcal{C} = 0$ is alternately greater or smaller than $2\alpha_V$.

would be satisfying to account for them numerically to improve the absolute precision of the technique, although we have no *a priori* reason to question the real value as deviating significantly from the theoretical value based on the two-transition line (D1 and D2) approximation made. We stress again that *neither* of these systematics impede precision nulling of the VLS, there still exists a θ that satisfies $\mathcal{C} = 0$.

Rotation stage precision

The precision of the technique can be improved with longer Ramsey times, although this requires the background magnetic field gradient to be reduced to allow for longer dephasing times. Also, a higher precision rotation stage is necessary: the angular increments used in Figure 7.15 are 2.5 arcmin (0.7 mrad) and result in coarse sampling. We attempted to procure such a stage, but space constraints in the dipole trap setup precluded its immediate implementation. State-of-the-art digitally controlled rotation stages are able to achieve minimum incremental rotations on the order of 10 µrad, although absolute reproducibility can be poor. This is not expected to impede an empirical minimisation of the VLS, however.

Relative polarisation imperfections

One systematic that will affect the degree to which we can eliminate the VLS is the presence of optical-path dependent birefringence, which can be caused by spatially varying birefringent stress across the cell, or waveplate thickness variations. This induces a *relative polarisation imperfection*, which invalidates our assumption that high-quality linearly polarised light achieved using a differential polarisation measurement maps to similarly pure linearly polarised for a single beam (that follows approximately the same beam path). Despite this, the setup used for the VLS experiment is immediately applicable to in-trap magnetic gradiometry experiments, and our technique can be used to cancel the VLS at a range of different gradiometer baselines.

A relative polarisation imperfection would effectively result in a VLS that is impossible to completely remove with a differential measurement, although the discussion to follow suggests that this may be a very small correction. Differential interferometry of course offers a means to quantify the effect: we could perform measurements for a range of splitting separations and beam translations. This would allow the spatial dependence of the birefringence of the cell (and other optical components on the beam path) to be directly probed.

The increased scatter of points in $\Delta\phi$ vs. ΔP_{rf} plots (and corresponding increased uncertainty of gradients, Figure 7.15) in the vicinity of the nulling

angle is of some interest and may offer information about the magnitude of the relative polarisation imperfection. The balanced-power intersect point (Figure 7.15,top) is less well-defined than in the case of the 3 ms data, we assume here that the higher sensitivity of the 15 ms data reveals what appears to be an intensity difference induced by the waveplate rotation; although the results in Figure 7.16 suggest that this is small. There may indeed be a relative polarisation difference that is subject to the alignment of the dipole trapping beams, possibly induced by the waveplate itself. Such an effect could be easily investigated with *ex vacuo* polarimetry and would be interesting to pursue. Relative intensity noise, which may be of a transient nature[14] between the two split beams can also result in the observed noise of the intersect point, and could also be examined with standard polarimetry.

Common-mode rejection using balanced intensities

Referring to Figure 7.2(c), the spurious differential VLS signal in an in-trap magnetic gradiometry experiment can be eliminated by ensuring equal beam intensities. As part of the calibration we performed to convert rf drive power differences into intensity differences, the rf drive power imbalance that results in equal intensity in each beam was experimentally determined. Figure 7.16 shows the results of an imposed common-mode polarisation imperfection (from rotating the quarter-wave plate) with imbalanced intensities and when the rf drive power difference is set so that $I_A = I_B$. In this example, the Ramsey interrogation time is 15 ms and the quarter-wave plate is rotated over 10 arcmin: three ellipses at 5 arcmin increments are acquired and then whole dataset combined: when the power is imbalanced by $\sim 20\%$, the result is noise. With balanced intensities, there is no *differential* VLS imposed by rotating the waveplate, and so the points are tightly confined on a single ellipse. This scenario is hence a plausible alternative to eliminating residual elliptical polarisation for the purposes of in-trap magnetic gradiometry, free from the VLS.

Elliptical data reduction

We mentioned previously in this Chapter the limitations of elliptical data reduction, particularly in extracting a fit uncertainty that accounts for asymmetric sampling of the ellipse and typical scatter of points. The fit uncertainty associated with an ellipse depends only on the intrinsic uncertainty in each F_z measurement ($\sim 3\times$ SQL) and the number of phases sampled. For gradiome-

---

[14]Some noise sources appear and disappear on an irregular basis, which may be a result of ambient laboratory equipment and originate from temperature/humidity fluctuations in the building.
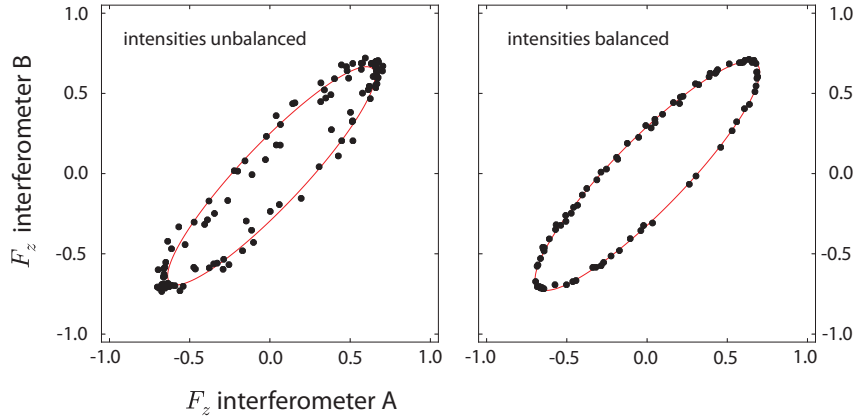
Figure 7.16: Balanced optical intensities in the two beams can be used for common-mode rejection of small, common mode polarisation imperfections. Left: Shown are the results of 9 ellipses, three each at quarter-wave plate angles differing by 5 arcmin. An intensity difference of $\sim 20\%$ between the two split beams results in a non-zero differential VLS, that manifests as scatter of points. Right: with the intensities in each beam equal, the differential VLS is zero, and so in a repeat of the previous measurement the points are constrained along a single ellipse. The measured $\Delta\phi$ is thus independent of the VLS, making this an ideal scenario for magnetic gradiometry.

try experiments at short interrogation times, an ellipse is generally uniformly sampled. At short interrogation times ($< 3\,\mathrm{ms}$), the azimuthal direction the second $\pi/2$-pulse rotates the spin vector about is well defined, and changing the pulse phase between 0-360° results in well-defined phase domain Ramsey fringes (Figure 6.3, top). The $\pi/2$-pulse phase can then be incremented with each shot to trace out an evenly sampled ellipse. At longer times, the rotation axis is completely scrambled by detuning perturbations and the common fringe phase sampled in each iteration of the Ramsey experiment is essentially random. The result of this is that an ellipse may not necessarily be evenly sampled by the normal 10-15 shots of different $\pi/2$-pulse phases.

Figure 7.17 shows the relative phase $\Delta\phi$ reported from an ellipse fit (with uncertainty) as a function of the number of shots ($\pi/2$-pulse phases) used in the fit. Ellipses with 12 points are used for most experiments in this work, which offers a good trade-off between fit uncertainty and duty cycle. Although the phase measured from 12 shots compares well with the relative phase from 24 shots (in this example), some ellipses are primarily sampled in a single quadrant for many shots, the reported $\Delta\phi$ may then be inaccurate from only 12 shots.[15] We speculate that some points near the

---

[15]Extreme examples of this problem are generally manually discarded.
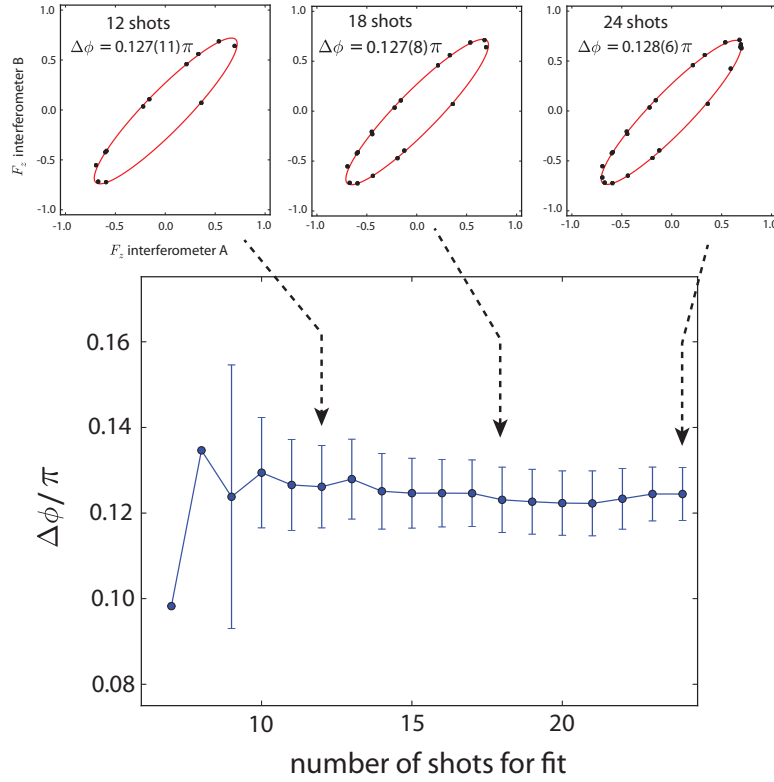
Figure 7.17: Relative phase $\Delta\phi$ and fitted uncertainty $\delta(\Delta\phi)$ reported from ellipse fits as a function of shots ($\pi/2$-pulse phases). Most analyses in this thesis use 12 points, a compromise between accuracy and duty cycle. The fitted $\Delta\phi$ uncertainty depends on the intrinsic uncertainty of each point and the number of shots used in the fit. Uneven sampling of an ellipse results in an inaccurate measurement of $\Delta\phi$, using more points in each ellipse fit also improves the accuracy of the fitted $\Delta\phi$ by sampling the ellipse more evenly.

nulling angle in the 15 ms data (outside the bounds of the sample error bar shown) may be subject to this effect and representative of the scatter of points for certain, poorly sampled ellipses where the reported $\Delta\phi$ uncertainty is an underestimate. This problem can be overcome in the short term by increasing the number of phases used to sample each ellipse.

Contribution from orthogonal beam

The greater scatter of $\Delta\phi$ data close to the nulling angle (present to some extent in the 3 ms data as well) still remains something of an enigma. We are yet to determine a satisfying explanation for this observation, even accounting for the effects described previously. A small relative polarisation imperfection that is time varying and random due to acoustic vibration of

optical components may be to blame. We also suspected it may be the result of residual VLS from the crossing beam (which we detect in quadrature) that is detected due to imperfect alignment of the magnetic bias field with the z' dipole beam. This may only be significant when the VLS from the z' beam is very small.

We performed a series of measurements to determine if this was plausible. We measure a differential VLS phase $\Delta\phi$ from the split z' beams, and then rotate the x' beam quarter-wave plate whilst a z'-bias field presumably renders the measurement immune to the x' VLS. We can thus determine to what extent the magnetic bias field is aligned with the z'-axis. We did not detect any substantial variation of the measured $\Delta\phi$ for x' waveplate deviations of up to $10°$ from the nulling angle. This observation strongly suggests that the spurious contribution of the x' dipole beam is not the cause of the additional scatter.

7.7 Absolute polarisation measurement with a spin echo sequence

Extremely high precision absolute polarisation measurement is possible using the vector light shift. While we have focused on the case of a relative polarisation measurement that well approximates an absolute polarisation measurement in this thesis, we can adapt the technique to function as an absolute polarisation measurement. Consider the experimental setup immediately prior to performing the delayed drop measurements, *i.e.* two BECs axially positioned along the crossing beam in two cross dipole traps; we intend to measure the absolute polarisation of the crossing beam. We have seen that the two clouds will in general always feel an overall intensity difference from the crossing beam due to small misalignments and different trapping beam waists (Section 7.6.3), so dropping the clouds is not essential. Also, aligning a bias field along the direction of the crossing beam eliminates the VLS contribution from the intersecting split beams. However, simply rotating the quarter wave plate that controls the polarisation of the crossing beam will induce a VLS phase that is indistinguishable from a magnetic gradient contribution. For small enough rotations, the effect of intensity variation due to transmission of polarised light through the cell can be safely neglected. The relative phase induced per wave plate rotation angle (which can be very small, $\sim 10\,\mu$rad) can then be assumed to be purely from the VLS and hence a direct, absolute measure of the polarisation of the trapping beam.

However, to make such a measurement of the polarisation worthwhile it is essential to decouple the magnetic field gradient contribution from that of

the VLS, which is essentially equivalent to a nulling measurement since we seek to find the relative phase that corresponds to a magnetic-gradient only contribution. Simply changing the intensity of the beam will translate the clouds, and hence change the magnetic gradient phase. In order to separate the effect of an intensity dependent VLS shift from a gradient magnetic field, we can modulate the trapping light and apply a spin echo pulse at the centre of the modulation cycle. In a normal spin-echo sequence applied to DRI, the phase developed by each interferometer will be zero, and thus the relative phase will be zero. Unlike background magnetic fields, we are able to change the strength of the VLS field with the intensity of the light. Before the spin-echo pulse, the relative phase of the system increases as $\Delta \phi = (\omega_1 - \omega_2)t$, with $\omega_i = \gamma B_i$. After the spin-echo pulse, the phase decreases at exactly the same rate, since the magnetic field at each BEC is the same before and after the spin-echo pulse. However, if we apply a step modulation of light intensity δI (a VLS field) either side of the spin-echo pulse (applied halfway through the Ramsey sequence), such that

$$\Delta \phi_{\text{before}} = [(\gamma B_1 + \omega_{1,\text{VLS}} + \delta \omega_{1,\text{VLS}}) - (\gamma B_2 + \omega_{2,\text{VLS}} + \delta \omega_{2,\text{VLS}})] \frac{T}{2},$$

$$\Delta \phi_{\text{after}} = -[(\gamma B_1 + \omega_{1,\text{VLS}} - \delta \omega_{1,\text{VLS}}) - (\gamma B_2 + \omega_{2,\text{VLS}} - \delta \omega_{2,\text{VLS}})] \frac{T}{2},$$

$$\Delta \phi_{\text{total}} = 2(\delta \omega_{1,\text{VLS}} - \delta \omega_{2,\text{VLS}})T, \tag{7.42}$$

where the first two lines are before and after the spin-echo pulse respectively, $\omega_{i,\text{VLS}} \propto I_i \sin 2(\theta - \theta_N)$ is the baseline VLS shift at each BEC and $\delta \omega_{i,\text{VLS}} \propto \delta I_i \sin 2(\theta - \theta_N)$ the VLS imparted with the intensity modulation, and θ_N the nulling angle. Naturally, for $\theta = \theta_N$ the VLS shift vanishes and the light is linearly polarised, the target result now being to minimise the relative phase, rather than eliminate the dependence on intensity. For long Ramsey times the modulation depth can be made quite small, minimising the perturbative effect on the trapping potential. The spin-echo sequence will also result in narrow-band sensitivity to AC magnetic field gradients, although we can choose the number of spin-echo pulses so that these effects may be minimised.

The cross-beam, in-trap experiments described in this thesis (Section 7.6) serve to substantially eliminate the VLS shift from the orthogonal beams before undertaking the more precise spin echo measurements. If we assume we are able to align the bias field with the crossing beam wavevector to a precision of 1 part in 10^3, the results of this thesis suggest that the VLS from the intersecting split beams can be eliminated to less than 1 part in 10^6. Additionally, balancing the intensities of the split beams offers a third means of eliminating any extraneous VLS from the intersecting beams. If we now consider the specifics of the current experimental setup and measuring

the VLS from the x'-oriented dipole beam, we have seen that the effective background gradient can be reduced to as low as $\sim 4\,\mathrm{mG/cm}$ for an x'-bias field, allowing for Ramsey interrogation times of 50-100 ms.

Take the case of an overall beam intensity of $10^4\,\mathrm{W\,cm^{-2}}$, the two BECs sampling a 20% intensity difference, and a modulation depth of 10 %. Assuming we are able to measure phases from well-sampled elliptical data sets to a precision of $\delta(\Delta\phi) \approx 0.003\pi$ (typical fitted uncertainty from well sampled ellipses) and a Ramsey time of 100 ms, the polarisation of the light can be made linear to within 50 μrad.

7.8   Conclusions

Our measurements constitute the first high-precision interferometric measurement of the vector light shift and allow it to be suppressed to 0.04(1)% of its maximum value. We have demonstrated variations of the technique, as well as discussed sources of systematic error. These results make possible the realisation of in-trap magnetic gradiometry using differential Ramsey interferometry. The technique can be applied to measure and eliminate the VLS from other beams, not just dipole trapping beams as we have demonstrated here. Further applications and extensions of the technique will be discussed in Chapter 8.

Future and conclusions

<div style="text-align: right">8</div>

In this thesis we have seen how spinor BECs behave in noisy, inhomogeneous magnetic field environments and that they can still be used as precision sensors. In this concluding Chapter, we summarise what we have learnt and describe what extensions and applications the findings of this thesis engender.

The initial aims of this thesis were to develop methods to coherently control spin-mixing interactions in spinor BECs. Towards that aim, we built and commissioned a BEC apparatus that reliably produces condensates of 3×10^5 atoms every $\sim 25\,\mathrm{s}$. We explored spin-mixing dynamics and found the magnetic environment our BEC exists in to be inhomogeneous and unsuitable for long-timescale observations of spinor physics. We then undertook a series of experiments to precisely characterise the magnetic field landscape of the BEC, due to real magnetic field gradients and 'fictitious' optically-induced magnetic fields arising from vector light shifts. These techniques have spawned a range of ideas and extensions that will be discussed further in this Chapter.

8.0.1 Spin-mixing experiments

The combined results of this thesis constitute progress toward ensuring a clean magnetic environment for a spinor BEC, and as such the future of spin-mixing experiments appears bright. Much remains to be explored in the area of spin collisional dynamics and we note here a few avenues of future investigation.

The results in Chapter 5 offer strong evidence that magnetic field gradients play a key role in the equilibration of an evolving spinor BEC. A number of interesting theoretical models in the literature may assist identifying the mechanism for this, which will require further experimental work. It would be particularly interesting to rigorously explore how the relaxation time scales with the strength of the gradient, in both the high-gradient regime ex-

plored in this thesis and the low and intermediate-gradient regime, as well as at different quadratic shifts. In this way we could compare the results of this thesis and the future work it engenders with the equilibration experiments of the Berkeley group [66] and the observation of dissipative relaxation by the NIST group [51].

The efforts towards eliminating gradient dephasing with spin-echo pulses heralded several interesting results, none more so than the suppression of ground state relaxation. The results of Section 5.3 have described in detail the effects of refocusing pulses in the strong rf coupling regime. With these lessons in mind, future work could explore refocusing pulses within the rotating-wave approximation, *i.e.* at higher Larmor frequencies, and use microwave fields to dress q to levels where spin dynamics are appreciable. In the RWA regime, where the sensitivity to Larmor phase is minimal and refocusing pulses can be applied (theoretically) at any evolution time, the limit of almost continuous decoupling could be explored. There are many other extensions to the general scheme we envisioned, such as composite adiabatic pulses [222] where the amplitude and frequency of the pulses are varied. Indeed, we envision applying refocusing pulses separated by $\sim 1\,\mathrm{ms}$ and employing strategies to deal with cumulative pulse errors; a vast array of NMR pulse sequences with ever more sophisticated acronyms exist exactly for this purpose.

Precisely characterising magnetic field gradients using the results of Chapter 6 offers the prospect of diagnosing gradient cancellation to unprecedented precision. Gradient coils can be specially designed to cancel dominant gradient terms and the use of bias coils aligned along the eigenaxes of the gradient tensor could potentially be explored: this reduces the problem of gradient elimination to only diagonal gradients. With the elimination of gradients one could explore coherent collisional control as well as weak equilibration mechanisms otherwise obfuscated by gradient-attributed effects.

8.0.2 Magnetic gradiometry

Microscale magnetic sensors are required where magnetic fields change over similarly microscopic length scales. The magnetic gradiometry experiments in this thesis constitute the first measurement of the magnetic field gradient tensor at the microscale. This is an important result, demonstrating high common-mode rejection of ambient magnetic field noise and allowing for precise measurement of magnetic fields in a noisy environment. The future for these measurements is particularly promising. With the elimination of vector light shifts making possible the realisation of contrast-limited interrogation times, the sensitivity of in-trap gradiometry measurements can

immediately be improved fivefold over the freefall measurements described in this thesis. Suppressing the background gradient allows for much longer interrogation times: the sensitivity limits of the technique could be probed, potentially surpassing the spatio-temporal sensitivity per unit bandwidth of existing magnetometers. A range of experiments can be foreseen to test and validate the predictive power of tensor measurements on the microscale, such as *in vacuo* characterisation of small magnetic structures or current distributions. Recently, we were able to use the gradiometer to determine the presence of a magnetic source near the atoms.[1]

A number of interesting extensions exist. For example, at what point does the interaction of the magnetic moment of each condensate with the other factor into the measurement of a magnetic field? Each condensate has its own magnetic field, generated by the $\sim 10^5$ magnetic dipoles it is composed of. Although this field is small, it may play an interesting role when the separation between the BECs is reduced to $\sim$ 10-20 µm, which is highly feasible with a tighter dipole trapping beam and an AOM with a higher number of resolvable spots. The magnetic field 25 µm from a 10^5 atom BEC is 6 pT, easily detectable using in-trap gradiometry at prospective interrogation times. Ultracold neutral atoms are promising candidates for realising atomic analogues of magnetic quantum systems, and the ability to perform incisive magnetic measurements on them enriches the prospective experiments considerably.

8.0.3 Vector light shifts

Optical dipole traps are only spin-independent when the trapping light is linearly polarised. Even a small elliptical polarisation imperfection gives rise to effective magnetic fields that contaminate magnetometry experiments (Figure 7.2), and in the case of a trapped cloud displaced from the position of peak intensity, strong spatial gradients of the VLS result in large effective magnetic field gradients (Figure 7.3).

Our measurements probed the VLS-induced effective field and its spatial variation. We were able to suppress the VLS by a factor of 2500 by adjusting the trapping beam polarisation. These measurements constitute a high-precision interferometric measurement of the trapping beam polarisation, and can be adapted to measure the polarisation of other optical fields.

Our measurement protocol could be extended in a number of ways. For instance, precisely characterising the VLS could supplement comprehensive elimination of the differential light shift of the clock transition. We could

---

[1] A magnetic ball driver unintentionally left near the science cell registers a measurable gradient, despite being hundreds of mm away from the atoms.

also extend the interrogation time to make a significantly more sensitive VLS measurement: with an x'-bias field, Ramsey interrogation times of between 50-100 ms should be possible with the current experimental configuration. Future demonstration of the spin-echo VLS measurement discussed in Section 7.7 will allow for improved VLS suppression. The work described in this thesis has already been applied in the experiments performed by Martijn Jasperse on Faraday MRI of a spinor BEC [95]: the VLS at 790 nm is 290 Hz W^{-1} cm^2, making the effective field induced by the Faraday beam significant.

Since the scalar light shift vanishes at 790 nm, light at this wavelength can be used as a minimally-perturbative means of engineering optically-induced magnetic fields. For example, one can even investigate the possibility of *using* the VLS gradient to cancel background magnetic field gradients. Measurement of the VLS could be used to map the spatial intensity profile of trapping or probe laser beams. Simply varying the timing of the 'delayed drop' method (Section 7.5) allows the falling BEC to traverse the intensity profile for a fixed beam polarisation. The variation of $\Delta\phi$ with drop time would then be a direct measure of the intensity profile. The trapped clouds could be translated axially along the length of the beam, allowing for full *in vacuo* spatially resolved beam characterisation. This could be applied to characterise other beams, such as the dispersive Faraday probe or even Raman coupling beams, where spatially varying Rabi frequencies must be accounted for.

Together with the results of Chapter 6, the VLS measurements described here constitute novel applications of a new atomic magnetometer, one with demonstrated precision sensitivity and a promising future.

List of Acronyms

AC   alternating current
AOM   acousto-optic modulator

BEC   Bose–Einstein condensate

CCD   charge-coupled device
CMOS   complementary metal-oxide-semiconductor
CP   Carr-Purcell
CPMG   Carr-Purcell-Meiboom-Gill

DDS   direct digital synthesis
DLS   differential scalar light shift
DRI   differential Ramsey interferometry

ECDL   external-cavity diode laser
EOM   electro-optic modulator

IQ   in phase and in quadrature

LO   local oscillator

MOPA   master oscillator-power amplifier
MOT   magneto-optical trap
MT   magnetic trap
MTS   modulation transfer spectroscopy
MW   microwave

OD   optical depth

PGC   polarisation gradient cooling
PLC   programmable logic controller
PM   polarisation maintaining
PSD   phase-space density

RF   radiofrequency
RGA   residual gas analyser
RWA   rotating-wave approximation

sCMOS   Scientific CMOS
SI trap   semi-isotropic trap
SMA   single-mode approximation

TEC   thermoelectric cooler
TSP   titanium sublimation pump

UHV   ultra high vacuum

VLS   vector light shift

Bibliography

[1]   M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. *Observation of Bose–Einstein condensation in a dilute atomic vapor*. Science **269**, 198 (1995). DOI: 10.1126/science.269.5221.198. [pp 1 and 29]

[2]   K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. *Bose–Einstein condensation in a gas of sodium atoms*. Physical Review Letters **75**, 3969 (1995). DOI: 10.1103/PhysRevLett.75.3969. [pp 1, 29, and 66]

[3]   C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. *Evidence of Bose–Einstein condensation in an atomic gas with attractive interactions*. Physical Review Letters **75**, 1687 (1995). DOI: 10.1103/PhysRevLett.75.1687. [pp 1 and 29]

[4]   Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. *Quantum simulations with ultracold quantum gases*. Nature Physics **8**, 267 (2012). DOI: 10.1038/nphys2259. [p 2]

[5]   Y.-J. Lin, K. Jiménez-García, and I. B. Spielman. *Spin-orbit-coupled Bose–Einstein condensates*. Nature **471**, 83 (2011). DOI: 10.1038/nature09887. [pp 2 and 156]

[6]   Victor Galitski and Ian B. Spielman. *Spin-orbit coupling in quantum gases*. Nature **494**, 49 (2013). DOI: 10.1038/nature11841. [p 2]

[7]   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). DOI: 10.1103/PhysRevA.65.033608. [p 2]

[8]   J. E. Debs, P. A. Altin, T. H. Barter, D. Döring, G. R. Dennis, G. McDonald, R. P. Anderson, J. D. Close, and N. P. Robins. *Cold-atom gravimetry with a Bose–Einstein condensate*. Physical Review A **84**, 033610 (2011). DOI: 10.1103/PhysRevA.84.033610. [pp 2 and 5]

[9]   Stephan Wildermuth, Sebastian Hofferberth, Igor Lesanovsky, Elmar Haller, L. Mauritz Andersson, Sönke Groth, Israel Bar-Joseph, Peter

Krüger, and Jörg Schmiedmayer. *Bose–Einstein condensates: Microscopic magnetic-field imaging*. Nature **435**, 440 (2005). DOI: 10.1038/435440a. [pp 2, 5, and 125]

[10] M. Vengalattore, J. M. Higbie, S. R. Leslie, J. Guzman, L. E. Sadler, and D. M. Stamper-Kurn. *High-resolution magnetometry with a spinor Bose–Einstein condensate*. Physical Review Letters **98**, 200801 (2007). DOI: 10.1103/PhysRevLett.98.200801. [pp 2, 3, 4, 5, 9, 10, 125, and 156]

[11] M. Koschorreck, M. Napolitano, B. Dubost, and M. W. Mitchell. *High resolution magnetic vector-field imaging with cold atomic ensembles*. Applied Physics Letters **98**, 074101 (2011). DOI: 10.1063/1.3555459. [pp 2 and 125]

[12] H. Schmaljohann, M. Erhard, J. Kronjäger, M. Kottke, S. van Staa, L. Cacciapuoti, J. J. Arlt, K. Bongs, and K. Sengstock. *Dynamics of $F = 2$ spinor Bose–Einstein condensates*. Physical Review Letters **92**, 040402 (2004). DOI: 10.1103/PhysRevLett.92.040402. [pp 2 and 10]

[13] Axel Griesmaier, Jörg Werner, Sven Hensler, Jürgen Stuhler, and Tilman Pfau. *Bose–Einstein condensation of chromium*. Physical Review Letters **94**, 160401 (2005). DOI: 10.1103/PhysRevLett.94.160401. [p 2]

[14] J. Stenger, S. Inouye, D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, and W. Ketterle. *Spin domains in ground-state Bose–Einstein condensates*. Nature **396**, 345 (1998). DOI: 10.1038/24567. [pp 2, 9, 20, 22, 23, and 100]

[15] Tin-Lun Ho. *Spinor Bose condensates in optical traps*. Physical Review Letters **81**, 742 (1998). DOI: 10.1103/PhysRevLett.81.742. [pp 2, 10, 17, and 20]

[16] L. E. Sadler, J. M. Higbie, S. R. Leslie, M. Vengalattore, and D. M. Stamper-Kurn. *Spontaneous symmetry breaking in a quenched ferromagnetic spinor Bose–Einstein condensate*. Nature **443**, 312 (2006). DOI: 10.1038/nature05094. [pp 2, 10, and 11]

[17] E. M. Bookjans, A. Vinit, and C. Raman. *Quantum phase transition in an antiferromagnetic spinor Bose–Einstein condensate*. Physical Review Letters **107**, 195306 (2011). DOI: 10.1103/PhysRevLett.107.195306. [pp 2, 11, and 23]

[18] Ming-Shien Chang, Qishu Qin, Wenxian Zhang, Li You, and Michael S. Chapman. *Coherent spinor dynamics in a spin-1 Bose condensate*. Nature Physics **1**, 111 (2005). DOI: 10.1038/nphys153. [pp 3, 10, and 26]

[19] M. V. Romalis and E. N. Fortson. *Zeeman frequency shifts in an optical dipole trap used to search for an electric-dipole moment*. Physical Review A **59**, 4547 (1999). DOI: 10.1103/PhysRevA.59.4547. [pp 3, 121, 150, and 156]

[20] Dmitry Budker and Michael Romalis. *Optical magnetometry*. Nature Physics **3**, 227 (2007). DOI: 10.1038/nphys566. [pp 3, 124, and 156]

[21] A. M. Kaufman, B. J. Lester, and C. A. Regal. *Cooling a single atom in an optical tweezer to its quantum ground state*. Physical Review X **2**, 041014 (2012). DOI: 10.1103/PhysRevX.2.041014. [p 3]

[22] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletić, and M. D. Lukin. *Coherence and Raman sideband cooling of a single atom in an optical tweezer*. Physical Review Letters **110**, 133001 (2013). DOI: 10.1103/PhysRevLett.110.133001. [p 3]

[23] Yuki Kawaguchi and Masahito Ueda. *Spinor Bose–Einstein condensates*. Physics Reports **520**, 253 (2012). DOI: 10.1016/j.physrep.2012.07.005. [pp 3, 14, 18, 19, 22, and 80]

[24] Dan Stamper-Kurn and Masahito Ueda. *Spinor Bose gases: Symmetries, magnetism, and quantum dynamics*. Reviews of Modern Physics **85**, 1191 (2013). DOI: 10.1103/RevModPhys.85.1191. [pp 4, 11, 14, 19, 20, 21, 22, and 97]

[25] D. M. Stamper-Kurn and W. Ketterle. *Spinor condensates and light scattering from Bose–Einstein condensates*. In R. Kaiser, C. Westbrook, and F. David (editors), *Coherent atomic matter waves*, volume 72 of *Les Houches - Ecole d'Ete de Physique Theorique*, pages 139–217. Springer Berlin Heidelberg (2001). DOI: 10.1007/3-540-45338-5_2. [pp 4, 14, 17, and 20]

[26] M. W. Ray, E. Ruokokoski, S. Kandel, M. Möttönen, and D. S. Hall. *Observation of Dirac monopoles in a synthetic magnetic field*. Nature **505**, 657 (2014). DOI: 10.1038/nature12954. [pp 4 and 11]

[27] Magnus O. Borgh and Janne Ruostekoski. *Topological interface engineering and defect crossing in ultracold atomic gases*. Physical Review Letters **109**, 015302 (2012). DOI: 10.1103/PhysRevLett.109.015302. [pp 4 and 11]

[28] H. Jing, Y. Jiang, Weiping Zhang, and P. Meystre. *Laser-catalyzed spin-exchange process in a Bose–Einstein condensate.* Physical Review A **81**, 031603 (2010). DOI: 10.1103/PhysRevA.81.031603. [p 4]

[29] C. D. Hamley, E. M. Bookjans, G. Behin-Aein, P. Ahmadi, and M. S. Chapman. *Photoassociation spectroscopy of a spin-1 Bose–Einstein condensate.* Physical Review A **79**, 023401 (2009). DOI: 10.1103/PhysRevA.79.023401. [pp 4 and 93]

[30] I. K. Kominis, T. W. Kornack, J. C. Allred, and M. V. Romalis. *A subfemtotesla multichannel atomic magnetometer.* Nature **422**, 596 (2003). DOI: 10.1038/nature01484. [pp 4, 123, and 124]

[31] W. Clark Griffith, Svenja Knappe, and John Kitching. *Femtotesla atomic magnetometry in a microfabricated vapor cell.* Optics Express **18**, 27167 (2010). DOI: 10.1364/OE.18.027167. [pp 4, 124, and 151]

[32] J. R. Kirtley, M. B. Ketchen, K. G. Stawiasz, J. Z. Sun, W. J. Gallagher, S. H. Blanton, and S. J. Wind. *High-resolution scanning SQUID microscope.* Applied Physics Letters **66**, 1138 (1995). DOI: 10.1063/1.113838. [pp 4 and 124]

[33] J. M. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. R. Hemmer, A. Yacoby, R. Walsworth, and M. D. Lukin. *High-sensitivity diamond magnetometer with nanoscale resolution.* Nature Physics **4**, 810 (2008). DOI: 10.1038/nphys1075. [pp 4 and 124]

[34] Romana Schirhagl, Kevin Chang, Michael Loretz, and Christian L. Degen. *Nitrogen-vacancy centers in diamond: Nanoscale sensors for physics and biology.* Annual Review of Physical Chemistry **65**, 83 (2014). DOI: 10.1146/annurev-physchem-040513-103659. [pp 4 and 123]

[35] J. M. Higbie, L. E. Sadler, S. Inouye, A. P. Chikkatur, S. R. Leslie, K. L. Moore, V. Savalli, and D. M. Stamper-Kurn. *Direct nondestructive imaging of magnetization in a spin-1 Bose–Einstein gas.* Physical Review Letters **95**, 050401 (2005). DOI: 10.1103/PhysRevLett.95.050401. [pp 5, 10, 129, and 150]

[36] Dan M. Stamper-Kurn. *Seeing spin dynamics in atomic gases.* (2014). arXiv: 1501.00057. [p 5]

[37] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto. *Rapid production of ^{87}Rb Bose–Einstein condensates in a combined magnetic and optical potential.* Physical Review A **79**, 063631 (2009). DOI: 10.1103/PhysRevA.79.063631. [pp 6, 32, 33, 37, and 66]

[38] S. R. Leslie, J. Guzman, M. Vengalattore, Jay D. Sau, Marvin L. Cohen, and D. M. Stamper-Kurn. *Amplification of fluctuations in a spinor Bose–Einstein condensate*. Physical Review A **79**, 043631 (2009). DOI: 10.1103/PhysRevA.79.043631. [p 9]

[39] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle. *Optical confinement of a Bose–Einstein condensate*. Physical Review Letters **80**, 2027 (1998). DOI: 10.1103/PhysRevLett.80.2027. [p 9]

[40] Tetsuo Ohmi and Kazushige Machida. *Bose–Einstein condensation with internal degrees of freedom in alkali atom gases*. Journal of the Physical Society of Japan **67**, 1822 (1998). DOI: 10.1143/JPSJ.67.1822. [pp 10 and 20]

[41] C. K. Law, H. Pu, and N. P. Bigelow. *Quantum spins mixing in spinor Bose–Einstein condensates*. Physical Review Letters **81**, 5257 (1998). DOI: 10.1103/PhysRevLett.81.5257. [pp 10 and 20]

[42] H. Pu, C. K. Law, S. Raghavan, J. H. Eberly, and N. P. Bigelow. *Spin-mixing dynamics of a spinor Bose–Einstein condensate*. Physical Review A **60**, 1463 (1999). DOI: 10.1103/PhysRevA.60.1463. [pp 10, 19, and 20]

[43] H. Pu, S. Raghavan, and N.P. Bigelow. *Manipulating spinor condensates with magnetic fields: Stochastization, metastability, and dynamical spin localization*. Physical Review A **61**, 023602 (2000). DOI: 10.1103/PhysRevA.61.023602. [p 10]

[44] J. Kronjager, C. Becker, M. Brinkmann, R. Walser, P. Navez, K. Bongs, and K. Sengstock. *Evolution of a spinor condensate: Coherent dynamics, dephasing, and revivals*. Physical Review A **72**, 063619 (2005). DOI: 10.1103/PhysRevA.72.063619. [pp 10, 99, and 100]

[45] M.-S. Chang, C. D. Hamley, M. D. Barrett, J. A. Sauer, K. M. Fortier, W. Zhang, L. You, and M. S. Chapman. *Observation of spinor dynamics in optically trapped ^{87}Rb Bose–Einstein condensates*. Physical Review Letters **92**, 140403 (2004). DOI: 10.1103/PhysRevLett.92.140403. [p 10]

[46] Wenxian Zhang, D. L. Zhou, M.-S. Chang, M. S. Chapman, and L. You. *Coherent spin mixing dynamics in a spin-1 atomic condensate*. Physical Review A **72**, 013602 (2005). DOI: 10.1103/PhysRevA.72.013602. [pp 10, 24, and 25]

[47] Artur Widera, Fabrice Gerbier, Simon Fölling, Tatjana Gericke, Olaf Mandel, and Immanuel Bloch. *Coherent collisional spin dynam-

ics in optical lattices. Physical Review Letters **95**, 190405 (2005). DOI: 10.1103/PhysRevLett.95.190405. [p 10]

[48] Fabrice Gerbier, Artur Widera, Simon Fölling, Olaf Mandel, and Immanuel Bloch. *Resonant control of spin dynamics in ultracold quantum gases by microwave dressing*. Physical Review A **73**, 041602 (2006). DOI: 10.1103/PhysRevA.73.041602. [p 10]

[49] A. T. Black, E. Gomez, L. D. Turner, S. Jung, and P. D. Lett. *Spinor dynamics in an antiferromagnetic spin-1 condensate*. Physical Review Letters **99**, 070403 (2007). DOI: 10.1103/PhysRevLett.99.070403. [p 10]

[50] Y. Liu, S. Jung, S. E. Maxwell, L. D. Turner, E. Tiesinga, and P. D. Lett. *Quantum phase transitions and continuous observation of spinor dynamics in an antiferromagnetic condensate*. Physical Review Letters **102**, 125301 (2009). DOI: 10.1103/PhysRevLett.102.125301. [pp 10 and 11]

[51] Y. Liu, E. Gomez, S. E. Maxwell, L. D. Turner, E. Tiesinga, and P. D. Lett. *Number fluctuations and energy dissipation in sodium spinor condensates*. Physical Review Letters **102**, 225301 (2009). DOI: 10.1103/PhysRevLett.102.225301. [pp 10, 11, 97, 98, and 202]

[52] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein, and T. Pfau. *The physics of dipolar bosonic quantum gases*. Reports on Progress in Physics **72**, 126401 (2009). DOI: 10.1088/0034-4885/72/12/126401. [p 10]

[53] Yuki Kawaguchi, Hiroki Saito, and Masahito Ueda. *Can spinor dipolar effects be observed in Bose–Einstein condensates?* Physical Review Letters **98**, 110406 (2007). DOI: 10.1103/PhysRevLett.98.110406. [p 10]

[54] Yuki Kawaguchi, Hiroki Saito, and Masahito Ueda. *Einstein–de Haas effect in dipolar Bose–Einstein condensates*. Physical Review Letters **96**, 080405 (2006). DOI: 10.1103/PhysRevLett.96.080405. [p 10]

[55] L. Santos and T. Pfau. *Spin-3 chromium Bose–Einstein condensates*. Physical Review Letters **96**, 190404 (2006). DOI: 10.1103/PhysRevLett.96.190404. [pp 10 and 20]

[56] M. Vengalattore, S. R. Leslie, J. Guzman, and D. M. Stamper-Kurn. *Spontaneously modulated spin textures in a dipolar spinor Bose–Einstein condensate*. Physical Review Letters **100**, 170403 (2008). DOI: 10.1103/PhysRevLett.100.170403. [p 10]

[57] Q. Beaufils, R. Chicireanu, T. Zanon, B. Laburthe-Tolra, E. Maréchal, L. Vernac, J.-C. Keller, and O. Gorceix. *All-optical production of*

chromium Bose–Einstein condensates. Physical Review A **77**, 061601 (2008). DOI: 10.1103/PhysRevA.77.061601. [p 11]

[58] B. Pasquiou, E. Maréchal, G. Bismut, P. Pedri, L. Vernac, O. Gorceix, and B. Laburthe-Tolra. *Spontaneous demagnetization of a dipolar spinor Bose gas in an ultralow magnetic field.* Physical Review Letters **106**, 255303 (2011). DOI: 10.1103/PhysRevLett.106.255303. [p 11]

[59] Yujiro Eto, Hiroki Saito, and Takuya Hirano. *Observation of dipole-induced spin texture in an ^{87}Rb Bose–Einstein condensate.* Physical Review Letters **112**, 185301 (2014). DOI: 10.1103/PhysRevLett.112.185301. [p 11]

[60] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm, and F. Ferlaino. *Bose–Einstein condensation of erbium.* Physical Review Letters **108**, 210401 (2012). DOI: 10.1103/PhysRevLett.108.210401. [p 11]

[61] Mingwu Lu, Nathaniel Q. Burdick, Seo Ho Youn, and Benjamin L. Lev. *Strongly dipolar Bose–Einstein condensate of dysprosium.* Physical Review Letters **107**, 190401 (2011). DOI: 10.1103/PhysRevLett.107.190401. [p 11]

[62] Jay D. Sau, Sabrina R. Leslie, Marvin L. Cohen, and Dan M. Stamper-Kurn. *Spin squeezing of high-spin, spatially extended quantum fields.* New Journal of Physics **12**, 085011 (2010). DOI: 10.1088/1367-2630/12/8/085011. [p 11]

[63] C. Klempt, O. Topic, G. Gebreyesus, M. Scherer, T. Henninger, P. Hyllus, W. Ertmer, L. Santos, and J. J. Arlt. *Parametric amplification of vacuum fluctuations in a spinor condensate.* Physical Review Letters **104**, 195303 (2010). DOI: 10.1103/PhysRevLett.104.195303. [p 11]

[64] C. D. Hamley, C. S. Gerving, T. M. Hoang, E. M. Bookjans, and M. S. Chapman. *Spin-nematic squeezed vacuum in a quantum gas.* Nature Physics **8**, 305 (2012). DOI: 10.1038/nphys2245. [p 11]

[65] A. Vinit, E. M. Bookjans, C. A. R. Sá de Melo, and C. Raman. *Antiferromagnetic spatial ordering in a quenched one-dimensional spinor gas.* Physical Review Letters **110**, 165301 (2013). DOI: 10.1103/PhysRevLett.110.165301. [pp 11 and 23]

[66] J. Guzman, G.-B. Jo, A. N. Wenz, K. W. Murch, C. K. Thomas, and D. M. Stamper-Kurn. *Long-time-scale dynamics of spin textures in a degenerate $F = 1$ ^{87}Rb spinor Bose gas.* Physical Review A **84**, 063625 (2011). DOI: 10.1103/PhysRevA.84.063625. [pp 11, 23, 98, and 202]

[67] David Jacob, Lingxuan Shao, Vincent Corre, Tilman Zibold, Luigi De Sarlo, Emmanuel Mimoun, Jean Dalibard, and Fabrice Gerbier. *Phase diagram of spin-1 antiferromagnetic Bose–Einstein condensates*. Physical Review A **86**, 061601 (2012). DOI: 10.1103/PhysRevA.86.061601. [p 11]

[68] J. Jiang, L. Zhao, M. Webb, and Y. Liu. *Mapping the phase diagram of spinor condensates via adiabatic quantum phase transitions*. Physical Review A **90**, 023610 (2014). DOI: 10.1103/PhysRevA.90.023610. [pp 11, 98, and 104]

[69] Jae-yoon Choi, Woo Jin Kwon, Moonjoo Lee, Hyunseok Jeong, Kyung-won An, and Yong-il Shin. *Imprinting skyrmion spin textures in spinor Bose–Einstein condensates*. New Journal of Physics **14**, 053013 (2012). DOI: 10.1088/1367-2630/14/5/053013. [p 11]

[70] Jae-yoon Choi, Woo Jin Kwon, and Yong-il Shin. *Observation of topologically stable 2D skyrmions in an antiferromagnetic spinor Bose–Einstein condensate*. Physical Review Letters **108**, 035301 (2012). DOI: 10.1103/PhysRevLett.108.035301. [p 11]

[71] M. W. Ray, E. Ruokokoski, K. Tiurev, M. Möttönen, and D. S. Hall. *Observation of isolated monopoles in a quantum field*. Science **348**, 544 (2015). DOI: 10.1126/science.1258289. [p 11]

[72] Sang Won Seo, Seji Kang, Woo Jin Kwon, and Yong-il Shin. *Observation of half-quantum vortices in an antiferromagnetic spinor Bose–Einstein condensate*. arXiv:1503.05648 [cond-mat] (2015). ARXIV: 1503.05648. [p 11]

[73] Franco Dalfovo, Stefano Giorgini, Lev P. Pitaevskii, and Sandro Stringari. *Theory of Bose–Einstein condensation in trapped gases*. Reviews of Modern Physics **71**, 463 (1999). DOI: 10.1103/RevModPhys.71.463. [pp 12, 30, and 73]

[74] C. Pethick and H. Smith. *Bose–Einstein condensation in dilute gases*. Cambridge University Press (2002). [pp 12 and 30]

[75] C. J. Foot. *Atomic physics*. Oxford University Press (2005). [pp 12, 55, and 56]

[76] Ming-Shien Chang. *Coherent dynamics of a spin-1 Bose–Einstein condensate*. PhD thesis, Georgia Institute of Technology (2006). [pp 14, 16, 24, and 96]

[77] Wenxian Zhang. *Spin-1 atomic condensates in magnetic fields*. PhD thesis, Georgia Institute of Technology (2005). [p 14]

[78] Norman Ramsey. *Molecular beams*. Clarendon Press (1963). [pp 14 and 36]

[79] Daniel A. Steck. Rubidium 87 D line data, revision 2.1.4. (2010). [pp 14, 161, and 163]

[80] R. P. Anderson. *Non-equilibrium dynamics and relative phase evolution of two-component Bose–Einstein condensates*. PhD thesis, Swinburne University of Technology (2010). [pp 14, 76, 77, and 79]

[81] S. Yi, Ö. E. Müstecaplıoğlu, C. P. Sun, and L. You. *Single-mode approximation in a spinor-1 atomic condensate*. Physical Review A **66**, 011601 (2002). DOI: 10.1103/PhysRevA.66.011601. [p 19]

[82] H.-J. Miesner, D. M. Stamper-Kurn, J. Stenger, S. Inouye, A. P. Chikkatur, and W. Ketterle. *Observation of metastable states in spinor Bose–Einstein condensates*. Physical Review Letters **82**, 2228 (1999). DOI: 10.1103/PhysRevLett.82.2228. [pp 22 and 23]

[83] C. S. Gerving, T. M. Hoang, B. J. Land, M. Anquez, C. D. Hamley, and M. S. Chapman. *Non-equilibrium dynamics of an unstable quantum pendulum explored in a spin-1 Bose–Einstein condensate*. Nature Communications **3**, 1169 (2012). DOI: 10.1038/ncomms2179. [p 26]

[84] T. M. Hoang, C. S. Gerving, B. J. Land, M. Anquez, C. D. Hamley, and M. S. Chapman. *Dynamic stabilization of a quantum manybody spin system*. Physical Review Letters **111**, 090403 (2013). DOI: 10.1103/PhysRevLett.111.090403. [p 26]

[85] H. K. Pechkis, J. P. Wrubel, A. Schwettmann, P. F. Griffin, R. Barnett, E. Tiesinga, and P. D. Lett. *Spinor dynamics in an antiferromagnetic spin-1 thermal Bose gas*. Physical Review Letters **111**, 025301 (2013). DOI: 10.1103/PhysRevLett.111.025301. [p 26]

[86] Xiaodong He, Bing Zhu, Xiaoke Li, Fudong Wang, Zhi-Fang Xu, and Dajun Wang. *Coherent spin-mixing dynamics in thermal ^{87}Rb spin-1 and spin-2 gases*. Physical Review A **91**, 033635 (2015). DOI: 10.1103/PhysRevA.91.033635. [p 26]

[87] Wolfgang Ketterle. *Nobel lecture: When atoms behave as waves: Bose–Einstein condensation and the atom laser*. Reviews of Modern Physics **74**, 1131 (2002). DOI: 10.1103/RevModPhys.74.1131. [p 29]

[88] J. Walraven. *Thermodynamic and Collisional Properties of Trapped Atomic Gases*. Les Houches - predoc lectures 2013 (Vienna version - unpublished), University of Vienna (2013). [p 30]

[89] A. A. Wood. *Design of a collimated rubidium-87 effusive oven and beam line system*. Undergraduate Project Report, Monash University (2009). [p 32]

[90] Alexander Wood. *A rubidium cold atom beam source*. Honours thesis, Monash University (2010). [pp 32 and 35]

[91] S. C. Bell, M. Junker, M. Jasperse, L. D. Turner, Y.-J. Lin, I. B. Spielman, and R. E. Scholten. *A slow atom source using a collimated effusive oven and a single-layer variable pitch coil Zeeman slower*. Review of Scientific Instruments **81**, 013105 (2010). DOI: 10.1063/1.3276712. [pp 32 and 37]

[92] L. M. Bennie. *Design of a Single-Layer Zeeman Slower for ^{85}Rb and ^{87}Rb*. Undergraduate project report, Monash University (2009). [p 32]

[93] Lene Vestergaard Hau, J. A. Golovchenko, and Michael M. Burns. *A new atomic beam source: The ''candlestick''*. Review of Scientific Instruments **65**, 3746 (1994). DOI: 10.1063/1.1145246. [p 36]

[94] M. R. Walkiewicz, P. J. Fox, and R. E. Scholten. *Candlestick rubidium beam source*. Review of Scientific Instruments **71**, 3342 (2000). DOI: 10.1063/1.1288261. [p 36]

[95] M. Jasperse. *Faraday Magnetic Resonance Imaging Of Bose–Einstein Condensates*. PhD thesis, Monash University (2015). [pp 37, 38, 43, 46, 91, 151, and 204]

[96] K. M. Birnbaum. *Cavity QED with Multilevel Atoms*. PhD thesis, California Institute of Technology (2005). Appendix B: Ultra-High vacuum chambers. [p 40]

[97] CSRDVS Group. *Vacuum preparation procedures*. Technical note, Daresbury Laboratory, Daresbury UK (1996). [p 40]

[98] D. Edwards. *Methane outgassing from a Ti sublimation pump*. Journal of Vacuum Science and Technology **17**, 279 (1980). DOI: 10.1116/1.570412. [p 43]

[99] H. J. Halama. *Behavior of titanium sublimation and sputter-ion pumps in the 10-11 Torr range*. Journal of Vacuum Science & Technology **14**, 524 (1977). DOI: 10.1116/1.569295. [p 43]

[100] D. J. McCarron, S. A. King, and S. L. Cornish. *Modulation transfer spectroscopy in atomic rubidium*. Measurement Science and Technology **19**, 105601 (2008). DOI: 10.1088/0957-0233/19/10/105601. [p 44]

[101] V. Negnevitsky and L. D. Turner. *Wideband laser locking to an atomic reference with modulation transfer spectroscopy*. Optics Express **21**, 3103 (2013). DOI: 10.1364/OE.21.003103. [p 44]

[102] Jürgen Appel, Andrew MacRae, and A. I. Lvovsky. *A versatile digital GHz phase lock for external cavity diode lasers*. Measurement Science and Technology **20**, 055302 (2009). DOI: 10.1088/0957-0233/20/5/055302. [p 44]

[103] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. *Double-pass acousto-optic modulator system*. Review of Scientific Instruments **76**, 063112 (2005). DOI: 10.1063/1.1930095. [p 45]

[104] L. P. Maguire, S. Szilagyi, and R. E. Scholten. *High performance laser shutter using a hard disk drive voice-coil actuator*. Review of Scientific Instruments **75**, 3077 (2004). DOI: 10.1063/1.1786331. [p 45]

[105] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. *Making, probing and understanding Bose–Einstein condensates*. In M. Inguscio, S. Stringari, and C. E. Wieman (editors), *Bose–Einstein Condensation in Atomic Gases*, Proceedings of the International School of Physics. IOS Press (1999). [p 45]

[106] Jason Merrill. *Characterization of Inelastic Losses from Bose–Einstein Condensates in the $|2, 1\rangle$ State of ^{87}Rb*. Senior thesis, Amherst College (2006). [p 49]

[107] Walther Gerlach and Otto Stern. *Das magnetische moment des silberatoms*. Zeitschrift für Physik **9**, 353 (1922). DOI: 10.1007/BF01326984. [p 51]

[108] P. A. Altin, M. T. Johnsson, V. Negnevitsky, G. R. Dennis, R. P. Anderson, J. E. Debs, S. S. Szigeti, K. S. Hardman, S. Bennetts, G. D. McDonald, L. D. Turner, J. D. Close, and N. P. Robins. *Precision atomic gravimeter based on Bragg diffraction*. New Journal of Physics **15**, 023009 (2013). DOI: 10.1088/1367-2630/15/2/023009. [p 52]

[109] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson. *A scripted control system for autonomous hardware-timed experiments*. Review of Scientific Instruments **84**, 085111 (2013). DOI: 10.1063/1.4817213. [p 53]

[110] L. M. Bennie. *A large atom number magneto-optical trap for BEC production*. Honours thesis, Monash University (2010). [p 53]

[111] 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). DOI: 10.1364/JOSAB.6.002023. [p 56]

[112] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. *Optical molasses*. Journal of the Optical Society of America B **6**, 2084 (1989). DOI: 10.1364/JOSAB.6.002084. [p 56]

[113] Simon Stellmer, Benjamin Pasquiou, Rudolf Grimm, and Florian Schreck. *Laser cooling to quantum degeneracy*. Physical Review Letters **110**, 263003 (2013). DOI: 10.1103/PhysRevLett.110.263003. [p 60]

[114] Wolfgang Ketterle and N. J. Van Druten. *Evaporative cooling of trapped atoms*. In Benjamin Bederson and Herbert Walther (editor), *Advances In Atomic, Molecular, and Optical Physics*, volume 37, pages 181–236. Academic Press (1996). [pp 61 and 66]

[115] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. *Optical dipole traps for neutral atoms*. In Benjamin Bederson and Herbert Walther (editors), *Advances In Atomic, Molecular, and Optical Physics*, volume 42, pages 95–170. Academic Press (2000). DOI: 10.1016/S1049-250X(08)60186-X. [pp 62 and 160]

[116] M. Zaiser, J. Hartwig, D. Schlippert, U. Velte, N. Winter, V. Lebedev, W. Ertmer, and E. M. Rasel. *Simple method for generating Bose–Einstein condensates in a weak hybrid trap*. Physical Review A **83**, 035601 (2011). DOI: 10.1103/PhysRevA.83.035601. [p 63]

[117] K. J. Arnold and M. D. Barrett. *All-optical Bose–Einstein condensation in a 1.06 μm dipole trap*. Optics Communications **284**, 3288 (2011). DOI: 10.1016/j.optcom.2011.03.008. [pp 63 and 67]

[118] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer. *All-optical runaway evaporation to Bose–Einstein condensation*. Physical Review A **79**, 061406 (2009). DOI: 10.1103/PhysRevA.79.061406. [pp 63 and 151]

[119] D. L. Jenkin, D. J. McCarron, M. P. Köppinger, H. W. Cho, S. A. Hopkins, and S. L. Cornish. *Bose–Einstein condensation of ^{87}Rb in a levitated crossed dipole trap*. The European Physical Journal D **65**, 11 (2011). DOI: 10.1140/epjd/e2011-10720-5. [p 63]

[120] D. H. Guest. *A cross-beam far off-resonance optical trap for ^{87}Rb Bose–Einstein condensates*. Undergraduate thesis, Amherst College (2007). [p 64]

[121] D. S. Naik and C. Raman. *Optically plugged quadrupole trap for Bose–Einstein condensates*. Physical Review A **71**, 033617 (2005). DOI: 10.1103/PhysRevA.71.033617. [p 66]

[122] Daniel McCarron. *A quantum degenerate mixture of ^{87}Rb and ^{133}Cs*. PhD thesis, University of Durham (2011). [p 67]

[123] S. Friebel, C. D'Andrea, J. Walz, M. Weitz, and T. W. Hänsch. *CO_2-laser optical lattice with cold rubidium atoms*. Physical Review A **57**, R20 (1998). DOI: 10.1103/PhysRevA.57.R20. [p 70]

[124] L. D. Landau and E. M. Lifschitz. *Mechanics*, volume 1 of *Course of theoretical physics*. Pergamon Press, 2 edition (1969). [p 70]

[125] I. I. Rabi, N. F. Ramsey, and J. Schwinger. *Use of rotating coordinates in magnetic resonance problems*. Reviews of Modern Physics **26**, 167 (1954). DOI: 10.1103/RevModPhys.26.167. [p 76]

[126] Norman Ramsey. *A molecular beam resonance method with separated oscillating fields*. Physical Review **78**, 695 (1950). DOI: 10.1103/PhysRev.78.695. [p 81]

[127] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. *Optics and interferometry with atoms and molecules*. Reviews of Modern Physics **81**, 1051 (2009). DOI: 10.1103/RevModPhys.81.1051. [p 81]

[128] A. Smith, B. E. Anderson, S. Chaudhury, and P. S. Jessen. *Three-axis measurement and cancellation of background magnetic fields to less than 50 µG in a cold atom experiment*. Journal of Physics B: Atomic, Molecular and Optical Physics **44**, 205002 (2011). DOI: 10.1088/0953-4075/44/20/205002. [p 86]

[129] *LM339 quad comparator*. Technical datasheet Rev. 1.0.5, Fairchild Semiconductor Corporation (2012). [p 86]

[130] *74HC423; 74HCT423 dual retriggerable monostable multivibrator with reset*. Technical datasheet Rev. 6, NXP Semiconductor (2011). [p 86]

[131] S. Leslie. *On Spinor Condensates as Amplifiers, Sensors and Tunable Quantum Playgrounds for Studies of Spin*. PhD thesis, University of California, Berkeley (2008). [p 88]

[132] A. Vinit and C. Raman. *Precise measurements on a quantum phase transition in antiferromagnetic spinor Bose–Einstein condensates*. (2014). ARXIV: 1410.1938. [p 90]

[133] L. Santos, M. Fattori, J. Stuhler, and T. Pfau. *Spinor condensates with a laser-induced quadratic Zeeman effect*. Physical Review A **75**, 053606 (2007). DOI: 10.1103/PhysRevA.75.053606. [p 93]

[134] Cheng Chin, Rudolf Grimm, Paul Julienne, and Eite Tiesinga. *Feshbach resonances in ultracold gases*. Reviews of Modern Physics **82**, 1225 (2010). DOI: 10.1103/RevModPhys.82.1225. [p 93]

[135] F. K. Fatemi, K. M. Jones, and P. D. Lett. *Observation of optically induced Feshbach resonances in collisions of cold atoms*. Physical Review Letters **85**, 4462 (2000). DOI: 10.1103/PhysRevLett.85.4462. [p 93]

[136] M. Theis, G. Thalhammer, K. Winkler, M. Hellwig, G. Ruff, R. Grimm, and J. Hecker Denschlag. *Tuning the scattering length with an optically induced Feshbach resonance*. Physical Review Letters **93**, 123001 (2004). DOI: 10.1103/PhysRevLett.93.123001. [p 93]

[137] Kevin M. Jones, Eite Tiesinga, Paul D. Lett, and Paul S. Julienne. *Ultracold photoassociation spectroscopy: Long-range molecules and atomic scattering*. Reviews of Modern Physics **78**, 483 (2006). DOI: 10.1103/RevModPhys.78.483. [p 93]

[138] D. J. Papoular, G. V. Shlyapnikov, and J. Dalibard. *Microwave-induced Fano-Feshbach resonances*. Physical Review A **81**, 041603 (2010). DOI: 10.1103/PhysRevA.81.041603. [p 93]

[139] J. Kronjager. *Coherent Dynamics of Spinor Bose–Einstein condensates*. PhD thesis, University of Hamburg (2007). [p 94]

[140] Yuki Endo and Tetsuro Nikuni. *Kinetic theory of a spin-1 Bose condensed gas at finite temperatures*. Journal of Low Temperature Physics **163**, 92 (2011). DOI: 10.1007/s10909-010-0340-y. [p 97]

[141] J. Guzman. *Explorations of Magnetic Phases in $F = 1$ ^{87}Rb Spinor Condensates*. PhD thesis, University of California, Berkeley (2012). [pp 99 and 100]

[142] E. L. Hahn. *Spin echoes*. Physical Review **80**, 580 (1950). DOI: 10.1103/PhysRev.80.580. [p 101]

[143] M. F. Andersen, A. Kaplan, and N. Davidson. *Echo spectroscopy and quantum stability of trapped atoms*. Physical Review Letters **90**, 023001 (2003). DOI: 10.1103/PhysRevLett.90.023001. [pp 101 and 156]

[144] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. *Analysis of dephasing mechanisms in a standing-wave dipole trap*. Physical Review A **72**, 023406 (2005). DOI: 10.1103/PhysRevA.72.023406. [pp 101 and 156]

[145] M. Egorov, R. P. Anderson, V. Ivannikov, B. Opanchuk, P. Drummond, B. V. Hall, and A. I. Sidorov. *Long-lived periodic revivals of coherence in an interacting Bose–Einstein condensate*. Physical Review A **84**, 021605 (2011). DOI: 10.1103/PhysRevA.84.021605. [p 101]

[146] H. Y. Carr and E. M. Purcell. *Effects of diffusion on free precession in nuclear magnetic resonance experiments*. Physical Review **94**, 630 (1954). DOI: 10.1103/PhysRev.94.630. [p 101]

[147] S. Meiboom and D. Gill. *Modified spin-echo method for measuring nuclear relaxation Times*. Review of Scientific Instruments **29**, 688 (1958). DOI: 10.1063/1.1716296. [p 101]

[148] Yujiro Eto, Mark Sadgrove, Sho Hasegawa, Hiroki Saito, and Takuya Hirano. *Control of spin current in a Bose gas by periodic application of π pulses*. Physical Review A **90**, 013626 (2014). DOI: 10.1103/PhysRevA.90.013626. [p 120]

[149] B. S. Mathur, H. Tang, and W. Happer. *Light shifts in the alkali atoms*. Physical Review **171**, 11 (1968). DOI: 10.1103/PhysRev.171.11. [p 121]

[150] Claude Cohen-Tannoudji and Jacques Dupont-Roc. *Experimental study of Zeeman light shifts in weak magnetic fields*. Physical Review A **5**, 968 (1972). DOI: 10.1103/PhysRevA.5.968. [p 121]

[151] J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. Gurudev Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin. *Nanoscale magnetic sensing with an individual electronic spin in diamond*. Nature **455**, 644 (2008). DOI: 10.1038/nature07279. [pp 123 and 124]

[152] Gopalakrishnan Balasubramanian, I. Y. Chan, Roman Kolesov, Mohannad Al-Hmoud, Julia Tisler, Chang Shin, Changdong Kim, Aleksander Wojcik, Philip R. Hemmer, Anke Krueger, Tobias Hanke, Alfred Leitenstorfer, Rudolf Bratschitsch, Fedor Jelezko, and Jörg Wrachtrup.

Nanoscale imaging magnetometry with diamond spins under ambient conditions. Nature **455**, 648 (2008). DOI: 10.1038/nature07278. [p 124]

[153] S. D. Billings. *Discrimination and classification of buried unexploded ordnance using magnetometry.* IEEE Transactions on Geoscience and Remote Sensing **42**, 1241 (2004). DOI: 10.1109/TGRS.2004.826803. [p 124]

[154] Riitta Hari and Riitta Salmelin. *Magnetoencephalography: From SQUIDs to neuroscience: Neuroimage 20th anniversary special edition.* NeuroImage **61**, 386 (2012). DOI: 10.1016/j.neuroimage.2011.11.074. [p 124]

[155] Dmitry Budker and Derek F. Jackson Kimball. *Optical Magnetometry.* Cambridge University Press (2013). [pp 124 and 149]

[156] H. B. Dang, A. C. Maloof, and M. V. Romalis. *Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer.* Applied Physics Letters **97**, 151110 (2010). DOI: 10.1063/1.3491215. [pp 124 and 151]

[157] Jan Jeske, Jared H. Cole, and Andrew D. Greentree. *Laser threshold magnetometer reaching attotesla precision.* (2014). ARXIV: 1410.6239. [p 124]

[158] D. M. Harber, H. J. Lewandowski, J. M. McGuirk, and E. A. Cornell. *Effect of cold collisions on spin coherence and resonance shifts in a magnetically trapped ultracold gas.* Physical Review A **66**, 053616 (2002). DOI: 10.1103/PhysRevA.66.053616. [p 125]

[159] Matthew L. Terraciano, Mark Bashkansky, and Fredrik K. Fatemi. *Faraday spectroscopy of atoms confined in a dark optical trap.* Physical Review A **77**, 063417 (2008). DOI: 10.1103/PhysRevA.77.063417. [p 125]

[160] Fredrik K. Fatemi and Mark Bashkansky. *Spatially resolved magnetometry using cold atoms in dark optical tweezers.* Optics Express **18**, 2190 (2010). DOI: 10.1364/OE.18.002190. [p 125]

[161] M. Koschorreck, M. Napolitano, B. Dubost, and M. W. Mitchell. *Subprojection-noise sensitivity in broadband atomic magnetometry.* Physical Review Letters **104**, 093602 (2010). DOI: 10.1103/PhysRevLett.104.093602. [p 125]

[162] Amita B. Deb, Bianca J. Sawyer, and Niels Kjærgaard. *Dispersive probing of driven pseudospin dynamics in a gradient field.* Physical Review A **88**, 063607 (2013). DOI: 10.1103/PhysRevA.88.063607. [p 125]

[163] P. Schmidt, D. Clark, K. Leslie, M. Bick, D. Tilbrook, and C. Foley. *GETMAG — a SQUID magnetic tensor gradiometer for mineral and oil exploration*. Exploration Geophysics **35**, 297 (2004). DOI: 10.1071/EG04297. [p 126]

[164] R. Stolz, V. Zakosarenko, M. Schulz, A. Chwala, L. Fritzsch, H. Meyer, and E. Köstlin. *Magnetic full-tensor SQUID gradiometer system for geophysical applications*. The Leading Edge **25**, 178 (2006). DOI: 10.1190/1.2172308. [p 126]

[165] S. T. Keenan, J. A. Young, C. P. Foley, and J. Du. *A high-tc flip-chip SQUID gradiometer for mobile underwater magnetic sensing*. Superconductor Science and Technology **23**, 025029 (2010). DOI: 10.1088/0953-2048/23/2/025029. [p 126]

[166] R. H. Koch, G. A. Keefe, and G. Allen. *Room temperature three sensor magnetic field gradiometer*. Review of Scientific Instruments **67**, 230 (1996). DOI: 10.1063/1.1146576. [p 126]

[167] Yangyi Sui, Guang Li, Shilong Wang, and Jun Lin. *Compact fluxgate magnetic full-tensor gradiometer with spherical feedback coil*. Review of Scientific Instruments **85**, 014701 (2014). DOI: 10.1063/1.4856675. [p 126]

[168] T. R. Clem, G. J. Kekelis, J. D. Lathrop, D. J. Overway, and W. M. Wynn. *Superconducting magnetic gradiometers for mobile applications with an emphasis on ordnance detection*. In Harold Weinstock (editor), *SQUID Sensors: Fundamentals, Fabrication and Applications*, number 329 in NATO ASI Series, pages 517–568. Springer Netherlands (1996). [p 126]

[169] A. C. Fraser-Smith. *The magnetic field gradiometer*. Technical report E723-1, Radioscience Laboratory, Stanford Electronics Laboratories, Stanford University (1983). Govt. Accession Number: ADA131477. [p 126]

[170] R. H. Koch, J. R. Rozen, J. Z. Sun, and W. J. Gallagher. *Three SQUID gradiometer*. Applied Physics Letters **63**, 403 (1993). DOI: 10.1063/1.110032. [p 126]

[171] L. B. Pedersen and T. M. Rasmussen. *The gradient tensor of potential field anomalies: Some implications on data collection and data processing of maps*. Geophysics **55**, 1558 (1990). DOI: 10.1190/1.1442807. [pp 126 and 147]

[172] Michael S. Zhdanov, Hongzhu Cai, and Glenn A. Wilson. *3D inversion of SQUID magnetic tensor data*. Journal of Geology & Geosciences **1**, 1000104 (2012). DOI: 10.4172/2329-6755.1000104. [p 126]

[173] David A. Clark. *New methods for interpretation of magnetic vector and gradient tensor data i: eigenvector analysis and the normalised source strength*. Exploration Geophysics **43**, 267 (2012). [pp 126 and 147]

[174] *T877 tensor gradiometer*. Technical datasheet, Tristan Technologies (2004). [p 128]

[175] Zhong-Kun Hu, Xiao-Chun Duan, Min-Kang Zhou, Bu-Liang Sun, Jin-Bo Zhao, Mao-Mao Huang, and Jun Luo. *Simultaneous differential measurement of a magnetic-field gradient by atom interferometry using double fountains*. Physical Review A **84**, 013620 (2011). DOI: 10.1103/PhysRevA.84.013620. [p 129]

[176] Yujiro Eto, Hayato Ikeda, Hirosuke Suzuki, Sho Hasegawa, Yasushi Tomiyama, Sawako Sekine, Mark Sadgrove, and Takuya Hirano. *Spin-echo-based magnetometry with spinor Bose–Einstein condensates*. Physical Review A **88**, 031602 (2013). DOI: 10.1103/PhysRevA.88.031602. [p 129]

[177] C. Affolderbach, M. Stähler, S. Knappe, and R. Wynands. *An all-optical, high-sensitivity magnetic gradiometer*. Applied Physics B **75**, 605 (2002). DOI: 10.1007/s00340-002-0959-8. [p 130]

[178] K. O. Roberts, T. McKellar, J. Fekete, A. Rakonjac, A. B. Deb, and N. Kjærgaard. *Steerable optical tweezers for ultracold atom studies*. Optics Letters **39**, 2012 (2014). DOI: 10.1364/OL.39.002012. [p 130]

[179] G. T. Foster, J. B. Fixler, J. M. McGuirk, and M. A. Kasevich. *Method of phase extraction between coupled atom interferometers using ellipse-specific fitting*. Optics Letters **27**, 951 (2002). DOI: 10.1364/OL.27.000951. [pp 133 and 134]

[180] 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). DOI: 10.1103/PhysRevA.88.053620. [p 133]

[181] A. Fitzgibbon, M. Pilu, and R. B. Fisher. *Direct least square fitting of ellipses*. IEEE Transactions on Pattern Analysis and Machine Intelligence **21**, 476 (1999). DOI: 10.1109/34.765658. [p 134]

[182] Zygmunt L. Szpak, Wojciech Chojnacki, and Anton van den Hengel. *Guaranteed ellipse fitting with the Sampson distance*. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid

(editors), *Computer Vision – ECCV 2012*, number 7576 in Lecture Notes in Computer Science, pages 87–100. Springer Berlin Heidelberg (2012). [p 135]

[183] N. I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press (1996). [p 135]

[184] Majid Beiki and Laust B. Pedersen. *Eigenvector analysis of gravity gradient tensor to locate geologic bodies*. Geophysics **75**, I37 (2010). DOI: 10.1190/1.3484098. [p 147]

[185] Andreas Steffen, Wolfgang Alt, Maximilian Genske, Dieter Meschede, Carsten Robens, and Andrea Alberti. *Note: In situ measurement of vacuum window birefringence by atomic spectroscopy*. Review of Scientific Instruments **84**, 126103 (2013). DOI: 10.1063/1.4847075. [pp 150 and 166]

[186] Peter D. D. Schwindt, Svenja Knappe, Vishal Shah, Leo Hollberg, John Kitching, Li-Anne Liew, and John Moreland. *Chip-scale atomic magnetometer*. Applied Physics Letters **85**, 6409 (2004). DOI: 10.1063/1.1839274. [p 151]

[187] Vishal Shah, Svenja Knappe, Peter D. D. Schwindt, and John Kitching. *Subpicotesla atomic magnetometry with a microfabricated vapour cell*. Nature Photonics **1**, 649 (2007). DOI: 10.1038/nphoton.2007.201. [p 151]

[188] D. Robbes. *Highly sensitive magnetometers—a review*. Sensors and Actuators A: Physical **129**, 86 (2006). DOI: 10.1016/j.sna.2005.11.023. [p 151]

[189] W. Muessel, H. Strobel, D. Linnemann, D. B. Hume, and M. K. Oberthaler. *Scalable spin squeezing for quantum-enhanced magnetometry with Bose–Einstein condensates*. Physical Review Letters **113**, 103004 (2014). DOI: 10.1103/PhysRevLett.113.103004. [p 152]

[190] J. P. Barrat and C. Cohen-Tannoudji. *Optical pumping and the density matrix formalism*. Journal de Physique et Le Radium **22**, 329 (1961). DOI: 10.1051/jphysrad:01961002206032900. [p 154]

[191] W. Happer and B. S. Mathur. *Effective operator formalism in optical pumping*. Physical Review **163**, 12 (1967). DOI: 10.1103/PhysRev.163.12. [p 154]

[192] N. L. Manakov and V. D. Ovsyannikov. *Stark effect in hyperfine structure sublevels and splitting of $n^2 S_{1/2}$ states of alkali atoms in a nonresonant optical field*. Journal of Experimental and Theoretical Physics **48**, 406 (1978). [p 154]

[193] N. L. Manakov, V. D. Ovsiannikov, and L. P. Rapoport. *Atoms in a laser field*. Physics Reports **141**, 320 (1986). DOI: 10.1016/S0370-1573(86)80001-1. [p 154]

[194] J. C. Camparo, R. P. Frueholz, and C. H. Volk. *Inhomogeneous light shift in alkali-metal atoms*. Physical Review A **27**, 1914 (1983). DOI: 10.1103/PhysRevA.27.1914. [p 154]

[195] Masao Takamoto, Feng-Lei Hong, Ryoichi Higashi, and Hidetoshi Katori. *An optical lattice clock*. Nature **435**, 321 (2005). DOI: 10.1038/nature03541. [p 155]

[196] Jun Ye, H. J. Kimble, and Hidetoshi Katori. *Quantum state engineering and precision metrology using state-insensitive light traps*. Science **320**, 1734 (2008). DOI: 10.1126/science.1148259. [p 155]

[197] R. Chicireanu, K. D. Nelson, S. Olmschenk, N. Lundblad, A. Derevianko, and J. V. Porto. *Differential light-shift cancellation in a magnetic-field-insensitive transition of ^{87}Rb*. Physical Review Letters **106**, 063002 (2011). DOI: 10.1103/PhysRevLett.106.063002. [pp 155 and 190]

[198] Huidong Kim, Hyok Sang Han, and D. Cho. *Magic polarization for optical trapping of atoms without stark-induced dephasing*. Physical Review Letters **111**, 243004 (2013). DOI: 10.1103/PhysRevLett.111.243004. [pp 155 and 156]

[199] Y. O. Dudin, R. Zhao, T. A. B. Kennedy, and A. Kuzmich. *Light storage in a magnetically dressed optical lattice*. Physical Review A **81**, 041805 (2010). DOI: 10.1103/PhysRevA.81.041805. [p 155]

[200] Andrei Derevianko. *Theory of magic optical traps for Zeeman-insensitive clock transitions in alkali-metal atoms*. Physical Review A **81**, 051606 (2010). DOI: 10.1103/PhysRevA.81.051606. [p 155]

[201] V. V. Flambaum, V. A. Dzuba, and A. Derevianko. *Magic frequencies for cesium primary-frequency standard*. Physical Review Letters **101**, 220801 (2008). DOI: 10.1103/PhysRevLett.101.220801. [pp 155 and 164]

[202] P. Rosenbusch, S. Ghezali, V. A. Dzuba, V. V. Flambaum, K. Beloy, and A. Derevianko. AC *stark shift of the Cs microwave atomic clock transitions*. Physical Review A **79**, 013404 (2009). DOI: 10.1103/PhysRevA.79.013404. [pp 155, 159, and 161]

[203] Jai Min Choi and D. Cho. *Elimination of inhomogeneous broadening for a ground-state hyperfine transition in an optical trap*. Journal of Physics: Conference Series **80**, 012037 (2007). DOI: 10.1088/1742-6596/80/1/012037. [p 155]

[204] N. Lundblad, M. Schlosser, and J. V. Porto. *Experimental observation of magic-wavelength behavior of ^{87}Rb atoms in an optical lattice*. Physical Review A **81**, 031611 (2010). DOI: 10.1103/PhysRevA.81.031611. [p 155]

[205] Andrei Derevianko. *"Doubly magic" conditions in magic-wavelength trapping of ultracold alkali-metal atoms*. Physical Review Letters **105**, 033002 (2010). DOI: 10.1103/PhysRevLett.105.033002. [p 155]

[206] A. Kaplan, M. F. Andersen, T. Grünzweig, and N. Davidson. *Hyperfine spectroscopy of optically trapped atoms*. Journal of Optics B: Quantum and Semiclassical Optics **7**, R103 (2005). DOI: 10.1088/1464-4266/7/8/R01. [p 156]

[207] I. Novikova, A. B. Matsko, V. L. Velichansky, M. O. Scully, and G. R. Welch. *Compensation of AC stark shifts in optical magnetometry*. Physical Review A **63**, 063802 (2001). DOI: 10.1103/PhysRevA.63.063802. [p 156]

[208] Kunyan Zhu, Neal Solmeyer, Cheng Tang, and David S. Weiss. *Absolute polarization measurement using a vector light shift*. Physical Review Letters **111**, 243006 (2013). DOI: 10.1103/PhysRevLett.111.243006. [p 156]

[209] K. L. Corwin, S. J. M. Kuppens, D. Cho, and C. E. Wieman. *Spin-polarized atoms in a circularly polarized optical dipole trap*. Physical Review Letters **83**, 1311 (1999). DOI: 10.1103/PhysRevLett.83.1311. [p 156]

[210] Ivan H. Deutsch and Poul S. Jessen. *Quantum-state control in optical lattices*. Physical Review A **57**, 1972 (1998). DOI: 10.1103/PhysRevA.57.1972. [p 156]

[211] P. J. Lee, M. Anderlini, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto. *Sublattice addressing and spin-dependent motion of atoms in a double-well lattice*. Physical Review Letters **99**, 020402 (2007). DOI: 10.1103/PhysRevLett.99.020402. [p 156]

[212] N. Lundblad, J. M. Obrecht, I. B. Spielman, and J. V. Porto. *Field-sensitive addressing and control of field-insensitive neutral-atom qubits*. Nature Physics **5**, 575 (2009). DOI: 10.1038/nphys1330. [p 156]

[213] Y.-J. Lin, R. L. Compton, K. Jiménez-García, J. V. Porto, and I. B. Spielman. *Synthetic magnetic fields for ultracold neutral atoms*. Nature **462**, 628 (2009). DOI: 10.1038/nature08609. [p 156]

[214] G. Edward Marti, Andrew MacRae, Ryan Olf, Sean Lourette, Fang Fang, and Dan M. Stamper-Kurn. *Coherent magnon optics in a ferromagnetic spinor Bose–Einstein condensate*. Physical Review Letters **113**, 155302 (2014). DOI: 10.1103/PhysRevLett.113.155302. [p 157]

[215] J. Mitroy, M. S. Safronova, and Charles W. Clark. *Theory and applications of atomic and ionic polarizabilities*. Journal of Physics B: Atomic, Molecular and Optical Physics **43**, 202001 (2010). DOI: 10.1088/0953-4075/43/20/202001. [pp 159, 161, 163, 165, and 190]

[216] Fam Le Kien, Philipp Schneeweiss, and Arno Rauschenbeutel. *Dynamical polarizability of atoms in arbitrary light fields: general theory and application to cesium*. The European Physical Journal D **67**, 1 (2013). DOI: 10.1140/epjd/e2013-30729-x. [pp 159, 160, 161, 162, 163, and 164]

[217] Bindiya Arora and B. K. Sahoo. *State-insensitive trapping of Rb atoms: Linearly versus circularly polarized light*. Physical Review A **86**, 033416 (2012). DOI: 10.1103/PhysRevA.86.033416. [pp 159, 161, 162, 163, 164, 165, and 166]

[218] Keith D. Bonin and Michael A. Kadar-Kallen. *Theory of the light-force technique for measuring polarizabilities*. Physical Review A **47**, 944 (1993). DOI: 10.1103/PhysRevA.47.944. [p 166]

[219] Eugene Hecht. *Optics*. Addison-Wesley, 4th edition (2002). [pp 166 and 167]

[220] C. D. Herold, V. D. Vaidya, X. Li, S. L. Rolston, J. V. Porto, and M. S. Safronova. *Precision measurement of transition matrix elements via light shift cancellation*. Physical Review Letters **109**, 243003 (2012). DOI: 10.1103/PhysRevLett.109.243003. [p 175]

[221] Jie Zhang, Baoguo Yang, and Yunbo Zhang. *Magnetization oscillation of a spinor condensate induced by a magnetic field gradient*. Physical Review A **83**, 053634 (2011). DOI: 10.1103/PhysRevA.83.053634. [p 180]

[222] M. D. Hürlimann. *Carr–Purcell sequences with composite pulses*. Journal of Magnetic Resonance **152**, 109 (2001). DOI: 10.1006/jmre.2001.2370. [p 202]