

Faraday Magnetic Resonance Imaging of Bose-Einstein Condensates

Martijn Jasperse

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

Supervisory committee:

Dr Lincoln Turner

Dr Russell Anderson

Prof Kris Helmerson

School of Physics and Astronomy

Monash University

December 2015

Copyright Notices

Notice 1

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

Notice 2

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Abstract

This thesis presents a novel magnetic resonance imaging (MRI) technique for spinor Bose-Einstein condensates (BEC) with the paramagnetic Faraday effect. This quantum-photon interface couples the spin of the BEC to the polarisation of an off-resonant laser, which in a magnetic field gradient allows a 1D density profile of the BEC to be reconstructed.

As opposed to conventional imaging techniques that are destructive or diffraction limited, this technique is minimally destructive to the BEC and not diffraction limited. Multishot *in situ* imaging of a single condensate is therefore possible, paving the way for time-resolved studies and 2D/3D reconstructions.

A theoretical model for Faraday imaging is developed using the tensor polarisability formalism for the atom-light interaction, and the signal-to-noise ratio is derived. Resolution limits induced by Stern-Gerlach separation are considered, encouraging rapid measurements with strong magnetic field gradients.

Analysing the trapping potential reveals that the ‘magic-zero’ wavelength 790 nm produces no dipole force and enables a tightly focused probe laser to be used without perturbing the trap. Polarisation-maintaining fibers were seen to impart large fluctuations depending on the incident polarisation, which were minimised. The broad diode laser emission background is seen to reduce the lifetime of the BEC, which was corrected with an interference filter.

Short-time Fourier transforms are used to process the photodetected signal, showing the evolution of the polarisation rotation frequency and amplitude. Frequency-modulation is observed at the power line frequency, corresponding to magnetic field fluctuations induced by nearby equipment. Birefringence of optical elements results in elliptical polarisation of the probe and an effective magnetic field, which is cancelled using a quarter-waveplate. The quadratic Zeeman effect causes rapid amplitude modulation, and is eliminated by applying a microwave dressing field.

Radiative spin echo is achieved in a magnetic field gradient, proving that dephasing is coherent, and that multishot imaging possible. Magnetic resonance imaging (MRI) is performed on a partially-evaporated atom cloud, distinguishing the thermal and condensed fractions. A split dipole trap is used to create two spatially separated BECs, which are individually resolved during MRI.

The resolution is only limited by the strength of the applied magnetic field gradient, and is not subject to the diffraction limit. Custom coils will enable MRI of condensates at the sub-micron scale.

Declaration

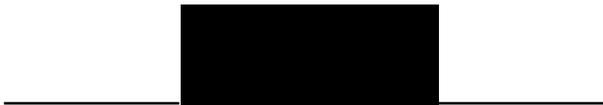
This thesis contains no material that has been accepted for the award of any other degree or diploma in any university or other institution. To the best of my knowledge the thesis contains no material previously published or written by another person, except where due reference is made in the text of 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.



Dr Lincoln Turner



Dr Russell Anderson



Prof Kristian Helmerson



Acknowledgements

Research is a community effort: the work of many individuals striving to the common goal of pushing the boundaries of science. There are so many people who contributed to this project – with knowledge, influence, encouragement or inebriation – that helped bring it to a successful conclusion. Collaboration, camaraderie and caffeine were the crucial components that culminated in the critical conditions for condensation.

I have been lucky in the course of my research career to have met a large number of amazing scientists across a diverse range of disciplines who continue to inspire me. I don't know how to begin to thank everyone who contributed in one way or another over the years I have been at Monash, and I do not attempt an exhaustive list. I can only hope that the friends I made, the mentors who guided me, and the professionals who assisted me are aware of my gratitude.

In the beginning there was an empty lab, an optics table, and a whole lot of dreams. It has been the stewardship of Lincoln, Russ and Kris that got us where we are today. The trajectory may have exhibited Brownian motion at times, but I'm proud of what we accomplished together. I especially want to thank my fellow warriors of the spinor BEC lab – Alex and Lisa – for the machine we collectively built and the hours we collectively spent debugging it. Thanks also to the wider BEC group at Monash – members past and present – because BEC physics is hard, and because 'borrowing' things is an essential part of experimental science.

Physicists are interesting characters, and none moreso than my partners-in-crime, Katie, Jeremy, Elliot, Wan, Bopper and Jack. Thanks for being distractingly sexy and endlessly entertaining. Thanks also to my fellow former dungeon-dwellers, the coffee-train, all of the tearoom banditos, the glorious pod machine, Tuesday for being Tuesday and Wednesday for sometimes also being Tuesday. A special note of thanks to Jean, for yelling at people who deserved it, and without whom the department would simply fall apart. Thanks as well to the New Horizons building, without which this project would have been in danger of finishing on time.

I'd especially like to thank my family for their unwavering support and encouragement. When it rains, it pours, and there were some really hard times, especially when my health was failing. Thanks for the understanding, the strong encouragement, and for believing in me when I didn't believe in myself. As ever, this is for you guys.

And finally, to my darling Peta, thanks for sharing this journey with me. No matter what you say, it is still you.

Table of Contents

Abstract	i
Declaration	iii
Acknowledgements	v
1 Introduction	1
1.1 Cold atoms beyond the diffraction limit	1
1.2 Time-of-flight imaging	3
1.3 Other common imaging methods	5
1.4 The Faraday effect as a probe	7
1.5 Existing Faraday imaging techniques	8
1.6 NMR and MRI	9
1.7 MRI in atomic vapours	11
1.8 What’s new in this work	12
1.9 Outline of this thesis	12
2 The off-resonant dipole interaction	15
2.1 Review of the atomic polarisability expansion	15
2.2 The far-detuned limit	18
2.3 The off-resonant trapping force	21
2.4 The ‘magic’ wavelength	21
2.5 The Faraday effect	23
2.6 The vector light-shift and effective magnetic field	27
2.7 The tensor light-shift	28
2.8 Summary	29
3 Faraday magnetic resonance imaging	31
3.1 Measuring Faraday rotation	31
3.2 Faraday imaging in 1D	33
3.3 Signal-to-noise ratio	35
3.4 Choice of detuning	37
3.5 Coherent Faraday signal dephasing	39
3.6 Rephasing with spin echo	42
3.7 Slice selection and higher dimensional images	44
3.8 Stern-Gerlach blurring	48
3.9 Choice of magnetic field gradient	49
3.10 Summary	51

4	Spinor BEC apparatus	53
4.1	The Bose-Einstein condensate factory	53
4.2	Beatnote microwave offset lock	57
4.3	Industrial reliability with microcontroller interlocks	62
4.4	BIAS: A modular image acquisition system	67
4.5	Summary	69
5	Faraday probe beam	71
5.1	Probe laser and locking scheme	71
5.2	Pre-cell optics	73
5.3	Polarisation stability and fiber axis alignment	74
5.4	Post-cell optics and imaging system	76
5.5	Striking the BEC bulls-eye	79
5.6	Pointing stability of the Faraday beam	83
5.7	Trapped BEC lifetime	85
5.8	Anomalous resonant scattering	87
5.9	Characterisation of off-resonant scattering	89
5.10	Summary	91
6	Photodetection and signal processing	93
6.1	High-precision polarimetry	93
6.2	Photodetector design considerations	94
6.3	Shot-noise-limited photodetection	97
6.4	Electro-optic BEC calibration	99
6.5	Revised photodetector design	102
6.6	Data acquisition	103
6.7	Structure of the Faraday signal	104
6.8	Short-time Fourier transform analysis	107
6.9	Overlapping and oversampling	109
6.10	Role of window size	109
6.11	Summary	112
7	Signal calibration and optimisation	113
7.1	Magnetic field calibration	113
7.2	Measuring and correcting power-line modulation	116
7.3	Blurring from power line modulation	120
7.4	Effect of probe polarisation	121
7.5	Eliminating the probe vector light-shift	124
7.6	Quadratic Zeeman effect	126
7.7	Microwave dressing	129
7.8	Measuring the microwave Rabi frequency	131
7.9	BEC component separation	134

7.10 Summary	137
8 Magnetic resonance imaging of BEC	139
8.1 Coordinate systems	139
8.2 Imaging axis	141
8.3 Magnetic field gradient synthesis	142
8.4 Faraday gradiometry	145
8.5 Bias coil gradients	146
8.6 Quadrupole coil gradients	149
8.7 Radiative spin echo	151
8.8 Considerations for MRI of BEC	154
8.9 MRI of a partially evaporated BEC	155
8.10 MRI of twin BECs	158
8.11 Summary	160
9 Conclusions and outlook	163
9.1 Conclusions	163
9.2 Outlook	166
References	171
A Oven controller implementation	187
A.1 Inputs and outputs	187
A.2 Code listing	189
B Loss mechanisms	195
B.1 Thomas-Fermi profile	196
B.2 Photon scattering rate	197
B.2.1 Far-detuned limit	200
B.3 Photoassociation resonances	202

List of Figures

1.1	Example of a background-subtracted absorption image.	4
1.2	Polarisation rotation of an initially vertically linearly polarised beam.	7
2.1	Hyperfine energy-level scheme for ^{87}Rb	20
2.2	Wavelength dependence of scalar polarisability for ^{87}Rb	22
2.3	Larmor precession for an off-axis field, and a partial spin-tip.	24
2.4	Wavelength dependence of the Faraday effect.	26
3.1	Schematic of the balanced polarimeter.	31
3.2	A magnetic field gradient causes Larmor precession at different rates.	33
3.3	Dependence of signal-to-noise ratio on wavelength.	37
3.4	Dependence of probe power on wavelength for a fixed scattering rate.	38
3.5	Schematic of coherent spin dephasing.	40
3.6	Expected Faraday signal for different cloud shapes.	41
3.7	Comparison of the expected signal envelope for clouds with dark features.	42
3.8	Schematic of gradient recall echo.	43
3.9	Schematic of a phase-encoding spin echo sequence.	46
3.10	Construction of a 'k-space' image.	46
3.11	Stern-Gerlach limit on resolution for varying magnetic field gradients.	50
4.1	A visualisation of the vacuum system.	54
4.2	Laser detunings used to perform cooling and trapping of ^{87}Rb atoms.	55
4.3	Scheme to generate the MOT, Zeeman slower and imaging beams.	56
4.4	Design and arrangement of the magnetic field bias coils.	57
4.5	The ground-state hyperfine splitting of ^{87}Rb , and the D2 line.	58
4.6	Schematic of the digital offset-lock.	59
4.7	The signal processing circuit for the digital offset lock.	61
4.8	The FET-modulation circuit to enable fast-feedback on laser current.	63
4.9	The PLL-synthesiser produces a step response at the desired frequency.	63
4.10	Rf beatnote spectrum of laser locked using digital offset-lock at 6.8 GHz.	64
4.11	The controller state-machine and error-detection routine.	66
4.12	Interaction between modules of the Labscrip suite.	67
4.13	Overview of BIAS and its integration with LABSCRIPT.	68
4.14	Screenshot of BIAS in operation.	69
5.1	Experimental setup to produce Faraday imaging light.	72
5.2	Conditioning optics to generate the probe beam.	74
5.3	Experimental schematic for Faraday beam analysis.	77

5.4	Reimaging of the BEC using relay lenses.	78
5.5	Image of the BEC captured through the Faraday beam path.	83
5.6	Number of atoms remaining in BEC after various hold times.	86
5.7	Expected transmission spectrum of the SEMROCK LL01-808.	88
5.8	Optical spectrum of the Faraday ECDL before and after the filter.	88
5.9	Comparison of lifetime measurement with and without filter.	89
5.10	Loss rate dependence on different probe powers.	90
6.1	Schematic of the polarimeter.	93
6.2	Circuit diagram for the differential photodetector.	95
6.3	Calibration of the photodetector gain.	97
6.4	Comparison of single beam noise spectrum to balanced measurement.	98
6.5	Measured noise power for balanced photodetection.	99
6.6	Optical schematic for the ‘electro-optic BEC’.	100
6.7	Example of the calibration possible using the electro-optic BEC.	101
6.8	Measurement of photodetector bandwidth using the electro-optic BEC.	101
6.9	A typical polarimeter measurement of the raw Faraday signal.	105
6.10	Fourier-filtered Faraday measurement.	106
6.11	Schematic of the STFT algorithm.	108
6.12	Spectrogram showing the STFT of a typical measurement.	108
6.13	Close-up of the STFT spectrogram about the Larmor frequency.	108
6.14	Time domain oversampling of the STFT.	110
6.15	Frequency domain oversampling of the STFT.	110
6.16	Comparison of a standard STFT with one 10x oversampled.	110
6.17	Comparison of spectrograms for STFTs with varying window size.	111
7.1	Calibration of the bias coils using the Larmor frequency.	115
7.2	Calibration of bias coils without requiring nulling of other components.	117
7.3	Residuals of the quadratic coil fit showing the accuracy of the fit.	117
7.4	Faraday signal in a y -bias with fitted frequency modulation.	119
7.5	Time dependence of the Faraday signal with an x' -bias.	119
7.6	Worst-case power line blurring for varying magnetic field gradients.	121
7.7	Effect of probe beam polarisation ellipticity on the Faraday signal.	123
7.8	The measured Larmor frequency allows the VLS to be mapped out.	126
7.9	The quadratic Zeeman shift shifts the energy level of the $ m_F = 0\rangle$ state.	127
7.10	Amplitude modulation arising from the quadratic Zeeman effect.	128
7.11	Microwave dressing shifts the $ F = 1, m_F = 0\rangle$ state.	130
7.12	Faraday signal with microwaves of varying detuning.	132
7.13	Separation of the two Larmor tones caused by the quadratic Zeeman shift.	133
7.14	Observation of component-separation in a y -bias but not an x' -bias.	135
7.15	Spectrograms comparing y -bias to x' -bias.	136

8.1	Schematic of the apparatus coordinate systems.	140
8.2	Generation of a linear magnetic field gradient along z	144
8.3	The dipole trap beams can be laterally translated, or split in two.	146
8.4	Faraday signal dephasing rate dependence on applied bias gradient.	148
8.5	Faraday signal with background gradient cancelled.	148
8.6	Larmor frequency dependence on quadrupole field gradient.	150
8.7	Faraday dephasing rate dependence on quadrupole field gradient.	150
8.8	Gradiometry of the quadrupole coil along the z -axis.	151
8.9	Single-shot observation of spin echo.	152
8.10	Multiple sequential spin echoes fit simultaneously.	153
8.11	Observation of 50 sequential spin echoes.	153
8.12	Fit of echo height amplitude over time.	154
8.13	Faraday signal and reconstruction of a partially evaporated BEC.	157
8.14	Comparison of MRIs for different truncation parameters.	158
8.15	Spectrogram of twin BECs.	159
B.1	Known $1g$ -photoassociation resonances for ^{87}Rb	203

List of Tables

2.1	Properties of the D1 and D2 lines in ^{87}Rb	19
2.2	Normalised polarisability coefficient sums in the far-detuned limit.	19
2.3	Normalised coefficients of the quadratic expansion of polarisability.	20
2.4	Comparison of the nematic interaction strength for ^{87}Rb and ^{133}Cs	29
3.1	Magnetic field gradients required to achieve a range of resolutions.	49
4.1	Register settings for Zeeman repump PFD offset-lock.	60
5.1	Comparison of flipper mirror and translation stage.	85
6.1	Comparison of first- and second-generation photodetectors.	103
7.1	Bias coil calibration factors as measured using the quadratic fit method.	117
7.2	Comparison of fit parameters quantifying power line modulation	118
7.3	Fitted parameter values for the vector light-shift measurement.	125
8.1	Bias coil control settings for different bias field orientations.	147
B.1	Fit parameters for photoassociation model.	203

1

Introduction

In this chapter I introduce Bose-Einstein condensates (BECs) and their applications in quantum simulation. I consider the problem of imaging features within the BEC, the most interesting of which are smaller than the optical diffraction limit and require destructive expansion to be resolved. I introduce the paramagnetic Faraday effect as a probe for the BEC spin state, and propose an imaging technique that uses a magnetic field gradient to obtain a profile of the BEC density. This magnetic resonance imaging (MRI) technique is not subject to the diffraction limit, permitting *in situ* observation of the BEC and the finest features within it as they evolve over time.

1.1: Cold atoms beyond the diffraction limit

Bose-Einstein condensates (BECs) are an exquisite tool for understanding fundamental physics. Formed when a cloud of bosonic atoms is cooled to a fraction of a degree above absolute zero [1], they constitute the coldest matter in the known Universe [2]. Confined by a trapping potential, at these extremely low temperatures the wave-like nature of the atoms dominates and they collapse into the quantum mechanical ground state of the potential. The atoms then behave as if they were a single macroscopic quantum object with an overall wavefunction that describes their collective behaviour, and that object is called a ‘condensate’. This phenomenon was predicted by Einstein in 1925 [3, 4], used to explain superfluidity in ^4He by London in 1938 [5], and achieved experimentally with alkali atoms in 1995 [6, 7].

The condensate brings the normally tiny, delicate world of quantum mechanics into the regime where the wavefunction can be studied and manipulated directly. This permits investigation of quantum phenomena such as vortices [8, 9, 10, 11], dark solitons [12, 13, 14], bright solitons [15, 16], and control of the atom-atom interaction through Feshbach resonances [17, 18, 19, 20].

These extremely cold and fragile condensates are only created under the most tightly controlled experimental conditions [21]. This level of precise control and sensitivity to the environment makes BECs perfectly suited to precision measurement through interferometry [22], including inertial sensing [23], gravity gradiometry [24, 25, 26] and magnetometry [27, 28].

Having such precise control over the environment also allows the effective Hamiltonian of the system to be manipulated with lasers and microwaves to mimic other systems. This enables condensates to quantum-mechanically ‘simulate’ other systems [29] by using lattices and synthetic potentials to shed light on their behaviour, such as in the study of turbulence [30, 31, 32] and superconductivity [5, 33].

Initial experiments created condensates in a magnetic trap, which is only capable of trapping weak-field-seeking spin states. This largely eliminates the spin degree of freedom, and restricts the magnetic properties of the system. However, by transferring the BEC to an optical dipole trap during preparation, a *spinor condensate* is formed, which permits confinement of any of the Zeeman substates [34]. These substates then behave as multiple interacting condensates occupying the same volume of space, which experience a spin-dependent contact interaction (collisions). This gives rise to a wide range of rich dynamics [35, 36, 37], including the spontaneous formation of spin domains [38], spin mixing [39], and induced ferromagnetic-to-antiferromagnetic transitions using Feshbach resonances [40, 41].

Persistent features such as spin textures can be created, which are patterns in the local spin orientation of the cloud [42, 43, 36]. Recently these spin textures have been analysed as magnetic quasiparticles (‘magnons’) [44, 28] opening up a new field of ultra-cold ‘magnonics’. Development of a ‘spinful’ imaging technique capable of studying such spin textures is one of the primary motivations for this project.

In the first instance we consider *in situ* imaging of small-scale structures in the density profile. Such features are of interest in both scalar (single-component) and spinor BECs, and minimally-destructive 3D imaging of their dynamics will be an extremely valuable experimental technique in its own right. In the future applications of this work, extensions to enable spin-sensitive imaging will be considered.

The mean-field interaction in a BEC causes density perturbations to be smoothed out over a characteristic distance called the ‘healing length’ [45],

$$\xi = \frac{1}{\sqrt{8\pi n_0 a_s}}, \quad (1.1)$$

where n_0 is the peak condensate density and a_s the s -wave scattering length, which for ^{87}Rb is $a_s = 5.39 \text{ nm}$ [46]. For a typical density of $n_0 \sim 10^{14} \text{ atoms/cm}^3$, the BEC healing length is $\xi \sim 270 \text{ nm}$.

Persistent features of interest such as vortices exist on this length scale, and observing them *in situ* requires an even finer imaging resolution. Typically BECs are imaged using (near) resonant light, which for rubidium is $\lambda = 780 \text{ nm}$, with an associated Abbe diffraction limit [47] of¹ $\lambda/2 = 390 \text{ nm}$. Limited optical access around the vacuum chamber typically restricts the collection numerical aperture (NA), and further

¹ Taking NA = 1 to give a conservative estimate on the smallest achievable resolution.

increases the diffraction limit. In principle shorter wavelength light could be used to image the BEC for an improved diffraction limit, but the quantum efficiency of common photodetectors (such as silicon photodiodes) is typically reduced at short wavelengths.

Direct *in situ* imaging of the features within a BEC is therefore a difficult prospect, as any diffraction-based imaging technique will be incapable of resolving features at the healing length scale. I will now summarise the existing imaging techniques and their advantages and disadvantages in attempting to overcome this, before introducing an imaging technique that is not diffraction-based and therefore not subject to the same limits.

1.2: Time-of-flight imaging

The reigning technique for collecting information from BECs is an absorption-imaging method called time-of-flight (TOF) imaging. In TOF imaging, the atom cloud is released from its trap and allowed to fall under gravity. Without the trapping potential to confine the atoms, the repulsive mean-field interaction causes rapid free expansion,² which becomes ballistic as the BEC grows and the density decreases [48, 49, 50]. After a brief interval of expansion, the features within the cloud become large enough to resolve using resonant light. The BEC is then exposed to a resonant laser beam, which is absorbed by the atoms in accordance with the Beer-Lambert law.³ This absorption leaves a shadow in the probe beam that is recorded on a camera, which can be used to compute the optical depth of the BEC, and hence its column density.

The remarkable property of ballistic expansion is that it occurs coherently: the structure of the BEC is preserved. Features such as vortices which were present in the BEC remain after expansion, but are also magnified. In particular, as the density decreases during expansion, the healing length increases, so features at this scale grow in size [52, 53]. Once they grow beyond the diffraction limit, those features can be clearly resolved by absorption imaging.

An absorption image is typically produced from a series of three exposures called the ‘flat’, ‘atoms’, and ‘dark’ frames (Figure 1.1). The flat-field frame is an image of the probe beam without any atoms present, quantifying the intensity profile of the probe. The atoms frame contains the shadow of the BEC produced by atoms scattering the probe beam. The dark field is taken without any illumination, and is a measure of the camera’s pedestal reading and dark noise level. The optical depth (OD) is then

$$\text{OD} = -\log \left(\frac{I_{\text{atoms}} - I_{\text{dark}}}{I_{\text{flat}} - I_{\text{dark}}} \right). \quad (1.2)$$

² An immediate consequence is that in an anisotropic trap the most tightly confined direction expands most rapidly.

³ Typically a correction factor is required to account for saturation of the atomic transition [51].

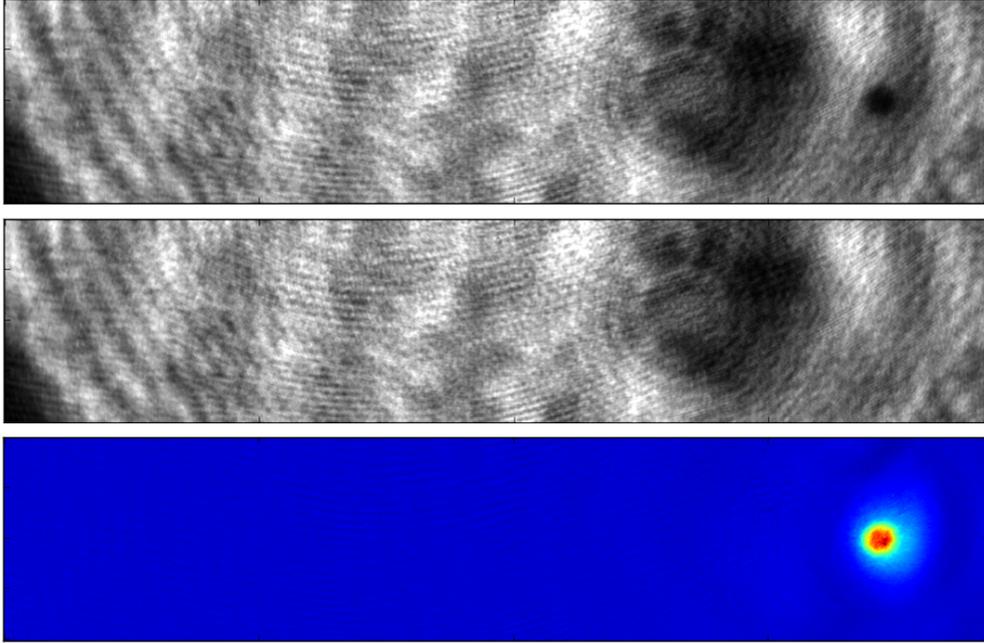


Figure 1.1: Example of a background-subtracted absorption image. Most of the structure in the absorption image (top) is common to the flat-field image (middle) and is removed from the computed optical depth (bottom), leaving a well-defined image of the cloud's column density with some weak residual fringes.

The coherent probe beam often exhibits fringes from the vacuum system windows and other optical components, which are susceptible to vibration. It is therefore desirable to take the flat-field and atoms exposures as close together as possible. One technique that accomplishes this is the 'back-to-back' procedure, whereby the atoms are held in a dark state (which does not absorb the imaging light) for the flat-field exposure, then are optically pumped into a state resonant with the imaging light, after which the atoms frame is captured. Interline charge-coupled device (CCD) cameras using an 'overlapped' readout mode can digitise the frame during the acquisition of the subsequent exposure, allowing the time between frames to be tens of microseconds. This limits the motion of the probe beam fringe pattern and improves the quality of the background cancellation.

There are a number of other constraints and limitations on this imaging process. The exposure must be brief to prevent blurring from photon recoil, falling under gravity, and ongoing ballistic expansion. The probe beam must be weak to prevent saturation of the resonant transition, which competes with the desire to have 'full wells' in the flat-field image and achieve the highest dynamic range in the computed optical depth. In the extreme case, a weak probe will produce only a few counts per pixel, resulting in images dominated by dark counts and the readout noise of the camera. Resonant absorption is therefore rarely used for quantitative measurement *in situ* as very high ODs (> 100) are routine.

Furthermore, in order to observe the different spin states of a spinor BEC, it is necessary to physically separate its spin components using Stern-Gerlach separation before imaging it. This must be done quickly, because the local density changes when the spin-states separate, which causes the spin-dependent interaction to change.

Clearly, this imaging method is destructive: the cloud is released from the trap, split into spin components, ballistically expanded, and then imaged with resonant light that scrambles its state and/or heats it above the critical temperature of condensation. Performing time-resolved studies typically requires the formation and subsequent imaging of many BECs, each with a different evolution time. However, this requires reliable recreation of the initial state, which may not always be possible – particularly for turbulent or chaotic systems. Time-resolved studies therefore require the collection of large amounts of data to gain statistical insight into the processes under investigation.

1.3: Other common imaging methods

One solution to achieve minimally-destructive imaging is using phase contrast, which probes the BEC *in situ* with an off-resonant laser beam. The BEC behaves like a weak phase object, and instead of reducing the probe intensity through absorption, the condensate imparts a phase shift on the probe that is proportional to the column density. Measuring the spatially varying phase shift of the probe then allows reconstruction of the cloud profile.

There are a number of techniques to convert the phase shift into an intensity pattern that can be recorded on a CCD camera.⁴ The earliest method applied to imaging BECs *in situ* is ‘dark-ground’ imaging (also known as the ‘schlieren’ technique [55, 56]), which collects the light diffracted by the BEC and then blocks the unscattered component with an opaque mask in the Fourier plane of the imaging system. Only the scattered light propagates through to the camera, producing an image of the cloud.

This technique has long been applied to obtain multiple sequential images of a single BEC [57]. However, the image typically suffers from a low signal-to-noise ratio as the low intensity of light reaching the camera means that the camera registers few counts per pixel. The image is then dominated by dark noise, which is a problem common to all ‘dark-field’ imaging techniques. Furthermore the mask that blocks transmission of the undiffracted beam in the Fourier plane acts as a spatial filter on the image and introduces distortions at low spatial frequencies. The mask should therefore be as small as possible, although a lower limit is set by the waist size of the transmitted beam at the position of the mask, which is increased by lensing induced by refraction through the BEC.

⁴ See [54] for an in-depth comparison of techniques and associated signal-to-noise ratios.

The resolution of the image is also inherently diffraction limited, although ‘super-resolution’ images of vortex lattices have been obtained experimentally [58]. Despite the fact that the vortex cores (400 nm diameter) were smaller than the diffraction limit, distinct intensity minima were observed in the image at positions corresponding to the vortex centres. This is only possible because the separation between vortices in the lattice ($\sim 9 \mu\text{m}$) was much greater than the diffraction limit, and the measured size of the vortex cores (2.4(5) μm) is six times greater than their predicted size.⁵ This technique is effective for a regular structure such as a vortex lattice with well-spaced vortices, but is not applicable to arbitrary structures (such as turbulent vortex tangles [31]). The authors of Ref. [58] note that the low signal-to-noise ratio necessitated small detunings (1–5 linewidths), which was highly destructive to the BEC and prevented multiple images being acquired.

Improvements to the signal-to-noise ratio can be obtained by replacing the opaque mask with a quarter-wave plate [59]. Other techniques such as spatial heterodyne [60], holography [61] and Zernike phase-contrast [62] have also been applied to image cold atom clouds. However, the BEC is optically thin to a far off-resonance probe, and achieving a large numerical aperture at the large working distances necessitated by the vacuum system is technically challenging.

A completely different approach is partial transfer absorption imaging (PTAI) [63, 64] which can be used for magnetically trapped BECs, and has been applied to observe the trajectory of a single vortex in a scalar BEC [63]. The technique works by transferring a fraction of the BEC population into an untrapped state [65, 66, 67], which then undergoes time-of-flight expansion and is subsequently absorption imaged. This retains the benefits of absorption imaging such as the high optical depth for improved signal-to-noise ratios, as well as magnification as a result of ballistic expansion.

However, in applying this technique to multi-shot imaging, there is a limit to the repetition rate, as the outcoupled cloud must have fallen clear of the BEC before the process is repeated. It is also unclear how the technique can be applied to a spinor BEC, as in an optical trap there is no untrapped state that can be transferred to. Potentially Raman coupling to a high-momentum state (see e.g. [68]) could be used to transfer some of the population into a state with enough momentum to escape the dipole trap, but this has not yet been demonstrated.

It would be advantageous to have an imaging technique that is capable of minimally destructive 3D reconstruction of a spinor BEC *in situ*, whose resolution is not limited by diffraction. Such a measurement cannot be based on the existing techniques of absorption or phase-contrast, so a different interaction will now be considered.

⁵ The measured size is limited by the point-spread function of the imaging system, which is at best the diffraction limit (estimated to be 1.9 μm for Ref. [58]).

1.4: The Faraday effect as a probe

As first described by Faraday in 1845, light propagating through matter subject to a magnetic field along the direction of propagation experiences polarisation rotation (Figure 1.2), with the degree of rotation proportional to the magnetic field strength [69, 70]. Physically, this phenomenon arises from the Zeeman splitting of the energy levels inside the medium, causing the σ^\pm polarisation components to experience different detunings, and hence different refractive indices – a process known as circular birefringence. This results in a relative phase lag between the polarisation components, which has the effect of rotating the linear polarisation vector [71]. Measuring the degree of rotation then provides a means of determining properties of the magnetic field, the material, or both.

In a quantum mechanical description, the Faraday effect can be described as a direct coupling of the atom spin in the magnetic material with the spin (polarisation) of the light passing through [72]. This interaction has been studied in cold atoms for some time [73] and is of the right form for quantum non-demolition measurement [74, 75, 76] as well as the creation of polarisation- [77] and spin-squeezed states [78, 79]. Notably, this concept has been extended to macroscopic entanglement of atoms in physically separated vapour cells using the Faraday effect [80].

Applying an orthogonal magnetic field to a spin-polarised medium will cause the spins to precess about the external magnetic field at the Larmor frequency [81]. When Faraday probing a precessing spin, the probe beam then encounters a spin projection along its propagation axis that varies sinusoidally in time, causing the induced Faraday rotation to also oscillate at the Larmor frequency [82]. Sometimes referred to as ‘Faraday spectroscopy’, this pushes the Faraday signal measurement frequency well above baseband where the small induced polarisation rotation is difficult to measure, allowing shot-noise limited measurements to be made.

Long used in warm atomic magnetometers, off-resonant Faraday spectroscopy was first performed in a laser-cooled atomic gas by Isayama *et al.*, published in 1999 [83]. Taking advantage of reduced decoherence effects in a cold atom cloud, precise mea-

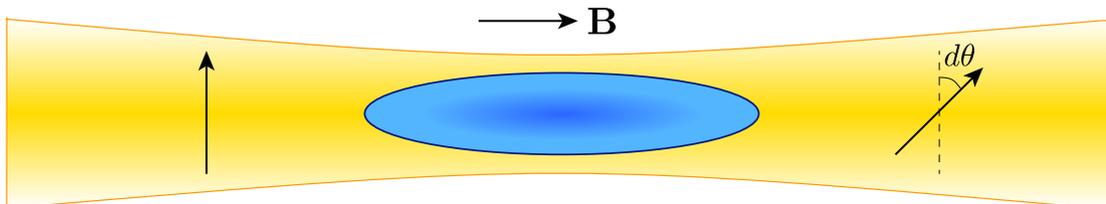


Figure 1.2: A beam with initially vertical linear polarisation passes through a medium subject to a magnetic field B as shown. As a result of the Faraday effect, the polarisation state of the beam is rotated by an amount $d\theta \propto |B|$ at the output.

measurements of the Larmor precession rate were used to determine the external magnetic field strength to within 18 pT. This formed the basis of later cold atom magnetometry experiments measuring Larmor precession of cold atoms in optical lattices [84] and dark optical traps [85].

Developments in spin-squeezing have enabled magnetometry beyond the projection-noise limit through Faraday measurement in cold rubidium [86], demonstrating the potential for squeezed vapour magnetometers based on the Faraday effect. Similarly, performing Faraday spectroscopy using polarisation squeezed light generated from an optical-parametric oscillator increased the measurement sensitivity by 3.2 dB beyond the shot-noise limit [87]. As these two techniques are complementary, Faraday spectroscopy is a promising contender for high-precision magnetometry [88].

1.5: Existing Faraday imaging techniques

The Faraday effect can be applied to the diffraction-contrast imaging techniques described in §1.3 by using linearly polarised light and placing a linear polariser in the imaging path. Adjusting the relative orientation θ of the polariser with respect to the polarisation of the incident beam, the recorded intensity pattern [89] is

$$I_s = I_0 \left[\cos^2 \theta + \frac{\sqrt{3}}{4} \phi \sin(2\theta) - \frac{3}{16} \phi^2 \cos(2\theta) \right], \quad (1.3)$$

where I_0 is the incident intensity and ϕ is the phase shift. Taking $\theta = 90^\circ$ results in dark-field Faraday imaging (DFFI) as only the polarisation-rotated component is detected, and $\theta = 45^\circ$ gives linear phase-contrast imaging (PCI). Other combinations are possible, and $\theta = \pm 70^\circ$ was found to maximise the signal-to-noise ratio in Ref. [89].

Multi-shot imaging with a high repetition rate has been successfully achieved with both DFFI and PCI by using CCD cameras in ‘kinetics’ mode, which does not require complete readout between exposures. PCI was successfully applied to record 40 sequential images of a single BEC at a rate of 20 kHz [44], enabling direct observation of Larmor precession of the spin.⁶

In particular, coherent winding of the transverse magnetisation across the cloud was observed over long timescales (~ 2 ms) in a magnetic field gradient. The spins behave as if they were ‘frozen in place’ and precessing at the value of the local magnetic field. This is significant as the cloud exhibits motional dynamics on this timescale, sampling different magnetic field strengths – yet the magnetisation demonstrates no ‘memory’ of this. Instead, a coherent ‘corkscrew’ winding of the magnetisation is observed over the duration of the experiment⁷ (12 complete phase windings).

⁶ The Larmor frequency in Ref. [44] was 38 kHz, which combined with the 20 kHz acquisition rate results in an aliased Larmor frequency of 2 kHz.

⁷ This observation is crucial to developing magnetic resonance imaging in spinor BEC, which relies on imprinting phase windings that persist over long timescales.

This technique has also been applied to magnetometry [28], and enabled the observation of spin textures [43]. The minimally-destructive nature of Faraday measurement has been demonstrated with DFFI, which was used to record 2000 sequential images of a single spinor BEC with an electron-multiplying CCD [90]. Using a beam-splitter instead of a polariser enables both polarisation components to be simultaneously recorded. This ‘dual-port’ Faraday imaging (DPFI) [91], this allows common-mode subtraction and correction for distortions such as refraction by the BEC.

While these imaging techniques have been demonstrated to be minimally-destructive and enable spatially-resolved imaging of the spin distribution, they are inherently two-dimensional images with diffraction-limited resolution. Imaging of features beyond the diffraction limit, or reconstruction of the full three-dimensional profile of the BEC therefore requires a different technique – one not limited by diffraction.

Magnetic resonance imaging (MRI) is a promising candidate: it is a completely different approach to image formation that circumvents the diffraction limit, but has not yet been applied to imaging BECs.

1.6: NMR and MRI

Medical MRI revolutionised diagnostic imaging by providing a new non-invasive modality for seeing inside the human body. The operational principle uses nuclear magnetic resonance (NMR) to probe the chemical composition and local environment of nuclear spins within the body. Applying a sequence of magnetic field gradients in different directions enables spatial structure to be extracted, and a full three-dimensional profile of the patient to be reconstructed.

The simplest version of NMR is ‘continuous wave’ (CW) NMR, which places a sample in a strong bias magnetic field and exposes it to CW rf radiation. If the applied frequency matches the Zeeman splitting of the nucleons, they can absorb an rf photon and undergo a spin flip. This removes energy from the rf field and enables resonance to be observed by sweeping the frequency of the rf or the magnetic field strength.

The resonant frequency is proportional to the product of the gyromagnetic ratio of the nucleus and the local magnetic field strength. It is therefore sensitive to the composition of the nucleus,⁸ as well as the environment of the atom.⁹ NMR is therefore a powerful technique to determine not just the elemental composition of a sample but its chemical structure as well: a discovery for which the 1952 Nobel prize in Physics was awarded [92].

⁸ Both the gyromagnetic ratio and nuclear isospin vary between different elements and isotopes, changing the resonance condition.

⁹ For example, nuclear shielding by electrons, molecular bonds, and electronegativity all cause perturbations to the local magnetic field that result in a measurable change in the resonant frequency known as the ‘chemical shift’.

More accurately, the applied rf field drives Rabi flopping of the nucleons between the aligned and anti-aligned spin states. It is therefore possible to prepare a superposition of spin-up and spin-down states using an rf $\frac{\pi}{2}$ -pulse, resulting in a spin vector that points perpendicular to the bias field. The spin therefore undergoes Larmor precession, resulting in weak rf *emission* at the local Larmor frequency which can be detected with an induction pick-up coil. This enables NMR to be performed rapidly using a short, broad-band pulse to tip all the spins in the sample, and taking the Fourier transform of the measured signal to determine the sample composition.

This technique, called NMR spectroscopy (or FT-NMR), has massively improved sensitivity compared to CW-NMR, revolutionising chemical analysis and earning the 1991 Nobel prize in Chemistry [93, 94]. For example, NMR can discern the difference between cancerous cells and healthy cells [95, 96]. However, despite providing detailed information about an entire sample, this kind of NMR analysis contains no spatial information and cannot be directly used to form an image.

To use NMR for imaging, the necessary spatial information can be introduced by placing the sample in a linear magnetic field gradient. This alters the resonance condition at different locations in the sample. In CW-NMR, the frequency of the rf radiation can be swept, causing resonance at different locations in the sample and enabling a 1D profile to be built up. Tomographic reconstruction techniques can then be used to build up a two-dimensional image of the object from multiple one-dimensional slices through the sample. This was first achieved by Lauterbur, who reconstructed a pair of water-filled capillaries surrounded by heavy water [97] from a series of 1D projections in a gradient, in a technique he called ‘zeugmatography’.

The critical step that made MRI practical was the development of a rapid imaging technique by Mansfield [98]. He applied the magnetic field gradient idea to FT-NMR, where the spread of Larmor frequencies results in gradient-induced dephasing (discussed in §3.5) that can be reversed with spin echo [99, 100]. By changing the strength of the gradient between echoes, extra spatial information can be coupled into the signal that can be reconstructed into a 2D image. This technique, called ‘echo-planar imaging’, revolutionised MRI not just for its rapid acquisition time (reducing a typical acquisition from 10–20 min to 20–50 ms [101]), but because it provided a platform for developing new ‘spin sequences’ with gradients and rf pulses for improved contrast in different applications. This work led directly to the medical MRI in use today, and the 2003 Nobel prize in Medicine was jointly awarded for these two contributions [102].

The incredible success of MRI in the ‘living state’ begs the question as to whether the techniques can be adapted to ultracold atoms for rapid, minimally-destructive acquisition of three-dimensional profiles.

1.7: MRI in atomic vapours

In nuclear MRI, the constituent spins of the sample are coherently manipulated with rf radiation to extract information about its composition and density. The CW-NMR technique uses the weak absorption of the rf to detect resonance, whereas in FT-NMR the weak re-emission of radiation following the $\frac{\pi}{2}$ -pulse is captured by an rf antenna. Both of these methods exploit the fact that the ‘living-state’ sample has a large number of precessing spins (typically 10^{16} – 10^{20} per volume element) contributing to the signal. Spinor BECs have at most $\sim 10^6$ atoms, making direct detection of rf impossible via electromagnetic induction of a pick-up coil.

However, the purpose of measuring the re-emitted rf was simply to record the state of the spins in the system, and this interface is perfectly provided by the Faraday measurement. Faraday spectroscopy therefore takes the place of FT-NMR in any MRI sequence, allowing the coherent spin control techniques to be directly adapted. In the same way as FT-NMR, spatial information can be encoded into a Faraday measurement by placing the BEC in a magnetic field gradient and recording different Larmor precession at different rates in different locations,¹⁰ enabling magnetic resonance imaging to be performed.

The first application of MRI to a warm alkali vapour was by Skalla *et al.* [103], who obtained a two-dimensional image of the spin distribution. This was used to observe atomic diffusion by optically pumping the vapour with a spatially varying intensity profile and tracking the resulting spin dynamics with pulsed gradients [104], and extended to three dimensions by Young *et al.* [105]. However, the dominance of diffusion-induced decoherence and lack of persistent spatial features in thermal systems makes other applications limited.

A proposal to use MRI to determine the internal spin state of a BEC was first published by Toyoda *et al.* in 2002 [106]. It was demonstrated theoretically that the diffraction limit could be overcome by Faraday measurement with realistic experimental parameters, implying that sub-wavelength features such as dark solitons could be directly imaged. In 2008, Liu *et al.* [107] published Faraday rotation by a BEC in a uniform magnetic field, demonstrating minimally-destructive measurement of the relative spin populations and phase in the single-mode approximation.

These results, combined with the coherent spin winding of a spinor BEC in a magnetic field gradient observed by Higbie *et al.* [44], demonstrate the versatility of applying spin-manipulation techniques from nuclear MRI to minimally destructive imaging of cold quantum gases. In particular, the long coherence times of a BEC will also enable spin-echo techniques to be repeatedly applied to perform time-reversal of the BEC spin dynamics [108, 109] and periodically re-phase the spins for an MRI measurement.

¹⁰ In MRI nomenclature this is called ‘frequency encoding’.

This overcomes coherent dephasing of the different Larmor frequencies present in the MRI signal, and will potentially allow observation of the evolution of internal condensate dynamics over prolonged time-scales.

1.8: What's new in this work

This thesis is motivated by the following research question: can magnetic resonance imaging with the Faraday effect resolve the structure of a spinor condensate *in situ*?

I answer this question by constructing a Faraday spectrometer at the 'magic-zero' wavelength to probe a ^{87}Rb BEC with the Faraday effect for long interrogation times. I show that the measurement is highly sensitive to the local magnetic field, and cancel modulation from mains-power noise and the vector light-shift induced by the probe beam. I apply magnetic field gradients to induce coherent dephasing, which I rephase with spin echo. I then present one-dimensional magnetic resonance images of a partially evaporated BEC, and of twin BECs in a split optical dipole trap. I show that their overall spatial structure can be resolved, and discuss potential extensions to achieve resolution at the healing-length scale.

1.9: Outline of this thesis

In [Chapter 2](#), I introduce the off-resonant dipole interaction, derive the Faraday effect, and discuss its use as a probe for the BEC. I quantify undesirable effects induced by the probe beam, such as perturbations to the trapping potential, vector and tensor light-shifts, and derive conditions under which they vanish.

In [Chapter 3](#), I show how a magnetic field gradient can be used to extract the spatial profile of a spinor BEC using the Faraday effect, and derive a signal-to-noise ratio for the reconstruction. I show how techniques from MRI may be adapted to construct 2D and 3D profiles of the BEC, and discuss limits on the prospective resolution.

In [Chapter 4](#), I discuss our group's spinor BEC apparatus, which was designed and constructed over the course of this project. I describe in detail the specific parts of the system that I developed for reliable and independent operation of the apparatus.

In [Chapter 5](#), I outline practical considerations regarding the Faraday probe beam and discuss solutions to technical issues such as stability, alignment and imaging. I show that the diode laser's amplified spontaneous emission background causes resonant scattering in the BEC, which is prevented using an interference filter.

In [Chapter 6](#), I describe the differential photodetector built for high precision polarimetry of the Faraday rotation and show that it is shot-noise limited. I use Fourier filtering to process the Faraday signal to eliminate out-of-band noise, and construct a spectrogram to illustrate the amplitude and frequency modulation present in the signal.

In [Chapter 7](#), I characterise the Faraday signal and show that the long interrogation time permits precise measurement of the Larmor frequency for magnetometry. I measure magnetic field fluctuations caused by power supplies near the apparatus, which result in frequency-modulation. I observe how the signal is affected by the vector light-shifts and the quadratic Zeeman effect, and demonstrate how each can be eliminated.

In [Chapter 8](#), I explain how magnetic field gradients can be produced without specialised gradient coils, and demonstrate radiative spin echo to reverse coherent dephasing of the Faraday signal. I perform the first magnetic resonance imaging of BEC by reconstructing the 1D profiles of a bimodal partially evaporated BEC and twin BECs in a split trap.

In [Chapter 9](#), I summarise the work presented, describe the advantages and limitations of the technique, discuss how the limitations could be improved upon, and speculate on potential future applications.

The off-resonant dipole interaction

In this chapter I analyse the effective Hamiltonian for the dipole interaction between a multi-level atom and an off-resonant laser beam using an irreducible spherical decomposition of the polarisability tensor. Multiple fine-structure transitions are explicitly considered and the ‘magic-zero’ wavelengths at which the beam produces no trapping force are derived. Terms relevant for the Faraday effect are obtained and used to quantify the induced polarisation rotation. The vector and tensor light-shifts are considered, and conditions are found for which they vanish.

2.1: Review of the atomic polarisability expansion

An atom in a laser beam can be modelled as interacting with the electric field using the dipole potential [110, 111],

$$\hat{\mathcal{H}} = -\hat{\mathbf{d}} \cdot \hat{\mathbf{E}}, \quad (2.1)$$

where $\hat{\mathbf{E}}$ is the electric field operator and $\hat{\mathbf{d}} = -e\hat{\mathbf{r}}$ is the electric dipole moment operator of the atom, which for an alkali atom arises from the position operator $\hat{\mathbf{r}}$ of the valence electron. For monochromatic light with wavelength λ , angular frequency $\bar{\omega} = 2\pi c/\lambda$, and polarisation unit vector ϵ , the time-dependence of the electric field operator can be separated out, resulting in

$$\begin{aligned} \hat{\mathbf{E}} &= \hat{\mathbf{E}}^{(+)} e^{-i\bar{\omega}t} + \hat{\mathbf{E}}^{(-)} e^{+i\bar{\omega}t}, \\ \text{with } \hat{\mathbf{E}}^{(+)} &= \frac{1}{2}\epsilon\hat{E}_0 \quad \text{and} \quad \hat{\mathbf{E}}^{(-)} \equiv (\hat{\mathbf{E}}^{(+)})^\dagger = \frac{1}{2}\epsilon^*\hat{E}_0, \end{aligned} \quad (2.2)$$

where \hat{E}_0 is the electric field amplitude operator of the specific EM-field mode¹ with polarisation ϵ .

In a frame co-rotating with the electric field, the dipole operator can be written as $\hat{\mathbf{d}} = \hat{\mathbf{d}}^{(+)} e^{-i\bar{\omega}t} + \hat{\mathbf{d}}^{(-)} e^{i\bar{\omega}t}$, so the dipole interaction Hamiltonian becomes

$$\hat{\mathcal{H}} = -e^{-2i\bar{\omega}t} \hat{\mathbf{d}}^{(+)} \cdot \hat{\mathbf{E}}^{(+)} - \hat{\mathbf{d}}^{(-)} \cdot \hat{\mathbf{E}}^{(+)} - \hat{\mathbf{d}}^{(+)} \cdot \hat{\mathbf{E}}^{(-)} - e^{2i\bar{\omega}t} \hat{\mathbf{d}}^{(-)} \cdot \hat{\mathbf{E}}^{(-)}.$$

Assuming that in the Heisenberg picture the dipole operator evolves slowly in this co-rotating frame, the terms oscillating at $2\bar{\omega}$ are too rapid to contribute to dynamics. The

¹ The generalisation to polychromatic light and multiple polarisation modes is immediate and treated in standard texts.

rotating wave approximation can then be made and these terms can be time-averaged to zero.² The simplified Hamiltonian is then

$$\hat{\mathcal{H}} \approx -\hat{\mathbf{d}}^{(-)} \cdot \hat{\mathbf{E}}^{(+)} + \text{h.c.} \quad (2.3)$$

In the case of far-detuned laser light, a negligible fraction of atoms undergo direct transitions to an excited state. However, Raman transitions between ground states are still possible via an excited state. The quantum master equation for this system predicts these off-resonant transitions, and by adiabatically eliminating the excited state [113] an *effective* Hamiltonian can be obtained in the off-resonant limit:

$$\hat{\mathcal{H}} = \hat{\mathbf{E}}^{(-)} \cdot \left(\sum_{|i\rangle} \frac{\overleftrightarrow{\alpha}}{\hbar\Delta_i} \right) \cdot \hat{\mathbf{E}}^{(+)}, \quad (2.4)$$

where $\overleftrightarrow{\alpha} \equiv \hat{\mathbf{d}}\hat{\mathbf{d}}^\dagger$ is termed the *polarisability tensor*. The sum is over all excited states $|i\rangle = |J', F', m'\rangle$, with each term weighted by the detuning $\Delta_i \equiv \bar{\omega} - \omega_i$ of the light field above the energy of the transition $|J, F, m\rangle \rightarrow |J', F', m'\rangle$. The approximation is only valid when $\Delta_i \gg \Gamma_i$, where Γ_i is the natural linewidth of the transition.

Since the polarisability tensor $\overleftrightarrow{\alpha}$ is the product of two vectors (a ‘dyadic’ tensor), it can be decomposed into a sum of irreducible spherical tensor components [114, 115, 116],

$$\overleftrightarrow{\alpha} = \alpha^{(0)} + \boldsymbol{\alpha}^{(1)} + \overleftrightarrow{\alpha}^{(2)}, \quad (2.5)$$

where $\alpha^{(0)}$ is a scalar, $\boldsymbol{\alpha}^{(1)}$ a vector and $\overleftrightarrow{\alpha}^{(2)}$ a traceless symmetric tensor.

Through the Wigner-Eckart theorem [117, 118] and properties of the Clebsch-Gordan coefficients, each of the differently ranked irreducible spherical contributions $\alpha^{(i)}$ to the dyadic tensor $\overleftrightarrow{\alpha}$ can be expressed in terms of the Pauli spin operators $\hat{F}_x, \hat{F}_y, \hat{F}_z$. The derivation is a standard but algebraically involved affair [119, 120], therefore only the primary results will be summarised here.

Each of the contributions $\alpha^{(i)}$ to the irreducible spherical tensor decomposition has a corresponding Hamiltonian,

$$\hat{\mathcal{H}}^{(i)} = \hat{\mathbf{E}}^{(-)} \cdot \left(\sum_i \frac{\alpha^{(i)}}{\hbar\Delta_i} \right) \cdot \hat{\mathbf{E}}^{(+)}, \quad (2.6)$$

$$\text{such that } \hat{\mathcal{H}} = \hat{\mathcal{H}}^{(0)} + \hat{\mathcal{H}}^{(1)} + \hat{\mathcal{H}}^{(2)}. \quad (2.7)$$

The dot products in each component Hamiltonian can be evaluated by expressing the electric field vector in a compatible representation. The relevant quantised form of the electric field (2.2) is

$$\hat{\mathbf{E}}^{(+)} = \sqrt{\hbar g} \left(\hat{a}_- \boldsymbol{\epsilon}_- + \hat{a}_+ \boldsymbol{\epsilon}_+ \right), \quad \hat{\mathbf{E}}^{(-)} = \sqrt{\hbar g} \left(\hat{a}_-^\dagger \boldsymbol{\epsilon}_-^* + \hat{a}_+^\dagger \boldsymbol{\epsilon}_+^* \right),$$

² Discarding these terms causes a small perturbation known as the Bloch-Siegert shift [112].

where \hat{a}_\pm are the annihilation operators for photons with σ^\pm -polarisation, ϵ_\pm are the corresponding spherical basis vectors and $g = \bar{\omega}/2\epsilon_0 V$ with V being the quantisation volume of the mode.

The associated Stokes operators describing the polarisation state are

$$\begin{aligned}\hat{S}_0 &= \frac{1}{2}(\hat{a}_+^\dagger \hat{a}_+ + \hat{a}_-^\dagger \hat{a}_-), & \hat{S}_x &= \frac{1}{2}(\hat{a}_+^\dagger \hat{a}_- + \hat{a}_-^\dagger \hat{a}_+), \\ \hat{S}_y &= \frac{1}{2}(\hat{a}_-^\dagger \hat{a}_+ - \hat{a}_+^\dagger \hat{a}_-), & \text{and} & \quad \hat{S}_z = \frac{1}{2}(\hat{a}_+^\dagger \hat{a}_+ - \hat{a}_-^\dagger \hat{a}_-).\end{aligned}\quad (2.8)$$

They represent the total intensity of the beam (\hat{S}_0) and the polarisation projections in the linear x - y basis (\hat{S}_x), in a basis at 45° to that (\hat{S}_y), and in the circular σ^\pm basis (\hat{S}_z).

It can be shown [119] that the terms of the effective Hamiltonian couple the Stokes parameters of the light field to the Pauli spin operators as

$$\hat{\mathcal{H}}^{(0)} = g \sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\Delta_{J'F'}} \frac{2}{3} \hat{S}_0 \hat{1}_F, \quad (2.9)$$

$$\hat{\mathcal{H}}^{(1)} = g \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\Delta_{J'F'}} \hat{S}_z \hat{F}_z, \quad (2.10)$$

$$\begin{aligned}\hat{\mathcal{H}}^{(2)} &= g \sum_{J'F'} \frac{\alpha_{J'F'}^{(2)}}{\Delta_{J'F'}} \left(\hat{S}_x (\hat{F}_x^2 - \hat{F}_y^2) + \hat{S}_y (\hat{F}_x \hat{F}_y + \hat{F}_y \hat{F}_x) \right. \\ &\quad \left. + \hat{S}_0 [3\hat{F}_z^2 - F(F+1)\hat{1}_F]/3 \right),\end{aligned}\quad (2.11)$$

where the $\Delta_{J'F'}$ is the frequency detuning of the laser from the hyperfine transition $|J, F\rangle \rightarrow |J', F'\rangle$, and the sums are over all possible excited states $|J', F'\rangle$. Note that the Zeeman splitting of the hyperfine states is taken as negligible compared to the detunings involved.

The associated coupling strengths are³

$$\alpha_{J'F'}^{(0)} = \alpha_{JF}^{J'F'} \left((2F-1)\delta_{F-1}^{F'} + (2F+1)\delta_F^{F'} + (2F+3)\delta_{F+1}^{F'} \right), \quad (2.12)$$

$$\alpha_{J'F'}^{(1)} = \alpha_{JF}^{J'F'} \left(-\frac{2F-1}{F}\delta_{F-1}^{F'} - \frac{2F+1}{F(F+1)}\delta_F^{F'} + \frac{2F+3}{F+1}\delta_{F+1}^{F'} \right), \quad (2.13)$$

$$\alpha_{J'F'}^{(2)} = \alpha_{JF}^{J'F'} \left(\frac{1}{F}\delta_{F-1}^{F'} - \frac{2F+1}{F(F+1)}\delta_F^{F'} + \frac{1}{F+1}\delta_{F+1}^{F'} \right), \quad (2.14)$$

$$\alpha_{JF}^{J'F'} = \alpha_0 (2J'+1) \left| \begin{Bmatrix} 1 & J & J' \\ I_s & F' & F \end{Bmatrix} \right|^2, \quad (2.15)$$

$$\alpha_0 = |\langle J || \hat{\mathbf{d}} || J' \rangle|^2 \frac{2J+1}{2J'+1} = \frac{3\epsilon_0 \hbar \lambda_{J'}^3 \Gamma_{J'}}{8\pi^2} \quad (2.16)$$

³ The convention for the magnitude of the reduced dipole matrix element $\langle J || \hat{\mathbf{d}} || J' \rangle$ in (2.16) follows Refs. [119] and [121], although alternative conventions exist in the literature. The normalisation convention can be identified by the relationship between the reduced matrix element and $\Gamma_{J'}$, and notational differences should not be relied upon.

where $\delta_F^{F'}$ is the Kronecker delta, I_s is the nuclear isospin, $\lambda_{J'}$ is the wavelength of the fine structure transition $|J\rangle \rightarrow |J'\rangle$, $\Gamma_{J'}$ is the associated natural linewidth of the transition. Note that although the definition of α_0 appears to depend on J' , it does not as $\lambda_{D1}^3 \Gamma_{D1} = \lambda_{D2}^3 \Gamma_{D2}$ (see §B.2).⁴

Each of the component Hamiltonians $\hat{\mathcal{H}}^{(i)}$ involve a different coupling between the atomic spin state \hat{F} and the light field \hat{S} , and each result in different dynamics. The dynamics corresponding to each of these contributions are discussed below.

2.2: The far-detuned limit

The Hamiltonians (2.9)–(2.11) contain summations over all possible hyperfine transitions. In the large detuning limit, these expressions can be simplified as the excited state hyperfine splitting becomes negligible compared to the detuning.

Introducing the detuning from the fine structure line centre $\Delta_{J'}$, the summations are

$$\Delta_{J'F'} \approx \Delta_{J'} \quad \Rightarrow \quad \sum_{J'F'} \frac{\alpha_{J'F'}^{(i)}}{\hbar \Delta_{J'F'}} \approx \sum_{J'} \frac{1}{\hbar \Delta_{J'}} \sum_{F'} \alpha_{J'F'}^{(i)}. \quad (2.17)$$

These sums over F' contain no information about the light field, and can be evaluated for a given J' using the Kronecker deltas in (2.12)–(2.14). The result is a linear combination of Wigner 6- j symbols depending on the initial ground state through J and F , and the excited state through J' ,

$$\sum_{F'} \frac{\alpha_{J'F'}^{(0)}}{\alpha_0} = (2J' + 1) \left((2F - 1)|W_{F-1}|^2 + (2F + 1)|W_F|^2 + (2F + 3)|W_{F+1}|^2 \right), \quad (2.18)$$

$$\sum_{F'} \frac{\alpha_{J'F'}^{(1)}}{\alpha_0} = (2J' + 1) \left(-\frac{2F - 1}{F}|W_{F-1}|^2 - \frac{2F + 1}{F(F + 1)}|W_F|^2 + \frac{2F + 3}{F + 1}|W_{F+1}|^2 \right), \quad (2.19)$$

$$\sum_{F'} \frac{\alpha_{J'F'}^{(2)}}{\alpha_0} = (2J' + 1) \left(\frac{1}{F}|W_{F-1}|^2 - \frac{2F + 1}{F(F + 1)}|W_F|^2 + \frac{1}{F + 1}|W_{F+1}|^2 \right), \quad (2.20)$$

$$\text{where } W_{F'} = \begin{Bmatrix} 1 & J & J' \\ I_s & F' & F \end{Bmatrix}. \quad (2.21)$$

Considering ^{87}Rb in the $F = 1$ ground state, we have $J = 1/2$ and $I = 3/2$. There are two electric dipole transitions of interest, called the D1 and D2 lines (Table 2.1), for which the interaction constants can be evaluated (Table 2.2).

⁴ If transitions other than the D1 and D2 line were being considered, (2.15) could be generalised to include a Wigner-6 j to also factor out L' dependence in (B.22).

Line	J'	$\lambda_{J'}$ (nm)	$\Gamma_{J'}$ (MHz)	I_{sat} (mW/cm ²)
D1	1/2	794.979	$2\pi \times 5.75$	1.495
D2	3/2	780.241	$2\pi \times 6.07$	1.669

Table 2.1: Properties of the D1 and D2 transitions in ⁸⁷Rb, tabulated from Ref. [121].

	Line	Scalar ($i = 0$)	Vector ($i = 1$)	Tensor ($i = 2$)
$\sum_{F'} \frac{\alpha_{J'F'}^{(i)}}{\alpha_0}$	D1	1	1/3	0
	D2	2	-1/3	0

Table 2.2: Normalised polarisability coefficient sums in the far-detuned limit for ⁸⁷Rb in the $F = 1$ ground state, computed using (2.18)–(2.20).

Hence the interaction picture Hamiltonian contributions in the far-detuned limit for ⁸⁷Rb in $F = 1$ become

$$\hat{\mathcal{H}}^{(0)} = \frac{2g}{3\hbar} \left(\frac{\alpha_0}{\Delta_{D1}} + \frac{2\alpha_0}{\Delta_{D2}} \right) \hat{S}_0 \hat{I}_0, \quad (2.22)$$

$$\hat{\mathcal{H}}^{(1)} = \frac{g}{3\hbar} \left(\frac{\alpha_0}{\Delta_{D1}} - \frac{\alpha_0}{\Delta_{D2}} \right) \hat{S}_z \hat{F}_z, \quad (2.23)$$

$$\hat{\mathcal{H}}^{(2)} = 0. \quad (2.24)$$

In this approximation, the role of the nuclear spin is negligible and the hyperfine coupling vanishes. The polarisability of the atom therefore depends only on the spin of the valence electron, behaving as a spin- $\frac{1}{2}$ system. The symmetries of such a system mean the two contributions to the vector interaction have opposite sign, and the tensorial contribution vanishes.

Reintroducing the excited state splitting, the detunings are $\Delta_{J'F'} = \Delta_{J'} + \delta_{J'F'}$, where $\delta_{J'F'}$ is the splitting of $|J'F'\rangle$ relative to the line centre (Figure 2.1). Following Ref. [120], a power series expansion of (2.17) in $\delta_{J'F'}/\Delta_{J'}$ gives

$$\begin{aligned} \sum_{J'F'} \frac{\alpha_{J'F'}^{(i)}}{\hbar\Delta_{J'F'}} &= \sum_{J'} \frac{1}{\hbar\Delta_{J'}} \sum_{F'} \frac{\alpha_{J'F'}^{(i)}}{1 + \delta_{J'F'}/\Delta_{J'}} = \sum_{J',F'} \frac{\alpha_{J'F'}^{(i)}}{\hbar\Delta_{J'}} \sum_k \left(\frac{-\delta_{J'F'}}{\Delta_{J'}} \right)^k \\ &= \alpha_0 \sum_{J'} \left(\frac{1}{\hbar\Delta_{J'}} \sum_{F'} \frac{\alpha_{J'F'}^{(i)}}{\alpha_0} - \frac{\Gamma_{J'}}{\hbar\Delta_{J'}^2} \sum_{F'} \frac{\delta_{J'F'}}{\Gamma_{J'}} \frac{\alpha_{J'F'}^{(i)}}{\alpha_0} + \dots \right). \end{aligned} \quad (2.25)$$

Each sequential term is weighted by $\delta_{J'F'}/\Delta_{J'} \ll 1$, allowing the power series to be truncated at the second term. The sums over F' can then be computed numerically for the transitions of interest (Table 2.3), allowing the Hamiltonians to be computed in the far-detuned limit to a higher accuracy. The corrections to the scalar and vector Hamiltonians are small compared to the linear contributions, but the nematic interaction $\hat{\mathcal{H}}^{(2)}$ is now nonzero and potentially contributes to dynamics.

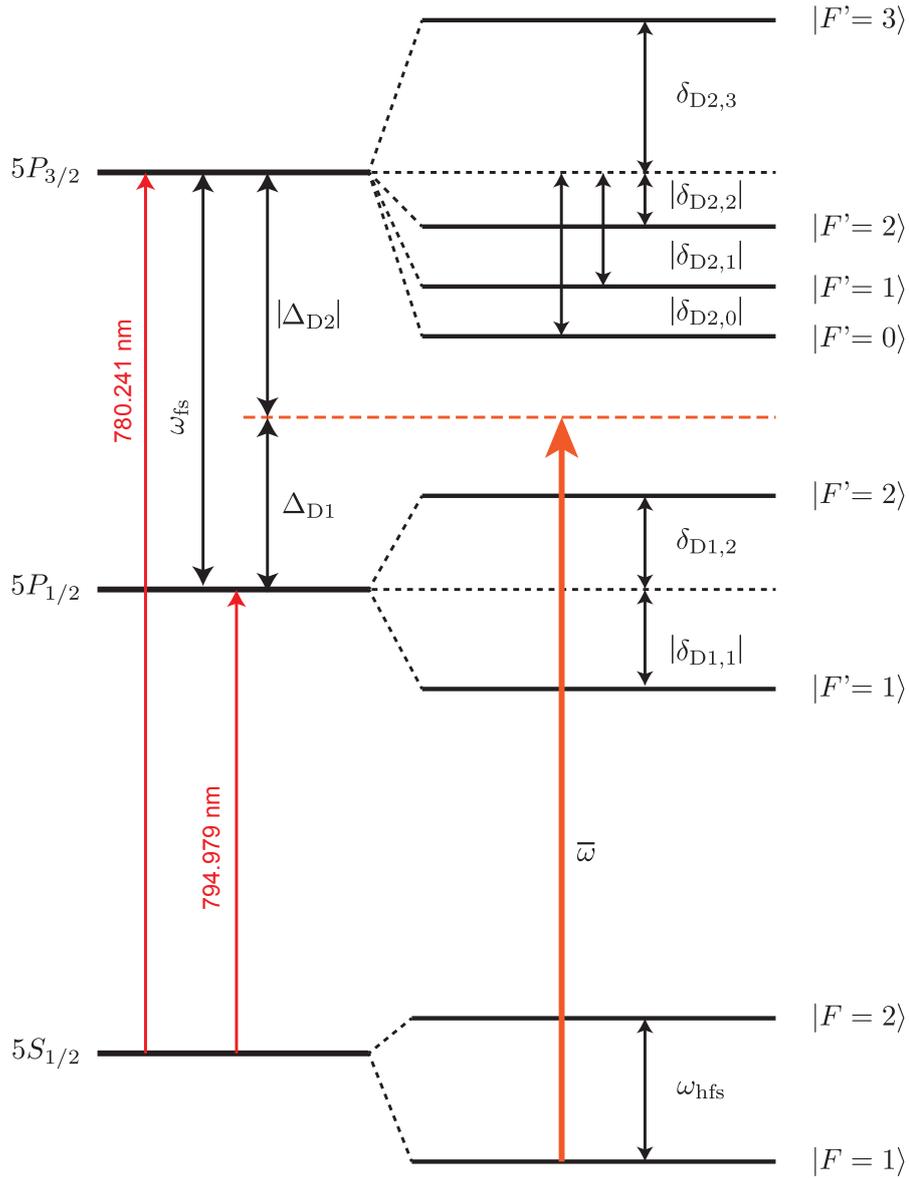


Figure 2.1: Hyperfine energy-level scheme for ^{87}Rb , defining symbols used in this thesis. The detuning of each hyperfine state is $\Delta_{J'F'} \equiv \bar{\omega} - \omega_{J'F'} = \Delta_{J'} + \delta_{J'F'} \approx \Delta_{J'}$ where $\bar{\omega}$ is the laser frequency and $\omega_{J'F'}$ is the resonant frequency, so $\Delta_{D1} > 0$ and $\Delta_{D2} < 0$.

	Line	Scalar ($i = 0$)	Vector ($i = 1$)	Tensor ($i = 2$)
$\sum_{F'} \frac{\delta_{J'F'}}{\Gamma_{J'}} \frac{\alpha_{J'F'}^{(i)}}{\alpha_0}$	D1	29.6	29.6	11.8
	D2	-58.2	27.4	-1.81

Table 2.3: Normalised coefficients of the quadratic term in the expansion of the polarisability in the far-detuned limit (2.25) for ^{87}Rb in the $F = 1$ ground state, computed using (2.12)–(2.14) and the known hyperfine structure of the excited states.

2.3: The off-resonant trapping force

The scalar contribution $\alpha^{(0)}$ to the Hamiltonian depends only on the total intensity of the beam and is unaffected by the spin state of the atoms. This contribution therefore results in a state-independent energy-level shift that is proportional to the intensity of the laser beam, termed the *dipole potential*. Considering a cylindrically symmetric beam with intensity $I(r)$, the potential is

$$V_{\text{dip}}(r) = \mathbf{E}^{(-)} \cdot \sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\hbar\Delta_{J'F'}} \cdot \mathbf{E}^{(+)} = \frac{2I(r)}{\varepsilon_0 c} \left(\sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\hbar\Delta_{J'F'}} \right). \quad (2.26)$$

Because the laser beam has a spatially varying profile, it imparts a net mechanical force on the atoms, $\mathbf{F}_{\text{dip}} = -\nabla V_{\text{dip}}$, that draws them towards the region of highest intensity for red-detuned ($\Delta < 0$) lasers. This creates an effective harmonic potential that causes transverse confinement.

The potential is approximated as harmonic about the minimum, with the trap strength quantified by the radial trapping frequency ω_r defined by $V_{\text{dip}}(r) = \frac{1}{2}m\omega_r^2 r^2$. Taking the second derivative of (2.26), the trapping frequency is given by

$$\omega_r^2 = \frac{1}{m} \left. \frac{\partial^2 V_{\text{dip}}(r)}{\partial r^2} \right|_{r=0} = \frac{2}{m\varepsilon_0 c} \left(\sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\hbar\Delta_{J'F'}} \right) \left. \frac{\partial^2 I(r)}{\partial r^2} \right|_{r=0}. \quad (2.27)$$

For a Gaussian laser beam with power P_0 and $1/e^2$ beam radius σ_r , the intensity profile is

$$I(r) = \frac{2P_0}{\pi\sigma_r^2} \exp\left(-\frac{2r^2}{\sigma_r^2}\right), \quad (2.28)$$

which results in a dipole potential with trapping frequency

$$\omega_r^2 = \frac{-16P_0}{m\pi\sigma_r^4\varepsilon_0 c} \left(\sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\hbar\Delta} \right). \quad (2.29)$$

Note that for $V_{\text{dip}} > 0$ this gives $\omega_r^2 < 0$ and the dipole force is repulsive, resulting in an anti-trapped state.

2.4: The ‘magic’ wavelength

The dipole potential (2.29) contains contributions that scale as $1/\Delta$. For a given detuning between two resonances, one line will appear to be blue detuned and one will be red detuned. By the intermediate-value theorem, there is a wavelength between any two adjacent transitions at which these contributions cancel exactly, resulting in no net scalar light shift. These are termed ‘magic’ or ‘tune-out’ wavelengths as at that particular wavelength a probe beam exhibits no trapping force on the atoms [122, 123].

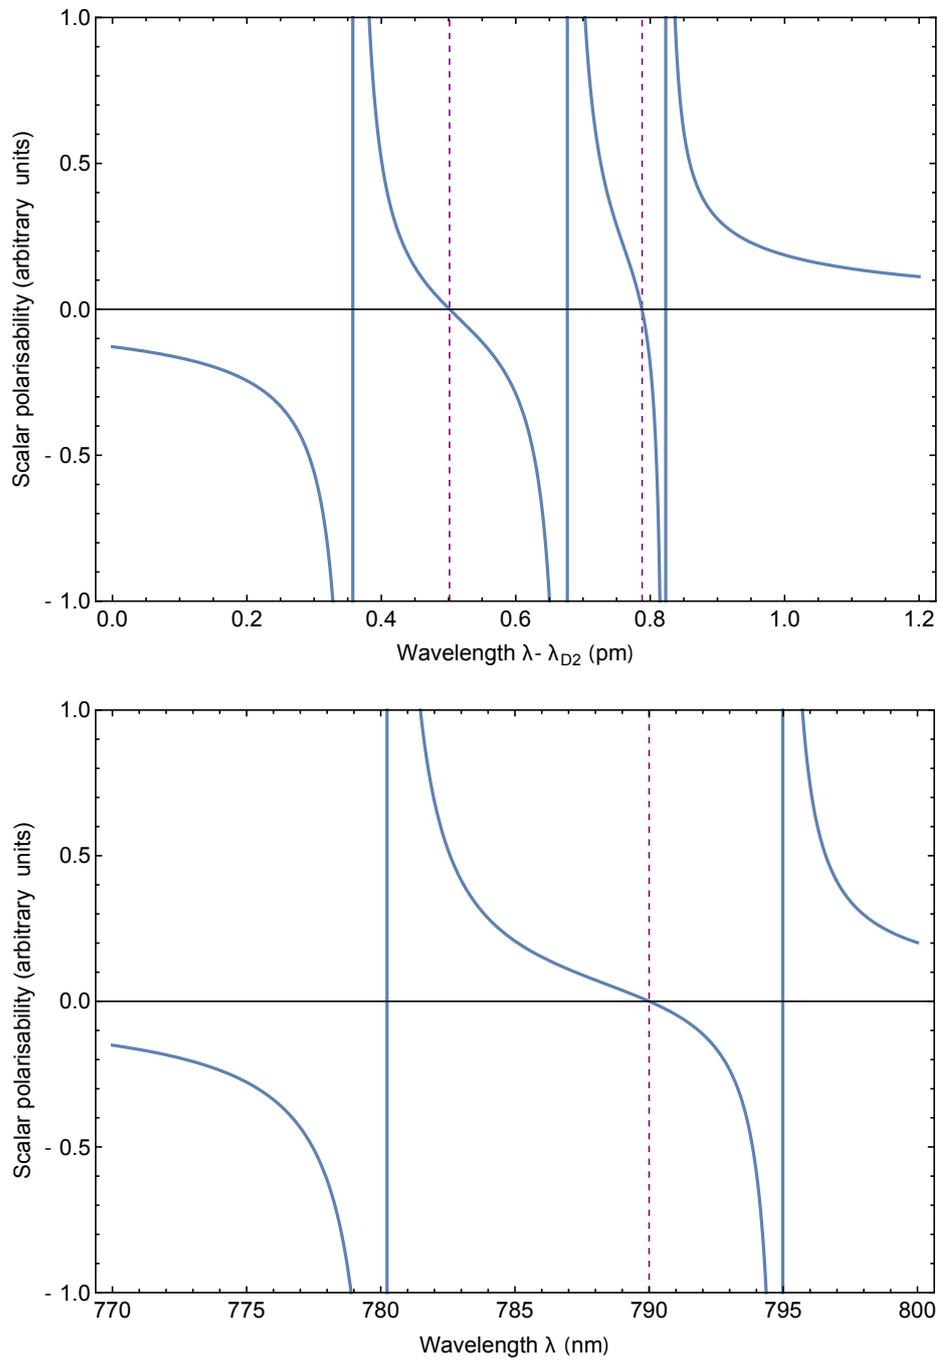


Figure 2.2: Wavelength dependence of scalar polarisability for ^{87}Rb showing the zeroes (dashed lines) corresponding to magic-zero wavelengths between accessible hyperfine excited states on the D2 line (top), and between the D1 and D2 lines (bottom).

As the BEC will be exposed to the Faraday beam while imaging is occurring, operating at a magic wavelength will prevent the probe beam from applying an additional trapping force and stops the dipole potential from being perturbed. There are a number of potential magic wavelengths to choose from (Figure 2.2), corresponding to either a small detuning between the hyperfine levels, or large detuning between the fine structure lines.

Illuminating the BEC with the Faraday beam induces scattering, causing atoms to be lost and destroying the BEC. The scattering rate can be kept constant by trading off between detuning and laser beam intensity, so the choice of magic wavelengths reduces to the choice between a weak probe beam ($\sim 100 \mu\text{W}$) with small detuning, or a bright probe beam ($\sim 10 \text{mW}$) with large detuning.

However, performing a shot-noise limited measurement with a weak probe is a significant technical challenge, as the shot-noise level approaches the dark noise level, where technical electronic noise begins to dominate the measurement. The bright-beam limit makes the magic wavelength between D1 and D2 fine structure transitions an excellent candidate for making shot-noise limited measurements.

Since this is far detuned from both resonances, the value of the magic wavelength can be estimated by finding the zero of (2.22),

$$\Delta_{D2} = -2\Delta_{D1} \quad \Rightarrow \quad \lambda_{\text{magic}} = \frac{3\lambda_{D1}\lambda_{D2}}{\lambda_{D1} + 2\lambda_{D2}}, \quad (2.30)$$

since $\Delta_{D2} = \omega_{D1} + \Delta_{D1} - \omega_{D2}$.

For ^{87}Rb this gives $\lambda_{\text{magic}} = 790.00 \text{ nm}$, compared to the value obtained by numerically finding the root of (2.9), which is $\lambda_{\text{magic}} = 790.01 \text{ nm}$. This shall be taken as the probe wavelength for the remainder of the present work.⁵

2.5: The Faraday effect

In the interaction picture polarisability expansion (§2.1) of the atom-light interaction, the vector term is

$$\hat{\mathcal{H}}^{(1)} = g \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\Delta_{J'F'}} \hat{S}_z \hat{F}_z, \quad (2.31)$$

which directly couples the spin of the atoms (\hat{F}_z) to the spin of the light passing through (\hat{S}_z). This coupling is sensitive to the spin projection of the atoms and the polarisation state of the light. It introduces a state-dependent rotation of the polarisation vector, as will be shown below.

⁵ Following the submission of this thesis, a precise measurement of the magic-zero wavelength of ^{87}Rb in $F = 2$ was reported, finding $\lambda_{\text{magic}} = 790.03235(3) \text{ nm}$ [124].

Light transits through the 30 μm cloud in 100 fs, which is a much faster timescale than the evolution of the atomic spin vector. The state of the atoms will effectively remain constant during the transit, allowing the quasi-static approximation to be made. The atomic spin state can therefore be replaced with its expectation value, which for a spin-polarised media is

$$\hat{F}_z \rightarrow \langle \hat{F}_z \rangle = F \cos \theta, \quad (2.32)$$

where θ is the angle the spin vector makes with the z -axis.

The spin projection $\langle \hat{F}_z \rangle$ may be time-dependent, as is the case with Larmor precession. Provided this time-dependence occurs on a significantly slower timescale than the transit time of the light across the cloud, the quasi-static regime still applies. For the Larmor frequencies of interest (ranging from order 100 kHz to a few MHz), this is a good approximation.

However, if atoms are undergoing Larmor precession in a plane inclined at an angle ϕ to the z -axis, the spin projection is reduced (Figure 2.3A). This occurs if the net magnetic field vector has a component $B_z \neq 0$ which inclines the precession plane by $\phi = \sin^{-1}(B_z/|B|)$. Taking the Larmor frequency as ω_L , the expectation value is

$$\langle \hat{F}_z \rangle = F \cos(\phi) \cos(\omega_L t). \quad (2.33)$$

As choice of the magnetic field bias direction is arbitrary, where possible it should be chosen to be perpendicular to the z -axis ($\phi = 0$). The Faraday effect will be similarly reduced if the local magnetisation of the BEC is non-zero (Figure 2.3B). This could occur if Larmor precession is initiated in a polarised BEC by an imperfect $\frac{\pi}{2}$ -pulse that does not tip the spin completely into the x - z plane.

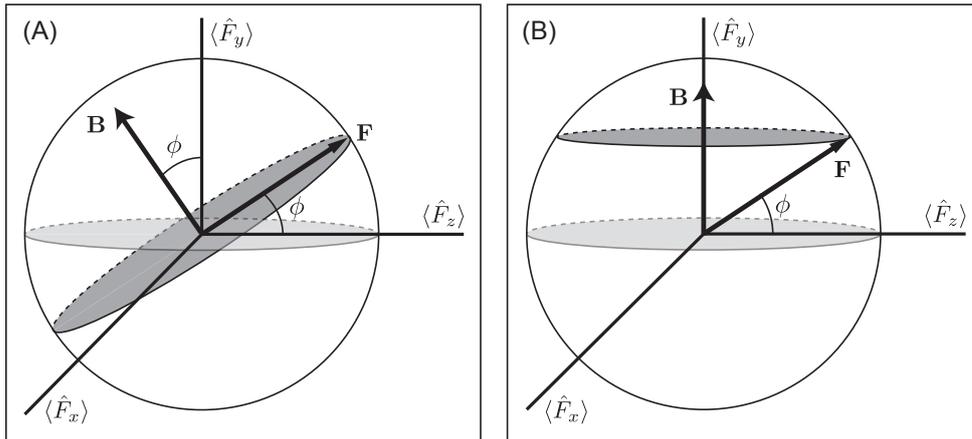


Figure 2.3: Larmor precession of the spin vector F about the magnetic field B in the cases of an off-axis magnetic field (A) and a non-zero magnetisation (B). The spin F traces out the dark shaded circle at the Larmor frequency, resulting in an oscillating projection $\langle \hat{F}_z \rangle$. For $\phi \neq 0$, the amplitude of the projection is reduced in each case, and the associated Faraday effect is weaker.

The quasi-static Hamiltonian and corresponding unitary time evolution operator are:

$$\hat{\mathcal{H}}^{(1)} = \left(g \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\Delta_{J'F'}} F \cos \theta \right) \hat{S}_z \propto \hat{S}_z, \quad (2.34)$$

$$\Rightarrow U(\delta t) = \exp \left(-i \hat{\mathcal{H}}^{(1)} \delta t / \hbar \right) = \exp \left(-i \varphi_z \hat{S}_z \right), \quad (2.35)$$

which is the form of the rotation operator for \hat{S} about the z -axis by an amount φ_z , where δt is the transit time of the light through the cloud.

Hence as light propagates through the sample, it experiences a rotation of its Stokes vector about the z -direction, by an amount proportional to the atomic density and net spin projection. This rotation preserves ellipticity, so for a probe beam that is initially linearly polarised, this is equivalent to a rotation of its polarisation axis. Taking the beam as being initially polarised along the x axis, its polarisation state before and after passing through the atom cloud is

$$\begin{aligned} \langle S_x(0) \rangle &= \langle S_0(0) \rangle, & \langle S_y(0) \rangle &= 0, \\ \langle S_x(\delta t) \rangle &= \langle S_x(0) \rangle \cos \varphi_z, & \text{and} & \quad \langle S_y(\delta t) \rangle = \langle S_x(0) \rangle \sin \varphi_z, \end{aligned} \quad (2.36)$$

Measuring the polarisation projection $\langle S_y(\delta t) \rangle$ of the beam at the output therefore provides an interface to measure the spin state of the cloud, which forms the basis of the Faraday imaging concept.

For a uniform density cloud containing N_a atoms and having cross-sectional area A , the interaction volume is $V = A c \delta t$, and the rotation can be written as

$$\varphi_z = \varphi_0 F \cos \theta \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\alpha_0 \Delta_{J'F'}}, \quad (2.37)$$

$$\text{where } \varphi_0 = \frac{N_a g \delta t}{\hbar} \alpha_0 = \frac{\pi}{\epsilon_0 \hbar \lambda} \left(\frac{N_a}{A} \right) \alpha_0, \quad (2.38)$$

The interaction strength φ_0 can then be expressed in terms of the natural linewidth and on-resonant optical-depth (OD) of the $|J\rangle \rightarrow |J'\rangle$ transition as

$$\varphi_0 = \left(\frac{\lambda_{J'} \Gamma_{J'}}{\lambda \frac{\Gamma_{J'}}{4}} \right) \text{OD} \quad \text{where} \quad \text{OD} = \frac{N_a}{A} \sigma_{0,J'} \quad \text{and} \quad \sigma_{0,J'} = \frac{3\lambda_{J'}^2}{2\pi}. \quad (2.39)$$

It should be noted that since $\lambda_{D1}^3 \Gamma_{D1} = \lambda_{D2}^3 \Gamma_{D2}$ (see [§B.2](#)), φ_0 is independent of J' even though the resonant cross section $\sigma_{0,J'}$ is not.⁶ This is a convenient representation as the optical-depth is directly measured by resonant absorption imaging ([§1.2](#)).

This result can be immediately generalised to a cloud with spatially-varying number density $\rho(x, y, z)$ as

$$\text{OD} = \tilde{\rho} \sigma_{0,J'} \quad (2.40)$$

⁶ Typically the result is simplified further by taking $\lambda \approx \lambda_{J'}$ (e.g. [\[119\]](#)).

where the *column density* $\tilde{\rho}$ is

$$\tilde{\rho}(x, y) = \int_{-\infty}^{\infty} \rho(x, y, z) dz, \quad (2.41)$$

having units of atoms per unit area.

The Faraday rotation can therefore be written as

$$\varphi_z = \xi_F \tilde{\rho} F \cos \theta, \quad (2.42)$$

where the strength of the atom-light Faraday coupling is

$$\xi_F = \left(\frac{\pi}{\epsilon_0 \hbar \lambda} \right) \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\Delta_{J'F'}}, \quad (2.43)$$

and is solely determined by the laser detuning and properties of the atomic transition. In the far-detuned limit (which applies at the magic wavelength) this becomes

$$\xi_F \approx \left(\frac{\pi \alpha_0}{\epsilon_0 \hbar c} \right) \sum_{J'} \frac{\omega}{\Delta_{J'}} \sum_{F'} \frac{\alpha_{J'F'}^{(1)}}{\alpha_0} = \frac{\pi \alpha_0}{3 \epsilon_0 \hbar c} \left(\frac{\omega}{\Delta_{D1}} - \frac{\omega}{\Delta_{D2}} \right). \quad (2.44)$$

Hence the Faraday effect is strongest near the D1 and D2 line centres (where one of Δ_{D1} or Δ_{D2} is small), constructively interferes for detunings between the lines (when the detunings have opposite sign), and destructively interferes for detunings outside the two lines (Figure 2.4).

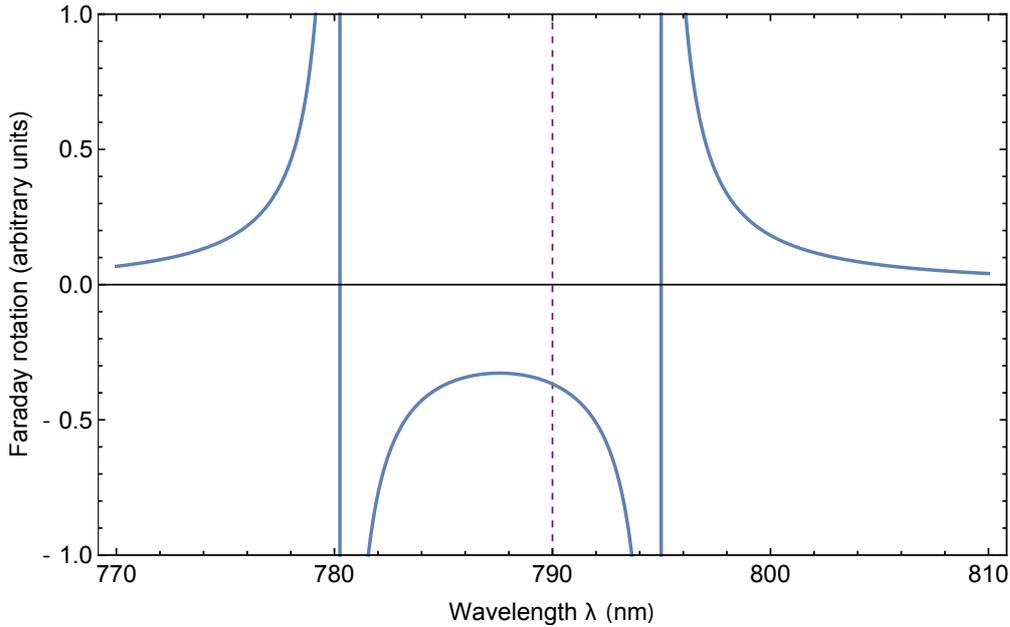


Figure 2.4: Wavelength dependence of the Faraday effect. The dashed line denotes the magic wavelength.

In particular, this can be seen by introducing the detuning from the midpoint of the lines,

$$\Delta \equiv \bar{\omega} - \frac{1}{2}(\omega_{D1} + \omega_{D2}) = \Delta_{D1} - \frac{1}{2}\omega_{fs} = \frac{1}{2}(\Delta_{D1} + \Delta_{D2}), \quad (2.45)$$

the coupling strength is then

$$\xi_F \approx \left(\frac{4\pi\alpha_0}{3\epsilon_0\hbar\lambda} \right) \frac{\omega_{fs}\Gamma}{4\Delta^2 - \omega_{fs}^2} \propto \begin{cases} \omega_{fs}\Gamma/\Delta^2 & \text{for } \Delta \gg \omega_{fs}, \\ 1 + (2\Delta/\omega_{fs})^2 & \text{for } \Delta \ll \omega_{fs}. \end{cases} \quad (2.46)$$

For very large detunings, it is therefore not possible to consider only the nearer line, as the destructive interference of the two lines makes the net Faraday rotation fall off as Δ^{-2} , which is much faster than the Δ^{-1} dependence in (2.43). Furthermore, although the Faraday interaction is weak when Δ is small, the interaction strength always remains non-zero due to the constructive interference.

2.6: The vector light-shift and effective magnetic field

An alternate interpretation of the vector interaction Hamiltonian is to consider a single atom and approximate the polarisation rotation due to that atom as negligible. The state of the light field can then be approximated as

$$\hat{S}_z \rightarrow \langle \hat{S}_z \rangle = \frac{1}{2} \left[\langle \hat{a}_+^\dagger \hat{a}_+ \rangle - \langle \hat{a}_-^\dagger \hat{a}_- \rangle \right], \quad (2.47)$$

which is a measure of the ellipticity of the probe beam. This leads to the ‘vector light-shift’ (VLS) Hamiltonian,

$$\hat{\mathcal{H}}^{(1)} = \left(g \sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\Delta_{J'F'}} \langle \hat{S}_z \rangle \right) \hat{F}_z. \quad (2.48)$$

This Hamiltonian couples directly to the magnetic quantum number m_F through \hat{F}_z , and looks like the usual linear Zeeman effect Hamiltonian,

$$\hat{\mathcal{H}}_z = \left(\frac{\mu_B g_F}{\hbar} B_{\text{vls}} \right) \hat{F}_z, \quad (2.49)$$

corresponding to some artificial magnetic field B_{vls} along the propagation direction z , whose strength depends on $\langle \hat{S}_z \rangle$.

Therefore the vector component Hamiltonian $\hat{\mathcal{H}}^{(1)}$ acts like an effective magnetic field for the atoms. The resulting energy level splitting experienced by the atoms is referred to as the ‘vector light-shift’ [125, 113].

Provided the incident probe light is perfectly linearly polarised, then $\langle \hat{S}_z \rangle = 0$ and there is no effective magnetic field. However, for a purely circular polarisation the effect is maximised, making the evolution of the atomic spin vector sensitive to the ellipticity of the probe beam.

This interpretation is dual to the Faraday effect, and occurs simultaneously. In fact, the form of the operator $\hat{S}_z \hat{F}_z$ acts to couple the atoms and light together such that they are no longer separable. However, as the light passes through the atom cloud quickly relative to the timescale of the evolution of the atomic spin state, the quasi-static approximation still applies and it is possible to treat them separately.

In the Faraday effect, \hat{F}_z causes rotation of the probe beam polarisation $\langle S_y \rangle$, whereas in the vector-light shift case the effective magnetic field induced by \hat{S}_z causes Larmor precession about the probe beam propagation direction. This result implies that it is extremely important to have precise control of the probe beam polarisation, as any ellipticity in the probe will contribute to undesirable non-trivial dynamics in the evolution of the spin state of the cloud being observed.

Moreover, since $\langle S_z \rangle$ depends on the local intensity of the laser, the spatially-varying intensity profile of the probe beam results in an effective magnetic field *gradient* across the BEC. This perturbs the condensate in a spin-dependent way, which prevents the measurement being minimally-destructive.

2.7: The tensor light-shift

The final Hamiltonian component is the nematic interaction (2.11), which contains a non-linear dependence on the atomic spin operators \hat{F}_i . The nematic interaction causes the atoms to induce a small rotation of the probe's polarisation state that does not preserve ellipticity [119]. This rotation vanishes to first-order for large detunings, and remains weak to second-order for atoms in the $F = 1$ ground state of interest (§2.2). The Faraday effect therefore dominates the evolution of the probe beam polarisation, allowing the nematic evolution of the probe beam to be neglected.

Considering the effect of light on atoms, a linearly polarised probe beam causes the vector light-shift to vanish exactly, so the nematic interaction becomes the dominant term of the interaction Hamiltonian. This interaction results in complicated evolution of the atomic spin state, resulting in coherent collapse and revival of the Faraday signal [126].

Taking the magnetic field direction along y , the linear polarisation state of the probe beam can be expressed as

$$\epsilon = \sin \theta e_x + \cos \theta e_y, \quad (2.50)$$

where e_x and e_y are unit vectors, and θ is the polarisation angle measured clockwise from the y -axis. The resulting dynamics depend on the orientation of the probe beam's linear polarisation through expectation values of the Stokes operators,

$$\langle \hat{S}_x \rangle = \frac{1}{2} (\cos^2 \theta - \sin^2 \theta) \langle \hat{S}_0 \rangle \quad \text{and} \quad \langle \hat{S}_y \rangle = -\sin \theta \cos \theta \langle \hat{S}_0 \rangle. \quad (2.51)$$

Transforming into a frame that is co-rotating with the Larmor precession, and taking the bias field as strong ($B_y \gg \hbar\Gamma/g_F\mu_B$), the rotating wave approximation can be made by discarding the counter-rotating term. The nematic Hamiltonian (2.11) becomes [126]

$$\hat{\mathcal{H}}_{\text{RWA}}^{(2)} \propto \left(-\frac{1}{2} \sin^2 \theta + \cos^2 \theta\right) \hat{F}_y^2. \quad (2.52)$$

Therefore, different dynamics occur whether the polarisation is aligned with ($\theta = 0$) or orthogonal to ($\theta = 90^\circ$) the bias magnetic field. Furthermore, polarising the incident probe at the ‘magic’ angle of $\theta = \arctan(\sqrt{2}) = 54.7^\circ$ to the magnetic field cancels out this non-linear contribution in the rotating wave approximation, eliminating the collapse and revival of the signal, and extending the lifetime of the measurements [126].

In principle, the nematic interaction can therefore also be eliminated by setting the linear polarisation axis of the probe beam to coincide with the magic angle. This does not affect the Faraday measurement, as the polarisation remains linear and only the polarisation rotation is measured. However, in practice this is not necessary as the interaction vanishes to first-order in the large detuning limit and the second-order contribution scales as $\Gamma_{J'}/\Delta \sim 10^{-6}$ at $\lambda = 790$ nm.

It should be noted that the observations of nematic dynamics presented in Ref. [126] studied ^{133}Cs in the $F = 4$ ground state, which experiences a significantly greater tensor light-shift. Not only is the strength of the nematic interaction enhanced by a factor F^2 , but the relatively small detuning⁷ $\Delta/2\pi \sim 50$ GHz gives $\Gamma_{J'}/\Delta \sim 10^{-4}$. The strength of the nematic interaction for our system is four orders of magnitude weaker (Table 2.4), so it will be neglected in the remainder of this thesis.

Atom	F	Detuning	$F^2 \sum \frac{\alpha_{J'F'}^{(2)}}{\alpha_0} \frac{\Gamma_{J'}}{\Delta_{J'F'}}$
^{87}Rb	1	$\lambda = 790.0$ nm	-6.6×10^{-11}
^{133}Cs	4	$\Delta/2\pi = 50$ GHz	-1.6×10^{-7}

Table 2.4: Comparison of the normalised nematic interaction strength for the ^{87}Rb at the magic wavelength and the D2 line of ^{133}Cs . Caesium transition data taken from [127].

2.8: Summary

This chapter considered the atom-light interaction for a multi-level atom in an off-resonant laser beam. The Hamiltonians corresponding to the different ranks of the irreducible spherical tensor decomposition of the polarisability tensor were studied and their effect on the probe beam’s polarisation and the atomic spin state considered.

⁷ The apparatus is described in [84] and uses detunings between 10 GHz and 100 GHz.

The rank-0 contribution was shown to result in a trapping potential that depends on the intensity gradient of the probe beam. The trapping potential was seen to vanish at specific ‘magic’ wavelengths where the contributions from neighbouring transitions cancel each other. A magic wavelength was predicted at $\lambda = 790$ nm, which presents a promising candidate for performing Faraday measurement, because the large detuning permits bright-beam polarimetry measurements without imparting a dipole force.

The rank-1 contribution produced the Faraday effect, resulting in density-dependent polarisation rotation of a linearly polarised probe beam. An elliptically polarised probe beam was seen to generate a vector light-shift, which acts as an effective magnetic field for the atoms and causes evolution of the atomic spin vector. This demonstrated the importance of the purity of the probe’s linear polarisation.

The rank-2 contribution vanished to first-order in the far-detuned limit, comprising a weak contribution to the dynamics of the system. The evolution of the probe’s polarisation is dominated by the Faraday effect, making this nematic contribution negligible to the probe. Conversely, as a linear polarisation yields no vector light-shift, the nematic interaction becomes the dominant effect in the evolution of the atomic spin state, resulting in decay and revival of the Faraday signal. The effect is small in the regime of interest however, and vanishes altogether for the specific polarisation angle of 54.7° .

This characterisation of the interaction between atomic spin and probe polarisation forms the basis of the Faraday measurement technique. I have demonstrated how the different contributions can be controlled, and how the Faraday effect can be used to interrogate an atomic system without significantly perturbing its spin state. I will now consider how to apply this minimally destructive technique to obtain an image of a condensate.

Faraday magnetic resonance imaging

In this chapter I develop the theory of the Faraday effect in a magnetic field gradient into an imaging method for atomic clouds. I derive an expression for the expected signal-to-noise ratio, and discuss the role of experimental parameters. I explain coherent dephasing of the signal, and propose methods for rephasing it. I quantify the blurring induced by Stern-Gerlach separation of the spin components of the condensate, and quantify the relationship between spatial resolution and the applied magnetic field gradient.

3.1: Measuring Faraday rotation

The Faraday effect provides a coupling between atom cloud density and the strength of the polarisation rotation experienced by the probe light, which can be measured with a balanced polarimeter (Figure 3.1). The condensate is taken to be quasi-static with density $\rho(r)$ that is effectively constant in time, and is measured by an optical system with transmission efficiency κ , photodiodes with quantum efficiency η , and an amplifier with transimpedance gain G . An aperture of radius a blocks the light that did not pass through the cloud, which limits the detected shot-noise to only that of light contributing to the signal from the cloud.

Let Φ be the total polarisation rotation experienced by the probe beam. A Wollaston prism is used as a polarising beam-splitter (PBS) to separate the polarisation components of the probe. A $\lambda/2$ waveplate before the prism is used to balance the light

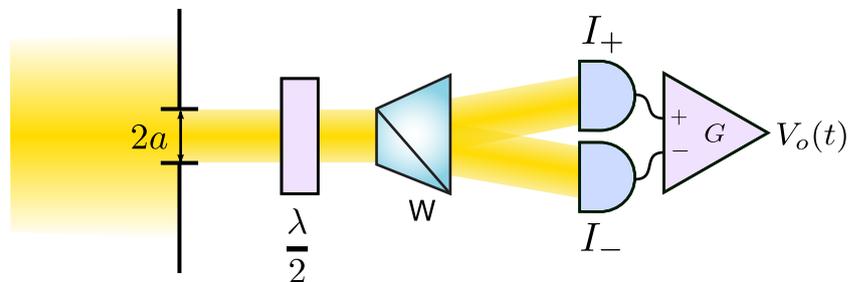


Figure 3.1: The balanced polarimeter consists of an aperture of radius a to block light not travelling through the condensate, followed by a $\lambda/2$ waveplate and Wollaston prism (W) to split the beam into its two polarisation components I_+ and I_- .

received by each detector when $\Phi = 0$. In this configuration, the two beams after the prism have intensities

$$I_{\pm} = I_0 \cos^2(\Phi \pm \frac{\pi}{4}). \quad (3.1)$$

Hence, for small rotations, the difference between the two beam intensities is

$$\Delta I = I_0 \cos^2(\Phi + \frac{\pi}{4}) - I_0 \cos^2(\Phi - \frac{\pi}{4}) = I_0 \sin(2\Phi) \simeq 2I_0\Phi. \quad (3.2)$$

Assuming that the photodetectors are neither saturated nor overfilled, the measured power difference is

$$\Delta P = \iint \Delta I \, dx \, dy \simeq 2I_0 \iint \Phi \, dx \, dy. \quad (3.3)$$

The expected degree of Faraday rotation (2.42) is small, with the rotation from an entire cloud of $N = 2 \times 10^5$ atoms with radius $r = 15 \mu\text{m}$ being

$$\Phi = \int \varphi_z dz = \frac{N}{\pi r^2} \frac{\xi_F}{4} = 25 \mu\text{rad}. \quad (3.4)$$

For a 5 mW beam this corresponds to a power difference of $\Delta P = 250 \text{ nW}$ at the detector.

Although it is possible to measure this power difference directly, it will be highly susceptible to noise and easily be overwhelmed by polarisation fluctuations in the probe beam. An oscillating Faraday signal would enable selective filtration and amplification, reducing sensitivity to drift. Choosing the oscillation frequency to be significantly faster than any random fluctuations in the probe beam filters out technical noise, and reduces the effect of flicker ($1/f$) noise.

The Faraday effect can be made time-dependent by applying a bias magnetic field perpendicular to the direction of propagation of the beam (z). Applying an rf $\frac{\pi}{2}$ -pulse then tips the spin of the atoms into the plane perpendicular to the bias field, which causes the spins to undergo Larmor precession. Applying the linear Zeeman approximation,¹ the spin projection (and hence the Faraday rotation) then oscillates at the Larmor frequency, with

$$\langle F_z \rangle = F \cos(\omega_L t) \quad \text{where} \quad \omega_L = \gamma B \quad (3.5)$$

and $\gamma = 2\pi \times 702 \text{ kHz/G}$ is the gyromagnetic ratio.

Hence the frequency can be tuned by adjusting the bias field, which is typically in the region of $B \sim 1 \text{ G}$ to achieve a Larmor frequency of $f_L = \omega_L/2\pi \sim 700 \text{ kHz}$. This field is easily produced by our magnetic bias coils, and matches the bandwidth of our photodetection (§6.2) and digitisation apparatus (§6.6). This is therefore a convenient magnetic field bias to operate at, with some flexibility on the exact value used.

¹ The quadratic Zeeman correction does not affect the Larmor frequency, as discussed in §7.6.

3.2: Faraday imaging in 1D

Measuring the Faraday rotation produced by the entire cloud provides information about the atom number and spin state of the whole cloud, but contains no spatial information about the cloud's structure. However, spatial information about the cloud can be recovered by placing the cloud in a magnetic field gradient b , which makes the Larmor frequency position-dependent (Figure 3.2).

The measured Faraday signal then contains a range of frequency components corresponding to the different Larmor frequencies, linking spatial position to measurement frequency. The weighting of each frequency component contains the local density and spin projection of the cloud, so the Fourier power spectrum of the signal provides a one-dimensional spatial map of the cloud.

I now derive an expression for the spatial profile of the cloud by considering the measured voltage on the polarimeter for a cloud placed in a linear magnetic field gradient in the z -direction. The local Larmor frequency can be expressed as

$$\omega_L(z) = \omega_{L0} + \gamma bz, \quad (3.6)$$

where $\omega_{L0} = \gamma B_0$ is the 'carrier' Larmor frequency in the centre of the cloud corresponding to the bias field B_0 (typically of order 1 G), and b is the strength of the magnetic field gradient (of order 0.1–100 G/cm).

Consider the Faraday rotation due to an infinitesimal slice of thickness dz of the cloud at position z in the magnetic field. The column density for just this slice is $d\tilde{\rho} = \rho(z) dz$, so by (2.42), the degree of Faraday rotation is

$$d\varphi_z(x, y, z) = F \xi_F \cos \theta(z, t) \rho(x, y, z) dz. \quad (3.7)$$

The angle $\theta(z)$ between the net spin vector and the propagation direction oscillates at the local Larmor frequency $\omega_L(z)$, where $\omega_L(z)$ is taken to be constant within this infinitesimal slice.

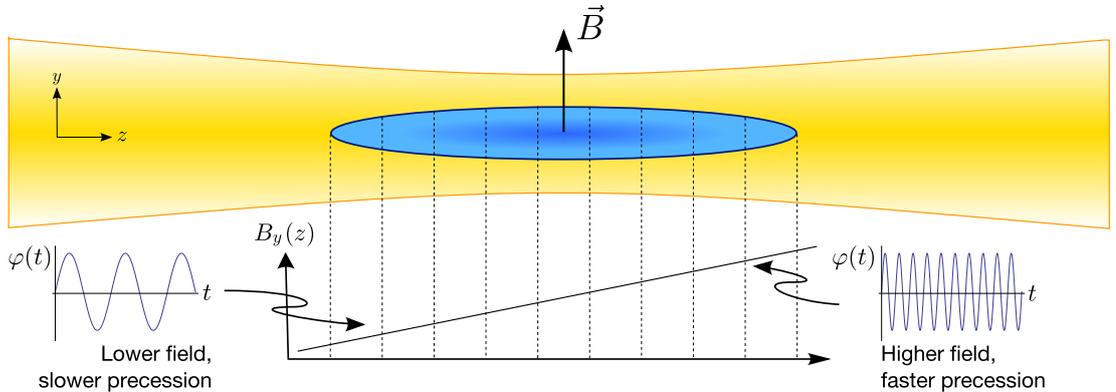


Figure 3.2: A magnetic field gradient causes Larmor precession at different rates in a BEC, coupling spatial information into the frequency domain.

Summing up the infinitesimal contributions to the polarisation rotation as the beam transits the cloud, the net rotation is

$$\Phi(x, y) = \int_{-\infty}^{\infty} \frac{d\varphi_z(z)}{dz} dz = F\xi_F \int_{-\infty}^{\infty} \cos \theta(z, t) \rho(x, y, z) dz. \quad (3.8)$$

Integrating across the spatial profile of the beam, the net power differential measured by the polarimeter is

$$\begin{aligned} \Delta P &\simeq 2I_0 \iint_{-\infty}^{\infty} \Phi(x, y) dx dy \\ &= 2I_0 \iint_{-\infty}^{\infty} \frac{1}{4} F\xi_F \int_{-\infty}^{\infty} \cos \theta(z, t) \rho(x, y, z) dz dx dy \\ &= 2FI_0\xi_F \iiint_{-\infty}^{\infty} \cos \theta(z, t) \rho(x, y, z) dx dy dz \\ &= 2FI_0\xi_F \int_{-\infty}^{\infty} \cos \theta(z, t) \bar{\rho}(z) dz, \end{aligned} \quad (3.9)$$

where $\bar{\rho}(z)$ is called the *line density*, and is given by

$$\bar{\rho}(z) \equiv \iint \rho(x, y, z) dx dy. \quad (3.10)$$

The BEC is created in the $|m_F = -1\rangle$ state (see §4.1), so Larmor precession is initiated by applying an rf $\frac{\pi}{2}$ -pulse. This results in a superposition state with fractional populations $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$, whose spins are polarised in the x - z plane. As the spins are all tipped simultaneously they begin Larmor precession in phase, with

$$\begin{aligned} \theta(z, t) &= \omega_L(z)t, \\ \therefore \cos \theta(z, t) &= \cos((\omega_{L0} + \gamma bz)t), \\ \mathcal{F} [\cos \theta(z, t)] &= \frac{1}{2} \left\{ \delta(\omega - (\omega_{L0} + \gamma bz)) + \delta(\omega + (\omega_{L0} + \gamma bz)) \right\} \\ &= \frac{1}{2\gamma|b|} \left\{ \delta\left(z - \frac{\omega - \omega_{L0}}{\gamma b}\right) + \delta\left(z + \frac{\omega + \omega_{L0}}{\gamma b}\right) \right\}, \end{aligned} \quad (3.11)$$

where δ is the Dirac delta distribution. Therefore the Fourier transform of the measured power differential is

$$\begin{aligned} \mathcal{F} [\Delta P] &= 2FI_0\xi_F \int \mathcal{F} [\cos \theta(z, t)] \bar{\rho}(z) dz \\ &= 2FI_0\xi_F \frac{1}{2\gamma|b|} \int \left\{ \delta\left(z - \frac{\omega - \omega_{L0}}{\gamma b}\right) + \delta\left(z + \frac{\omega + \omega_{L0}}{\gamma b}\right) \right\} \bar{\rho}(z) dz \\ &= \frac{FI_0\xi_F}{\gamma|b|} \bar{\rho}\left(\frac{|\omega| - \omega_{L0}}{\gamma b}\right). \end{aligned} \quad (3.12)$$

Hence for a given gain G , responsivity R , and optical collection efficiency κ , the voltage out of the differential photodetector is

$$V_{\text{out}} = GR\kappa \Delta P \quad \Rightarrow \quad \mathcal{F}[V_{\text{out}}] = \frac{GR\kappa FI_0 \xi_F}{\gamma|b|} \bar{\rho} \left(\frac{|\omega| - \omega_{L0}}{\gamma b} \right). \quad (3.13)$$

Therefore the frequency components of the measured voltage correspond to the integrated line density at each point along the cloud in the z -direction, and taking the Fourier transform inverts the signal for the one-dimensional cloud profile $\bar{\rho}(z)$:

$$\bar{\rho}(z) = \frac{\gamma|b|}{GR\kappa FI_0 \xi_F} \mathcal{F}[V_{\text{out}}] \Big|_{\omega=\omega_{L0}+\gamma bz}. \quad (3.14)$$

The spectrum will be obtained by measuring the voltage for some finite measurement time τ_f and computing its Fast Fourier Transform (FFT). The minimum length-scale Δz of the image is obtained by equating the spacing of samples in this Fourier representation Δf to the corresponding frequency spread induced by the gradient,

$$\Delta f = \frac{1}{\tau_f} = \frac{\gamma b}{2\pi} \Delta z \quad \Rightarrow \quad \Delta z = \frac{2\pi}{\gamma b \tau_f}. \quad (3.15)$$

For a cloud with radius L_z , the number of points ('bins') in the profile corresponds to how many resolution elements are contained in the size of the cloud $2L_z$. Hence to obtain an image with N_b bins, the required resolution is $\Delta z = 2L_z/N_b$ and hence the required magnetic field gradient is

$$b = \frac{\pi N_b}{\gamma L_z \tau_f}. \quad (3.16)$$

3.3: Signal-to-noise ratio

There are a number of trade-offs to be made in obtaining a 1D profile of the cloud by measuring the Faraday signal and inverting it with the Fourier transform to achieve a 1D image. In this section I consider the signal-to-noise of a reconstruction with a fixed measurement time and scattering rate.

Consider a measurement where b is chosen such that the Fourier transform of the signal divides it into N_b Fourier bins (3.16). The *average* contribution of each bin to the signal is therefore

$$\begin{aligned} \tilde{V}_{\text{sig}} &= \frac{1}{N_b} \int_{\omega_L - \gamma b L_z}^{\omega_L + \gamma b L_z} \mathcal{F}[V_{\text{out}}](\omega) d\omega \\ &= \frac{\gamma|b|}{N_b} \int_{-L_z}^{L_z} \mathcal{F}[V_{\text{out}}](\omega_L + \gamma bz) dz \\ &= \frac{GR\kappa FI_0 \xi_F}{N_b} \int_{-L_z}^{L_z} \bar{\rho}(z) dz \\ &= GR\kappa FI_0 \xi_F \frac{N_a}{N_b}. \end{aligned} \quad (3.17)$$

The width of each Fourier bin is $\Delta f = 1/\tau_f$, where τ_f is the measurement time, so the shot noise associated with each bin is

$$\tilde{V}_{\text{shot}} = G\sqrt{eR\kappa P/\tau_f} = G\sqrt{eR\kappa\pi a^2 I_0/\tau_f}. \quad (3.18)$$

The scattering rate for a given intensity (derived in §B.2) can also be written

$$\gamma_s = \frac{4\omega\xi_S^2}{hc^2} I_0 \quad (3.19)$$

$$\text{where } \xi_S^2 = \left(\frac{\lambda_{J'F'}^3 \Gamma_{J'}}{16\pi^2 c} \right)^2 \sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\alpha_0} \left(\frac{\omega}{\Delta_{J'F'}} \right)^2, \quad (3.20)$$

which in the far-detuned limit (§2.2) becomes

$$\xi_S^2 \approx \left(\frac{\lambda_{J'F'}^3 \Gamma_{J'}}{16\pi^2 c} \right)^2 \left(\frac{\omega^2}{\Delta_{D1}^2} + \frac{2\omega^2}{\Delta_{D2}^2} \right). \quad (3.21)$$

Assuming the detector is shot-noise limited (see §6.3), the signal-to-noise ratio (SNR) of the measurement is therefore

$$\text{SNR} = \frac{\tilde{V}_{\text{sig}}}{\tilde{V}_{\text{shot}}} = F \frac{\sqrt{\eta\kappa}}{\sqrt{\pi a^2}} \frac{N_a}{N_b} \sqrt{\frac{I_0 \tau_F}{\hbar\omega}} \xi_F = F \sqrt{\eta\kappa} \frac{\lambda}{2\sqrt{2}\pi a} \frac{N_a}{N_b} \frac{\xi_F}{\xi_S} \sqrt{\frac{\tau_f}{\tau_s}}, \quad (3.22)$$

where the quantum efficiency of the photodiodes is $\eta = \hbar\omega R/e$ and the photon-scattering lifetime is $\tau_s \equiv 1/\gamma_s$.

The form of the signal-to-noise ratio expression is easily interpreted as arising from several contributions. In particular, there is a linear dependence on the ratio of the measurement wavelength λ to the aperture size a , the average number of atoms contributing to the signal in each bin N_a/N_b , with the structure of the transitions entirely captured in the scaling factor ξ_F/ξ_S .

This dependence on detuning can be seen by expressing

$$\frac{\xi_F}{\xi_S} \equiv \frac{\sum \alpha_{J'F'}^{(1)}/\Delta_{J'F'}}{\sqrt{\sum \alpha_{J'F'}^{(0)}/\Delta_{J'F'}^2}} \approx \frac{\Delta_{D1}^{-1} - \Delta_{D2}^{-1}}{\sqrt{\Delta_{D1}^{-2} + 2\Delta_{D2}^{-2}}} = \frac{\Delta_{D1} - \Delta_{D2}}{\sqrt{\Delta_{D2}^2 + 2\Delta_{D1}^2}}. \quad (3.23)$$

Considering detuning close to the D2 line only ($\Delta_{D2} \rightarrow 0$), the ratio is

$$\lambda \approx \lambda_{D2} \quad \Rightarrow \quad \frac{\xi_F}{\xi_S} \approx \frac{1}{\sqrt{2}}. \quad (3.24)$$

Hence for a ‘whole cloud’ measurement ($N_b = 1$) on the D2 line, where the measurement is performed with no magnetic field gradient, the expected SNR is

$$\text{SNR} = \frac{\lambda N_a \sqrt{\eta\kappa}}{4\pi a} \sqrt{\frac{\tau_f}{\tau_s}}, \quad (3.25)$$

which agrees with the expression presented in Ref. [84].

3.4: Choice of detuning

The wavelength dependence of the signal-to-noise ratio is captured in the ratio ξ_F/ξ_S , which includes contributions from all associated fine-structure transitions (Figure 3.3). In the regime of interest, the D1 and D2 transitions both contribute to the Faraday effect, but have polarisability coefficients of opposite sign (§2.2). Between the two resonances, the detunings have opposite sign so the resulting Faraday rotations add constructively. The signal contributions ξ_F therefore add coherently by (2.43) but the noise contributions ξ_S add in quadrature by (3.20).

As in (2.45), this can be seen by introducing the detuning from the midpoint of the lines, $\Delta = \frac{1}{2}(\Delta_{D1} + \Delta_{D2})$. The ‘normalised’ SNR is

$$\text{SNR}_{\text{norm}} \equiv \frac{\xi_F}{\xi_S} = \frac{2\sqrt{3} \omega_{\text{fs}}}{\sqrt{(6\Delta + \omega_{\text{fs}})^2 + 8\omega_{\text{fs}}^2}}. \quad (3.26)$$

The SNR is therefore maximised at $\Delta = -\omega_{\text{fs}}/6$, which is equivalent to $\Delta_{D2} = -2\Delta_{D1}$. This is exactly the magic wavelength, resulting in an SNR that is enhanced by factor $\sqrt{3/2} \approx 1.22$.

For a fixed cloud size (both in number of atoms N_a and cloud size a) the remaining free parameters in the signal-to-noise ratio estimation are the number of bins in the measurement N_b , the measurement time τ_f and the scattering lifetime τ_s . Typically the measurement time and the number of bins will be set by the phenomena of interest to be observed, which leaves only the scattering rate as free in (3.22). Therefore the scattering rate is fixed by the desired signal-to-noise ratio.

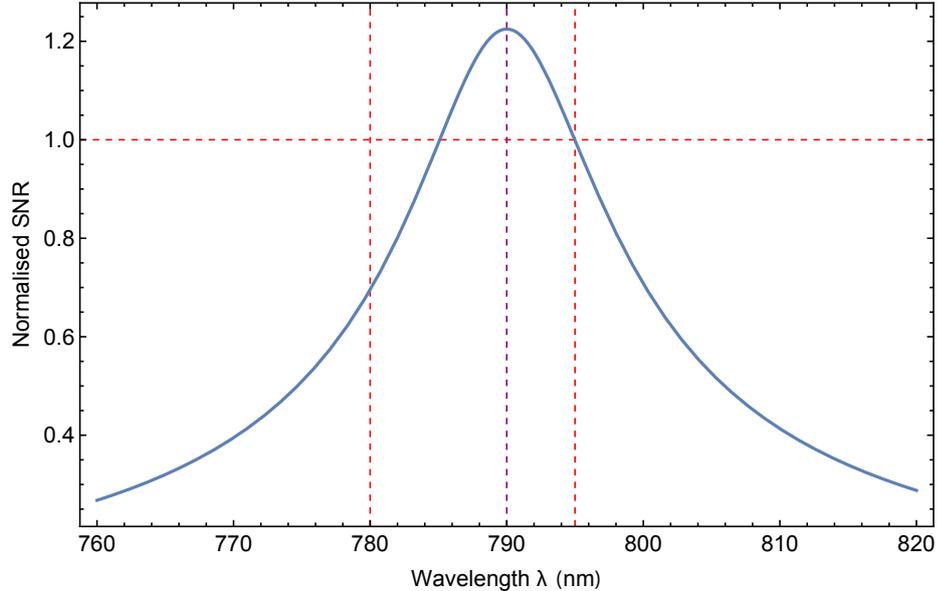


Figure 3.3: Dependence of signal-to-noise ratio (SNR) on wavelength, showing a clear optimum at the magic wavelength $\lambda = 790$ nm (purple), compared to the D1 and D2 transitions (red).

To gain the highest signal-to-noise ratio, it is preferable for the scattering rate to be high and the lifetime low, as the scattered photons contribute to the signal at the expense of ejecting atoms from the condensate. However, for a ‘minimally destructive’ measurement it is undesirable for a significant fraction of the condensate to be destroyed during the measurement, so $\tau_s \gg \tau_f$, setting a limit on the maximum achievable SNR.

Finally, as the scattering rate (3.19) is controlled by the detuning from resonance and the intensity of the beam, there is a choice to be made between using a weak probe beam detuned close to resonance, or a bright probe beam detuned far from resonance (Figure 3.4). For a given scattering rate, the total amount of laser power available introduces a practical limitation on the maximum detuning available.

However, the benefit of using far off-resonant bright light is that it is straightforward to build shot-noise-limited photodetectors for bright beams. Large area photodiodes can be made to work at megahertz frequencies [128] with the photon shot noise dominating the electrical noise. The magic wavelength is therefore an optimal wavelength choice as it is far detuned from both the D1 and D2 transition lines, and the required beam power to achieve a nominal scattering rate of $\sim 1 \text{ s}^{-1}$ is $\sim 1 \text{ mW}$ per photodiode, which provides a generous photocurrent to stay above the electronic noise floor.

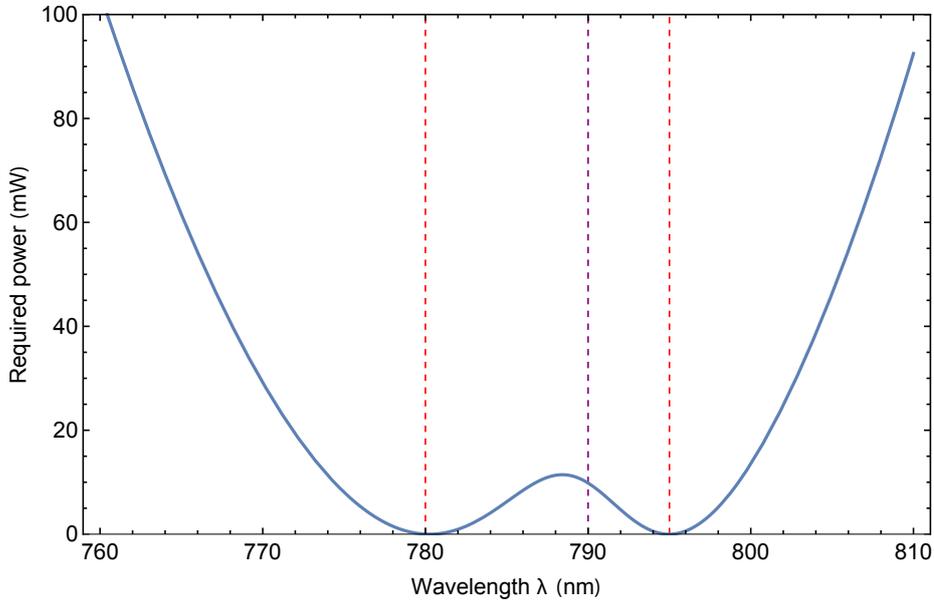


Figure 3.4: Dependence of probe power on wavelength for a fixed scattering rate, taking the beam radius as $75 \mu\text{m}$ and scattering lifetime as $\tau_s = 1 \text{ s}$. Detuning close to the D1 or D2 lines (red) increases the scattering rate for constant power, so the power has to be reduced to keep the scattering rate fixed.

3.5: Coherent Faraday signal dephasing

A magnetic field gradient is applied to couple spatial information into the Faraday signal by causing the Larmor frequency to become spatially dependent. However, having different frequency components mean that the spins will Larmor-precess in the x - z plane at different rates. Spins in a high-field part of the gradient will evolve rapidly and begin to wind from spins located in low-field parts of the gradient (Figure 3.5). As the spin vectors fan out in the x - z plane, the length of the net spin vector vanishes, causing $\langle \hat{F}_z \rangle \rightarrow 0$ and hence the induced Faraday rotation to vanish as well, $V_{\text{out}} \rightarrow 0$. This is known in the MRI community as ‘free induction decay’ (FID) and the characteristic decay time of the signal is called the ‘ T_2^* time’.²

This can be seen to be a consequence of (3.14), since the measured signal in the time domain looks like the inverse Fourier transform of the density profile. The density profile of the cloud therefore describes both the rate of dephasing, and the shape of the measured signal. Since the measurement is performed in the presence of noise, it is important to consider what contribution the interesting features of the target atom cloud make to the signal, and whether or not they can be distinguished above the noise.

A thermal cloud has a Gaussian profile, and since the Fourier transform of a Gaussian is also a Gaussian, the measured Faraday signal measured for a thermal cloud is expected to be contained within a Gaussian envelope. For a spherically symmetric cloud with $1/e$ radius σ , the expected signal is

$$\begin{aligned} \rho(r, \phi, \theta) \propto e^{-r^2/\sigma^2} &\Rightarrow \bar{\rho}(z) \propto e^{-z^2/\sigma^2} \\ &\Rightarrow V(t) \propto e^{-\gamma^2 b^2 \sigma^2 t^2 / 4} \cos(\omega_{L0} t), \end{aligned} \quad (3.27)$$

which decays with characteristic time constant $t_c = 2/\gamma b \sigma$. Importantly, the envelope of this measurement monotonically decays towards zero and shows no revivals (Figure 3.6A).

However, BECs do not have a Gaussian profile; they are more closely modelled as having a Thomas-Fermi shape. For a cylindrically symmetric condensate with chemical potential μ and radii L_z in the axial and L_r in the radial directions, the density in cylindrical coordinates is

$$\rho(r, \phi, z) = \frac{\mu}{g} \max\left(1 - \frac{r^2}{L_r^2} - \frac{z^2}{L_z^2}, 0\right). \quad (3.28)$$

The corresponding integrated line density for $|z| \leq L_z$ is

$$\bar{\rho}(z) = \int_0^{L_r \sqrt{1-z^2/L_z^2}} 2\pi r \rho(r, \phi, z) dr = \frac{\mu \pi L_r^2}{g} \frac{1}{2} \left(1 - \left(\frac{z}{L_z}\right)^2\right)^2. \quad (3.29)$$

² FID is a combination of coherent (gradient-induced) dephasing, decoherence (spin-spin relaxation), and remagnetisation (spin-lattice relaxation). However, strong magnetic field gradients are preferable for MRI (see §3.9) so gradient-induced dephasing dominates.

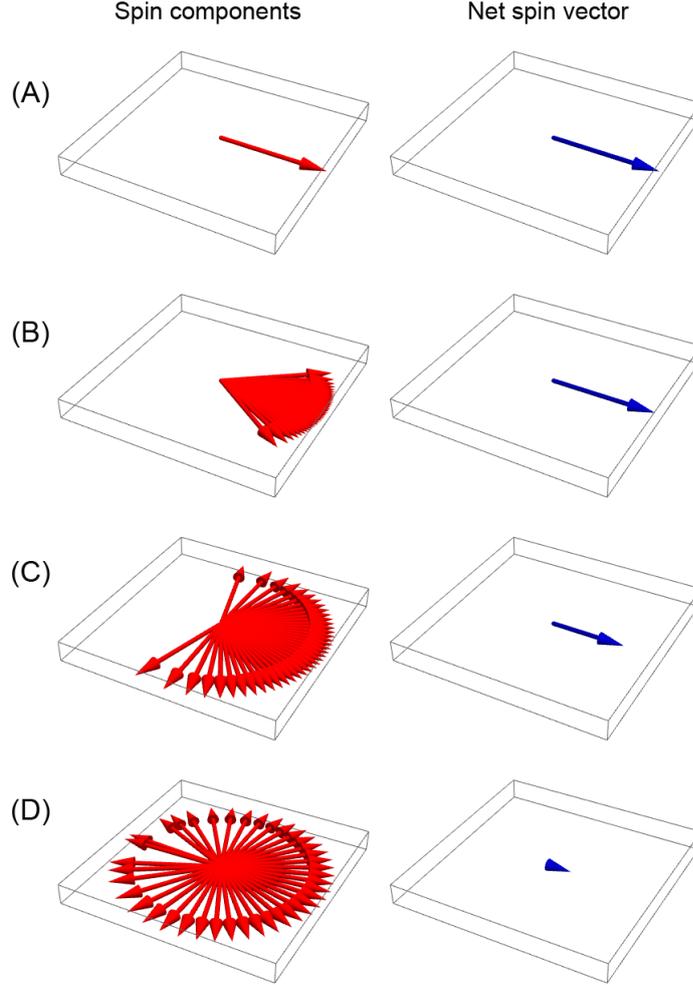


Figure 3.5: Schematic of coherent spin dephasing. Initially all spin components are aligned following the spin tip (A), but as the spin components evolve (B)–(C) they begin to spread out in the x – z plane, causing the net spin vector to decrease in magnitude. Eventually the spins wrap the plane and the net spin averages to zero (D).

The overall shape of the expected Faraday signal corresponding to the integrated Thomas-Fermi profile is very similar to the Gaussian profile, but shows several small revivals after the initial decay (Figure 3.6B). These are a result of the small-scale differences between the Thomas-Fermi and Gaussian line densities that become evident at longer evolution times, demonstrating that the method is sensitive to small changes in the overall profile.

Since spins physically separated by distance δz Larmor-precess at different rates, they develop a π -phase shift between them in time

$$t_c(\delta z) = \frac{\pi}{\gamma b \delta z}. \quad (3.30)$$

Therefore regardless of the cloud profile, the overall signal is expected to be contained in an envelope that decays over the characteristic timescale $t_c(2L_z) = \pi/2\gamma b L_z$.

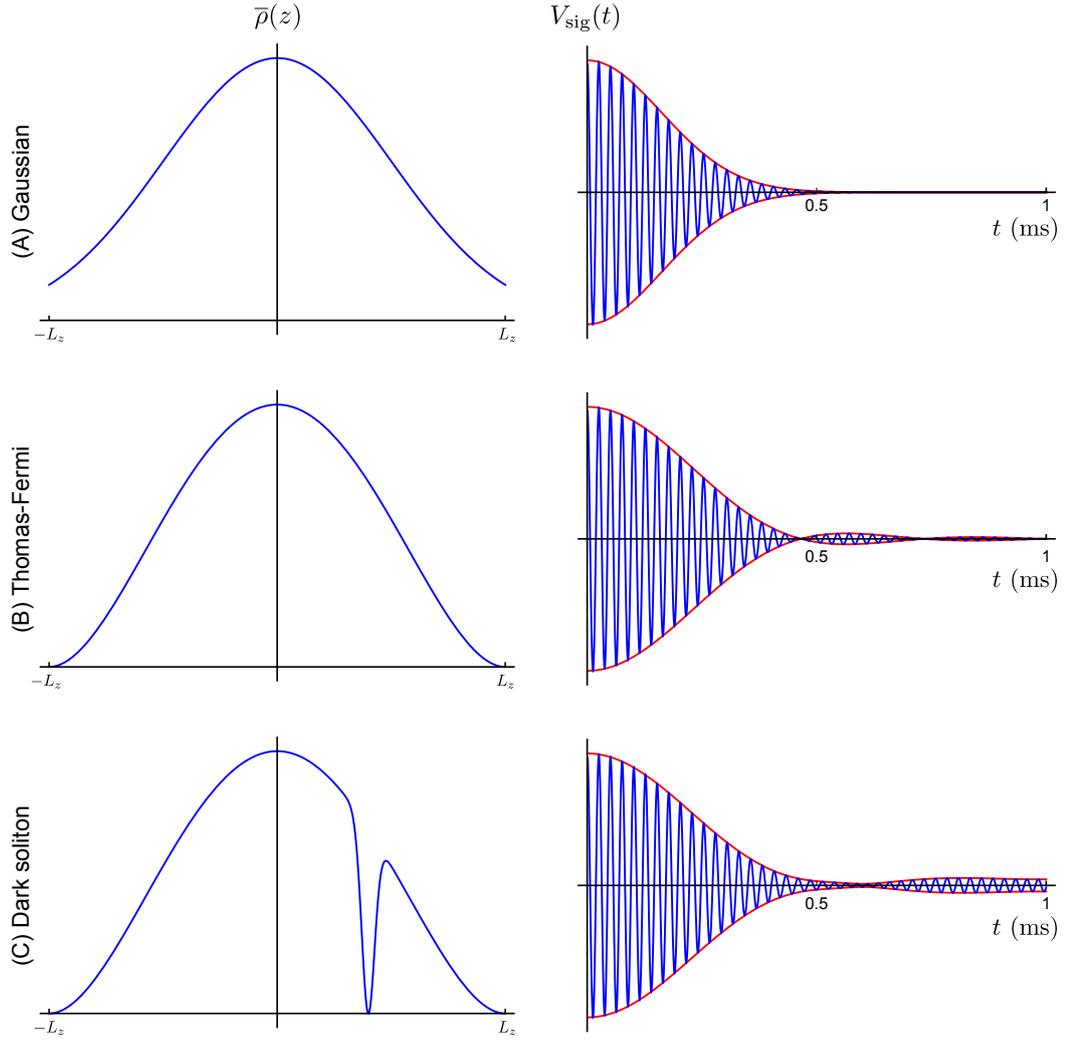


Figure 3.6: Schematic density profiles $\bar{\rho}(z)$ and expected signal $V_{\text{sig}}(t)$ for three different cloud shapes. As the spins dephase, the envelope of the signal (red) decays exponentially for a Gaussian profile (A), but a Thomas-Fermi profile (B) demonstrates small revivals. Small structures such as dark-solitons (C) lead to persistence on a longer timescale.

Small features within the cloud change the cloud density at that location, change the weighting of a small spread of corresponding frequencies in Fourier space. These require more time to develop a π phase difference across them to destructively interfere, and therefore persist for longer (Figure 3.6C). The smaller the feature, the longer it persists but the smaller the amplitude is (Figure 3.7).

For example, for a $L_z = 15 \mu\text{m}$ radius cloud in a 1 G/cm gradient the first null in the Faraday signal is expected at $t_c(2L_z) = 0.24 \text{ ms}$, which is the timescale of the initial FID of the signal. However, if there is a $\delta z = 1 \mu\text{m}$ feature within the cloud, the signal is expected to show low-amplitude revivals that persist on the timescale of $t_c(\delta z) = 7.1 \text{ ms}$.

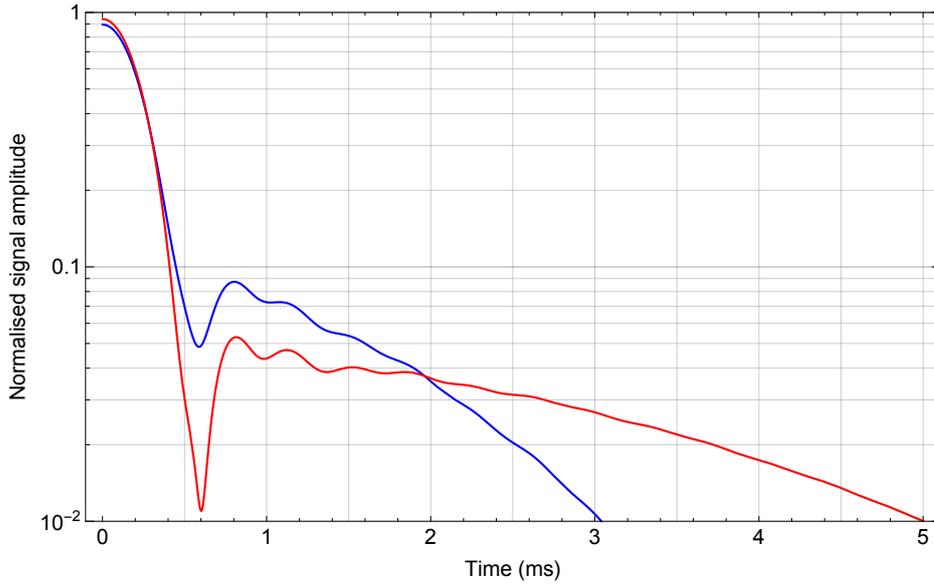


Figure 3.7: Comparison of the expected signal envelope for Thomas-Fermi profile clouds with dark features of width $1\ \mu\text{m}$ (red) and $2\ \mu\text{m}$ (blue). The smaller feature has a revival with smaller signal amplitude than the larger feature, but persists longer.

3.6: Rephasing with spin echo

The dephasing process described in the previous section is entirely coherent, and can therefore be reversed. Since the dephasing is due to Larmor precession occurring at different rates, by reversing the sign of the gradient the previously fast components become slow and vice-versa. The previously slow components then evolve rapidly and ‘catch up’ to what were previously the fast components, causing the ‘spreading out’ to be reversed and bringing the spins back into alignment. As the net spin vector becomes nonzero again, a revival of the Faraday signal is observed. This is an example of a spin echo technique called ‘gradient recall echo’.

Once the net spin has rephased, the now-fast components will then outrun the slow components, leading to dephasing in the opposite direction. Reversing the gradient at a later time reverses the process again, and causes another spin echo when the spins rephase. In this way, multiple Faraday measurements (‘images’) can be attained sequentially (Figure 3.8).

However, gradient-induced dephasing is not the only cause for decay of the Faraday signal, with fundamental effects such as quantum phase diffusion [129] and technical limitations on the spin echo prevent perfect rephasing of the spin. Decoherence of the constituent spins will result in gradual decay of the amplitude, quantified by the ‘ T_2 time’ and is typically much longer than the free-induction decay timescale (‘ T_2^* time’). The T_2^* time is the timescale of an individual measurement, the T_2 time limits the number of echoes that can be meaningfully extracted.

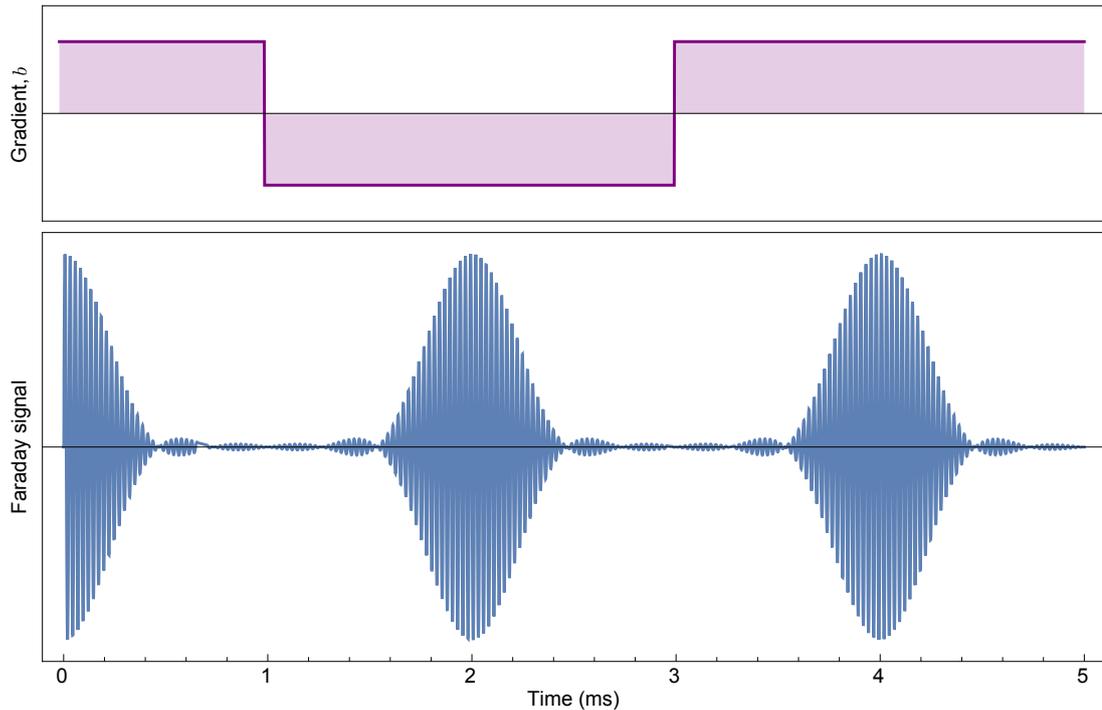


Figure 3.8: Schematic of gradient recall echo. After the initial spin tip at $t = 0$, the spins dephase and the envelope of the Faraday signal decays. When the magnetic field gradient b is reversed at $t = 1$ ms, the spins begin to rephase and an ‘echo’ is produced at $t = 2$ ms. Repeating the process generates further echoes, each of which can be individually reconstructed to form time-resolved profiles of the BEC. This example corresponds to five such profiles.

In nuclear MRI, the measurement is also limited by the rate at which the component spins realign with the bias field (the ‘longitudinal relaxation’ or ‘ T_1 ’ time). This process is generally called ‘spin-lattice’ relaxation, as the remagnetisation occurs through spins interacting with the environment. However, there is no analog to this effect in a BEC as magnetisation is conserved, so the T_1 timescale does not apply.³

The dephasing process can also be reversed with an rf π -pulse (‘radiative spin echo’), which inverts the spin states. Where the rapidly-evolving phases are phase-advanced before the pulse, they become phase-retarded after the pulse. However, since they are still rapidly-evolving, they catch up to the slowly-evolving phases (which became phase-advanced after the pulse), resulting in a non-zero magnetisation and hence producing a spin echo. Rephasing the spins in this way has the added advantage of undoing any phase-evolution from parasitic background gradients, which gradient-recall cannot reverse.⁴

³ It should be noted that unlike a condensed matter system, atoms are continually lost from the trapped BEC (see §5.7) resulting in a total magnetisation that decays over time. The net effect is a reduction in signal strength with an effective T_1 timescale that cannot be reversed.

⁴ In principle a gradient could be applied to cancel the background gradient, but the background gradient has components in all directions, all of which contribute to gradient-induced dephasing.

3.7: Slice selection and higher dimensional images

The imaging technique discussed thus far creates a time-resolved sequence of 1D images of a BEC along the ‘imaging axis’. While this enables observation of the dynamics of features in a quasi-1D BEC (e.g. a dark soliton propagating along the imaging axis), the 1D profile obtained ‘blurs’ out features in a higher-dimensional BEC by averaging over the transverse axes. This makes it difficult to observe features that do not extend across the cloud (e.g. vortices in 2D) or have no fixed orientation (e.g. vortex lines in 3D). However, further techniques may be able to be imported from medical MRI to use sequential magnetic field gradients along multiple axes.⁵

This is also a problem in absorption imaging, such as the imaging of vortices cores, as the cores would be blurred unless they perfectly aligned along the imaging axis. The solution is ‘slice-selection’ [130], which images only a 2D slice of the sample. Typically this is achieved in a BEC by holding it in a dark state and optically pumping a slice into another hyperfine state that is resonant with the imaging beam. The principle was first applied to dark-field imaging [131], and later applied to PTAI, enabling the motion of vortex cores in a 2D vortex lattice to be tracked [132].

The problem is more significant in Faraday measurement as the blurring occurs across the two transverse dimensions, instead of only the propagation axis as in absorption imaging. Any atoms which are not Larmor precessing are ‘dark’ to the measurement, so slice-selection can be performed in Faraday measurement by applying a spatially selective $\frac{\pi}{2}$ -pulse.⁶ The pulse is chosen such that it only effectively tips the spins in a fraction of the sample corresponding to a region of interest. The resulting Faraday signal only contains contributions from those atoms that were selected, resulting in slower dephasing than if the spins of the entire cloud were tipped.⁷ Performing slice-selection with a magnetic field gradient applied along the imaging axis enables very strong magnetic field gradients for enhanced resolution in that region. Alternatively, slice-selecting using a gradient perpendicular to the imaging axis allows a 2D image to be constructed by imaging a sequence of slices and combining the profiles.

Slice-selection can be performed with a shaped rf pulse such as a sinc or hyperbolic secant pulse [133, 134]. For a sample in a linear magnetic field gradient, the time-varying intensity and detuning selectively tip the spins in a tunable region. Choosing the centre-frequency of the pulse defines the location of the slice to be selected, and the width of the pulse (the pulse bandwidth) combined with the applied gradient strength

⁵ Because of the strong bias field, the relevant gradients are directional derivatives of the bias component of the net magnetic field (see §8.3), which for a y -bias are $\frac{\partial B_y}{\partial x}$, $\frac{\partial B_y}{\partial y}$ and $\frac{\partial B_y}{\partial z}$.

⁶ The $\frac{\pi}{2}$ -pulses are not currently spatially selective as they have a high rf Rabi frequency and have a short duration, making them spectrally broad compared to the spread of Larmor frequencies in the range of magnetic field gradient strengths applied.

⁷ In principle. If the dephasing occurs too rapidly then the signal will be undersampled because of the finite acquisition rate of the digitisation apparatus, leading to poor reconstruction.

defines the width of the slice being selected. While standard procedure in MRI, this technique has only recently been applied to ultracold atom clouds in a process called magnetic resonance control (MRC) [135, 136].

Typically the T_1 time is a measure of how rapidly sequential $\frac{\pi}{2}$ -pulses can be applied, as the pulse will not rotate the spin into the transverse plane if there is residual magnetisation. As described in §3.6, remagnetisation does not occur in a spinor BEC, which means that long sequences of rephasing pulses can be used, but sequential slice-selection is more difficult. One potential solution would be to apply a $\frac{-\pi}{2}$ -pulse at the exact time of a spin echo, which would tip the spin back into alignment with the external field ('forced remagnetisation') effectively 'deselecting' the slice. A new slice could then be selected using a spatially-selective $\frac{\pi}{2}$ -pulse for a different plane.

In principle, multiple reconstructed 1D profiles taken through different slices can be stacked together to form a 2D image. In practice however, a new slice can only be selected during the $\frac{\pi}{2}$ -pulse, and not varied between echos. While multiple echoes could be performed during this time to obtain multiple shots of a given slice and increase the signal-to-noise ratio by averaging, it would take a prohibitively long time to construct an image of the whole cloud in this way.

The primary method of producing higher-dimensional images using MRI is known as 'phase encoding' (PE). Unlike a conventional spin echo (CSE) sequence⁸ (Figure 3.9A), phase encoding in the simplest configuration is performed by simply applying a perpendicular gradient for a fixed duration before the echo (Figure 3.9B). This fixed period exposed to the gradient causes dephasing and couples the y -position of the spins into their phase, encoding spatial information into the free-induction decay of the subsequent echo. Varying the strength of the phase-encoding gradient between echoes enables this spatial direction to be mapped out and a ' k -space' image of the cloud to be produced by stacking together the echoes (Figure 3.10). A 2D image of the cloud is then reconstructed by taking the 2D inverse Fourier transform of the k -space image. The number of pixels in the reconstruction is given by the number of phase-encoding steps used, with the pixel size given by the strength of the gradient and the phase-encoding time used.

This method can be immediately extended to perform 3D reconstruction by replacing slice-selection with phase-encoding along the remaining axis. Stacking the resulting profiles into a 3D k -space then enables the 3D Fourier transform to be used to reconstruct the density of the BEC. Slice selection is often more beneficial in medical MRI because of spin-lattice relaxation, and the chosen slices are relevant in interpretation of the MRI. However, the phase coherence of a BEC makes phase-encoding better suited for 3D imaging.

⁸ Shown using radiative spin echo instead of gradient recall echo for comparison purposes.

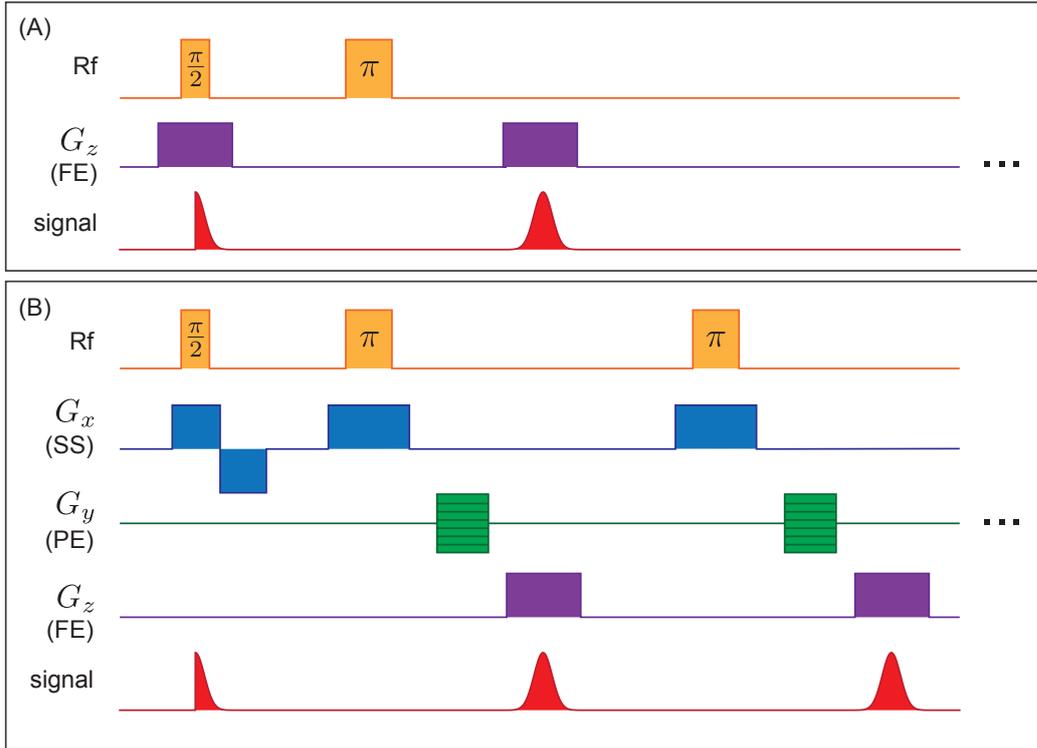


Figure 3.9: Schematic showing the magnetic field gradients required for a conventional sequence (A) using radiative π -pulses, and an extended sequence (B) including slice-selection (SS) along the x -axis and phase-encoding (PE) in the y -axis. The frequency-encoding (FE) gradient produces a 1D profile of the cloud along the z -axis, and the strength of the PE gradient is varied between echoes so that they can be combined to create a 2D image. Stacking images obtained in different slices (SS) then enables a full 3D reconstruction.

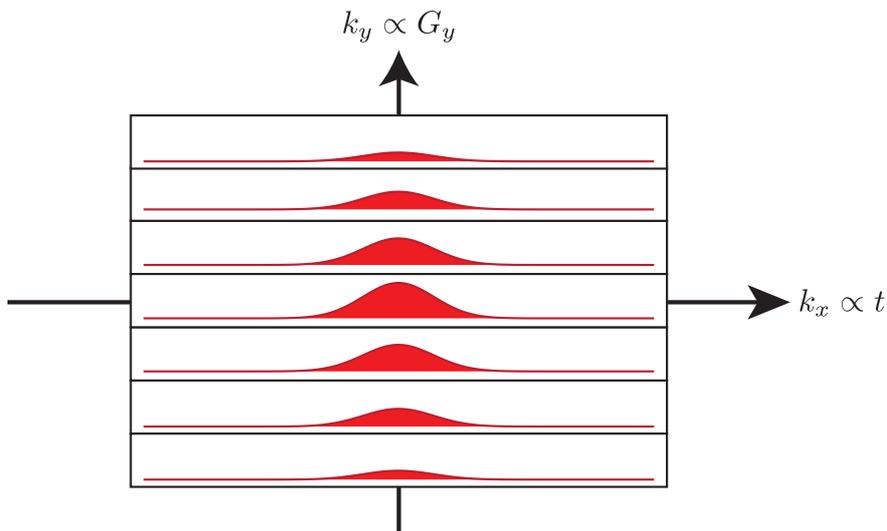


Figure 3.10: A 'k-space' image of the BEC is constructed in (k_x, k_y) -space by stacking the time-domain signal of echoes obtained with different phase-encoding gradients G_y . The image is reconstructed by taking the 2D inverse Fourier transform to reconstruct the 2D density profile of BEC as chosen during the slice-selection step.

Clearly there is a tradeoff between resolution of the reconstructed image and time-resolution, as echoes can either be used to improve the spatial resolution of an image, or contribute to the next image for improved temporal resolution. For a reconstruction with echo time t_e , and containing N_p phase-encoding steps repeated for N_s slice-selections, the resulting image comprises $(\frac{\tau_f}{f_s}, N_p, N_s)$ voxels⁹ and takes time $t = N_s N_p t_e$ to capture. Any dynamics which are to be observed must occur on a timescale significantly longer than this. In particular, dynamics on the timescale $t = N_p t_e$ will result in blurring whereas slower dynamics will result in the slices corresponding to different times.

The primary limitation on the timing of the echoes is the strength of the applied gradients, as stronger gradients enable faster encoding. In particular, independent control of all three gradient components of the bias field (B_y for a y -bias) is required, including the ‘off-diagonal’ components ($\frac{\partial B_y}{\partial x}$ and $\frac{\partial B_y}{\partial z}$) which are difficult to generate (see §8.3). The gradients must also be switched on and off rapidly to prevent limiting the duty cycle, and must be ramped symmetrically to prevent unintentional phase accumulation. Furthermore, eddy currents may be induced in the vacuum system and must also be taken into consideration. These are purely technical considerations however, and could be overcome with adequate design. For example, it may be possible to make a strong ‘effective’ magnetic field gradient using the vector light-shift (see §2.6) and a spatially-varying intensity profile.¹⁰

There is a wealth of knowledge behind the design of spin echo sequences, such as optimisation to minimise measurement time and compensation for apparatus imperfections. One example of this is using gradient reversal to prevent phase accumulation whereby applied gradients (SS and FE) induce unintended phase encoding. This is achieved during slice-selecting in Figure 3.9B, where the initial $\frac{\pi}{2}$ -pulse is followed by a reversal of the G_x gradient, to ‘unwind’ the phase-encoding it induces along the x -direction.

Similarly, it is important that the application of G_x before and after the π -pulses must be symmetric, so that any phase-winding before the π -pulse is undone by phase-winding after the pulse. Such reversals also reduce the influence of imperfect gradient synthesis (i.e. one set of coils generating a non-zero gradient along another axis) which would otherwise result in cumulative error. Typically, the frequency-encoding gradient is also reversed before the echo, which results in a gradient-recall echo that coincides with the radiative echo.

Many medical MRI techniques focus on generating contrast between regions of similar composition based on differing characteristic timescales (i.e. spatially varying $T_1, T_2,$

⁹ Where τ_f is the measurement time and f_s is the sample rate of the acquisition system.

¹⁰ The analysis is simplified by a linear gradient, but monotonic non-linear gradients are acceptable.

T_2^*), which is a different goal to BEC imaging. However, advances in generating and switching strong magnetic field gradients, and coherent control of spins are directly applicable. In general the much smaller physical size of the BEC makes the gradient generation and switching much less demanding than medical MRI systems, although it is still demanding by the standards of conventional BEC experiments.

3.8: Stern-Gerlach blurring

A side-effect of using a magnetic field gradient to encode spatial information in the Faraday measurement is that the gradient imparts a mechanical force on the different spin components of the cloud, causing them to become spatially separate. This separation can be reversed by either inverting the field gradient or applying a π -pulse, but can only be done following the acquisition of each image to prevent interfering with the measurement.

However, as separation occurs while a Faraday image is being captured, the atomic motion during the measurement interval will generate blurring. Provided this blurring is less than the resolution element of the image, it will not significantly affect the result. For a given gradient strength, this puts an upper limit on the maximum exposure time τ_f before blurring becomes significant.

The magnetic dipole potential can be written as

$$V_Z = m_F g_F \mu_B B = \hbar \gamma m_F B, \quad (3.31)$$

so the gradient-induced acceleration of the $m_F = \pm 1$ components along the z direction due to an applied gradient $b = \frac{\partial B_y}{\partial z}$ is

$$a = \frac{1}{m_{\text{Rb}}} \frac{\partial V_Z}{\partial z} = \pm \frac{\hbar \gamma}{m_{\text{Rb}}} b. \quad (3.32)$$

The magnetic field gradient is constant during the measurement, so the distance traveled during the measurement interval τ_f is

$$\Delta z_{\text{SG}} = \frac{1}{2} |a| \tau_f^2 = \frac{\hbar \gamma b}{2 m_{\text{Rb}}} \tau_f^2, \quad (3.33)$$

which is the degree of blurring introduced by Stern-Gerlach separation. To prevent limiting the resolution of the reconstructed image, this must be less than the size of a resolution element Δz , which requires that

$$\tau_f < \sqrt{\frac{2 m_{\text{Rb}}}{\hbar \gamma} \left(\frac{\Delta z}{b} \right)}. \quad (3.34)$$

This is an upper limit on the time that the gradient can be applied before reversal (either gradient inversion or spin echo) is required. However, this is not the only restriction on the resolution of the reconstruction, and (3.34) is insufficient to determine the gradient strength required to achieve a desired resolution.

3.9: Choice of magnetic field gradient

The density profile of the BEC is reconstructed by taking the Fourier transform of the time-domain Faraday signal (see §3.2). The resolution of the image is therefore given by the spacing of points in the discrete Fourier representation of the signal (the ‘Fourier resolution’), which by (3.15) is

$$\Delta z_F = \frac{2\pi}{\gamma b \tau_f}. \quad (3.35)$$

Increasing the measurement time τ_f improves the Fourier resolution at the expense of allowing more time for Stern-Gerlach separation to occur. These competing objectives are optimised when the degree of Stern-Gerlach blurring matches the resolution of the image, which is

$$\Delta z_{SG} = \Delta z_F \quad \Rightarrow \quad b = \sqrt{\frac{4\pi m_{\text{Rb}}}{\hbar \gamma^2} \frac{1}{\tau_f^3}}. \quad (3.36)$$

The field gradient and measurement interval required to achieve the desired resolution Δz are therefore

$$b = \frac{\pi \hbar}{m_{\text{Rb}} \gamma \Delta z^3}, \quad \tau_f = \frac{2m_{\text{Rb}} \Delta z^2}{\hbar}. \quad (3.37)$$

The primary limitation on improving the resolution is the strength of the magnetic field gradient, which requires specialised coil configurations to extend to gradients past ~ 10 G/cm. A given coil configuration will be capable of producing a particular maximum achievable gradient b_{max} , for which Stern-Gerlach blurring therefore limits the resolution to

$$\Delta z \geq \sqrt[3]{\frac{\pi \hbar}{m_{\text{Rb}} \gamma b_{\text{max}}}}. \quad (3.38)$$

The magnetic field gradient required to achieve a desired resolution (Figure 3.11) scales unfavourably as $b \propto 1/\Delta z^3$, requiring strong gradients to image below the diffraction limit without being dominated by Stern-Gerlach blurring (Table 3.1).

It should also be noted that the finite acquisition rate f_s of the digitisation apparatus limits the magnetic field gradient, as increasing the gradient results in faster dephasing, potentially causing the envelope to be undersampled. From the characteristic

Resolution, Δz (μm)	Gradient, b (G/cm)	Interval, τ_f (ms)
10	0.033	43
5	0.26	11
1	33	0.43
0.5	260	0.11
0.1	33,000	0.004

Table 3.1: Magnetic field gradients required to achieve a range of resolutions.

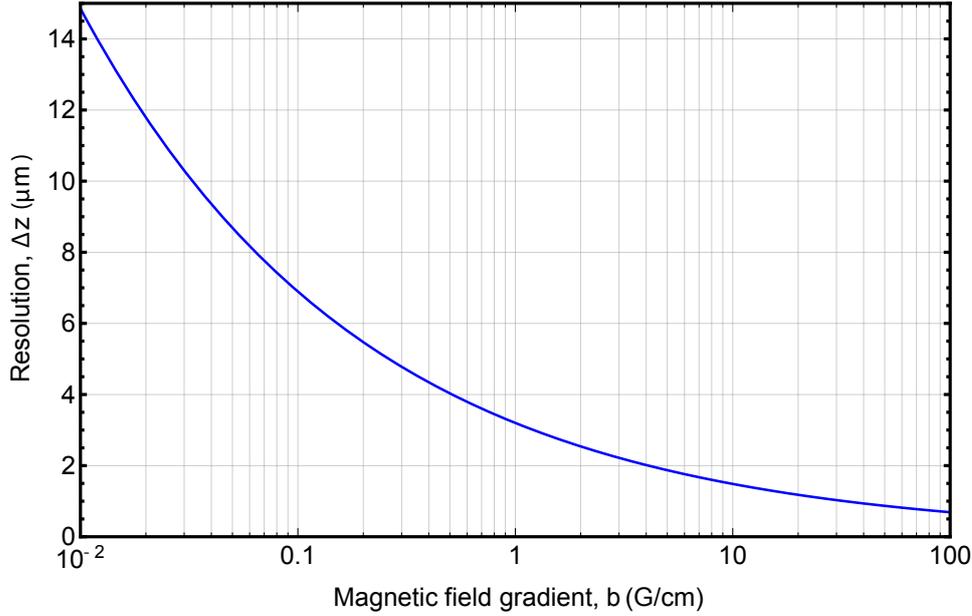


Figure 3.11: Stern-Gerlach limit on resolution for varying magnetic field gradients.

dephasing timescale (3.30), to oversample the envelope the sample rate must obey

$$\frac{1}{f_s} < \frac{\pi}{2\gamma b L_z} \quad \Rightarrow \quad b_{\max} < \frac{\pi f_s}{2\gamma L_z}. \quad (3.39)$$

Note that this condition is equivalent to requiring that the entire spread of Larmor frequencies fall within the Nyquist limit of the digitisation apparatus¹¹ (and also requires the bias to be set to give $f_L > f_s/4$ to avoid aliasing at DC).

Taking $f_s = 2\text{MS/s}$ and $L_z = 25\ \mu\text{m}$ gives $b_{\max} < 290\ \text{G/cm}$, corresponding to a best-case resolution $\Delta z = 0.49\ \mu\text{m}$. This can be improved with a faster digitiser, although the photodetection bandwidth and the bias field must be increased appropriately.

While these gradients are strong, they are achievable – especially in atom chip traps. An alternate method of generating strong gradients is using the vector light-shift, as a spatially-varying intensity can result in an effective magnetic field gradient (see §2.6). A VLS gradient generated using the Faraday probe beam would vary in strength radially, which is not useful for imaging. However, a second beam propagating orthogonal to the imaging axis could be used to produce the required gradient along that axis.

Note that static generation of strong gradients is much less problematic than rapid switching/inverting of the gradient, which encourages the use of radiative spin echo instead of gradient recall echo to rephase the spin (discussed further in §8.7).

¹¹ Even super-Nyquist sampling (see §6.6) requires the signal to fall within one Nyquist ‘window’.

3.10: Summary

In this chapter I showed how Faraday measurement might be used to obtain a minimally destructive image of a condensate by placing it in a magnetic field gradient. The resulting spatial dependence to the Larmor frequency should enable a 1D profile of the cloud to be recovered from the net Faraday signal using the Fourier transform. I derived an expression for the signal-to-noise ratio of the resulting profile, and discussed the influence of the free parameters.

The spread of Larmor frequencies contributing to the Faraday signal is predicted to result in ‘free induction decay’ as the frequency components dephase over time. Spin echo was considered as a method of rephasing the Faraday signal to generate multiple images and enable time-resolved studies to be carried out.

Stern-Gerlach separation was considered as a blurring mechanism, placing an upper limit on the duration of a single Faraday measurement. The magnetic field gradient required to obtain a desired imaging resolution Δz was derived, and shown to scale as $1/\Delta z^3$, with $\Delta z = 1 \mu\text{m}$ requiring $b = 33 \text{ G/cm}$. Despite being somewhat technically challenging to produce, these gradients are experimentally achievable.

The imaging technique as outlined in this chapter can therefore be applied to a real system, albeit with resolution that depends on the strength of the available magnetic field gradient. I shall therefore describe the experimental apparatus to create and study spinor BEC in the next chapter.

4

Spinor BEC apparatus

In this chapter I describe the apparatus used to create and study spinor condensates. Magneto-optical trapping and evaporation down to BEC is a standard procedure, covered in detail elsewhere [51, 137]. I therefore present only a brief overview of the experimental apparatus, followed by details on the unique contributions I have made to this apparatus. The interested reader is referred to the thesis of A. A. Wood [138] for further information on the design of the apparatus, choice of components, and the experiment staging required to produce a BEC.

4.1: The Bose-Einstein condensate factory

The layout of the condensate apparatus (Figure 4.1) follows the general design of Lin *et al.* [139], modified to include a large glass cell ‘science chamber’. Atoms are sourced from a 5 g rubidium-87 ampoule loaded into an oven, which is heated to 80°C and monitored by an interlock (see §4.3). This creates an effusive atom beam through a collimation tube towards a cold-cup [140], which both performs cyro-pumping and further collimation by capturing any atoms not travelling down the beam axis.¹

The atoms then enter a zero-crossing single-layer tapered Zeeman slower [141, 142], which creates a spatially varying magnetic field that counters the Doppler shift and enables a counter-propagating laser beam to slow atoms from 380 m/s to 20 m/s [143]. The atoms are then loaded into a six-beam Magneto-Optical Trap (MOT) based on variable-magnification beam-expanders [144] that create large-diameter MOT beams and increase the number of atoms caught in the trap. The MOT has a total power of ≈ 50 mW divided across its six beams, each with an approximately top-hat profile with 16 mm diameter. Typically the MOT captures $\sim 10^9$ atoms in a diameter of ~ 10 mm.

Following the MOT formation, the quadrupole field is deactivated and the atoms are polarisation-gradient cooled during an optical molasses cycle for 8 ms, reducing the cloud temperature to 31(3) μ K. The quadrupole gradient is then ramped up to 40 G/cm in 3 ms to catch the free-falling cloud in a magnetic trap, which is hybrid-loaded into a crossed-beam optical dipole trap formed by a KEOPSYS 1064 nm 20 W

¹ A pneumatic shutter blocks the atom beam unless the MOT is being loaded.

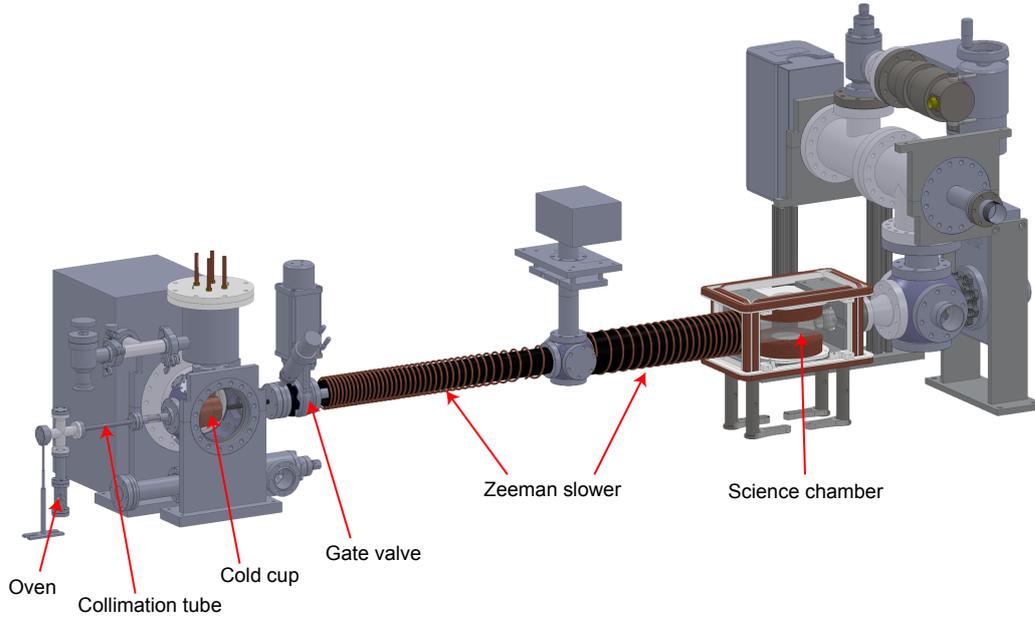


Figure 4.1: A CAD visualisation of the vacuum system showing the trajectory of the atoms from the oven (left) into the UHV section containing the science chamber (right). Optics and MOT assembly not shown.

single-mode, linearly polarised fiber laser. Finally, forced evaporation is carried out in the crossed-beam dipole trap by lowering the optical trap power to achieve condensation. This produces a BEC in the $|m_F = -1\rangle$ state, typically containing $\sim 3 \times 10^5$ atoms with trapping frequencies of $(\omega_x, \omega_y, \omega_z) = 2\pi \times (35, 60, 80)$ Hz.

Laser cooling and imaging requires beams at several different detunings (Figure 4.2). The ‘master’ laser is an EAGLEYARD RWE-0810 external-cavity diode laser (ECDL) in the Littrow configuration [145, 146] that generates seed light on the $|F = 2\rangle \rightarrow |F' = 3\rangle$ pump transition, which is locked using modulation transfer spectroscopy (MTS) [147, 148] for an extremely robust laser lock. The master laser seeds an M2K UC012 master oscillator power amplifier (MOPA) chip to produce ~ 1 W of 780 nm light [149], which is divided into several arms containing Acousto-Optic Modulators (AOMs) to frequency-shift each beam to the required detuning (Figure 4.3). Two repump lasers locked to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition are used to optically pump atoms out of the $|F = 1\rangle$ dark state. The Zeeman repump is locked using a digital offset-lock (§4.2) and the MOT repump is locked with saturated absorption spectroscopy [150].

Polarisation-maintaining fibers are used to transfer the light between different sections of the apparatus. This effectively decouples the different sections of the apparatus, preventing perturbations and accidents from misaligning the entire system. The single-mode fibers also have the additional benefit of cleaning up the poor spatial mode of the MOPA [151] by transmitting only the TEM_{00} component at the expense of losing the power contained in the higher-order spatial mode components.

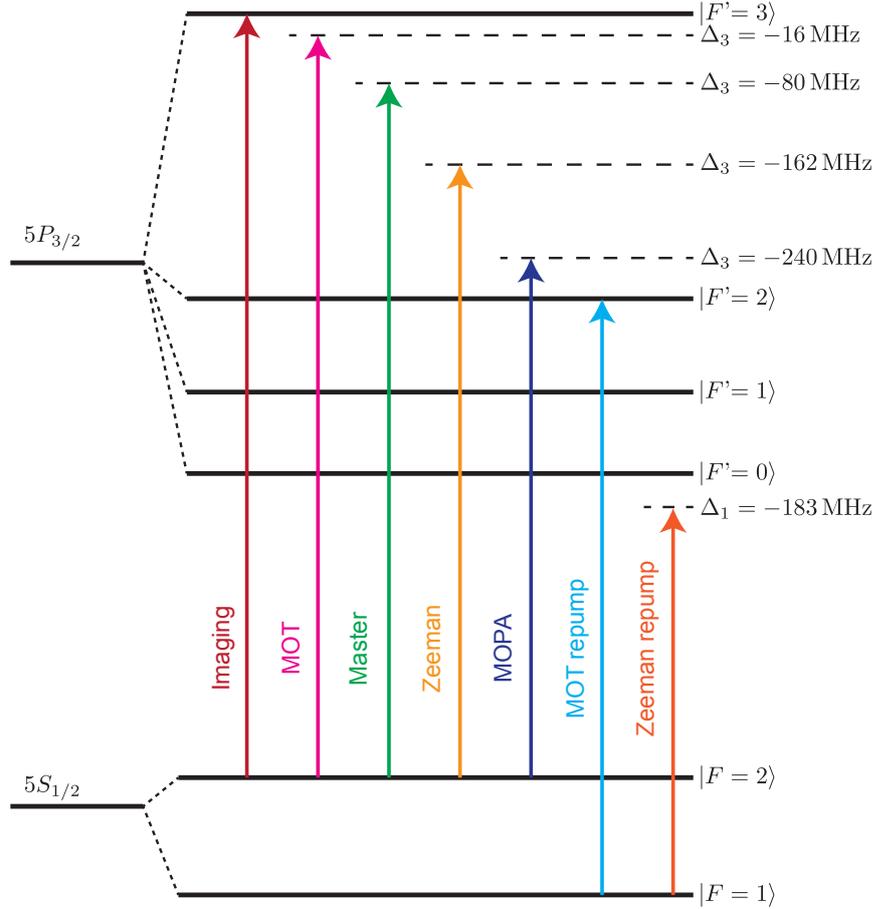


Figure 4.2: Laser detunings used to perform cooling and trapping of ^{87}Rb atoms, measured in hertz. The master laser seeds the laser amplifier (MOPA) from which the imaging, MOT and Zeeman are derived using AOMs. The repump beams are produced by independent lasers.

Magnetic field control is provided by a set of three orthogonal rectangle-based bias coil pairs, and a pair of quadrupole coils. The bias coils comprise 16 turns each in approximately Helmholtz configuration (Figure 4.4) and are operated by an in-house designed IGBT-based voltage-controlled current source called the ‘Mag-neat-o’. Each coil can be driven by up to 20 A, and steady-state bias fields of 10 G can be generated without requiring active cooling.

The quadrupole coils are wound from 42 turns of square-sectioned copper tubes with side length 4.76 mm and are internally water-cooled by a closed loop chiller.² Each coil is multi-layered, with inner/outer diameters of 60/126 mm and are separated by approximately 90 mm in anti-Helmholtz configuration. The coil pair is capable of generating up to 300 G/cm, though the MOT is typically formed in a gradient of 14.8 G/cm.

² A micropump is required to achieve the flow rate through both the quadrupole and Zeeman slower coils necessary to prevent significant heating; flow-rate sensors ensure the apparatus is not operated without active cooling.

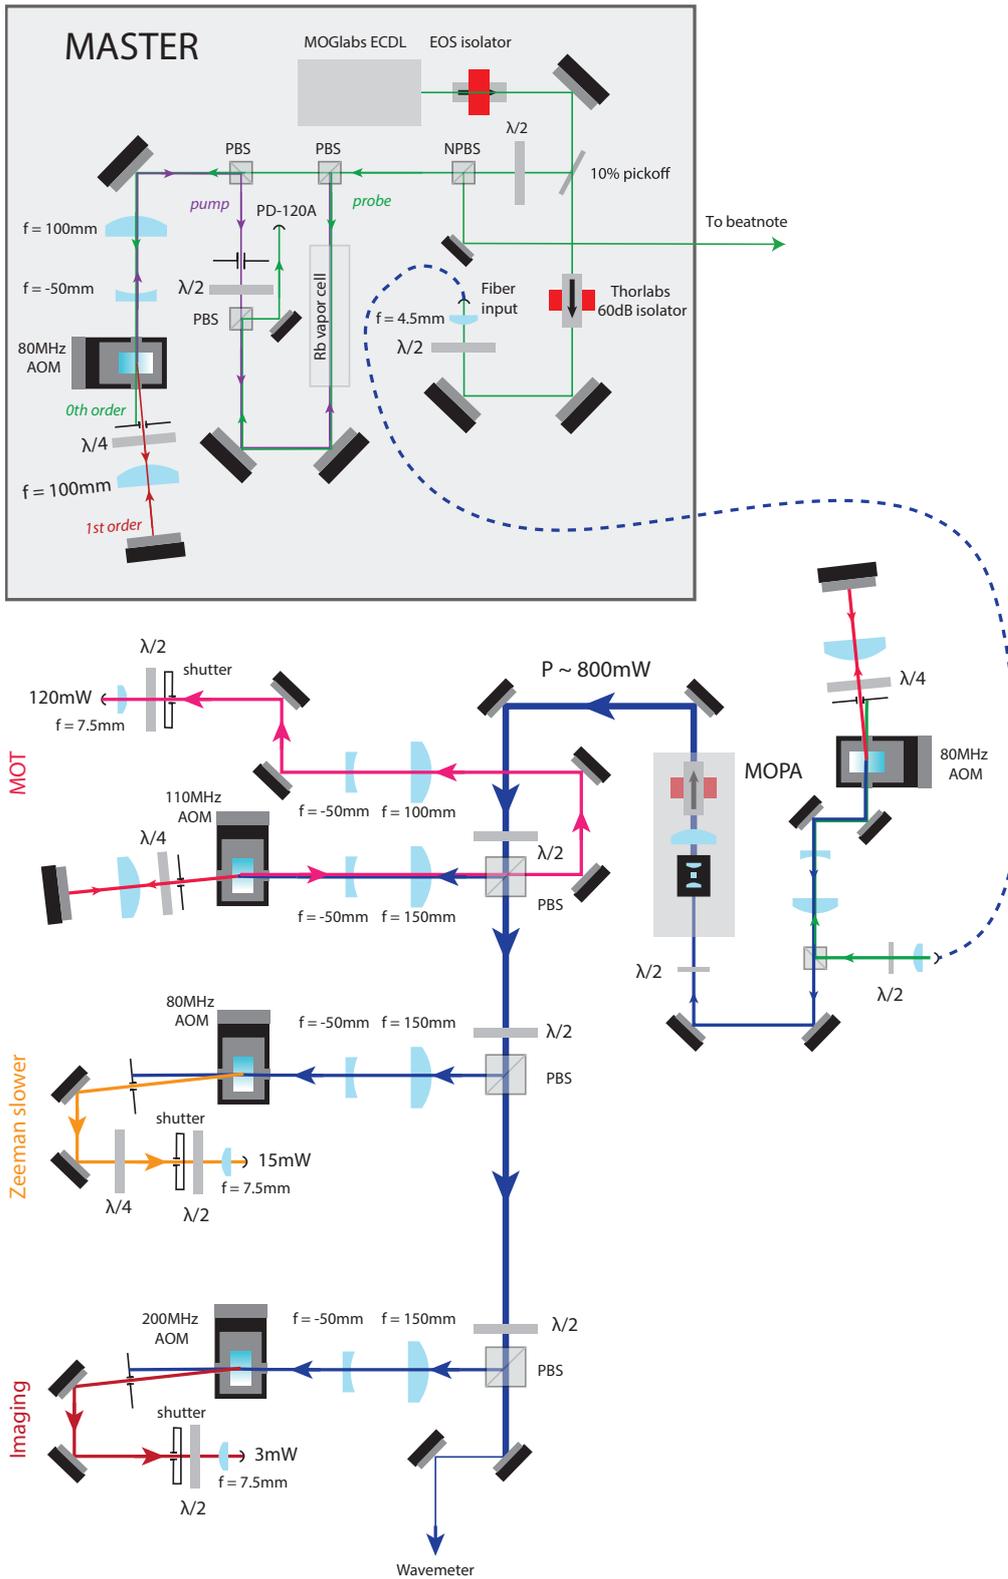


Figure 4.3: Scheme to generate the MOT, Zeeman slower and imaging beams using a single tapered amplifier. The master laser is locked using digital MTS and fiber-coupled to the MOPA. The MOPA output is divided into several arms, which use AOMs to frequency-shift the beams to obtain the required detunings. Adapted with permission from [138].

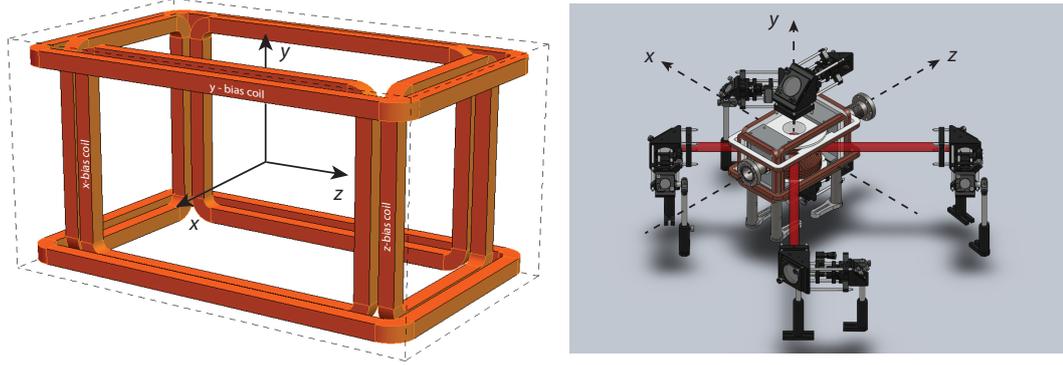


Figure 4.4: Design of the rectangular-base magnetic field bias coils (left), and arrangement of coils around the science chamber including MOT components (right). Inner coil dimensions are $(L_x, L_y, L_z) = (110, 157, 210)$ mm. See §8.1 for an explanation of the coordinate system. Reproduced with permission from [138].

Radio frequency pulses are generated by direct digital synthesis (DDS) with either a RfBLASTER³ or PULSEBLASTER DDS-II-300-AWG, and coupled to the atoms using a MINICIRCUITS LZY-22 amplifier connected to one of two radio antennae. The work presented in this thesis uses the ‘side’ antenna, which is a single-layer coil oriented in the y - z axis, comprised of 20 turns and having diameter 60 mm.

Microwaves are generated by a PHASEMATRIX FSW-0010 DDS, which operates at up to 10 GHz, and amplified by a MINICIRCUITS ZVE-3W-83X+ +35 dB rf-amplifier. The microwaves are coupled to the atoms with a custom half-wave dipole antenna designed for 6.834 GHz, located approximately 50 mm from the centre of the cell. Two series SKYWORKS SKY13298-360 switches allow the microwaves to be gated by a TTL trigger to prevent unintentional coupling between the hyperfine levels.

4.2: Beatnote microwave offset lock

Rubidium-87 has two hyperfine ground state energy levels, $|F = 1\rangle$ and $|F = 2\rangle$, with an energy level splitting between them of 6.834 GHz (Figure 4.5A). Typically only the $|F = 2\rangle$ state is used for cooling and trapping, but atoms promoted to an excited state can decay into either of the hyperfine ground states. Over time, atoms pumped by a cooling laser will randomly decay into the $|F = 1\rangle$ state, which is too far detuned to continue to interact with the pump laser, making it a ‘dark’ state.

A ‘repump’ laser is therefore required to perform the reverse process and optically pump atoms back into the $|F = 2\rangle$ state. However, since hyperfine splitting is large (6.8 GHz), the repump transition is not easily derived from the MOPA by modulation

³ The RfBlaster is a 2-channel DDS developed in-house, based on dual AD9910s controlled by a XILINX FPGA running PETALINUX and programmed in PYTHON using a custom compiler.

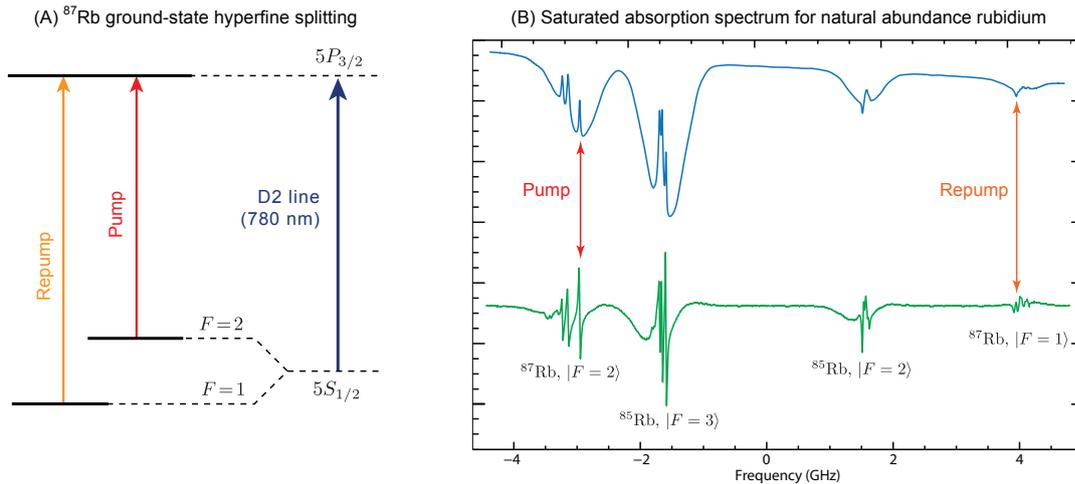


Figure 4.5: The ground-state hyperfine splitting of ^{87}Rb (A), and the D2 line of a natural abundance rubidium (B) showing the saturation absorption profile (blue) and derived error signal (green). Note the amplitudes of the error signal for the repump transition are significantly smaller than the pump transition.

techniques. Furthermore, the absorption signal from the repump transition is typically very weak (Figure 4.5B) so it is difficult to reliably frequency-lock a laser there using standard techniques like saturated-absorption spectroscopy [150]. Alternate techniques such as modulation transfer spectroscopy [148] have been demonstrated to provide a robust method of locking the pump laser [152], but are not as successful on the repump transition because of the lack of a cycling transition.

An alternate method to locking the pump and repump lasers independently is to lock the pump to the cooling transition using an atomic reference, and then lock the repump to the pump using a beatnote. The frequency separation between the pump and repump should match the known hyperfine splitting of the isotope, so a large single-featured error-signal can be derived by comparing the frequency of the beatnote to the desired separation [153].

As the beatnote frequency (6.8 GHz) is in the microwave regime, it is necessary to use a phase-locked loop (PLL) synthesiser to divide the beatnote frequency down to the rf-domain for signal processing. The divided signal can then be compared against a stable reference at known frequency (such as GPS clock), producing an error signal for reliable laser locking (Figure 4.6).

In this scheme, a small fraction of the each of the two lasers is extracted with a PBS cube, which are then combined on an NPBS to generate a beatnote on a HAMAMATSU G4176 fast photodetector, which has a 30 ps rise time at 7 V bias [154]. The photodetector is powered via a MINICIRCUITS ZX85-12G-S+ bias-tee with 9 V DC input, and the beatnote signal is amplified by five MINICIRCUITS ZX60-8008E-S+ +8 dB amplifiers to boost the rf power to ~ 0 dBm when 4 mW is incident on the photodetector. The re-

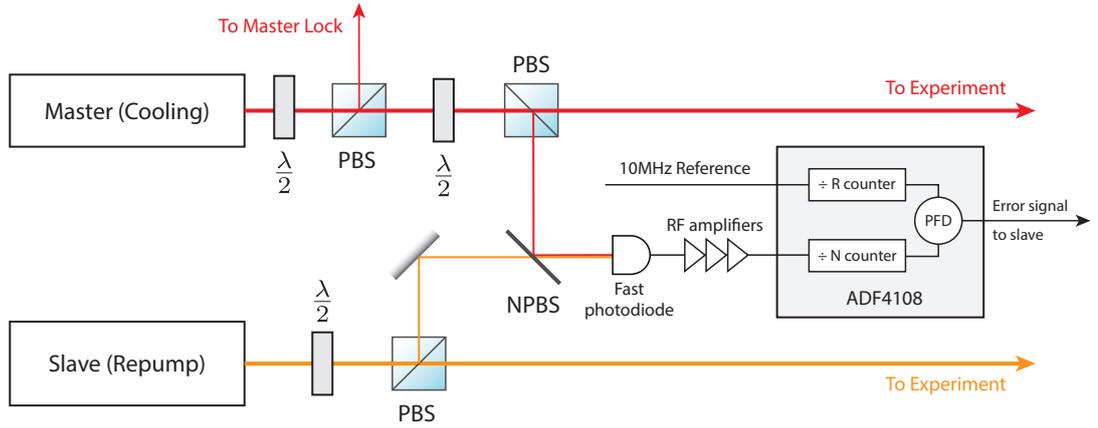


Figure 4.6: Schematic of the digital offset-lock. A fraction of the pump and repump lasers are separated with polarising beamsplitter (PBS) cubes to form a beatnote on the non-polarising beamsplitter (NPBS) as detected on the fast photodiode.

sulting beatnote is processed by the PLL and compared to a stable reference frequency by a phase-frequency detector (PFD), which produces an error signal that quantifies the difference between the measured (f_{in}) and desired (f_0) beatnote frequencies. The integrated circuit chosen for this task was the ANALOG DEVICES ADF4108 [155], for its high rf bandwidth (8 GHz) and availability in a premade evaluation-board (ANALOG EVAL-ADF4108EBZ1) to provide an interface to the chip.

The operational principle of the PLL device is to take the input f_{in} and reference f_{ref} signals and use counters to frequency-divide them by N and R respectively, where N and R are chosen such that $f_0 = N f_{ref}/R$. The resulting divided signals are compared, and if the frequencies are different then the PFD output is set to either 0 V or 5 V depending on which frequency is higher. However, if the divided frequencies are the same, the PFD outputs a voltage between 0 V and 5 V that is proportional to their difference in phase. Comparing the divided frequencies in this way is equivalent to comparing the input signal f_{in} to the target frequency f_0 and outputting a voltage proportional to their frequency difference that rails when the difference is more than the ‘channel spacing’ $f_{pfd} = f_{ref}/R$.

Because the beatnote frequency is so high, the N counter itself comprises a pre-scaler P and two programmable counters A and B that have the net effect of dividing the input by $N = BP + A$. There is a lot of freedom in selecting the counter values, but empirically it is best to choose the counters to limit phase noise, which is multiplied by the PFD at a rate of $20 \log N$ [156]. This implies that the counter values should be chosen to be as small as possible, subject to the operational requirements of the ADF4108 [155],

$$A < 64, \quad 2 < B < 8192, \quad f_{in}/P < 300 \text{ MHz} \quad \text{and} \quad N > P(P - 1).$$

A simple algorithm that chooses the smallest values for the counter parameters for a given desired frequency f_0 and reference frequency f_{ref} is

$$P \geq \left\lceil \frac{f_0}{300 \text{ MHz}} \right\rceil, \quad R = \left\lceil \frac{f_{\text{ref}}}{f_0} P(P-1) \right\rceil, \quad N = \left\lceil \frac{Rf_0}{f_{\text{ref}}} \right\rceil, \quad B = \left\lceil \frac{N}{P} \right\rceil, \quad A = N - BP,$$

where P is the smallest power of 2 that obeys the above inequality, and values are computed in order from left to right.

The possible values for the desired frequency f_0 are discretised into multiples of the channel spacing $f_{\text{pfd}} = f_{\text{ref}}/R$. In choosing R to be small to limit noise, the channel spacing is increased and the possible values of f_0 become more coarsely spaced. However, for the repump transition it is not critical to be exactly on resonance, and being detuned by a few MHz from the optimal frequency is acceptable. In particular, the offset lock was applied to the Zeeman repump laser which is off-resonant with the zero-velocity repump transition of the MOT, instead countering the Doppler shift for atoms exiting the oven. The beatnote frequency (Table 4.1) was optimised empirically by observing MOT load rates for a range of possible detunings. It should be noted that the offset lock has the unique advantage of being able to adjust the counter settings to arbitrarily set the lock point, and is not subject to the low efficiency and limited frequency range of modulation techniques (e.g. AOMs), or properties of the atomic transition (e.g. DAVLL [157]).

Parameter	f_0	P	R	A	B	N	f_{pfd}
Value	6305 MHz	32	2	13	39	1261	5 MHz

Table 4.1: Register settings required to lock the Zeeman repump laser using phase-frequency detection against a stable 10 MHz reference signal.

The PLL-synthesiser is controlled by settings stored in three 24-bit registers, which define the operation of the internal counters and the functionality of the PFD and are programmed using a 3-wire serial peripheral interface (SPI) through a DB-9 connector. Although programmable by computer via the parallel port using the manufacturer’s software, it is beneficial for the lock to operate independently of workstation computers. For this purpose, the ETHERNUT V2.1B micro-controller was chosen to form a bridge between the chip and user, constructing the register settings and programming the chip via a human-friendly interface over the local Ethernet. This required writing a custom TCP-server in the Ethernut’s NutOS operating system, calculating the required latch settings and ‘bit-bashing’ output over the board’s general-purpose I/O.

The range of the signal output by the PLL is 0–5 V, which is beyond the input specifications of the laser controller. To interface the error signal from the PFD with the laser controller, a signal processing circuit was constructed (Figure 4.7). The 0–5 V PFD output from the charge-pump (CP) pin is fed through a gain-2 ‘offset’ amplifier that level-shifts the input to be symmetric about zero. The subsequent variable-gain

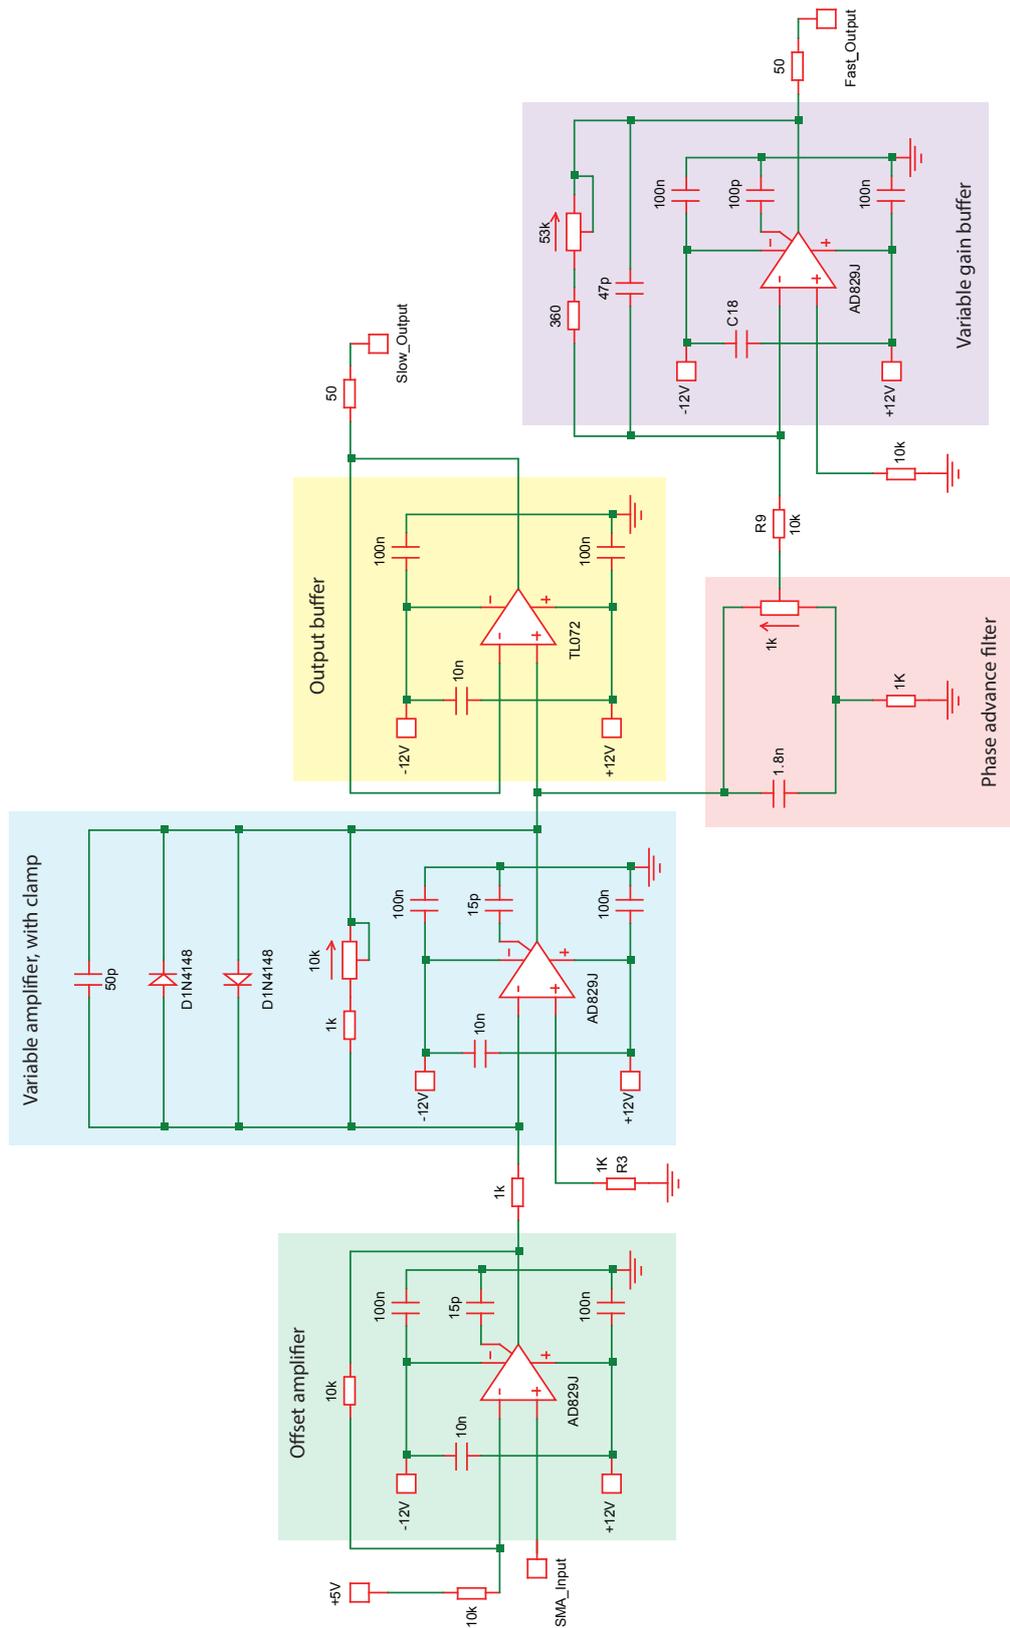


Figure 4.7: The signal processing circuit that level-shifts, amplifies and clamps the output from the PFD to produce a 'slow' feedback signal for the laser controller and a 'fast' feedback signal for direct injection into the laser head.

amplifier includes a protective clamp in the feedback arm to limit the range of the amplified error signal and prevent the laser entering oscillation. Two output signals are generated, one for feedback directly into the laser controller for PID locking ('slow' feedback), and the other for modulation of the laser diode current itself ('fast' feedback), for reasons which will be discussed presently.

The output of the phase-frequency detector rails when the frequencies being compared are different by more than the 'channel spacing' of the divider, typically 5 MHz. This is a tiny fraction of the beat frequency (6.8 GHz), or the frequency range of a laser scan (~500 MHz), so for a scanning laser the output changes rapidly when the laser passes the lock point, producing a square-wave error signal. Using this error signal directly therefore produces 'bang-bang' oscillations, whereby the extremely high gain of the error-signal causes the controller to continually overshoot the locking-point.

This is caused by insufficient servo bandwidth in our laser controller, making it necessary to construct a signal-processing board to produce a 'fast'-feedback signal for direct injection into the laser headboard to bypass the slow servo loop of the controller and obtain a lock. To prevent damage to the diode, the control signal was not added directly to the diode current; rather a FET-modulation technique [153] was used to drain a small amount of current away from the diode, in proportion to a control voltage derived from the PFD (Figure 4.8). As the laser scans, its frequency ramps up and the beat frequency increases. When the beatnote hits the desired value, the PFD output jumps, changing the FET base voltage and bleeding current away from the diode and reducing the laser frequency again. This has the effect of bringing the laser frequency closer to the desired detuning, linearising the step function across a scan sweep (Figure 4.9). Hence the overall gain of the frequency comparator is reduced, enabling the 'slow' locking servo to converge on the zero point.

Correctly choosing the degree of fast feedback is very important in this circuit; too much gain produces overshoot and causes the laser frequency to oscillate rapidly, as can be observed by monitoring the beatnote on an rf spectrum analyser. Optimal locking (with minimal linewidth) is obtained when the injection gain is maintained just above the point at which the fast-feedback enables the slow laser controller to lock (Figure 4.10).

4.3: Industrial reliability with microcontroller interlocks

The design philosophy of the apparatus was that it should be able to operate in the absence of human intervention for several hours at a time; ideally capable of performing experiments overnight without supervision. An important part of such independent operation is a system of interlocks that monitor the state of the apparatus, and *automatically* take action in the event of a systems failure to render the apparatus safe.

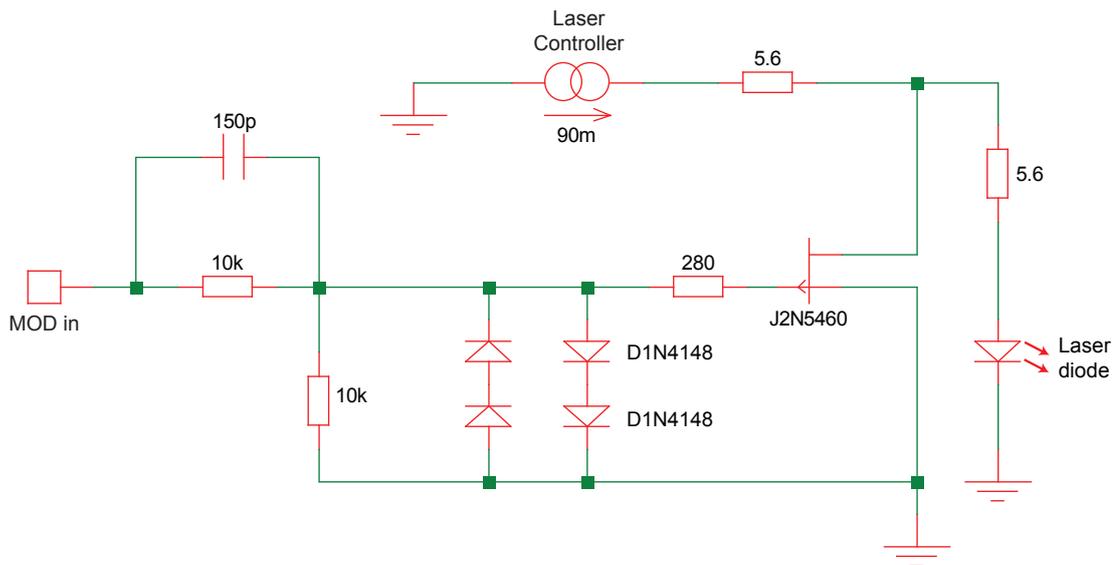


Figure 4.8: The FET-modulation circuit to enable fast-feedback by using a control voltage to reduce the laser current, adapted from [153]. Applying a positive voltage to 'MOD in' causes the JFET to bleed a small amount of current away from the diode laser, which modulates the lasing frequency and closes the offset-lock servo loop. This is termed 'fast' feedback as it bypasses the laser controller and has a high modulation bandwidth. Since the feedback only reduces the current through the diode, there is no danger of accidentally damaging the laser by overloading it.

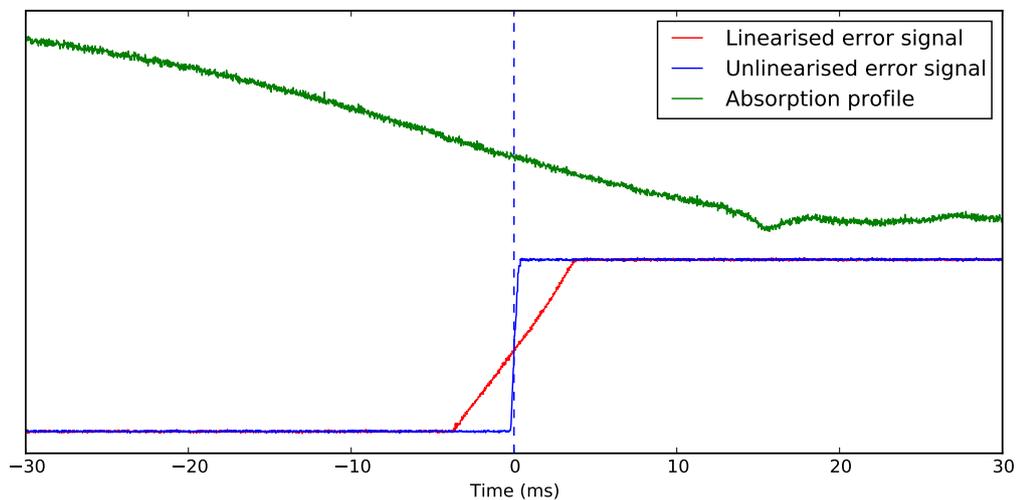


Figure 4.9: The PLL-synthesiser compares the beatnote frequency to a desired value, producing a sharp step response as the laser scans past. Engaging the fast-feedback circuit has the effect of linearising the error signal as the frequency is scanned, reducing the overall gain of the feedback loop and producing a slope that the laser controller can lock to. A saturated absorption spectrum is provided for reference.

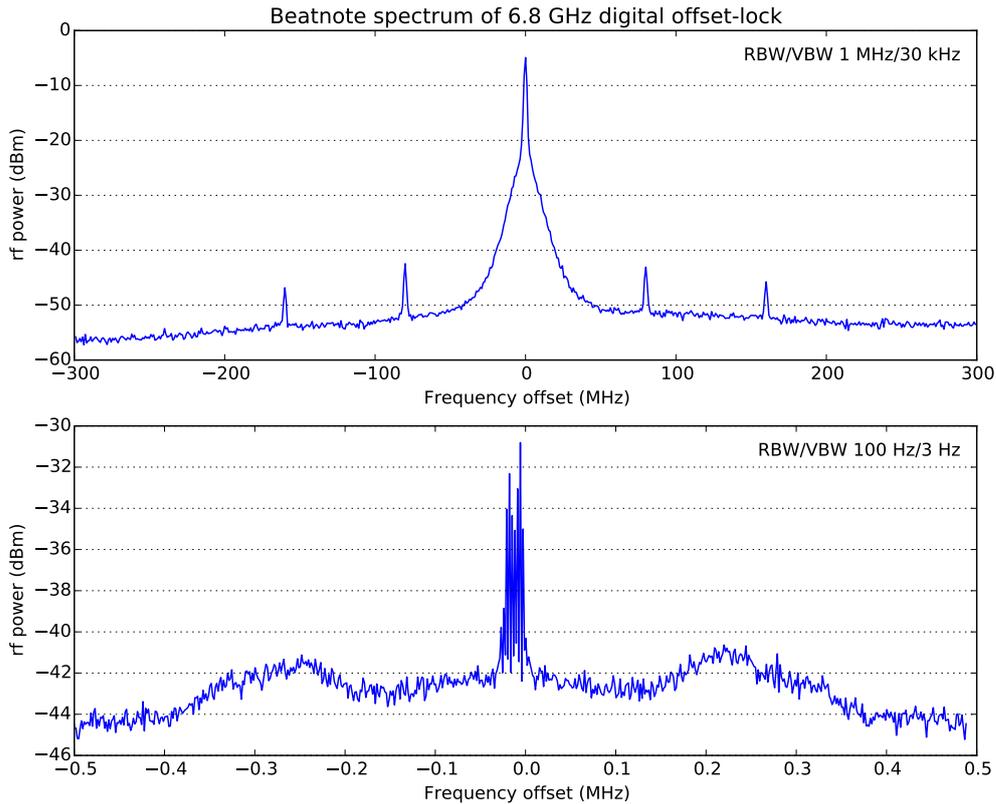


Figure 4.10: Rf beatnote spectrum of laser locked using digital offset-lock at 6.8 GHz. The width of the central feature is less than 20 kHz, and the ‘servo bump’ indicates a bandwidth of at least 100 kHz.

Rubidium is a highly reactive alkali metal, and as such the side of the vacuum system containing the rubidium ampoule is separated from the science chamber by a pneumatic gate valve, which can be sealed in an emergency. An independent controller operating as a watchdog is capable of continuously monitoring various sensors to determine when an emergency occurs and closing the gate valve to seal the source to protect the rest of the chamber. Careful control of the oven temperature also ensures that the atom flux remains constant, reducing atom number fluctuations in the BEC.

In particular, failure of either Peltier heat exchanger or the water cooling system would prevent the cold-cup collecting atoms, and instead cause previously trapped atoms to be released. Such a failure would prevent heat being removed from the hot side of the Peltier, leading to runaway heating of the Peltier and hence the cold-cup. This would generate a rubidium ‘fountain’: polluting the vacuum system, damaging sensitive vacuum equipment (such as the ion pumps) and necessitating an expensive and time-consuming rebuild of the vacuum system. Of the two other research groups around the world using this design, both have reported experiencing this problem to some degree because of a cooling system failure.

To automate control and monitoring of the oven-system, a custom control system based on the industrial GALIL RIO-47100 programmable-logic controller (PLC) was devised. The RIO controller was chosen for its large number of IO-ports, automatic and independent PID loops, and Ethernet connectivity. A finite state-machine ([Figure 4.11](#)) was written in the custom Galil ‘DMC’ language used by the RIO devices, to ensure that the oven controller remained in a well-defined state at all times, ensuring no dangerous actions were unintentionally initiated. The built-in PID loops were applied to control the oven heaters, enabling fast warm-up cycles and reliable temperature control to produce a consistent output flux from the oven. Details on the implementation are presented in [Appendix A](#).

Various failure modes were considered, ranging from minor (e.g. non-critical sensor disconnection) to extreme (e.g. vacuum failure, coolant failure, runaway PID loops). Upon detection of an error, the machine enters a ‘fail-safe’ mode, in which all potentially dangerous devices are deactivated by isolating the chamber (sealing the gate-valve), stopping the oven (heaters and Peltiers) and alerting the user (siren, email and SMS alerts). Physical user intervention is required to restart the machine after a failure, ensuring the machine never leaves a safe ‘shutdown’ state on its own without human intervention to ensure it is safe. Relays on the power supplies to the oven heaters mean that even in the event of sudden power-failure to the controller, the oven is rendered safe.

For monitoring purposes, the controller communicates with a logging server that emits SYSLOG-compatible messages. Log messages are sent whenever the machine changes state, as well as periodically to provide updates on the temperature of the oven and a number of attached gauges. These messages can then be processed by a number of standard SYSLOG tools for remote viewing, monitoring and graphing, as well as post-hoc diagnosis of emergency events. The latter has been essential for identifying unforeseen failure modes, which were then accounted for in the state machine.

A number of other interlock devices were also constructed; one to turn off the 120 A Zeeman slower power supplies in the event of water cooling failure, and another to prevent damage to the MOPA by being run at too high a current, or running without a seed beam. Either of these two disaster scenarios would cause significant damage to the apparatus and is enough to significantly set-back the project; they therefore warrant their own fail-safe watchdogs.

The interlocks successfully acted in multiple cases of coolant and power failure, rendering the apparatus safe, or alerting personnel that intervention was required. Multiple incidences were a result of malfunctions with the closed-loop water chiller, causing sharp rises in the temperature of the cold-cup, and hence pressure of the vacuum system. Damage was avoided in these cases thanks to the early warning provided by the interlocks.

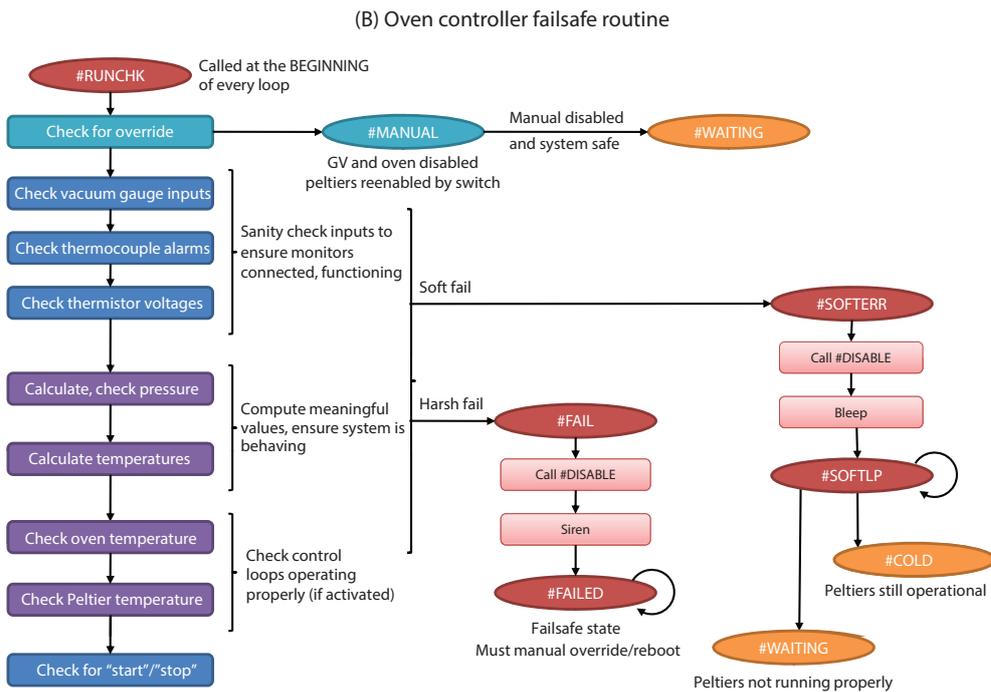
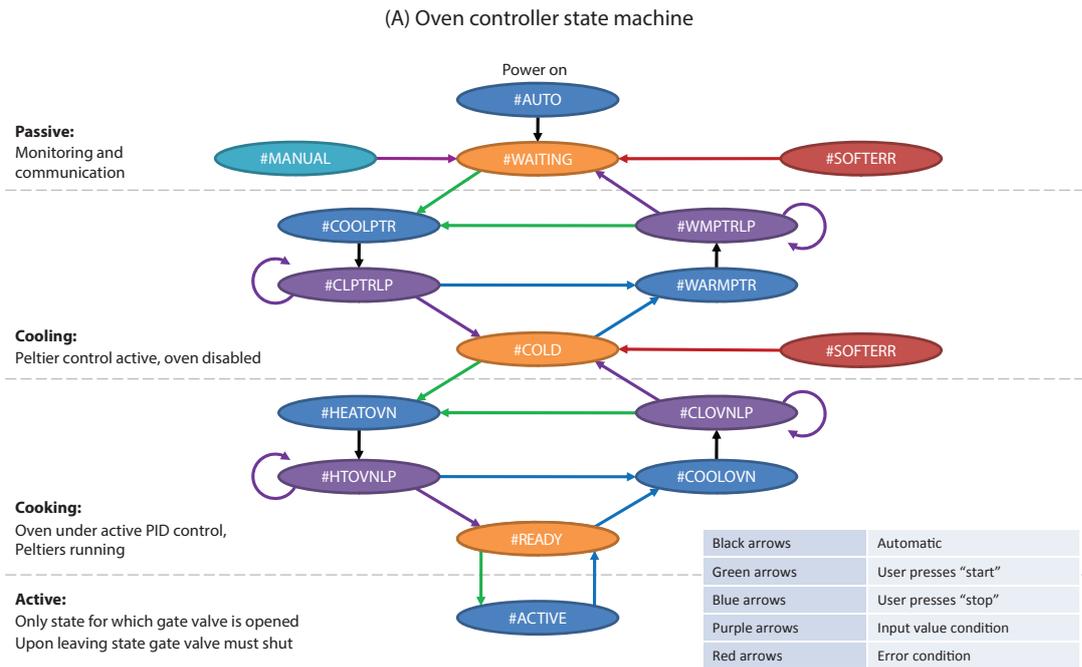


Figure 4.11: The controller state-machine (A) keeps the apparatus in a well-defined state. A variety of sensors are continually checked (B) to detect an error, entering one of two fail-safe modes depending on severity. User intervention is explicitly required to resume operation from a harsh fail, preventing the oven from entering a potentially dangerous state without diagnosis.

4.4: BIAS: A modular image acquisition system

The experimental apparatus is operated by a modular set of control system programs developed in-house, called the LABSCRIPT SUITE [158]. The labscript language provides a set of extensions to standard python that enables experiments to be described in terms of a series of events, abstracting away the low-level implementation details about how devices are programmed or connected. The script is compiled to hardware instructions that are stored in an HDF5 file [159] that encapsulates all data relating to a single realisation of the experiment (a ‘shot’). This hierarchical scientific data format enables experimental inputs (script, parameters, settings) to be stored with outputs (acquisitions, images, debug information) for each time the experiment is run, automatically keeping a detailed log of past experimental runs.

The suite is divided into several applications, each performing a different function (Figure 4.12). This modular design enhances separation of high-level experimental logic from low-level device implementation details by allowing modules to operate at different levels of abstraction. The individual components communicate over a network connection, allowing different applications (and different hardware) to be run from different computers and improving real-time diagnostic abilities.

Another advantage of the modularity is that it enables ‘secondary’ control programs to be used to communicate with specific devices when software to do so exists and has been written and debugged. In particular, this is valuable for software written in another programming language that would take significant time to rewrite and debug as part of the core suite. One such module is the program responsible for controlling camera hardware and providing interactive image capture, the BEC Image Acquisition System (BIAS).

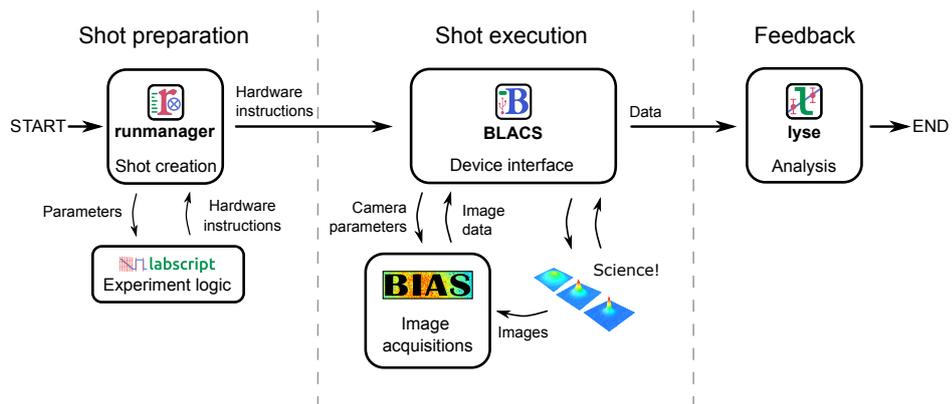


Figure 4.12: Interaction between modules of the Labscript suite. An experiment is defined in labscript logic, which is converted to hardware instructions by runmanager. BLACS coordinates experiment execution and manages hardware programming, interfacing with secondary programs like BIAS. Once a shot is complete it is post-processed and analysed by lyse. Adapted from [158].

The primary purpose of BIAS is to take hardware-agnostic instructions (such as exposure time, number of frames to capture, regions of interest) and to program attached cameras in an interchangeable way, translating instructions into commands accepted by the particular software development kit (SDK) in use. BIAS also performs image processing tasks such as background subtraction, saturation correction, optical depth calculation, and simple 2D fitting (Figure 4.13).

As per the rest of the LABSCRIPT SUITE, this is achieved with object orientation [158]. An abstract base class ('Camera') defines an interface to a generic device, and child classes extend this interface with SDK-specific implementation of instructions, and manufacturer-specific special cases where necessary. The net result is that once implemented, swapping between camera types and toolkits is seamless in the experimental script.

BIAS was written in LabVIEW to take advantage of the native image processing and display functionality, and availability of hardware drivers from many different manufacturers. Sequences of images can be captured and displayed, regions of interest highlighted, and calculations such as optical depth or profile fits performed (Figure 4.14). Regions of interest can be selected for capture or fitting, to reduce camera readout time or computation time.

The LABSCRIPT SUITE uses the HDF5 file format to encapsulate all instructions, data and metadata for a single shot in a single hierarchical file. At the time of writing the application, there was no native support for HDF5 in LabVIEW, and only a slow legacy library existed to provide an interface. Having determined that reading and writing large quantities of data was a bottleneck for operation, I created a fast new language bindings to the HDF5 file format. Termed h5labview [160], it was released as an open-source package and has been developed to provide a robust and fast interface

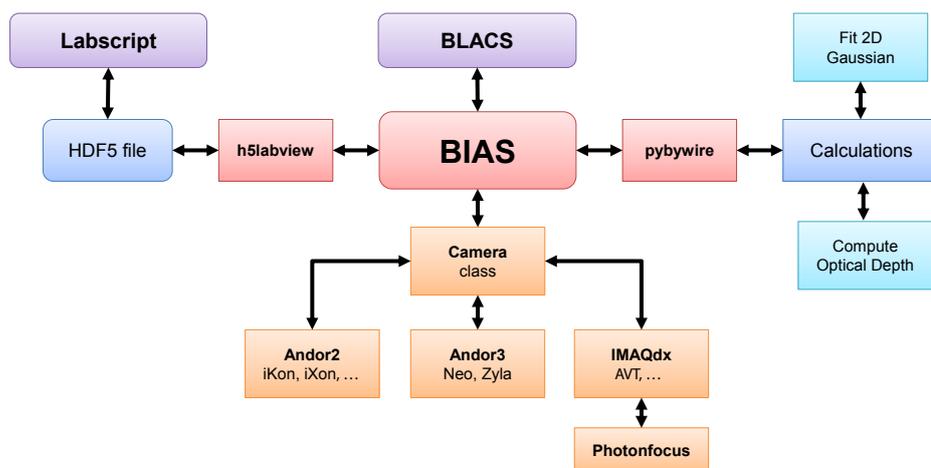


Figure 4.13: A simple overview of the components of BIAS and how it integrates with the rest of the LABSCRIPT SUITE of programs.

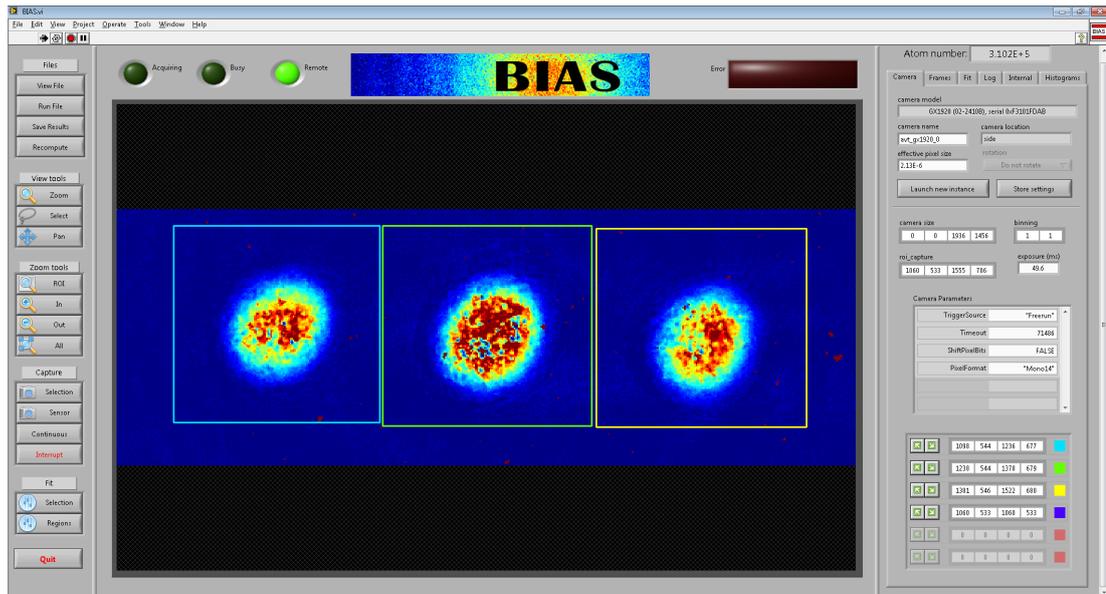


Figure 4.14: Screenshot of BIAS in operation, showing the three spin components of a BEC Stern-Gerlach separated. Multiple regions of interest have been defined to allow each component to be subsequently analysed.

to storing large hierarchical datasets in LabVIEW. It has since acquired a broad user base, and is used for scientific data storage by several groups worldwide.

Performing non-trivial calculations on large datasets (such as 2D curve-fitting) is also cumbersome in LabVIEW. The rest of the LABSCRIPT SUITE is written in python, which has a rapidly growing code-base of scientific and numerical code. To leverage this potential and enable code-sharing with the lyse analysis routines, I wrote the pybywire library [161] to enable the execution of arbitrary python code from LabVIEW. It uses ZeroMQ [162] for reliable data transfer and MessagePack [163] for serialisation, for which I also created LabVIEW interfaces [164, 165].

BIAS provides an extensible platform for image capture, display and analysis. Its object-oriented interface permits camera interoperability, and its data input/output routines enable it to both integrate with the wider LABSCRIPT SUITE and execute scripts in other languages.

4.5: Summary

In this chapter I described the apparatus we constructed for the creation and study of spinor BECs, with the goal of reliable and autonomous execution of experiments. An overview of the apparatus design was given, with specific details on the aspects of our implementation most relevant to the proposed experiments provided. Several components of the wider system were described in depth, comprising the parts of the system I developed.

A digital offset-locking technique was developed and implemented to lock the Zeeman repump laser at an arbitrary detuning using a phase-frequency detector to analyse the beatnote with another laser. The single-featured error signal generated this way is robust, and by digitally adjusting the counter settings the lock point can be changed. This enabled the detuning to be optimised to improve the efficiency of the Zeeman slower.

A fail-safe oven controller was designed and constructed using a programmable logic controller to operate the oven and continuously monitor its state, automatically taking action to render the oven safe in the event that a dangerous situation is detected. Water cooling failures have triggered the fail-safe on multiple occasions, and the controller has prevented damage to the apparatus by taking immediate action.

A modular imaging application was developed to capture and process data from a number of different camera types, presenting a unified interface and integrating with the rest of the LABSCRIPT control suite used to operate our apparatus and execute experiments. The techniques and interfaces developed to exchange data with HDF5, ZeroMQ and python were packaged into a number of open-source projects which have developed a wide user base.

The combination of independent watchdogs and the LABSCRIPT SUITE achieved the goal of automated operation, enabling long parameter-space scans to be conducted. For example, a sequence of 485 experiments were automatically executed to measure Rabi-cycling between the Zeeman substates of the BEC using absorption imaging. The sequence ran uninterrupted for 10 hours, with only 4 of the shots unsuccessful. The limiting factor on unattended data collection are the quality of the laser locks⁴ and attached devices becoming unresponsive.⁵ Work to address both of these issues is ongoing.

⁴ The MOT repump laser is presently locked using saturated absorption spectroscopy and typically unlocks every few hours. A digital-offset lock (as implemented on the Zeeman repump, see §4.2) would improve this behaviour.

⁵ Some devices such as the RfBlasters would enter an undefined state and required power-cycling to resume operation. This is detected by the hardware interface program BLACS, pausing operation until normal function can resume.

5

Faraday probe beam

In this chapter I describe experimental considerations regarding the Faraday probe beam: how it is generated, how the BEC is absorption imaged in-trap, how alignment is achieved in the presence of limited optical access, and the effect of the probe laser on the lifetime of the trapped BEC. The photodetection apparatus and signal processing methods used to capture and analyse the Faraday signal are discussed in the next chapter.

5.1: Probe laser and locking scheme

The light used as the Faraday probe beam was generated by a EAGLEYARD EYP-RWE-0810 laser diode in a Hawthorn external-cavity configuration [166]. The laser was controlled by a MOGLABS DLC-202 [167] and tuned to lase at 790 nm, generating 110 mW of light (before the optical isolator) with an injection current of 200 mA. The laser is located on a table separate from the vacuum system for modularity and mechanical isolation (Figure 5.1). A pick-off polarising beamsplitter (PBS) cube is used to sample some beam to measure its wavelength and provide an error signal to be fed back for locking, with the majority of the power fiber-coupled from the laser table to the experiment with a single-mode polarisation-maintaining fiber.

To prevent the atom cloud being exposed to the probe beam before the measurement begins, the beam is shuttered by both an AOM and a mechanical hard-drive shutter [168]. The mechanical shutter prevents the need to turn off the AOM for extended periods of time, which would cause the AOM to cool down and change diffraction efficiency over the timescale of the experiment, causing drift in the intensity of the probe beam.

A second beam containing resonant D2 (780 nm) light is simultaneously coupled into the orthogonal polarisation mode of the same fiber. This light derived from the MOPA output enables diagnostic absorption imaging to be carried out on the BEC using the same optical beam path as the Faraday beam, which aids in alignment (§5.5).

Multi-order 780 nm polarisation optics perform poorly with 790 nm light, so in the absence of specific 790 nm polarisation optics, zero-order 780 nm waveplates must be used. In particular, it is not possible to have perfectly linearly polarised light incident

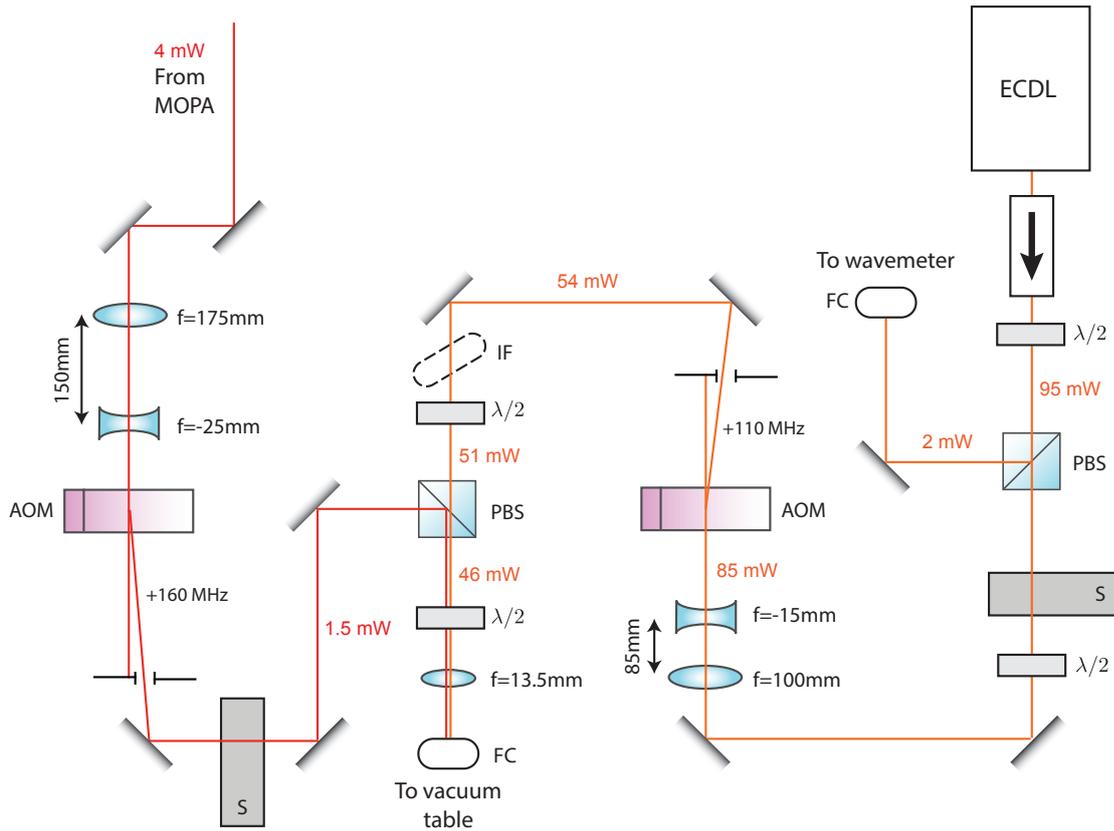


Figure 5.1: Experimental setup to produce Faraday imaging light, showing typical beam powers. 790nm light from the Faraday ECDL (orange) is shuttered by a mechanical hard-drive shutter (S) and acousto-optic modulator (AOM) before being combined on a polarising beamsplitter cube (PBS) with similarly shuttered resonant light derived from the MOPA (red) and fiber coupled (FC) to the vacuum table. The role of the interference filter (IF) is explained in §5.8.

on the fiber coupler whose polarisation axis aligns with the axis of the fiber, as the $\lambda/2$ plate induces some ellipticity to the polarisation. This has consequences for polarisation stability (§5.3) at the fiber output, resulting in long-term drift in the polarisation state of the output beam (on the timescale of minutes).

Laser locking techniques typically use an absorption line in a reference vapour cell to generate an error signal to servo the laser controller and fix the wavelength. Some techniques such as offset-locking (§4.2) can be used to lock a laser at some detuning from a reference. However, by choosing to operate at the inter fine-structure magic wavelength of 790 nm (§2.4), the nearest absorption feature is detuned by several nanometers and the beatnote frequency would be in the order of terahertz, making offset-locking impossible.

Instead, a HIGHFINESSE WAVELENGTH METER WSU-10 was used to measure the wavelength of the Faraday laser, at a repetition rate of a few Hz. The unit has a DAC output board which performs PID on the measurement and generates an analog error-

signal output for the laser controller. Although this is a low-bandwidth servo-loop, the atoms are relatively insensitive to the wavelength of the Faraday laser so it is not critical to eliminate the residual small fluctuations in laser wavelength.

The WSU-10 DAC control line was connected to the ‘sweep’ input of the MOGLABS laser controller with DIP13 set to ON, to bypass the PID built in to the laser controller and allow the wavemeter to directly control the piezo of the ECDL. The ‘Laser Control’ functionality of the HIGHFINESSE control program was configured to a sensitivity of 1 V/100 pm, positive polarity and gain settings of 0.1 proportional, 40 integral. The signal bounds were configured to output an error signal centred on 0 V, with limits at ± 1 V. A stable lock was achieved with an RMS wavelength fluctuation of 1 fm over 10 min, corresponding to a linewidth of 500 kHz. The lock point could be software-controlled by adjusting the ‘reference’ point by up to 1 pm, with the lock achieving convergence in < 1 s. Shifting the lock point by more than 1 pm typically caused mode hops and required adjustments to the laser current, but since the magic wavelength is known, the frequency does not need to be swept and this is not an issue.

5.2: Pre-cell optics

The Faraday signal is proportional to the total power of the probe beam that passes through the condensate. Any light not passing through the BEC contributes no signal to the measurement but adds further shot-noise. Not only does this light fail to contribute to the measurement, but it reduces the overall signal-to-noise ratio of the measurement by the inclusion of increased shot-noise. It is therefore important to focus the probe beam down to form a tight waist around the atom cloud and ‘shrink-wrap’ it to maximise the amount of light interacting with the atoms. Operating at the magic-zero wavelength means the highly focused, intense beam will not perturb the trap via a scalar light shift (§2.4), but the vector light-shift (§2.6) from the focused beam must be entirely eradicated.

The condensate is small, typically comprising 3×10^5 atoms in a Thomas-Fermi radius of 15 μm . The probe laser should have comparable waist to maximise the number of photons interacting with the atoms and generating the Faraday signal. Optical access limits how closely a lens can be placed to the science chamber: the closest optical access for a 2"-diameter focusing lens is 800 mm from the cloud. Based on available achromatic lenses, a 1000 mm lens was chosen and placed 1000 mm from the cloud position. To achieve a tight focus from such a long focal length lens, a large diameter beam must be incident on the lens, so a THORLABS BE10M beam expander was used to enlarge the probe beam immediately before the lens. This produced a focal spot size at the position of the BEC of $150 \times 125 \mu\text{m}$, as measured using a THORLABS BC106-VIS beam profiler and a fold mirror. The astigmatism is most likely introduced by entering the beam expander at a small angle of incidence.

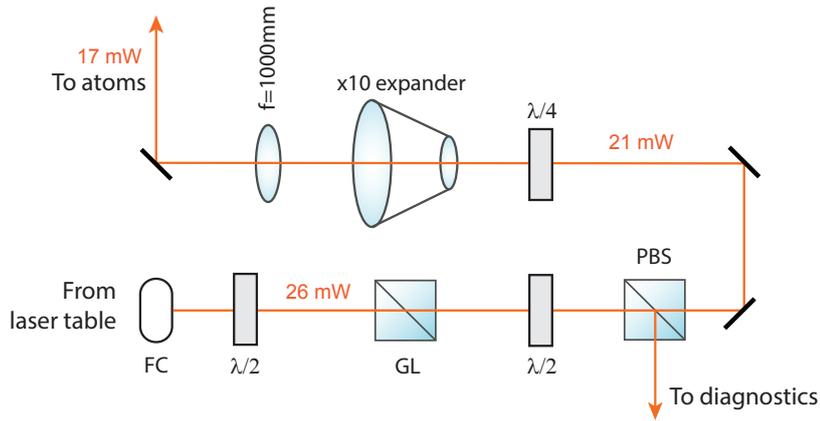


Figure 5.2: Conditioning optics to generate the probe beam, including typical beam powers. A Glan-Laser (GL) polariser purifies the polarisation state output by the fiber coupler (FC), and a polarising beam-splitter (PBS) picks off a small amount for diagnostics. The beam passes through a $\times 10$ beam expander then a 1000 mm lens to produce a focal spot inside the science chamber 1000 mm away.

Ideally the Faraday beam should co-propagate with one of the trap lasers to look at the cloud ‘end on’. This creates the most uniform intensity profile across the cloud and reduces the total optical power required to maintain a given probe beam intensity. In practice, the beam propagates at a small angle θ_F to the trap laser – which is limited by the optical access for steering mirrors before the cell and the pick-off mirrors after. The current configuration has $\theta_F \approx 3^\circ$.

The zero-order quarter-wave plate before the beam-expander allows control of the ellipticity of the probe beam polarisation. This is important as any circular component to the polarisation produces a vector-light shift on the atoms (§2.6). The polarising elements before the waveplate produce a linear polarisation state going into the science chamber, but the thick glass walls of the science cell itself induce significant birefringence on the beam as it travels through. Adjusting the polarisation with the waveplate allows the opposite ellipticity to be applied before the cell to produce linear polarisation at the position of the atoms, as characterised in §7.5.

5.3: Polarisation stability and fiber axis alignment

The Faraday beam is transmitted from the laser table to the vacuum table through a single-mode polarisation-maintaining (PM) fiber. The fiber has fast and slow linear polarisation axes, which prevent cross-coupling between polarisation modes as they travel through the fiber. However, unless the input polarisation is perfectly linear and polarised along one of the axes of the fiber, the two polarisation components will propagate at different speeds inside the fiber, resulting in an elliptical polarisation at the fiber output.

For a polarisation state that is not aligned with the fiber, the degree of ellipticity is sensitive to both the temperature of the fiber and any mechanical stress, both of which may vary during the course of an experiment. A misaligned input polarisation has been seen to result in a $> 50\%$ variation in PBS transmission in synchronisation with the laboratory air-conditioning cycle [151]. For polarisation-sensitive measurements it is therefore important to ensure the incident beam is linearly polarised with polarisation rotated to match one of the fiber axes.

Typically PM fibers are produced with a ‘key’ that defines the orientation of the fiber axes. However the fiber-couple mount itself is oriented in an arbitrary direction making it difficult to relate the polarisation of the incident light to the axes of the fiber. Furthermore, the width of the keyway reduces insertion repeatability enough to introduce large residual misalignment.

Axis alignment is achieved by using a PBS cube at the fiber input to purify the polarisation state, then a zero-order $\frac{\lambda}{2}$ -waveplate to rotate the polarisation axis to match the fiber. A $\frac{\lambda}{2}$ -waveplate followed by another PBS at the output then purifies the output state. The input waveplate is incrementally rotated, and at each stage the output waveplate is rotated to minimise the transmission through the cube. When a global minimum is found for the transmission through the cube, the ellipticity of the fiber output has been minimised, so the light is polarised ‘on-axis’ with the fiber. The result is checked by mechanically stressing the fiber gently and observing that the transmission does not change appreciably.

A change in transmission through the cube will be observed while pressure is applied, but should disappear again when the pressure is released. If not, the polarisation is not aligned with the fiber axis, but the net phase delay accumulated by the two modes makes the output state linear. However, since this phase is sensitive to pressure and temperature, it is not a stable solution and is sensitive to drift.

Ideally this alignment procedure is carried out with the input and output stages placed near each other and the optical fiber loosely coiled between, reducing strain on the fiber and producing a single intensity minimum that is easily located. However, applying the technique *in situ* to a fiber that is stretched across the lab is more difficult, as it will be under significant strain, causing $\gg 2\pi$ phase delay when the polarisation is misaligned with the axes of the fiber. This results in many ‘false minima’ as the waveplates are adjusted, depending on how many windings of the phase are occurring. In particular, if the output waveplate is adjusted to minimise leakage but the transmission is observed to jump to $\gtrsim 50\%$ of the beam power when strain is applied, then the polarisation is in the large delay regime. If possible, the fiber should be removed and coiled to reduce the accumulated phase, so that the above method can be used.

If alignment must be done *in situ*, sensitivity to applied pressure can be used as an alignment metric, since the accumulated phase is greater the further off-axis the input

polarisation is. The input waveplate should be slowly stepped through its range, the output waveplate adjusted to minimise transmission, and perturbation applied. When the jump is $\lesssim 5\%$ of the total probe power, the alignment can then be optimised by finding the global minimum in transmission as above.

This enables the alignment of the waveplate to within 1° , which is approximately the limit of the precision of a THORLABS CRM1 rotation stage. However, the residual polarisation fluctuations caused by being off-axis with the fiber by even this amount are not only significantly greater than the polarisation rotation due to the Faraday effect, but will cause fluctuations in the vector-light shift experienced by the atoms (§2.6). It is therefore critical to purify the polarisation state of the probe beam before the science chamber.

This is achieved using a high extinction-ratio Glan-Laser polarising beamsplitter cube (THORLABS GL5-CL26) as a ‘clean-up cube’ to project the polarisation onto a known axis, converting the polarisation fluctuations to intensity fluctuations in the transmitted beam. The polarisation rotation observed by the polarimeter is therefore entirely due to the Faraday effect, since the intensity fluctuations are common mode to the two inputs of the differential photodetector.

Ideally the polarisation optics would be located immediately before the science cell, so that the purity of the polarisation state is not affected by any subsequent optics. However, optical access restrictions around the cell and the large diameter of the beam after the beam-expander require that they be placed between the fiber and beam-expander, potentially introducing small fluctuations in polarisation. Although the polarimeter is sensitive, the signal being investigated occurs at the Larmor frequency, which is a significantly faster timescale than any such perturbations.

5.4: Post-cell optics and imaging system

The post-cell optics are designed to collect as many of the photons forward-scattered by the Faraday interaction as possible, create a magnified image of the BEC in trap, aperture the image to block light that did not pass through the BEC, then perform precise polarimetry on the resulting beam (Figure 5.3).

To have finer control over which light reaches the polarimeter, an adjustable iris instead of a pinhole was used to aperture the beam and control how much light reaches the detector. Shot-to-shot variations in pointing stability (see §5.6) result in perturbations to the centration of the BEC image, requiring the iris to be opened slightly further to prevent aperturing light that has passed through the BEC. A THORLABS SM1D12D ring iris was chosen for this task, which has a minimum aperture size of 0.8 mm. A high-magnification imaging system ($\sim 30\times$ net magnification) is therefore required to enlarge the image of the BEC (Thomas-Fermi radius $15\ \mu\text{m}$) to fill the aperture.

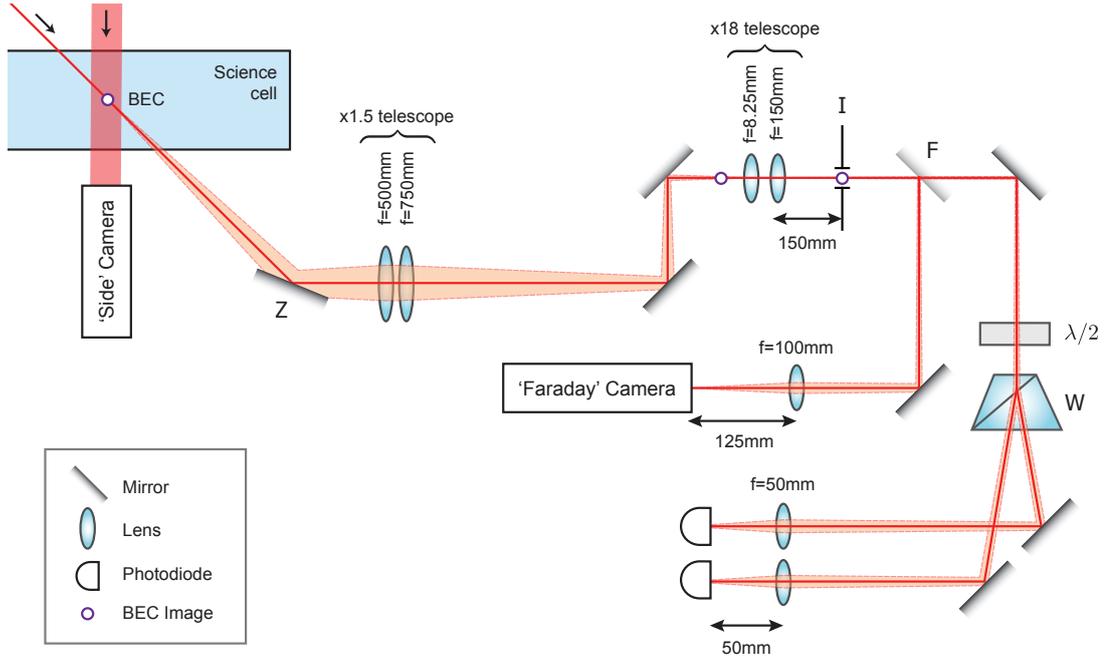


Figure 5.3: Experimental schematic for Faraday beam analysis, showing Abbe rays diffraction by the BEC. The transmitted beam is reflected off a 2" mirror on a ZABER translation stage (Z) into a pair of relay lenses, producing an image of the BEC (purple circle) in front of a microscope objective lens. The microscope magnifies the image and reimages it inside a translatable iris (I), which apertures the beam. The resulting beam is either reflected towards the camera by a mechanical flipper mirror (F) or analysed by the balanced polarimeter formed by the Wollaston prism (W) and dual-port differential photodiode. An independent 'side' imaging system captures TOF absorption images.

Ideally, a microscope objective could be placed directly after the cell to achieve high magnification with a high numerical aperture. Although there are designs for objectives with long enough working distances (> 25 mm) to observe the BEC inside the science chamber [169], the Faraday beam is approximately co-propagating with the dipole trap beam, so the objective cannot be closer to the cell than the dichroic mirror that dumps the dipole beam (200 mm from the BEC). Imaging the BEC is therefore achieved by using a pair of relay lenses to form an image at a more convenient distance from the cell, which can then be enlarged by an objective to achieve the required magnification (Figure 5.4). This lens combination magnifies the $15 \mu\text{m}$ radius BEC to a $200 \mu\text{m}$ radius image at the iris.

Optical access limitations are further complicated by the MOT beam which also co-propagates with the dipole beam. Placing the relay lenses close to the cell results in blocking the MOT beam, preventing a cold atom cloud from being formed. The solution is to either place the relay lenses sufficiently far from the cell so that they separate enough for the MOT beam to not be blocked, or to dynamically move a mirror in using a motorised flipper mirror (or translation stage) every shot, ensuring the MOT beam is only blocked after the MOT has been formed. While dynamically moving the

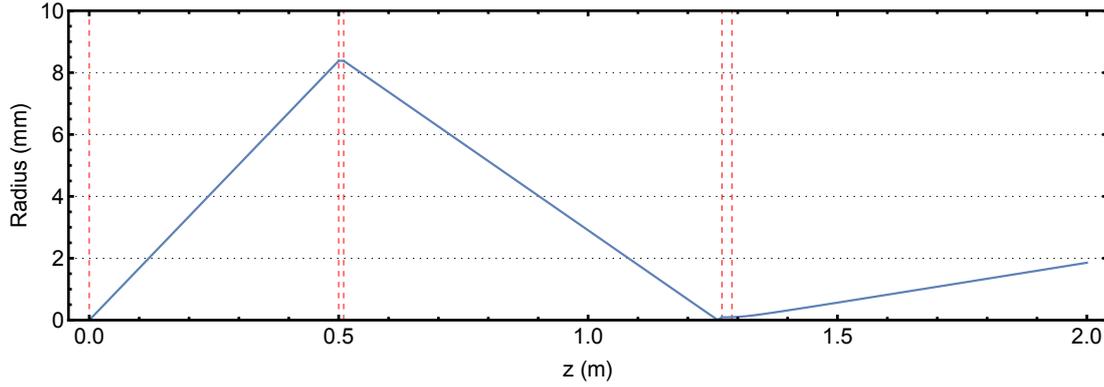


Figure 5.4: Beam waist of probe beam showing light scattered from the BEC at $z = 0$ being imaged by the relay lenses at $z = 0.5$ m, then reimaged by a microscope objective downstream at $z = 1.26$ m.

mirror enables the relay optics to be placed closer to the science cell, it introduces significant shot-to-shot fluctuations that are magnified by the microscope objective. This results in significant perturbations to the position of the BEC image at the iris (see §5.6), resulting in large variations in the amplitude of the measured Faraday signal.

Static collection optics were used to prevent shot-to-shot fluctuations, placed at a distance from the science cell that produced minimal clipping of the MOT beam. A 2" mirror was mounted to a ZABER T-LSR300D translation stage to have precise control of the mirror's position. An optimal position was located whereby the mirror mount clipped the MOT beam by a small amount, but no significant reduction in the MOT load rate (and therefore BEC size) was observed. The translation stage allowed optimisation of the mirror position, and allowed the mirror to be retracted for other experiments.

The collection mirror mounted on the ZABER stage is 380 mm from the centre of the cell, oriented with an angle of incidence of 67.5° to redirect the probe light into the relay lenses. The tilt of the mirror results in different numerical apertures in the x and y directions, with $NA = 0.026, 0.069$ respectively. Neglecting any aberrations in the relay lenses and microscope setup, the numerical aperture limits the resolution of diagnostic absorption images taken of the cloud. The lower bounds on the size of features that can be resolved by absorption imaging are given by $\Delta x = 15 \mu\text{m}$ and $\Delta y = 6 \mu\text{m}$.

It should be emphasised that the resolution of magnetic resonance images acquired with Faraday measurement is unaffected by this diffraction limit, as the size of a resolution element in the Faraday imaging technique is determined by the applied magnetic field gradient (see §3.2). In principle, diffractive blurring of the image of the BEC at the iris limits how tightly the aperture can be closed, but in practice a tolerance must be provided for shot-to-shot position fluctuations.

The relayed image of the BEC is then reimaged by a ZEISS 44-08-44 microscope objective (effective focal length 8.25 mm) and 150 mm tube lens. This generates a second image of the BEC within the iris, which is mounted in an XY-translation stage. The iris is translated to be centred on the image of the BEC (see §5.5), and partially closed so that only light that passed through the BEC is transmitted.

A mechanical flipper mirror then directs the transmitted light to either be reimaged onto a camera (AVT MANTA G609B) for absorption imaging, or balanced polarimeter to measure the Faraday signal for magnetic resonance imaging. Absorption imaging along the Faraday beam¹ is achieved by directing a small amount of resonant light into the optical fiber² to image along the same path using the same optics. Note that the tight focus of the probe beam means that extremely low levels of imaging light (\sim nW) are used to prevent saturation of the atomic transition used for resonant imaging.³

Because absorption imaging uses such low power, the reimaging lens in front of the camera is positioned to demagnify the image by a factor of 4. This spreads the image over fewer pixels, increasing the average count per pixel well above the dark noise level of the camera and increasing the SNR of the optical depth calculation. Demagnifying the image has the added advantage of being able to reduce the capture area, as the G609B has a 2752 x 2206 px CCD sensor with a full-frame readout time of 112 ms. This is the minimum interframe time between capturing the absorption and flat-field images, which should be kept as short as possible to reduce fringe movement. The demagnified image is contained within a 270 x 270 px window, which when used as the capture region reduces the readout time by nearly an order of magnitude to 16.4 ms.

5.5: Striking the BEC bulls-eye

It is a significant technical challenge to align the Faraday beam to hit the BEC (a 30 μ m target) as the probe beam has a small waist and optical access limitations mean the closest steering mirror is 50 cm away. Since the BEC is so small and its exact position is unknown, aligning the probe beam to overlap the trapping beams (MOT or dipole) only provides a coarse starting point, and the atoms themselves must be used as a fine alignment guide.

This fine alignment is achieved by coupling resonant light obtained from the MOPA into the Faraday fiber and exposing it to the atoms. The resonant light ejects atoms from the trap, reducing the number of atoms remaining after the exposure time. Initially the beam is poorly centred on the atoms, so the intensity they experience is

¹ As distinct from time-of-flight imaging using the 'side' camera.

² Note that the resonant light has orthogonal polarisation to the Faraday beam so it is mostly rejected by the Glan-Laser polariser at the fiber output. It is therefore necessary to rotate the waveplate after the fiber to transmit a small fraction. High losses are acceptable as the imaging beam is very weak.

³ The exact amount of imaging light used is not important as it is only used for alignment purposes.

weak, producing a low loss rate. As the alignment is improved to be centred on the atom cloud, the intensity increases and the loss rate is enhanced, resulting in more of the cloud being ‘destroyed’. Observing the number of atoms remaining after a fixed exposure time therefore provides a metric for alignment.

Rough initial alignment is achieved by targeting the MOT, which is substantially larger than the probe beam waist, making it easy to ‘hit’. By continuously loading the MOT from the Zeeman slower, the alignment can be adjusted in real time using fluorescence imaging (as observed with a THORLABS UC480 CMOS camera). Because the MOT is being continuously loaded, a significant amount of resonant light (~ 1 mW) is required to cause a noticeable perturbation. Initially the resonant light appears as a dark band through the atom cloud, but as losses increase closer to the centre of the cloud, it acts like a sinkhole, reducing the number of trapped atoms and the diameter of the cloud. When the beam strikes the centre of the cloud, the loss rate can overcome the load rate (depending on the probe intensity), and the MOT is seen to vanish altogether.

The magnetic field generated by the second Zeeman slower coil (which loads the MOT) perturbs the location of the magnetic field zero, and hence the position of the MOT. Switching off the Zeeman slower translates the MOT horizontally, and stops reloading it. The resonant beam can then be swept horizontally until it hits the new position of the MOT, obliterating it. Because the MOT is no longer being reloaded, the slower must be toggled on and off as the MOT will vanish over time even for a poorly aligned probe beam. The centre of the MOT can be determined by visually maximising the loss rate observed with the fluorescence camera.

With the centre of the MOT approximately located, the process is repeated for colder and smaller atom clouds until alignment on the final BEC is achieved. In our apparatus the magnetic field zero is vertically displaced from the centre of the crossed-beam dipole trap, so the MOT is centred on a different position to the BEC. Since the beam waist is small, alignment must be repeated for multiple stages of cooling between the MOT and Bose condensation so that the beam can be walked to follow the centre of the cloud. Adjustments must be made in a shot-based process (making a tweak then running a shot to check the result), which is laborious as the duty cycle of the experiment is ~ 30 s when evaporation is performed.

For these shot-based adjustments, the number of atoms remaining after the exposure time is measured using absorption imaging, captured using the ‘side’ imaging system.⁴ This independent imaging system captures background-subtracted absorption images along the x -axis using a AVT PROSILICA GX1920 with magnification 2.1 \times , and permits time-of-flight imaging with drop times of up to 25 ms. For unevaporated

⁴ See [138] for complete details of the ‘side’ imaging system.

atom clouds, the image of the atom cloud is larger than the camera's CCD, so the total atom number cannot be easily determined. However, in these cases the alignment of the beam can still be optimised by visual inspection of the optical depth following the resonant probe exposure with zero drop time.

Colder clouds are more sensitive to being perturbed by resonant light, so as the alignment improves, the power in the resonant beam needs to be decreased to prevent eliminating the cloud entirely.⁵ It is important to not completely annihilate the cloud, as the tails of the Gaussian beam profile can cause significant losses if the intensity has not been reduced far enough, preventing further refinement of the alignment. Therefore the alignment should be adjusted until the atom number reaches a definite minimum, then the experimental sequence should be adjusted to cool the cloud further, the intensity reduced, and alignment repeated.

The process is carried out for successively colder clouds, from magnetic trap down to BEC through to rf evaporation of increasing levels of truncation. When observing (nearly) condensed clouds, the absorption imaging procedure is modified to include some drop time (typically 23 ms in our apparatus) to permit ballistic expansion of the cloud and a 2D Gaussian fit of the optical depth to be performed to quantify changes to atom number.

Although the trapped BEC is smaller than the waist of the probe beam, the annihilation rate is maximised when the peak intensity is centred on the BEC, enabling the centre to be determined to within the precision of the kinematic mirror mount holding the steering mirror (THORLABS POLARIS K1). However, this alignment sensitivity combined with the long propagation distance from the steering mirror to the BEC means that a small amount of pressure on an Allen key in the kinematic mount is sufficient to sweep the probe across the BEC, so care must be taken.

The probe beam should be perfectly centred on the BEC so that most of the light interacts with the BEC to give the strongest Faraday signal. However it is not strictly required when performing Faraday measurement at the magic-zero wavelength, as the trap perturbation induced by the Faraday beam is negligible (see §2.4). Small drifts in alignment are acceptable, so the realignment procedure only needs to be repeated if the position of the BEC is perturbed (for example, if the crossed-beam dipole trap has been realigned).

The output stage optics can now be aligned, and an absorption image of the cloud produced using the Faraday beam path. First, the reimaging lens in front of the camera is positioned to produce a demagnified image (see §5.4) of the closed iris. The

⁵ Alternatively, the exposure time could be reduced, but beam power can be easily adjusted over several orders of magnitude by reducing the rf power into the AOM used to derive the resonant light from the MOPA. Reducing the exposure time substantially also increases the influence of switch-on transients from the shutter and/or AOM, increasing shot-to-shot fluctuations.

camera is mounted on a z -translation stage (THORLABS MT1), enabling its position to be adjusted until a sharp image of the blades of the iris is formed. As this imaging system is only used for diagnostics, aberrations introduced by using a single lens for demagnification are not a concern.

There are two sets of telescopes that require alignment: the relay lenses and the microscope. The relay lenses are relatively insensitive to alignment and can be positioned with a ruler, placing the centre of the science cell at the focus of the first lens. Perturbations to the relay telescope simply translate the position of the image, which can then be compensated for. The position of the microscope objective, however, must be carefully optimised to produce an image of the BEC at the position of the iris.

The objective, tube lens and iris must be concentric, so are held together in a THORLABS 30 mm cage system. The objective is held in a THORLABS SM1Z z -translation stage, and the iris in a THORLABS ST1XY-S xy -translation stage. Removing the objective lens, the probe beam can be aligned to travel down the central axis of the cage using a THORLABS CPA1 alignment tool, and through the centre of the iris. The objective is mounted in a cage plate using a C-mount to SM1 adapter, making it front-heavy and causing it to droop. Despite being on-axis with the cage, the beam therefore enters the objective at an angle, causing the beam path to deviate vertically. This must be corrected by vertically walking the input steering mirrors. The tight waist and short effective focal length of the objective prevents the probe from being collimated between the objective and tube lens, so the correct focusing condition can only be determined by looking at an image of the BEC using the diagnostic camera.

Since the camera is already independently focused on the iris, it should also produce an in-focus image of the BEC when the objective has been positioned correctly. When the BEC is out of focus, the image appears elongated or demonstrates spatial structure. It should be noted that since the NA is different in the x and y directions (§5.4) the cloud appears elongated. Once the objective is positioned, the iris is partially closed and centred to ‘shrink wrap’ the image of the BEC (Figure 5.5).

Now that the output stage optics are centred on the BEC, the input stage optics can be tweaked to similarly centre the illumination beam on the iris, and hence on the BEC itself. This final alignment does not require the atom cloud to be present and can therefore be done with the camera continuously imaging the iris. This instant feedback greatly facilitates precise adjustments to be made to the sensitive input stage optics, as opposed to the earlier alignment which had to be done in a shot-based method (~ 25 s per iteration).

Since the output stage optics have been aligned to capture the Abbe rays from the BEC, they are independent of the alignment of the input stage optics. In principle, the output stage only needs adjustment if the position of the BEC changes; any perturbations to the input stage that misalign the illumination beam only require realigning the

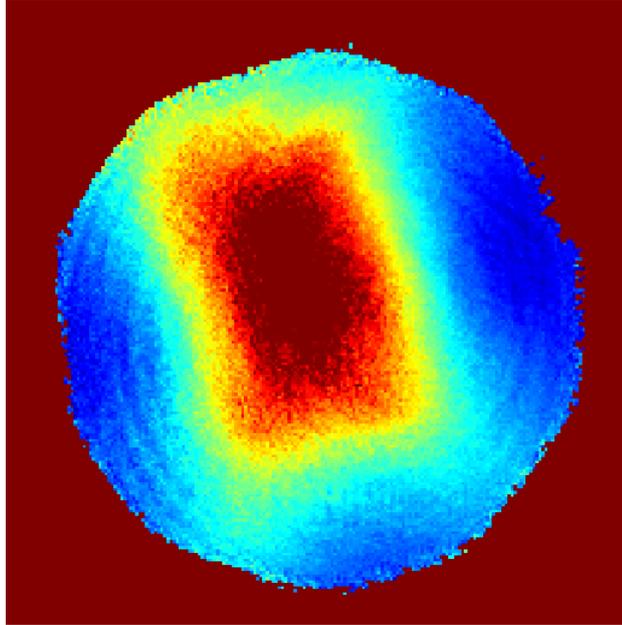


Figure 5.5: Resonant optical depth image of the BEC captured through the Faraday beam path using the MANTA 609B camera, showing typical centration of the iris and size of the aperture around the BEC image.

illumination beam back onto the iris. In practice however, the high magnification of the microscope objective means that slight drifts in the pointing stability of the input beam walk the image off the iris. These perturbations are more easily countered by making minor adjustments to the steering mirrors in front of the objective instead of the extremely sensitive input stage optics. This is preferable for dealing with small long-term drifts in position, though periodic realignment using the iris is required if the BEC is no longer approximately in the centre of the illumination beam.

5.6: Pointing stability of the Faraday beam

The high magnification of the objective (Figure 5.5) means that small displacements of the Faraday probe beam at the input translate into large displacements at the detector. Shot-to-shot fluctuations in the pointing stability of the probe beam translate the image of the BEC across the iris, clipping the image and reducing the signal contributed by that part of the beam. Although the iris could be opened further to compensate, this increases the amount of light reaching the detector that didn't pass through the BEC – which increases the amount of shot-noise measured by the detector, also reducing the overall SNR of the measurement.

Given the large magnifications and propagation distances involved, the sensitivity of the imaging system to perturbations in the probe beam will now be quantified. The system can be modelled with a linear ray optics matrix transformation formalism (ABCD matrices) as follows: a ray located a distance x_0 from the optic axis propa-

gating at angle θ_0 is represented by the vector $(x_0 \ \theta_0)^T$. Propagation by a distance d and passing through a thin lens with focal length f are respectively modelled by the transfer matrices [170]

$$\mathbf{P}(d) = \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{L}(f) = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}. \quad (5.1)$$

Consider a mechanical stage located a distance p from the cloud which laterally translates the beam an amount dx and perturbs its propagation direction by $d\theta$. Provided that dx and $d\theta$ are small enough to preserve the paraxial approximation, the position and direction of the probe beam at the iris is given by

$$\begin{pmatrix} x \\ \theta \end{pmatrix} = \underbrace{\mathbf{P}(f_t) \cdot \mathbf{L}(f_t) \cdot \mathbf{L}(f_o) \cdot \mathbf{P}(f_o)}_{\text{objective}} \cdot \underbrace{\mathbf{P}(f_2) \cdot \mathbf{L}(f_2) \cdot \mathbf{L}(f_1) \cdot \mathbf{P}(f_1 - p)}_{\text{relay}} \cdot \left[\mathbf{P}(p) \cdot \begin{pmatrix} 0 \\ \theta_0 \end{pmatrix} + \begin{pmatrix} dx \\ d\theta \end{pmatrix} \right].$$

The resulting transverse magnification of the image and its location are given by

$$M' \equiv \frac{2\theta_0}{\theta(\theta_0) + \theta(-\theta_0)} = \frac{f_t f_2}{f_o f_1} \equiv M_0 \quad \text{and} \quad x = M_0(dx - p d\theta). \quad (5.2)$$

Hence the magnification is unaffected but the position of the image is perturbed by an amount proportional to that magnification. The net magnification of the system is high ($\sim 30\times$) implying the detection apparatus is quite sensitive to perturbations dx , $d\theta$. Because the image is being formed on the iris, which tightly crops the BEC image, small perturbations to the position of the image result in clipping blocking some of the light that travelled through the BEC, leading to a reduction in the Faraday signal.

An initial design used a flipper mirror in close proximity to the dichroic mirror that deflects one of the dipole trap beams after the cell. Being close to the cell improved the numerical aperture of the collection optics but required flipping the fold mirror out of the way during loading of the MOT to prevent clipping the MOT beam. The flipper mirror was prone to introducing small perturbations that destabilised the position of the image at the iris.

For this reason the flipper mirror was exchanged with a mirror mounted on a ZABER T-LSR300D translation stage, which does not move between shots. This reduced the shot-to-shot standard deviation of the centroid position from $\sigma = 10.5$ px with the flipper mirror to $\sigma = 2.6$ px with the translation stage. However, to prevent blocking the MOT beam the mirror had to be located significantly further away from the cell, which increased the distance to the relay lenses. This reduced the overall magnification of the system, though led to a slight increase in the numerical aperture of the system (Table 5.1).

	Flipper mirror	Translation stage
Collection mirror diameter	1"	2"
Distance from cell	230 mm	370 mm
First relay lens	300 mm	500 mm
Net magnification	45	27
Numerical aperture (x, y)	(0.021, 0.055)	(0.026, 0.069)
Diffraction limit (x, y)	(19, 7) μm	(15, 6) μm
Centroid fluctuations	10.5 px	2.5 px

Table 5.1: Comparison of properties of the collection optics for a flipper mirror and a translation stage.

5.7: Trapped BEC lifetime

Exposing the BEC to the Faraday probe is expected to cause off-resonant scattering by (3.19), causing atoms to be ejected from the BEC. In order to quantify the loss rate due to the Faraday probe, it is necessary to develop a model for how the BEC atom number changes over time. In this section, I discuss the causes of number loss in the BEC, and measure the loss rate coefficients.

Collisions in a trapped condensate cause atoms to be ejected from the trap and lost, therefore reducing the atom number over time. The dominant loss mechanisms in an $F = 1$ condensate of ^{87}Rb are ‘one-body’ loss caused by collisions with the background gas in an imperfect vacuum, and ‘three-body’ collisions between the trapped atoms (see Appendix B). For a BEC with a Thomas-Fermi profile subject to these loss processes, the loss rate is given by

$$\frac{1}{N} \frac{dN}{dt} = -cN^{4/5} - b, \quad (5.3)$$

where b and c are the one- and three-body loss rates respectively. This differential equation must be solved numerically to determine the number of atoms remaining as a function of time, but approximating the exponent leads to an analytic solution that models the dynamics well for small cloud sizes,

$$\frac{1}{N} \frac{dN}{dt} \approx -aN - b \quad \Rightarrow \quad N(t) = \frac{be^{-bt}}{b + aN_0(1 - e^{-bt})} N_0 \quad (b \neq 0). \quad (5.4)$$

The two models are related by $c \approx a\langle N_0 \rangle^{1/5}$, where $\langle N_0 \rangle$ is the average initial BEC atom number.

Three-body loss dominates one-body loss for dense clouds, so the two loss mechanisms can be distinguished by observing how the number of atoms in BECs of varying initial atom number decays over time. Creating BECs from MOTs with different load times varies the initial atom number N_0 . Measuring the number of atoms remaining in the BEC after a range of dipole trap hold times therefore allows for simultaneous determination of both model parameters, a and b .

Each different MOT load time results in a different (unknown) value of N_0 , which is required to perform least-squares fitting of the model to obtain estimates of the parameters a and b . Initially N_0 can be estimated for each series from the measured atom number at zero hold time, but the parameter estimates are sensitive to errors in this value. An iterative algorithm is therefore used to take the estimates of a and b , then fit each series individually to better estimate N_0 , which enables a better estimate of a and b to be then obtained. This algorithm converges rapidly, as determined by the difference between consecutive estimates of the fit parameters, and results in good fits for observed decay rates (Figure 5.6).

The resulting values of the parameters are $a = 9.5(3) \times 10^{-7} \text{ s}^{-1}$ and $b = 2.6(4) \times 10^{-2} \text{ s}^{-1}$, with the vacuum lifetime given by $\tau = 1/b = 38(6) \text{ s}$. Fitting with the analytic model (5.3) yields $c = 1.21(4) \times 10^{-5} \text{ s}^{-1}$. Combined with the weighted average initial atom number $\langle N_0 \rangle = 2.27(5) \times 10^5$ this yields $a = 1.03(3) \times 10^{-6} \text{ s}^{-1}$, which agrees with the numerical model at the 2σ level. The discrepancy is attributed to the analytic model's value of b , whose uncertainties were too large for meaningful comparison, indicating the optimisation algorithm had stagnated.

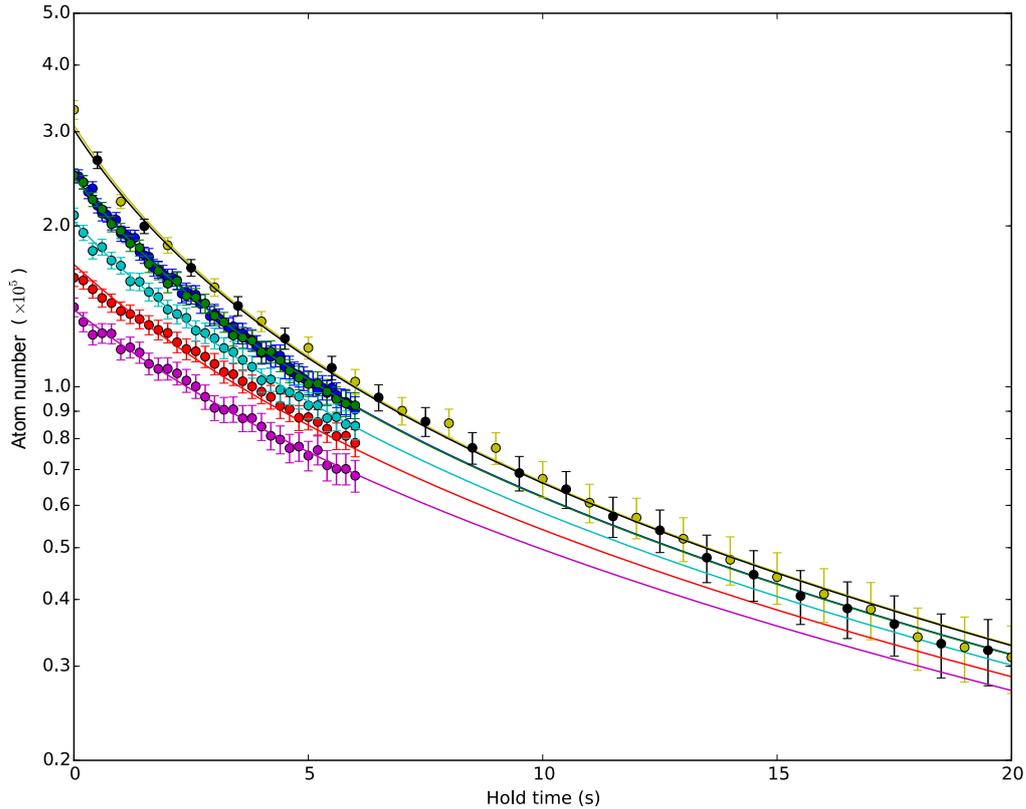


Figure 5.6: Number of atoms remaining in BEC after various hold times, as measured from 2D Gaussian fits to optical depth images, with each series corresponding to a different initial atom number N_0 . Self-consistent values of a , b and N_0 are obtained by iteratively fitting (5.4).

5.8: Anomalous resonant scattering

It was observed that exposing the BEC to the Faraday beam reduced the BEC lifetime by three orders of magnitude: 13 mW of probe light at $\lambda = 790$ nm reduced the lifetime from 38(6) s to 24(10) ms. The theoretical lifetime estimated using (3.19) was 900 ms, implying the presence of an unexpected loss mechanism. A SIRAH MATISSE-DS Ti:S laser was obtained to test the scattering rate at the same probe wavelength and intensity, and was found to result in a BEC lifetime that agreed with the theoretical calculation. The immediate conclusion is that the Faraday ECDL was not spectrally pure.

Despite lasing at 790 nm, the ECDL producing the Faraday beam has a broad amplified spontaneous emission (ASE) background that spans 770–810 nm [171]. This background emission is weak, but contains a small component which is resonant with the atomic transitions at 780 nm (D2 line) and 795 nm (D1 line). The BEC is extremely sensitive to small amounts of resonant light, and these small components are sufficient to cause significant scattering and substantially reduce the BEC lifetime.

A SEMROCK MAXLINE LL01-808 interference filter (IF) was introduced to the Faraday beam (§5.1), which reflects everything except a 4 nm band around $\lambda_0 = 808$ nm [172] for a normal incidence beam. Tilting the interference filter increases the optical path length through the filter, causing condition for constructive interference through the filter to shift to longer wavelength, according to

$$\lambda(\theta) = \lambda_0 \sqrt{1 - (\sin \theta / n_{\text{eff}})^2}, \quad (5.5)$$

where n_{eff} is the effective refractive index of the filter. Tilting the filter to $26.0(5)^\circ$ shifts the transmission window to 790 nm (Figure 5.7), heavily suppressing the ASE components at 780 nm and 795 nm as seen on an AGILENT 86140B Optical Spectrum Analyser (Figure 5.8). The observed suppression is > 15 dB, though this measurement is limited by the dynamic range of the spectrum analyser, which raises the measurement noise floor (-73 dB) above the dark noise level (-85 dB), so the true suppression is likely to be much higher (manufacturer specification is at least 50 dB [172]).

The lifetime of the condensate was measured with and without the interference filter (Figure 5.9), showing that introducing the filter improved the lifetime by a factor 40. So the observed reduction in trap lifetime as a result of exposing the cloud to the Faraday beam was indeed because the wide ASE component of the diode's optical spectrum caused resonant scattering in the condensate.

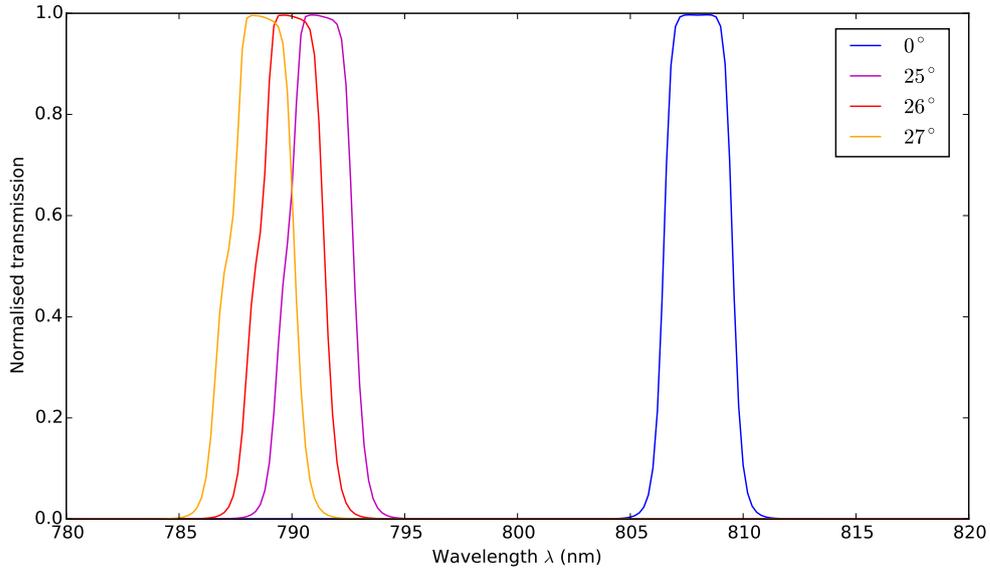


Figure 5.7: Expected transmission spectrum of the Semrock LL01-808 interference filter at different incidence angles, adapted from [172]. The transmission bandwidth is approximately 4 nm, and the window covers 790 nm at $26.0(5)^\circ$ incidence.

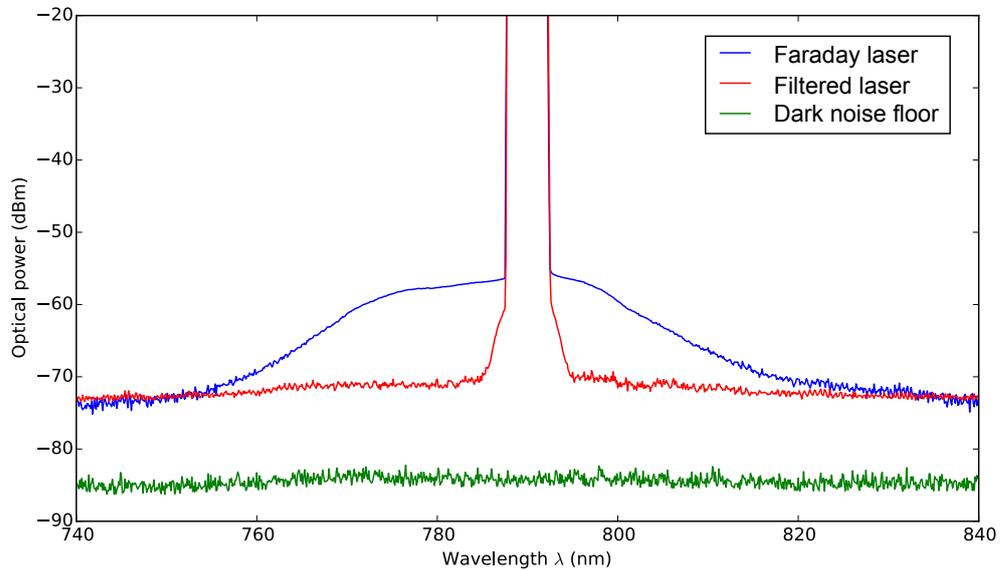


Figure 5.8: Comparison of the optical spectrum of the Faraday ECDL before and after introducing the filter. Note the strong suppression around the rubidium absorption lines (780 nm and 795 nm) which was previously reducing the lifetime of the atom cloud. Total beam power was 1 mW and the OSA measurement linewidth was 5 nm.

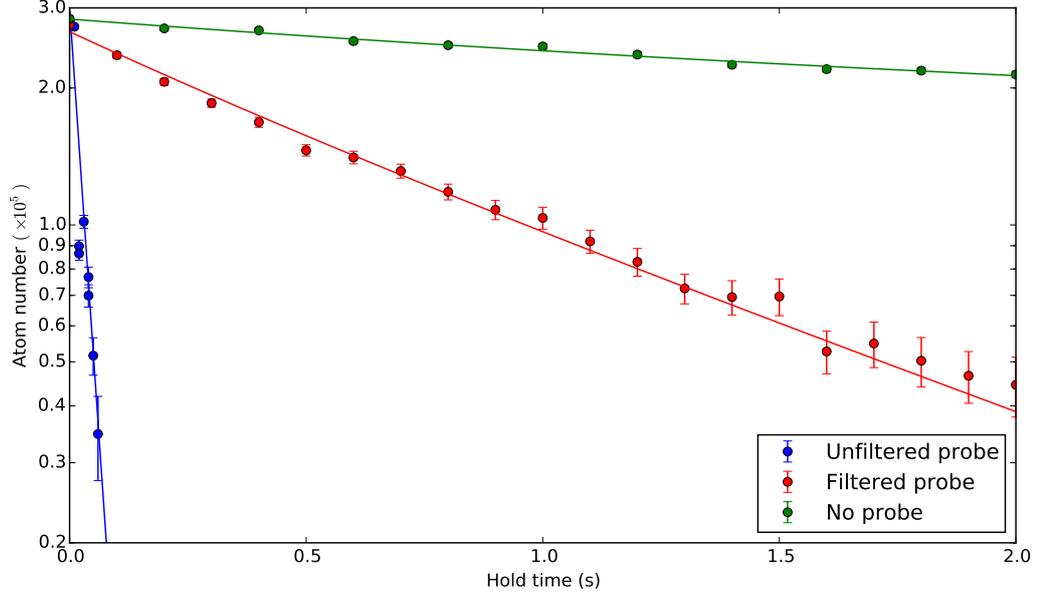


Figure 5.9: Exposing the BEC to the unfiltered Faraday beam (blue) results in a very high loss rate compared to the vacuum lifetime (green). Applying an interference filter to the Faraday beam (red) increases the lifetime (fitted decay constant) from 29(6) ms to 1.18(3) s.

5.9: Characterisation of off-resonant scattering

The interference filter prevents scattering from the resonant component of the Faraday laser light as described in the previous section. However, the Faraday probe laser still causes off-resonant scattering by (3.19), kicking atoms out of the trap and enhancing the one-body loss rate. The one-body loss rate b is expected to be linear in probe power P , so can be written as $b = b_0 + b_1 P$, where $\tau_0 = 1/b_0$ is the vacuum lifetime and b_1 quantifies scattering from the laser.

Considering a Gaussian beam of $1/e^2$ radius $r = 75 \mu\text{m}$, using (3.19) the scattering rate at the magic wavelength $\lambda = 790 \text{ nm}$ corresponding to the peak intensity is

$$I_0 = \frac{2P_0}{\pi r^2} \Rightarrow b_1 = \frac{2\xi_S^2}{\pi r^2} = 101.5 \text{ s}^{-1} \text{ W}^{-1}. \quad (5.6)$$

As previously, the loss rate was measured by varying the BEC hold time in the Faraday beam and measuring the remaining atom number (Figure 5.10A). An AOM was used to control the intensity of the probe beam and vary its power between sequences. The measured loss rates indeed show linear dependence in probe power (Figure 5.10B), with the fitted gradient $b_1 = 85(1) \text{ s}^{-1} \text{ W}^{-1}$.

The discrepancy between measured loss rate and the theoretical estimate is attributed to uncertainty in the probe beam intensity at the position of the atoms. The measured waist r corresponds to the *focal waist* of the probe beam, which is unlikely to be perfectly coincident with the position of the atom cloud.

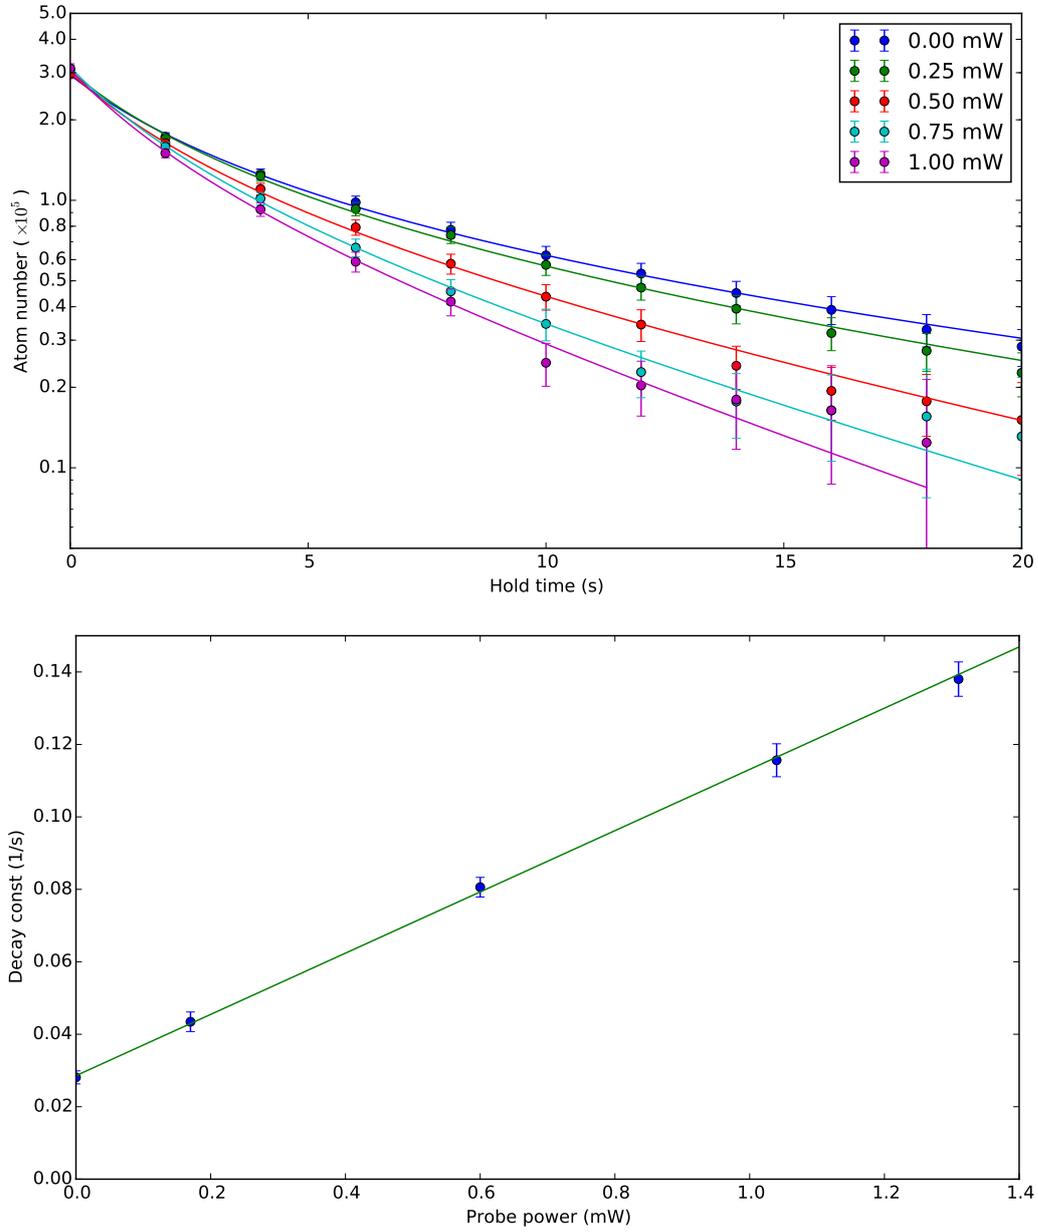


Figure 5.10: Number of atoms remaining in the trap after being exposed to different Faraday probe beam powers (top) and dependence of the resulting loss rates on beam power (bottom).

For example, lateral displacement of $0.3r = 23 \mu\text{m}$ or longitudinal displacement of $0.45z_R = 1 \text{ cm}$ from the focus would account for the difference. Although the lateral position of the focal point can be refined with absorption imaging (see §5.5), the longitudinal position is set by the position of the $f = 500 \text{ mm}$ focussing lens, which was not mounted on a translation stage. It is likely that a combination of imperfect alignment both laterally and horizontally resulted in the discrepancy.

Extrapolating the fitted regression to zero applied beam power then provides a measurement of the background one-body lifetime, $\tau = 1/b_0 = 35.0(7) \text{ s}$, which agrees with the earlier measurement (§5.7) but has much higher precision.

5.10: Summary

In this chapter I described the experimental considerations relevant to the generation, alignment and detection of the Faraday probe beam. A fiber-coupled ECDL tuned to lase at 790 nm was used to produce the probe beam, which was conditioned by purifying its polarisation and tightly focused at the position of the BEC.

A method to aligning the centre of the probe beam on the BEC in the presence of limited optical access was described, using resonant light coupled into the same fiber and aligning the beam to maximise the scattering rate. Absorption imaging was used to focus the microscope objective on the BEC, enabling the iris to be used to clip light that did not pass through the BEC from adding shot-noise to the Faraday measurement. Design considerations related to the collection optics on the other side of the cell were discussed, particularly the importance of pointing stability of the beam heading into the microscope objective.

The ECDL used to generate the Faraday beam was observed to display an anomalously high scattering rate, which was seen to be the result of the small resonant component of the diode's ASE background. Introducing an interference filter to transmit only a small window around the magic wavelength restored scattering to the expected rate, in excellent agreement with theory.

Our probe laser now generates a bright, tightly focused beam of light at the magic wavelength, which is centred on the BEC. I verified that we can expose the BEC to the Faraday beam for extended periods, without resulting in large atom loss. I will next discuss how polarimetry can be performed on the Faraday beam to measure properties of the BEC.

Photodetection and signal processing

In this chapter I discuss the photodetection apparatus used for polarimetry, demonstrate that it is shot-noise limited at operating probe powers, and calibrate it using a test target of terbium-glass. I undertake Faraday measurement of a BEC in a uniform applied magnetic field, and perform signal processing to extract the Faraday rotation signal.

A prototype high-gain balanced differential photodetector was designed to form the core of the polarimeter. This detector was used in the measurements and the characterisations are presented in this chapter. The design was further refined into a second-generation photodetector, whose response was characterised in detail elsewhere [173]. This revised detector was used to capture the Faraday measurements presented in the remainder of this thesis.

6.1: High-precision polarimetry

The Faraday measurement technique described in §3.1 requires continuous measurement of an oscillating polarisation rotation, with microradian amplitude and centre frequency in the hundreds of kilohertz. Capturing this Faraday signal therefore requires a sensitive, shot-noise limited polarimeter, containing high precision polarising optics and a photodetector with high bandwidth and transimpedance gain.

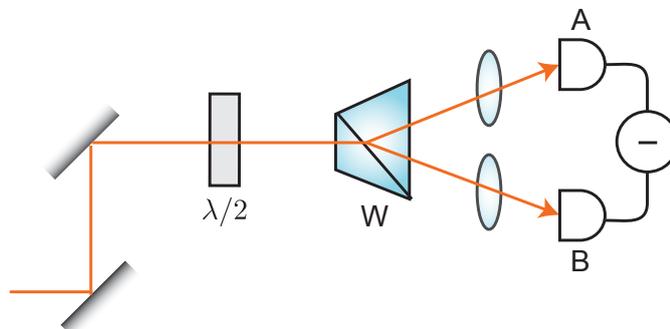


Figure 6.1: Polarimeter formed by a Wollaston prism (W) and differential photodetector ($A - B$). The $\lambda/2$ waveplate is adjusted to give equal power beams in the absence of Faraday rotation. The lenses focus the beams on the active area of the photodiodes.

The polarimeter constructed for this purpose was formed by a Wollaston prism and custom differential photodetector (Figure 6.1). The Wollaston splits the probe beam into its linear polarisation components with high fidelity ($> 10^5 : 1$). The two components are focused onto the ports of the photodetector, which outputs a voltage proportional to the difference in power between the two beams.

A half-waveplate before the Wollaston prism rotates the polarisation so that the two components have equal power in the absence of Faraday rotation. When the polarisation rotation θ experienced by the probe beam is small, the voltage output by the photodetector with the polarimeter 'balanced' in this way is (§3.1),

$$V = 2GR\kappa P\theta, \quad (6.1)$$

where κ is the collection efficiency, G the net transimpedance gain (V/A), R the responsivity of the photodetector (A/W) and P the total power in the incident probe.

6.2: Photodetector design considerations

The choice of electronic components has important consequences for the performance of the polarimeter, which shall now be discussed. The primary design considerations are a high bandwidth to prevent signal attenuation at the Larmor frequency, and high-gain/low-noise transimpedance amplification to detect small polarisation rotations.

The balanced differential photodetector design (Figure 6.2) is based on [128], and contains two photodiodes that drive a differential photocurrent through a high-gain transimpedance amplifier (gain $G = 5.1$ V/mA). The output voltage result is fed into a line driver (gain $2\times$) to isolate the transimpedance amplifier from the capacitance of the connected output cable. The design choices are selection of the photodiodes, operational amplifiers (op-amps), the transimpedance gain, plus associated capacitors.

The photodiode chosen should have high responsivity R to ensure efficient detection, and low terminal capacitance so that the photodiode bandwidth does not limit the detector response. This encourages the use of small area photodiodes and a high reverse bias voltage V_b . To prevent damage to the photodiode, the maximum probe power per photodiode is $P_{\text{opt}} = P_{e,\text{max}}/(V_b R)$ where $P_{e,\text{max}}$ is the maximum (electrical) power dissipation of the photodiode. Typically photodiodes with cutoff frequencies of ≥ 10 MHz are common at 15 V reverse bias, which present no difficulties for the Larmor frequencies of interest (~ 700 kHz).

The photodiode initially chosen was the HAMAMATSU S5971 [174], whose small active area (1.1 mm^2) and low capacitance (3 pF at 15 V reverse bias) results in a high photodiode bandwidth (100 MHz). However, short focal-length lenses ($f = 50$ mm) are required to focus the probe beam down into the active area and prevent overfilling. The high responsivity of the diodes, $R = 0.56$ A/W, corresponds to a quantum ef-

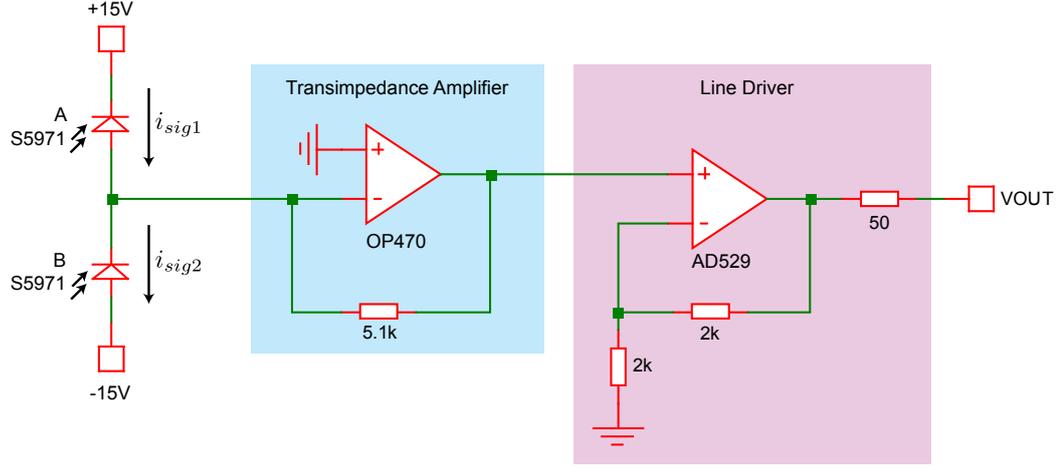


Figure 6.2: Circuit diagram for the differential photodetector. Photodiodes A and B generate a differential photocurrent that is fed through a two-stage amplifier. Bypassing arrangements and power supply connections not shown.

efficiency of $\eta = 88\%$ at the measurement wavelength $\lambda = 790$ nm. The damage threshold is 50 mW electrical power dissipation, which at 15 V reverse bias corresponds to an optical power of 6 mW per photodiode.

A detailed discussion of design considerations regarding transimpedance amplification can be found in Chapter 18 of [175], with a brief overview of pertinent details presented here.

The primary limitation on the photodetector bandwidth is the fixed ‘gain-bandwidth product’ of the transimpedance op-amp. Choosing an op-amp that has a high gain-bandwidth product allows the gain to be maximised, while ensuring the bandwidth remains greater than the Larmor frequency.¹ However, the op-amp introduces noise into the measurement which acts to obscure the Faraday signal, and its output can become unstable if incorrectly designed.

The voltage noise has contributions from Johnson noise in the feedback resistor (resistance $R_f = G$), input voltage noise, and input current noise from the op-amp flowing through the feedback resistor [176]. These respective contributions as voltage-referred noise levels are (in $V/\sqrt{\text{Hz}}$)

$$V_J^{(\text{rms})} = \sqrt{4k_B T G}, \quad V_{A_v}^{(\text{rms})} = e_n, \quad V_{A_i}^{(\text{rms})} = i_n G, \quad (6.2)$$

where e_n and i_n are the (manufacturer-specified) input noise voltage and current densities of the op-amp. These contributions should be similar in magnitude, as if one noise source dominated then a different amplifier could be chosen and the overall noise reduced. The gain should therefore be chosen to be of order $G \sim e_n/i_n$.

¹ The line driver op-amp can be used as a second-stage amplifier to increase the overall gain without reducing the bandwidth of the transimpedance amp.

The optical beam power P used in a Faraday measurement at the magic wavelength is fixed by the desired scattering rate (see §3.4). The electronic noise floor should be well below the shot-noise level at this power to ensure the measurement is not limited by technical noise. For optical power P , the shot-noise level (in $\text{V}/\sqrt{\text{Hz}}$) is

$$V_{\text{shot}}^{(\text{rms})} = G\sqrt{2eRP}. \quad (6.3)$$

The technical amplifier noise is therefore negligible compared to the shot-noise when

$$\begin{aligned} \left(V_{\text{shot}}^{(\text{rms})}\right)^2 &> \left(V_{\text{J}}^{(\text{rms})}\right)^2 + \left(V_{\text{Av}}^{(\text{rms})}\right)^2 + \left(V_{\text{Ai}}^{(\text{rms})}\right)^2 \\ \Rightarrow P &> \frac{4k_{\text{B}}TG + e_n^2 + i_n^2G^2}{2eRG^2}. \end{aligned} \quad (6.4)$$

The signal-to-noise ratio of the Faraday measurement (§3.3) depends on the scattering rate (3.19), so for a fixed scattering rate there is a trade-off between the detuning Δ and beam power P . A Faraday measurement can either be done with a small detuning and weak probe beam, or large detuning and bright probe. If a weak powers are being measured, it is important to select an op-amp and transimpedance gain such that the measurement is not dominated by electronic noise (see also [175, §18.4.3]).

Following [128], the chosen amplifier was the OP470, a low-noise quad op-amp.² The gain-bandwidth product of the OP470 is 6 MHz, so the transimpedance gain was chosen as 5.1 V/mA to maximise the gain while ensuring the bandwidth f_c remained well above the anticipated Larmor frequency $f_c > 1$ MHz. The noise spectrum of the OP470 is spectrally flat (white) above 1 kHz, and the datasheet [177] gives $e_n = 3.2 \text{ nV}/\sqrt{\text{Hz}}$ and $i_n = 0.4 \text{ pA}/\sqrt{\text{Hz}}$, so taking $R = 0.56 \text{ A/W}$ and $T = 300 \text{ K}$, the minimum measurement power is $P > 20 \text{ }\mu\text{W}$. Technical electrical noise in the photodetector therefore does not make a significant contribution to the measurement at anticipated Faraday measurement powers ($\sim \text{mW}$) and can be safely neglected.

The net transimpedance gain of the photodetector was calibrated by measuring the output voltage for a range of input beam powers, with one port of the photodetector blocked (Figure 6.3). The photocurrent was inferred from the known responsivity and the incident optical power (measured with a THORLABS PM100A power meter). The measured transimpedance gain is 10.16(3) V/mA, in good agreement with the theoretical value 10.2(2) V/mA – implying the photodiodes were not overfilled. The output voltage remains linear in photocurrent up to the power supply rails voltage of 15 V – implying that the photodiodes were not saturating at even the maximum available optical power of 2.5 mW per photodiode.

² The other op-amps in the quad package were used for other purposes in the original design, but the functionality was later removed to simplify the circuitry.

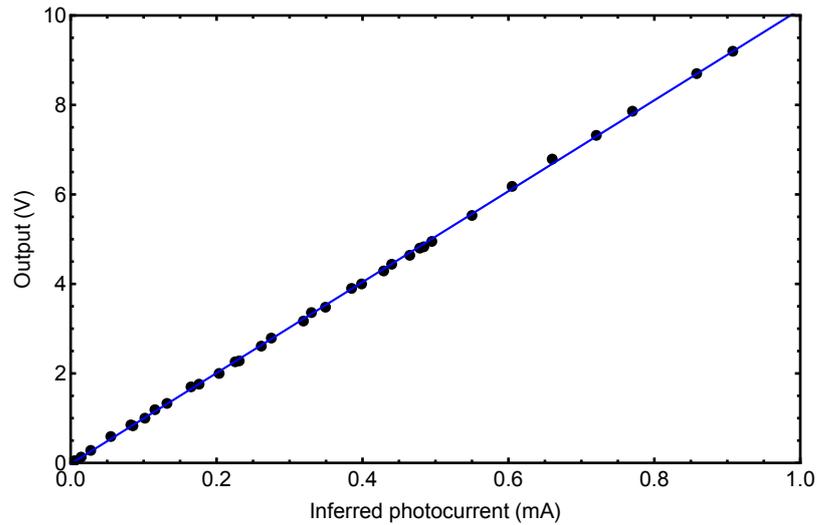


Figure 6.3: Calibration of the photodetector gain. The photodetector output is linear in the photocurrent, demonstrating the photodiodes are not saturated.

6.3: Shot-noise-limited photodetection

The intensity of a real light source, such as a laser, is not perfectly constant over time. Fluctuations in intensity act to obscure photodetection measurements and limit their precision, so the intensity noise should be eliminated where possible. Some noise sources are technical, and can be eliminated through careful selection of components, and some are fundamental. Contributions to the noise in a balanced photodetection apparatus will now be considered to quantify the noise level of our detector.

Classical intensity noise typically results from imperfections in the creation, propagation, or detection of the laser light. Examples of such imperfections are electronic noise in the laser current supply, mechanical instability of the lasing cavity, feedback from the frequency-locking servo loop, vibration in the optical components or etaloning between glass surfaces. A perfectly balanced polarimeter is insensitive to classical intensity noise, as the beamsplitter divides the beams evenly into two arms, each of which manifest the same classical fluctuations. These are then measured and subtracted by differential photodetection,³ cancelling out their contributions. Since the polarisation rotation induced by the Faraday effect is small, the polarimeter remains closely balanced throughout the measurement and classical noise can be eliminated.

Shot-noise, however, arises from the quantised detection of light as photons. Intensity measurements therefore follow Poisson statistics, and manifest a fundamental level of noise⁴ called the ‘standard quantum limit’ (SQL). Unlike classical noise, which is

³ Provided the arm lengths of the polarimeter are equal so fluctuations are observed simultaneously.

⁴ Squeezed light can increase sensitivity beyond the shot-noise limit [178], but producing squeezed light is technically challenging and beyond the scope of this project. Furthermore, the resulting sensitivity improvement is measured relative to the SQL, so characterisation of the SQL is still useful.

duplicated at a beamsplitter, the quantum noise on the resulting beams is uncorrelated [179, 180] and cannot be cancelled out by differential photodetection.⁵

The noise level was quantified by observing the power spectrum of the photodetector output using an HP 4395A vector network analyser (VNA) in spectrum analysis mode. Measuring the intensity noise with one port of the differential photodetector blocked shows the classical noise level, which demonstrates spectral structure and is up to 25 dB above the SQL (Figure 6.4). Unblocking the ports and balancing the polarimeter results in subtraction of the classical fluctuations and greatly suppresses the classical noise.

When the polarimeter is balanced, the classical fluctuations cancel exactly. So the polarimeter could be balanced by observing the noise level on the spectrum analyser and performing minor adjustments of the waveplate until the measured noise is minimised. The noise power was recorded as a function of incident beam power and observed to be linear (Figure 6.5). Since classical fluctuations would result in noise power that scales quadratically with beam power whereas shot-noise power scales linearly (6.3), this indicates the measurement was shot-noise limited.

⁵ In a fuzzy intuitive picture, photons are seen to be randomly deflected into one beam path or the other at the beamsplitter with fixed probability, so the number of photons in each arm is binomially distributed with uncorrelated fluctuations.

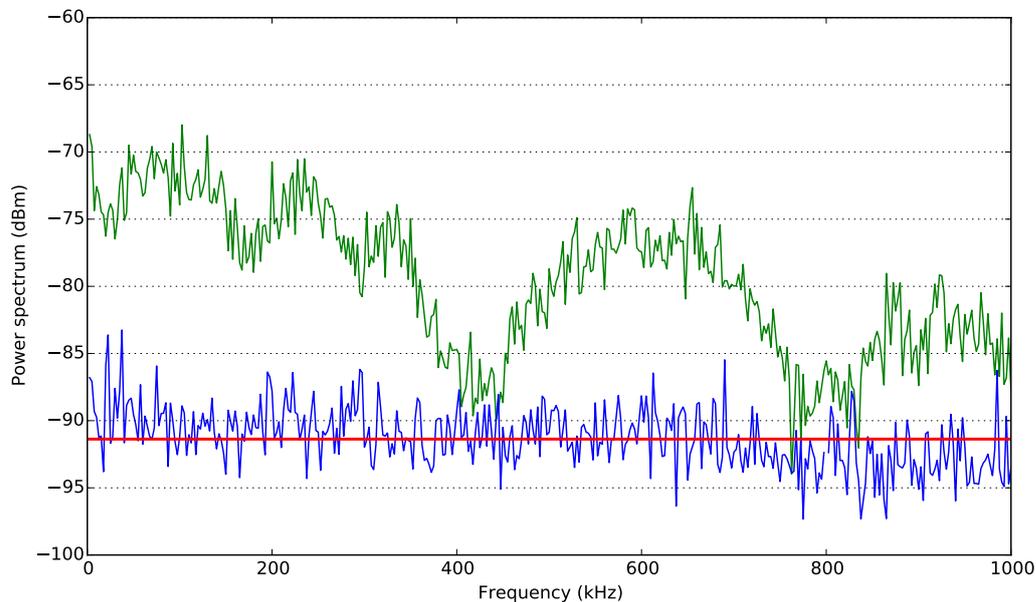


Figure 6.4: Comparison of the classical noise spectrum of a single beam (green) to the differential measurement (blue). Differential detection suppresses the classical noise by up to 20 dB, agreeing well with the independently measured shot-noise limit (red). Measurement RBW was 300 Hz, reference level was -60 dBm. Note that the SQL has not been corrected for gain roll-off at high frequencies.

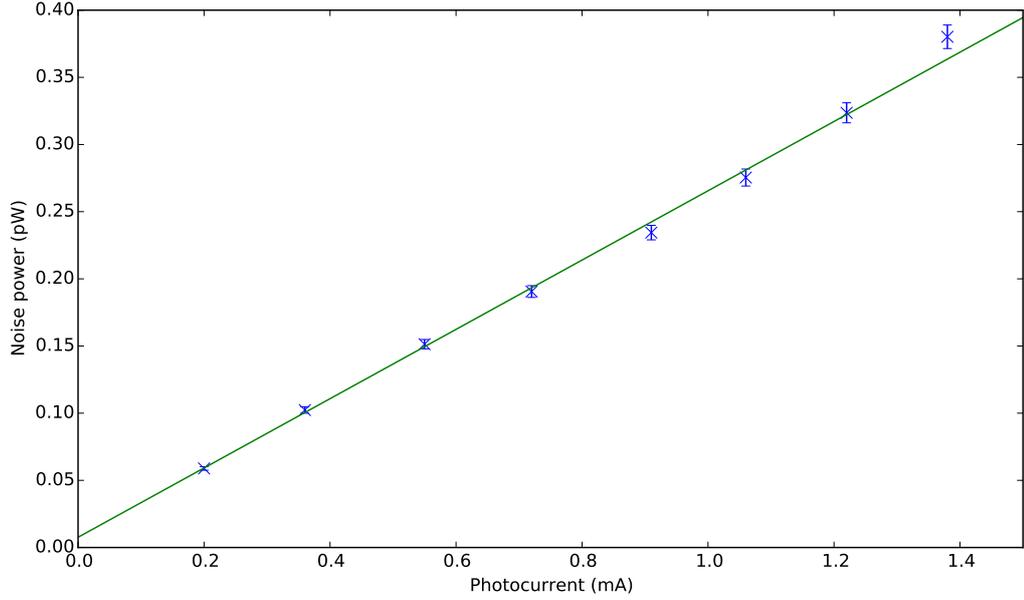


Figure 6.5: Measured noise power for balanced photodetection is observed to be linear in incident beam power, showing that the quadratic contribution from classical noise has been cancelled and the measurement is shot-noise limited as per (6.3).

6.4: Electro-optic BEC calibration

The polarimeter was calibrated with an ‘electro-optic BEC’, a cylinder of terbium-glass placed inside a solenoid. Terbium glass has a high Verdet constant, estimated to be $\mathcal{V} = -72 \text{ rad/T}\cdot\text{m}$ at $\lambda = 790 \text{ nm}$, enabling it to generate a large Faraday rotation upon light travelling through it. The polarisation rotation induced by a piece of such glass length l placed in a magnetic field of strength B is

$$\Delta\theta = \mathcal{V}Bl. \quad (6.5)$$

The terbium-glass cylinder was placed in a solenoid of diameter 6 mm constructed with $N = 140$ turns in a length of $l = 60 \text{ mm}$. The field well inside the solenoid is then

$$B = \frac{\mu_0 NI}{l}, \quad (6.6)$$

where I is the driving current. The coil was not impedance-matched to the current source, so a sense resistor ($R_s = 1.7 \Omega$) was used to measure the current in the coil by recording the voltage across the resistor.

The Wollaston prism split the beam into polarisation components so that the photodetector measured the polarisation rotation from the terbium-glass (Figure 6.6). The solenoid was driven with a function generator (TABOR WS8102) creating a sine-wave with adjustable frequency.⁶ The photodetector output was measured on an HP 4395A

⁶ The drive amplitude was kept to $100 \text{ mV}_{\text{rms}}$ to limit the peak current to 1 mA and prevent Joule heating from damaging the narrow-gauge coil wire.

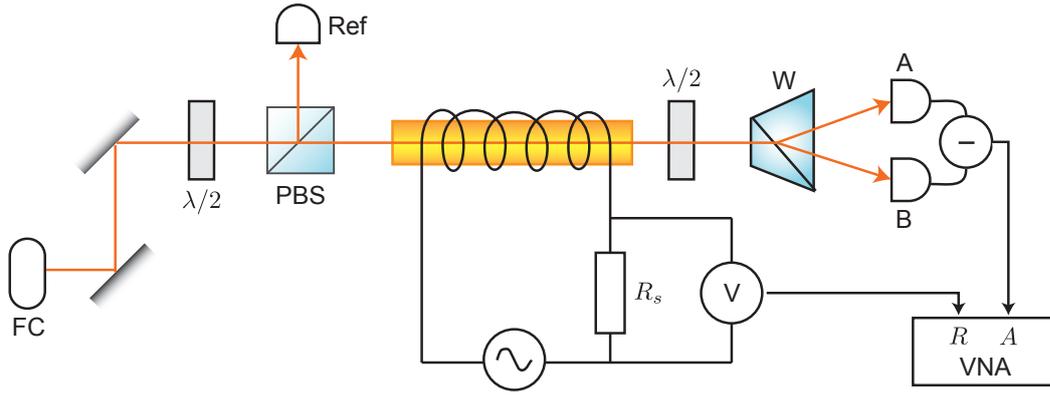


Figure 6.6: Optical schematic for the ‘electro-optic BEC’. A terbium-glass cylinder is wrapped in a solenoid which is driven by an AC current source, inducing a Faraday rotation on the transmitted light beam as measured by the differential photodetector. A VNA records the photodetector output (port A) and normalises to the current in the coil measure through the sense resistor R_s (port R).

vector network analyser (in spectrum analysis mode), with a clear peak observed at the drive frequency (Figure 6.7).

The frequency response of the photodetector was calibrated with the VNA (in ‘A/R’ network mode) to sweep the drive frequency while measuring the strength of the resulting Faraday signal (Figure 6.8). Network analysis simultaneously measures the response of the detector (A) and the voltage across the sense resistor (R) and takes the ratio of the two to normalise the measured Faraday rotation to the current flowing through the coil to correct for the frequency-dependence of the coil impedance.

The detector displays a linear response at low frequencies and rolling-off at higher frequencies, with a 3 dB point at 600 kHz.⁷ This reduced gain results in less signal at higher Larmor frequencies, but for a shot-noise-limited measurement the SNR remains the same as the shot-noise is also subject to reduced gain. However, if the gain is reduced too far the measurement will be dominated by the electronic noise floor.

To estimate the sensitivity of the polarimeter, we equate (6.1) and (6.3) to find that the smallest measurable rotation with a measurement bandwidth Δf is

$$\theta_{\min} = \sqrt{\frac{e\Delta f}{2R\kappa P}}. \quad (6.7)$$

For $P = 5 \text{ mW}$, $\Delta f = 1 \text{ kHz}$ and $\kappa = 0.9$ this gives $\theta_{\min} \sim 10^{-7}$ rad, comparing favourably to the estimated (whole cloud) Faraday rotation $\theta \sim 10^{-5}$ rad from (3.4).

⁷ The measured bandwidth is significantly lower than the expected 1.2 MHz bandwidth of the trans-impedance amplifier (see §6.2). Initial investigations did not resolve the discrepancy and were abandoned once construction of the new detector (see §6.5) was complete.

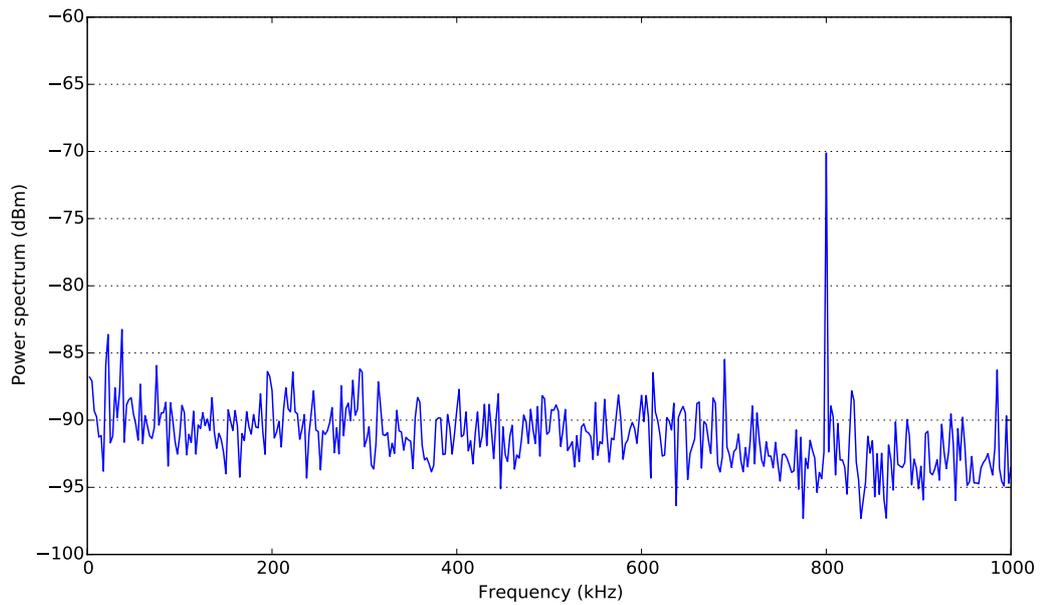


Figure 6.7: Example of the calibration possible using the electro-optic BEC. Driving the solenoid with a function generator at 800 kHz generates an oscillating polarisation rotation in the propagating beam which appears as a clear peak in the measured power spectrum, 20 dB above the noise floor. RBW was 300 Hz, reference level was -60 dBm.

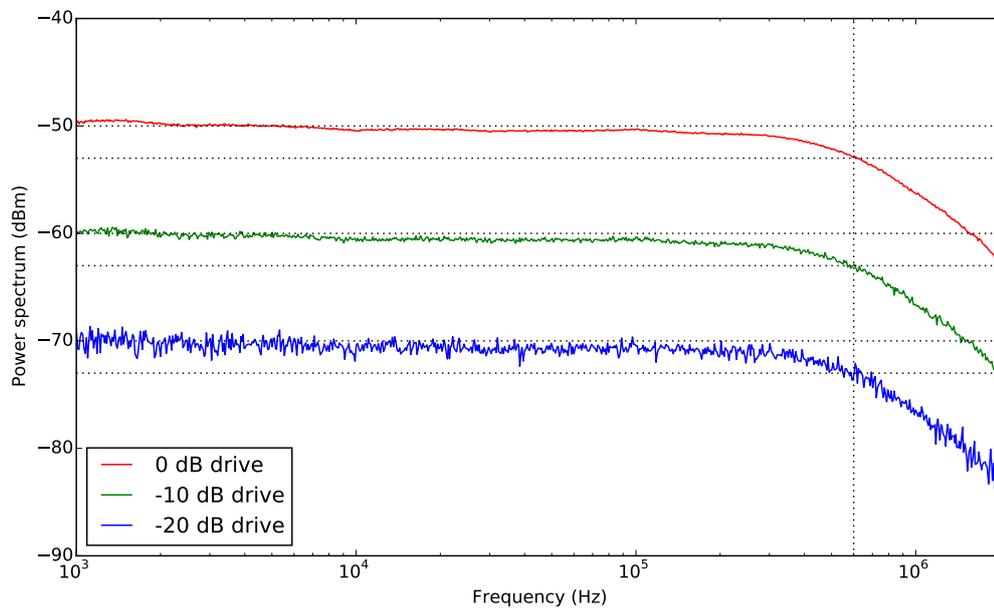


Figure 6.8: Measurement of photodetector bandwidth by driving the electro-optic BEC with a VNA, showing the response rolls off with a 3 dB point of 600 kHz (dashed lines). Driving the coil at different powers varies the magnetic field strength inside the coil, showing the response is linear across multiple orders of magnitude. RBW was 300 Hz.

6.5: Revised photodetector design

Prior to the selection of the magic wavelength (790 nm) as the detuning of the Faraday probe, the level of trap perturbation induced by the probe beam (see §2.3) was a primary concern for minimally destructive Faraday measurement. The dipole force induced by the probe would ‘squeeze’ the BEC and result in non-trivial time-evolution, so the beam power should be as low as possible to minimise the perturbation.⁸

Similarly, a Faraday beam with wide waist at the BEC reduces the intensity gradient across the cloud (and hence the dipole force),⁹ although only a small fraction of the probe passes through the BEC, wasting most of the beam power. The components of the initial photodetector constructed for this project were therefore chosen for their noise properties, to permit Faraday measurements at low probe powers [181].

However, detuning to the magic wavelength removes this restriction, with the only technical considerations on probe power arising from the photodiode damage threshold and total power available from the Faraday laser.¹⁰ This encourages measuring in the ‘bright beam’ regime using a modified photodetector design with larger-area photodiodes for increased saturation/damage threshold powers.

Although the larger area photodiodes have higher capacitance, the cutoff frequency remains well above the Larmor frequency. The limiting factor on the photodetection bandwidth is bandwidth of the transimpedance amplifier, which is set by the op-amp used and the chosen feedback resistor. Bright beam measurement enables a different op-amp to be chosen for transimpedance amplification with increased gain-bandwidth product, even increasing the photodetection bandwidth beyond that of the initial design. Furthermore, the second-stage amplifier which was intended as a line-driver can apply additional gain to match the amplitude of the Faraday signal to the dynamic range of the digital acquisition system (see §6.6).

A second-generation detector was designed and constructed by P. Pakkiam to achieve these goals, with the differences summarised in Table 6.1. Full characterisation of the new detector using the electro-optic BEC including details on design, construction, noise and common-mode rejection ratio (CMRR) analysis, is presented in the Honours thesis of P. Pakkiam [173]. The results in subsequent chapters were captured using this second-generation detector.

⁸ Not only is shot-noise-limited photodetection technically challenging at low light levels, but the detuning must also be reduced to maintain the same SNR, provided the far-detuned limit is obeyed for the polarisability expansion (§2.1) to be valid.

⁹ Beam-shaping techniques such as spatial-light modulation (SLM) or shaped beam diffusers (like the THORLABS ED1-C series) could also be used to produce a ‘top-hat’ intensity profile to reduce the intensity gradient, but these were not available at the time.

¹⁰ Note that since detuning is fixed by setting $\lambda = 790$ nm, adjusting the beam power is the only way to directly control the scattering rate and hence the measurement signal to noise (see §3.3). At present the scattering rate is low (see §5.9) and is limited by the available laser power and transmission efficiency.

	First generation	Second generation
Fabrication technique	Breadboard	PCB
Photodiode	Hamamatsu S5971	Hamamatsu S1223-01
Active area	1.1 mm ²	13 mm ²
Maximum optical power	6 mW	13 mW
Capacitance at -20 V bias	6 pF (300 MHz BW)	20 pF (20 MHz BW)
Transimpedance amplifier	OP470	ADA4898
Transimpedance gain	5.11 V/mA	1.0 V/mA
Second stage gain	2× (DC coupled)	100× (AC coupled)
Measured bandwidth	600 kHz	8 MHz
Noise-equivalent power	20 μW	140 μW

Table 6.1: Comparison of the properties of the first- and second-generation differential photodetectors used for Faraday measurement.

6.6: Data acquisition

The analog signal produced by the photodetector is digitised by an analog-digital converter (ADC) for post-processing. In accordance with the Nyquist theorem, the acquisition rate of the digitiser should be at least twice the Larmor frequency, $f_S > 2f_L$, to avoid aliasing and enable direct observation of the signal. The capabilities of the data acquisition card therefore limit the maximum strength of the bias magnetic field.

For these experiments a NATIONAL INSTRUMENTS PCIE-6363 DAQ was used, which has a maximum acquisition rate of $f_S = 2 \text{ MS/s}$ in single-channel mode.¹¹ The Nyquist frequency is therefore 1 MHz, which is ample for sampling at the desired bias field strength of 1 G ($f_L = 700 \text{ kHz}$).

Measurement at stronger bias fields, for which the Larmor frequency is greater than the Nyquist frequency, is also possible using super-Nyquist sampling [182, 183, 184]. The Faraday signal in this instance will be under-sampled (aliased), producing an ‘image’ signal that appears at the reduced frequency¹² $f_S - f_L$. Typically the ADC has a limited bandwidth (or an anti-aliasing filter) to avoid aliasing noise at higher frequencies into the measurement, limiting the maximum Larmor frequency. The small signal bandwidth of the PCIE-6363 has a -3 dB point of 1.7 MHz [185], permitting bias fields of up to 2.4 G to be used.

However, the roll-off of the photodetector’s transimpedance gain also produces an upper limit on the maximum possible bias field. The first-generation detector rolls off at 600 kHz, which limits the maximum bias field strength more than the ADC sam-

¹¹ In multi-channel mode, the maximum aggregate sample rate across all channels is 1 MHz, which inhibits the simultaneous capture of diagnostic measurements.

¹² For signals in the ‘second Nyquist zone’, $\frac{1}{2}f_S < f_L < f_S$, though similar relations exist for higher frequencies.

pling rate, whereas the second-generation detector extends the roll-off to 8 MHz [173]. Making use of the stronger bias fields permitted by this higher detector bandwidth (up to 11 G) would require a much faster ADC card,¹³ though the present work does not consider such high bias values.

It is also important that the signal be amplified to take full advantage of the dynamic range of the acquisition card, to prevent quantisation noise from affecting the measurement. A STANFORD SR560 variable gain pre-amplifier was used to apply a gain of $20\times$ to match the signal out of the photodetector with the ± 10 V input range¹⁴ of the PCIE-6363 and minimise the effect of quantisation noise.

It should be noted that polarisation drift can cause the input to clip the voltage range of the acquisition device over the course of the experiment, particularly if the post-amplification gain is too high. The polarimeter therefore needs to be periodically rebalanced to eliminate the DC offset of the signal by adjusting the $\frac{\lambda}{2}$ -waveplate. Typically the raw photodetector output can be zeroed to within 20 mV, though the waveplate typically needs adjusting every ~ 15 min to prevent drifting beyond 100 mV.

An autobalancing photodetector [128] is capable of automatically ‘rebalancing’ this measurement, by generating zero output signal even when the optical power on each photodiode is not exactly equal. It contains an electronic servo loop that eliminates slow polarisation drift in the polarimeter and further reduces the influence of common mode intensity fluctuations. Although such an autobalancing detector was constructed, it was not successfully integrated into the apparatus because of difficulty in reliably locking the servo loop. Manual rebalancing of the polarimeter by adjusting the $\frac{\lambda}{2}$ -waveplate proved successful enough to advance the project, though fixing the autobalancer is an extension to the present work.

With the required components to perform continuous Faraday measurement assembled, next I will describe the tools to analyse and process the measured signal.

6.7: Structure of the Faraday signal

The high-precision polarimeter is highly sensitive to long-term drifts in the polarisation of the probe beam, so the raw signal shows fluctuations over timescales on the order of hundreds of milliseconds (Figure 6.9). A Glan-Laser polariser in the probe beam before the cell purifies the polarisation state leading to the atoms, though small residual perturbations are detected by the polarimeter. These fluctuations are primarily due to ambient temperature variations affecting the PM-fiber (see §5.3). Further fluct-

¹³ An ALAZARTECH ATS9626 250 MS/s 16-bit DAQ was purchased for this purpose, and will be integrated into the control hardware in the near future.

¹⁴ Although the PCIE-6363 supports input voltage ranges down to ± 0.1 V, these are not currently implemented in the current generation of the experiment control system.

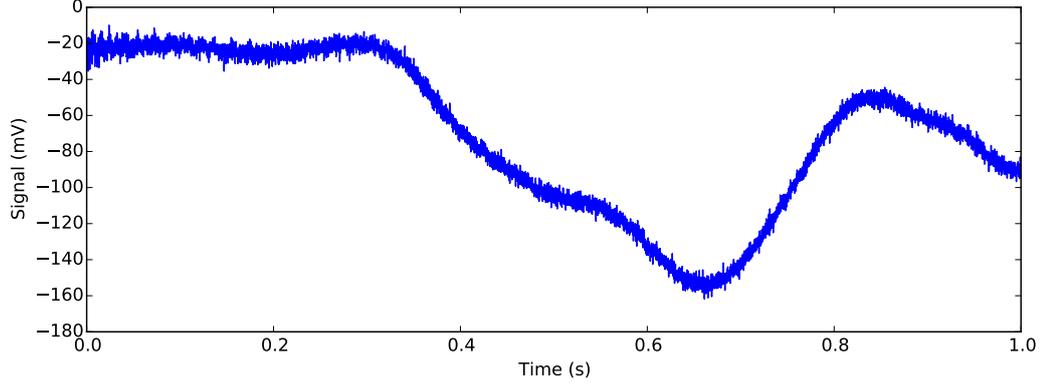


Figure 6.9: A typical polarimeter measurement of the raw Faraday signal, showing slow ($\lesssim 10$ Hz) drift due to probe beam polarisation fluctuations. The Faraday signal is buried under the shot-noise.

tuations arise from switching on the bright probe beam, which causes a small amount of heating in the polarisation optics between the polariser and atoms.

Fortunately, the slow timescale of these perturbations compared to the oscillation (Larmor) frequency means they are easily removed by filtering out the low-frequency components of the signal. The Faraday signal itself exists in a narrow frequency band (~ 10 kHz wide) around the Larmor frequency (~ 700 kHz), so a band-pass filter can be applied to eliminate the low-frequency drift,

$$V_F(t) = \mathcal{F}^{-1} [W(f - f_L) \cdot \mathcal{F} [V(t)]]$$

$$\text{with } W(f) = \begin{cases} H \left(\frac{1}{2} + \frac{f}{f_R} \right) & \text{for } |f| \leq \frac{f_R}{2} \\ 0 & \text{otherwise,} \end{cases} \quad (6.8)$$

where f_L is the Larmor frequency, $f_R = 10$ kHz is the filter bandwidth, and H is the window function (typically the Hanning window [186]). This has the added advantage of both eliminating narrow-band noise in other frequency bands (such as power supply noise), and reducing the contribution from spectrally broad noise sources (such as shot-noise and Johnson noise). The resulting signal is clearly visible above the noise floor (Figure 6.10A).

The BEC is exposed to the probe laser at $t = -20$ ms to ensure any switch-on transients in the polarimeter and the probe AOM do not affect the measurement. The polarimeter output shows white noise corresponding to the shot-noise on the beam for the first 20 ms, giving a measure of the shot-noise level. At $t = 20$ ms an rf $\frac{\pi}{2}$ -pulse is applied and the spin is tipped into the x - z plane, resulting in Larmor precession and the generation of a Faraday signal at the Larmor frequency. Coherent spin dephasing then occurs and the signal ‘decays’ over the first 50 ms (Figure 6.10B), after which non-trivial spin-dynamics dominate and the signal demonstrates multiple collapses and revivals (discussed later in §7.4 and §7.9).

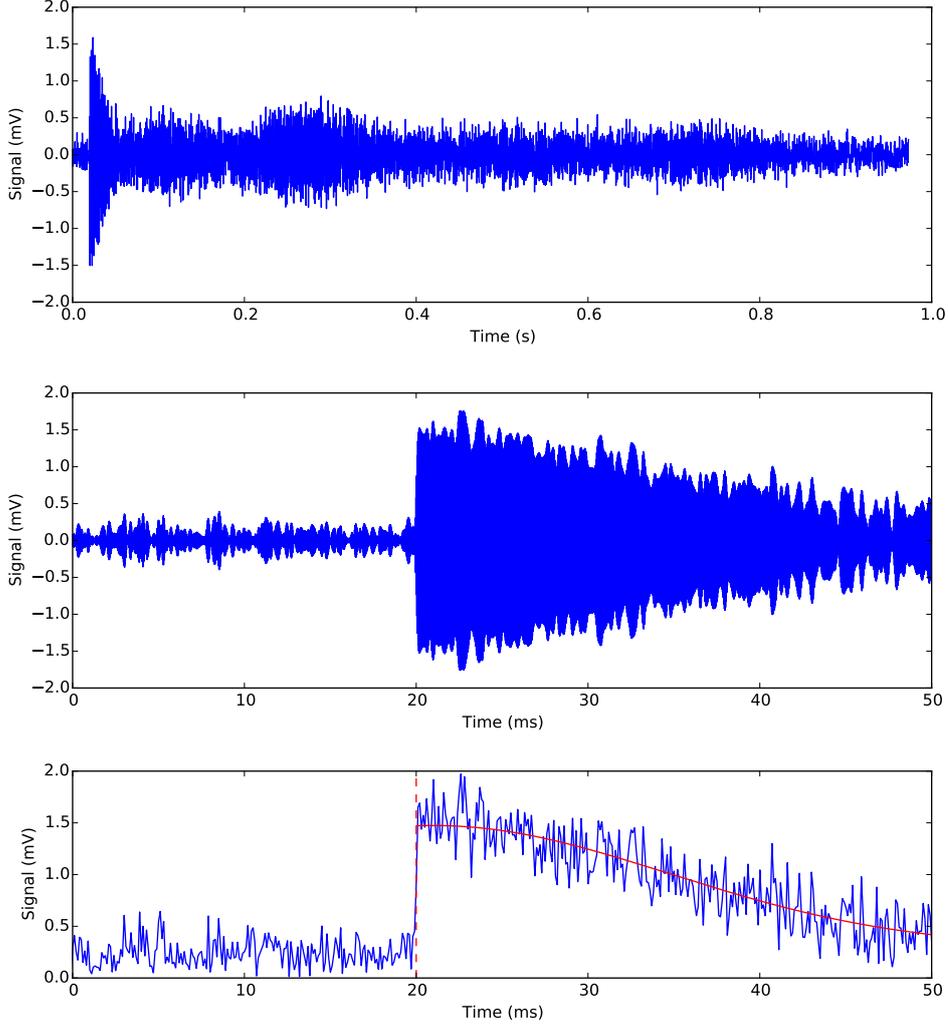


Figure 6.10: Fourier-filtered polarimeter measurement, with a band-pass filter of width 10 kHz applied around the Larmor frequency of 695 kHz (top), a close-up of the first 50 ms (middle), and the extracted envelope (bottom), giving a fitted decay constant of 19.4(4) ms. The shot-noise level is characterised by the signal envelope for $t < 20$ ms.

The Faraday signal oscillates at the Larmor frequency, with an amplitude describing the total magnetisation of the BEC. A loss of magnetisation can occur due to losing atoms from the trap, dephasing of the frequency components, spatial separation of the cloud, or other modulation of the transverse spin length. The envelope of the signal $V_E(t)$ can be extracted by applying the Fourier shifting theorem to the measured signal to remove the rapid Larmor oscillation,

$$V_E(t) = \mathcal{F}^{-1} \left[W(f) \cdot \mathcal{F} \left[e^{-2\pi i f_L t} V(t) \right] \right]. \quad (6.9)$$

This gives a measure of the amplitude of the Faraday signal over time, which can be fit (see §3.5) to give the dephasing rate (Figure 6.10C). Significantly, the amplitude remains greater than the shot-noise level for the duration of the measurement,

demonstrating Zeeman coherence and continuous spin projection measurement over the entire interval.

Extraction of the signal envelope also provides a method for averaging multiple shots together. The phase of the Larmor oscillation is sensitive to small fluctuations in the magnetic field during the experiment, so even identical experimental sequences will not give Faraday signals that remain in phase with each other. Combining the raw measurements will therefore average out to zero. In contrast, averaging the envelopes allows the signal amplitude to be determined more precisely.

6.8: Short-time Fourier transform analysis

The Faraday signal in a magnetic field gradient contains a range of different Larmor frequencies that correspond to different positions in the BEC. Each frequency component has a time-dependent amplitude that describes the spatial structure of the cloud and dynamics which are occurring. The filtered time-domain signal describes the cloud as a whole, but decomposing the Faraday signal into its Fourier components enables observation of the time-dependence of the spatial distribution.

Short-time Fourier transform (STFT) analysis is a useful way to study how the amplitude of the Fourier components of the signal vary over time. The STFT algorithm divides the raw signal into a number of segments called ‘windows’, and takes the Fourier transform each window to obtain its frequency spectrum. Stitching the power spectrum of each window together produces a spectrogram which shows how the spectral components change over time (Figure 6.11). Applying the STFT algorithm to the Faraday signal provides information about the time evolution of the Larmor signal, and separates it from the background noise at other frequencies (Figure 6.12).

The spectrogram enables a clear distinction to be made between noise sources. In particular, wide-band shot and flicker ($1/f$) noise can be separated from noise peaks due to power-supplies (primarily appearing between 200–400 kHz) which appear as horizontal lines in the spectrogram. The time-independence of these contributions delineates them from the signal at 695 kHz, which only appears at $t = 20$ ms when the $\frac{\pi}{2}$ -pulse occurs, and fades as the Faraday signal dephases.

Zooming in on the spectrogram about the Larmor frequency (Figure 6.13) shows a rich signal with non-trivial structure. In particular, peak brightness corresponds to the instantaneous Larmor frequency at a given time, and its amplitude gives the strength of the Faraday rotation. Immediately it can be seen that the Larmor frequency varies approximately sinusoidally at the power line frequency (50 Hz). This corresponds to magnetic field fluctuations at the position of the atoms introduced by power supplies located near the science chamber. These deviations are small ($< 0.1\%$) but can be easily measured and corrected to demodulate the signal (discussed in §7.2).

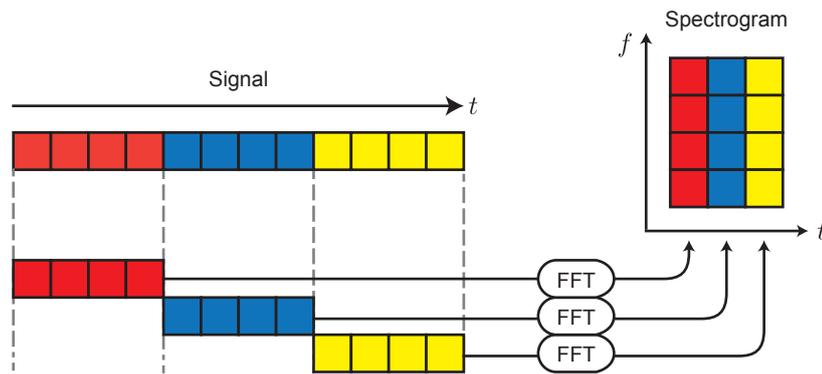


Figure 6.11: Schematic of the STFT algorithm. The signal is divided into sections of equal length called 'windows' which are individually Fourier-transformed and stacked together to form a spectrogram.

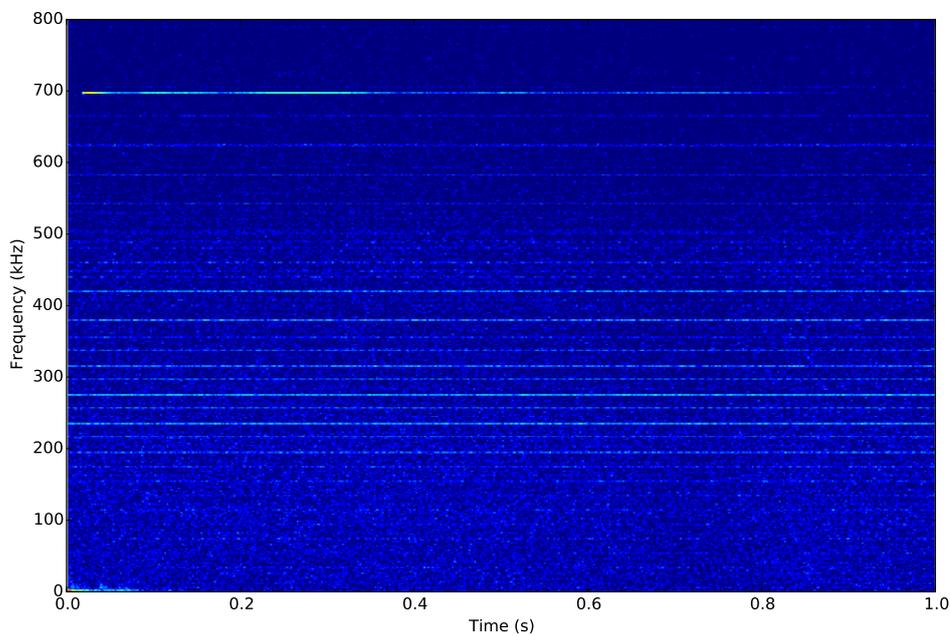


Figure 6.12: Spectrogram showing the STFT of a typical measurement. Flicker ($1/f$) noise is evident at low frequencies, with white noise dominating at high frequencies. The horizontal bands correspond to noise pick-up at particular frequencies, but the much brighter line near 700 kHz is the Faraday signal. Window size 5 ms.

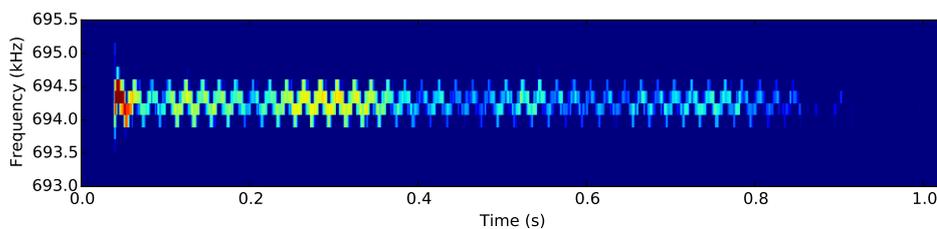


Figure 6.13: Close-up of the STFT spectrogram of [Figure 6.12](#) about the Larmor frequency, showing amplitude and frequency modulation of the signal.

6.9: Overlapping and oversampling

The spectrogram can be made easier to interpret by oversampling it in both the time and frequency domains, which interpolates the spectrogram and makes its features clearer. Oversampling in the time domain is achieved by dividing the signal into further windows that partially overlap with each other (Figure 6.14). This provides a finer-grained interpretation of the time-dependence of the associated Fourier spectra, although the windows are not independent so the resolution has not been increased.

Oversampling can be achieved in the frequency domain by zero-padding each window. This increases the number of data points in each window, which causes the FFT to interpolate the resulting spectrum (Figure 6.15). Adding zeros adds no new information to the window, so the resolution of the Fourier spectrum is still fundamentally limited by the window size.

As the fundamental resolution of the spectrogram is changed by neither overlapping nor zero padding, the resulting interpolated spectrogram will demonstrate blurring. However, it is still advantageous to interpolate to reduce pixellation of the spectrogram and track the signal in both time and frequency space (Figure 6.16).

6.10: Role of window size

In the STFT analysis, each vertical slice of the spectrogram is obtained by taking the Fast Fourier Transform (FFT) of a subset of the signal. The number of datapoints in each of these subsets is the ‘window size’ N_w , which is a free parameter. The fundamental (uninterpolated) resolution of the spectrogram in the frequency domain is

$$\Delta f = R_s/N_w,$$

where R_s is the sample rate of the original acquisition. Choosing a large window therefore increases the resolution in the Fourier domain but effectively time-averages more of the signal together, decreasing the resolution in the time domain. This behaviour is known as the Gabor limit [187], which is the signal-processing analog of the Heisenberg uncertainty relationship, stating that a signal cannot be localised in both time and frequency to arbitrary precision.

It is therefore important to choose a window size that permits observation of dynamics over the timescale of interest. In the case of frequency-modulated signals, the resolution in frequency space must be small enough to resolve the frequency deviation, while the resolution in the time domain must be small enough to resolve within the frequency modulation cycle. For the power line modulation observed in the Faraday signal, a window size of 5 ms balances this tradeoff (Figure 6.17).

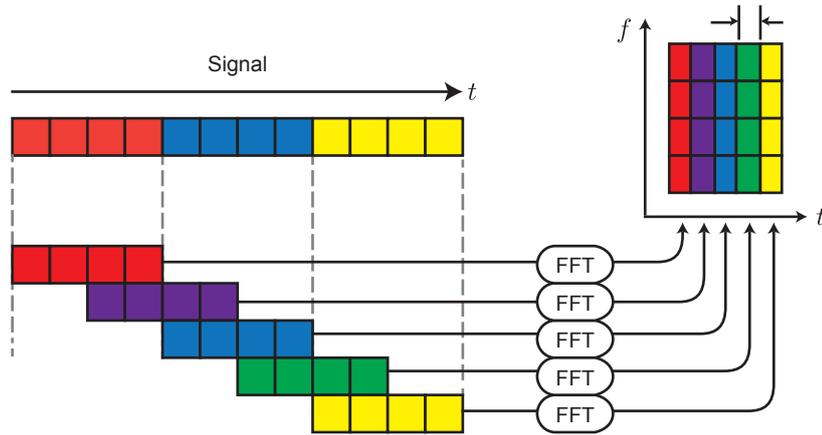


Figure 6.14: Time domain oversampling of the STFT is achieved by dividing the signal further into windows that overlap, resulting in Fourier transforms at more time points over the same interval.

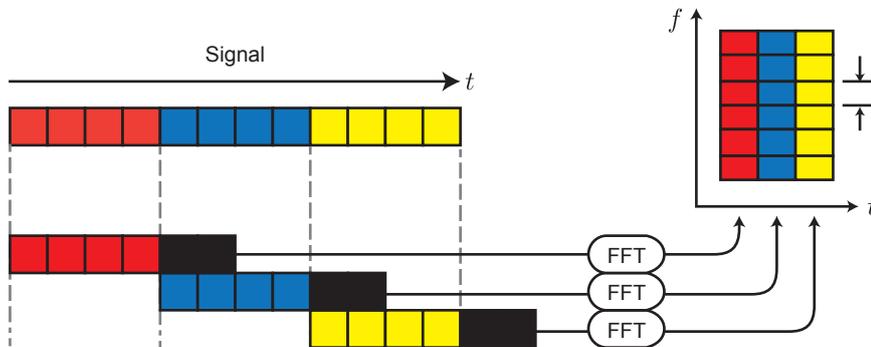


Figure 6.15: Frequency domain oversampling of the STFT is done by zero-padding (black boxes) each windowed subset of the signal, causing each Fourier transform to contain more points and resulting in interpolation.

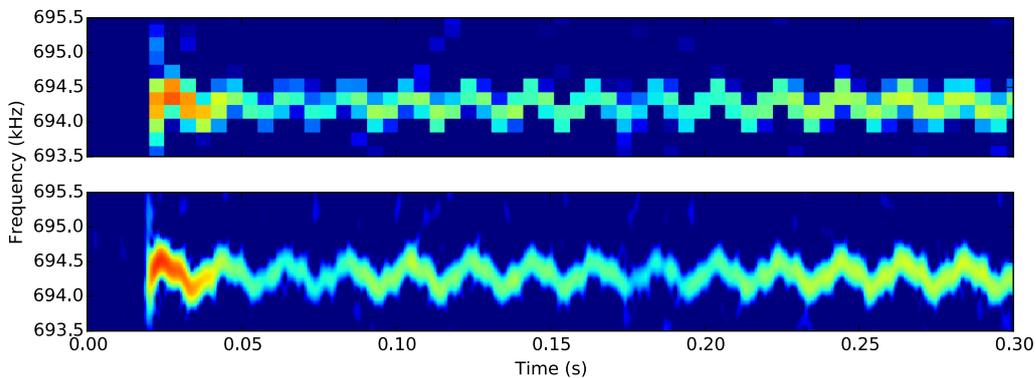


Figure 6.16: Comparison of the first 300 ms of a Faraday signal analysed using a standard STFT with window size 5 ms and sampled at 5 ms/200 Hz (top), against the same STFT 10x oversampled at 0.5 ms/20 Hz (bottom).

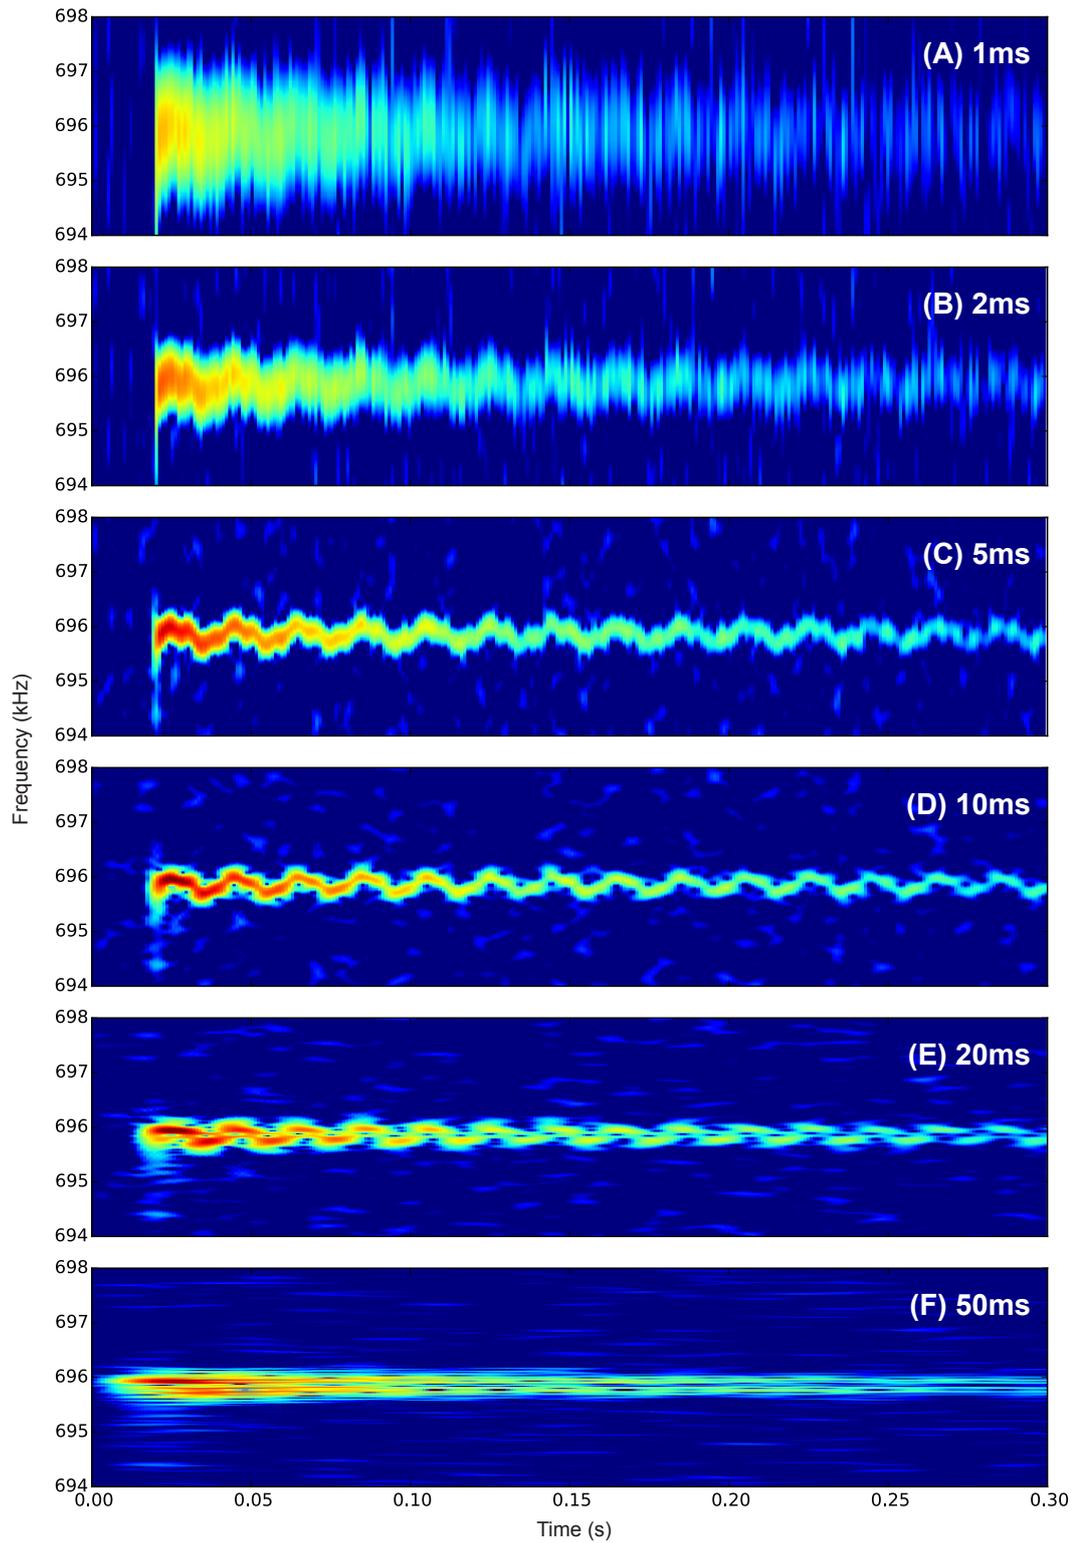


Figure 6.17: Comparison of spectrograms centred on the Larmor frequency for STFTs with varying window size (A)–(F). Note that windows that are too short/long lead to spectrograms that are ‘smeared’ vertically/horizontally. Spectrograms are oversampled at 1 ms/20 Hz intervals for consistency.

6.11: Summary

This chapter detailed the high-precision polarimeter used to measure polarisation rotation of the Faraday probe beam. Design considerations of the custom photodetector used to measure the Faraday signal were discussed, and the sensitivity quantified in terms of the measured noise floor.

A Faraday signal produced by a BEC in a uniform magnetic applied field was captured and methods to analyse it were considered. Fourier filtering in the time-domain provided a way to extract the signal envelope, showing collapse and revival of the Faraday rotation signal from a BEC. A spectrogram was constructed using short-time Fourier transforms, which showed the time-dependence of the frequency components of the signal, indicating the presence of both amplitude and frequency modulation. The roles of oversampling and window size in the STFT analysis were considered, and optimised to be representative of the structure observed in the Faraday signal.

The Faraday signal can now be used as a measurement tool to infer properties about the atoms, magnetic field, and probe beam. I will next characterise and optimise the signal in preparation for using it to image the BEC.



Signal calibration and optimisation

In this chapter I analyse the Faraday signal obtained from a BEC in further detail and discuss its features. The signal is then used as a diagnostic to determine a microwave Rabi frequency, the optimal microwave detuning to eliminate the quadratic Zeeman effect, and the optimal probe polarisation to eliminate the vector light shift.

7.1: Magnetic field calibration

The magnetic field coils used to generate the bias field for the Faraday measurement are powered by a multi-channel homemade voltage-controlled current source built in-house called the ‘Mag-neat-o’. The coils are rectangular-form solenoids in approximately Helmholtz configuration, producing a nearly uniform magnetic field at the position of the atoms. However, the net magnetic field at the position of the atoms is the vector sum of the generated field and the ‘ambient’, or ‘background’ magnetic field, which cannot be precisely determined from *ex vacuo* measurements.

Using the Faraday measurement to characterise the background and applied fields enables direct observation of their effect on the atoms. This calibration can be performed by measuring the Larmor frequency of the atoms at a range of control voltages, since the Larmor frequency depends solely on the net magnetic field at the position of the atoms.

In the case of a magnetic field bias in the y -direction with the control voltages for the other components set to null the background field, the net magnetic field strength is

$$|\mathbf{B}| \equiv B_{\text{net}} = \sqrt{B_x^2 + B_y^2 + B_z^2} \approx |B_y| + \frac{B_{\perp y}^2}{2B_y}, \quad (7.1)$$

where $B_{\perp y}^2 = B_x^2 + B_z^2 \ll B_y$ is the residual background field perpendicular to the bias B_y . Applying a voltage V_y to the transinductance amplifier driving the y -coils results in a net magnetic field primarily directed along the component B_y , where

$$B_y = \alpha_y V_y + B_{y0}. \quad (7.2)$$

Here α_y is the calibration of the y -coils¹ and B_{y0} is the y -component of the background field. Provided the field generated by the coils opposes the background field, the control voltage required to null this field component is therefore

$$V_{y,\text{null}} = -B_{y0}/\alpha_y. \quad (7.3)$$

In the linear Zeeman approximation, the Larmor frequency can be expressed in terms of the voltage applied to the bias coils $V_{y\text{-bias}}$ using (3.5) as

$$\omega_L \approx \gamma|B_y| = \gamma\alpha_y|V_y - V_{y,\text{null}}|. \quad (7.4)$$

The coil calibration can therefore be obtained by measuring the centre frequency of the Faraday signal for different field strengths as produced by varying the coil control voltage.

It should be noted however that varying the bias magnetic field strength in this way changes the Larmor frequency and causes the rf $\frac{\pi}{2}$ -pulse to become off-resonant. As the magnetic field strength is varied over a large range, the fidelity of the spin tip is greatly reduced, resulting in a Faraday signal with vanishing amplitude and making the centre Larmor frequency difficult to determine.

The pulse parameters can be optimised at any given control voltage by setting the frequency $f_{\pi/2}$ to the estimated/measured Larmor frequency and varying the pulse duration $t_{\pi/2}$ to achieve the spin tip. A successful $\frac{\pi}{2}$ -pulse is realised in a y -bias when

$$\langle F_y \rangle = \frac{N_{+1} - N_{-1}}{N} = 0, \quad (7.5)$$

as observed by measuring the Zeeman substate populations in the Stern-Gerlach separated absorption image. Typically our rf coil achieves a Rabi frequency² of $\Omega = 17.6$ kHz with a $\frac{\pi}{2}$ -pulse achieved when $t_{\pi/2} = 14.2$ μs and $f_{\pi/2} = 700$ kHz in a y -bias.

Provided there is an approximate existing calibration for the bias coils, the magnetic field strength can be estimated for any set of control voltages. This allows the frequency of the rf pulse to be estimated for each control voltage with the linear Zeeman effect (derived in §7.6) via (7.23) as

$$f_{\pi/2} \approx (702 \text{ kHz/G}) \cdot B_{\text{est}}. \quad (7.6)$$

Because the Larmor frequency is being swept over such a large range (50–700 kHz), the rf Rabi frequency does not remain constant. However, neither the pulse frequency nor duration need to be accurate as the centre Larmor frequency can be determined from the Faraday spectrogram even for an imperfect pulse.³

¹ The calibration factor α_i (units G/V) for each coil pair is the product of the transconductance gain of the Mag-neat-o (in A/V) and the magnetostatic factor (units G/A) that depends on the geometry of the coils.

² Note that the rf Rabi frequency depends on the coil alignment relative to the magnetic field bias.

³ Recall from §2.5 that the Faraday signal is proportional to the transverse spin length, so an imperfect $\frac{\pi}{2}$ -pulse results in a partially-tipped spin and reduced signal amplitude by (2.33).

Setting the x - and z -coils to null their corresponding components of the background field,⁴ the measured Larmor frequency as a function of y -coil control voltage shows a linear dependence (Figure 7.1). The quality of this fit demonstrates that the response of the Mag-neat-o and coils is indeed linear at these drive currents. From this measurement, the calibration of the y -coil was found to be $\alpha_y = 2.134(6)$ G/V and $B_{y0} = 352(2)$ mG.

It is also possible to perform the coil calibration without requiring that the perpendicular magnetic field components be nulled with the other coils. This removes the requirement to first perform vector magnetometry to determine the background field, and eliminates cross-coupling between calibrations. This is especially useful when the background field drifts intermittently (day-to-day), as is common to many ambient laboratory settings.

By taking the square of the Larmor frequency, the terms corresponding to the component of interest can be separated from the other components. For example, the x -coil can be calibrated by expressing the Larmor frequency as

$$\begin{aligned}
 (\omega_L/\gamma)^2 &= B_{\text{net}}^2 = B_x^2 + B_y^2 + B_z^2 \\
 &= (\alpha_x V_x + B_{x0})^2 + (\alpha_y V_y + B_{y0})^2 + (\alpha_z V_z + B_{z0})^2 \\
 &= \alpha_x^2 V_x^2 + 2\alpha_x B_{x0} V_x + [B_{x0}^2 + (\alpha_y V_y + B_{y0})^2 + (\alpha_z V_z + B_{z0})^2] \\
 &= aV_x^2 + bV_x + c,
 \end{aligned} \tag{7.7}$$

where a , b and c are constants obtained by fitting $B_{\text{net}}^2 = (\omega_L/\gamma)^2$ as a quadratic in V_x .

⁴ The background-nulling voltages were originally measured with Ramsey interferometry [188].

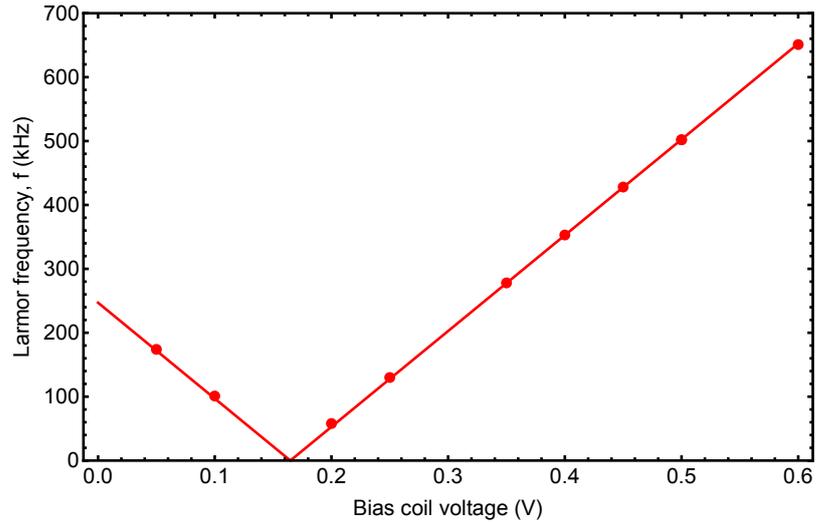


Figure 7.1: Measuring the Larmor frequency of the Faraday signal for different applied y -bias coil voltages provides a direct calibration of the magnetic field strength produced by a given coil voltage.

The calibration constants can then be computed as

$$\alpha_x = \sqrt{a}, \quad B_{x0} = \frac{b}{2\sqrt{a}}, \quad V_{x,\text{null}} = \frac{b}{2a}, \quad (7.8)$$

independent of the magnetic field strength in the y - or z -directions.

Performing this analysis on the three bias coil pairs (Figure 7.2) shows strong adherence to the model even when the perpendicular background fields are not nulled (demonstrated in the x - and z -coil measurements). Furthermore, the nulling voltage can still be determined even if it is beyond the range of control voltages available (as in the case of the z -coil). The fitting coefficients are then used to compute the coil calibration as per (7.8), to obtain the values in Table 7.1.

One primary advantage of this method is that the high quality of the quadratic fit results in small uncertainties on the fit coefficients, and hence provides precise estimates for the calibration parameters. It also significantly reduces the time taken to determine the calibrations α_i and the background fields B_{i0} as compared to the conventional approach of measuring Rabi spectra via Stern-Gerlach absorption imaging, which typically requires ~ 15 shots to measure B for one set of control voltages. Another benefit is it is calibration free and doesn't suffer from resonance shifts that may be present when using rf Rabi spectroscopy to determine ω_L and thus B .

7.2: Measuring and correcting power-line modulation

The spectrogram of the Faraday signal (§6.8) shows that the Larmor frequency appears to oscillate approximately sinusoidally at the power line frequency of 50 Hz. These small frequency deviations are produced by magnetic field fluctuations at the position of the atoms causing frequency modulation of the Faraday signal. The fluctuations in the magnetic field originate in electronic components (such as power supply transformers) that are located near the unshielded science chamber.

The corresponding magnetic field fluctuations are expected to have components at odd harmonics of the fundamental supply frequency [189] of 50 Hz, making it possible to fit and remove the variations in post-processing. The modulation is fitted by extracting the centre Larmor frequency at each time using a simple weighted estimator of the centroid,

$$f_L(t) = \frac{\sum_i f_i y_i^2}{\sum_i y_i^2}, \quad (7.9)$$

where $y_i = \text{STFT}\{t, i\}$ is the value of the i^{th} frequency bin of the STFT slice at time t , and f_i is the centre frequency of the i^{th} bin. This extracts the time-dependence of the Larmor frequency, which can then be fit with a sum of 50 Hz harmonics to model the frequency modulation. In practice it is sufficient to use only the first two terms in the harmonic expansion to model the behaviour,

$$f_L(t) = f_{L0} + A_{50} \sin(2\pi \times 50t - \phi_{50}) + A_{150} \sin(2\pi \times 150t - \phi_{150}). \quad (7.10)$$

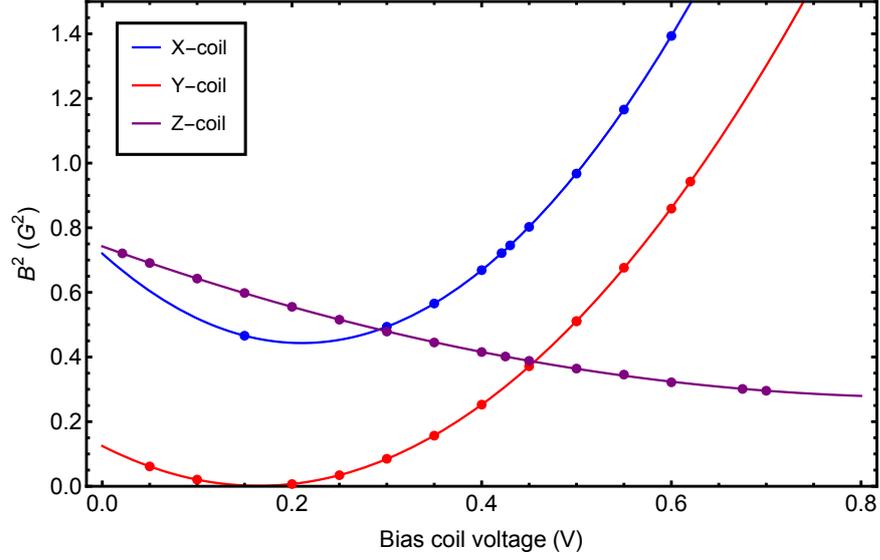


Figure 7.2: Calibration of each bias coil pair using the net magnetic field without requiring transverse components of the field to be nulled. The measured field is quadratically related to the control voltage, in good agreement with (7.7).

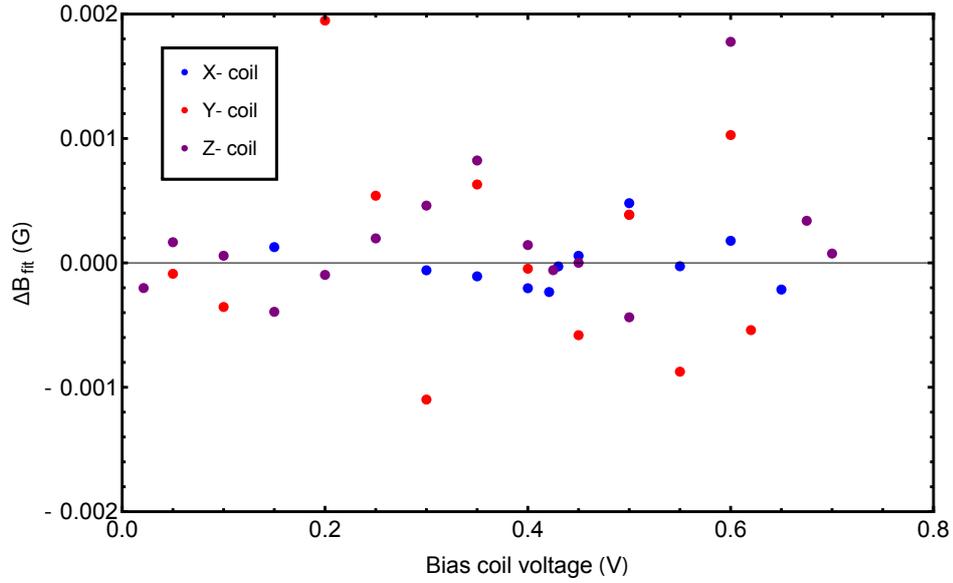


Figure 7.3: Residuals of the quadratic coil fit showing the accuracy of the fit, expressed in terms of the error in the predicted value $\Delta B_{\text{fit}} = (B_{\text{fit}}^2 - B_{\text{net}}^2)/2B_{\text{net}}$. The RMS error is $700 \mu\text{G}$.

Coil	Calibration, α_i (G/V)	Background field, B_{i0} (mG)	Nulling voltage $V_{i,\text{null}}$ (mV)
x	2.502(1)	526.3(8)	210.3(4)
y	2.130(2)	351(1)	164.8(7)
z	0.772(5)	683(1)	885(7)

Table 7.1: Bias coil calibration factors as measured using the quadratic fit method.

Extracting the time dependence of the Larmor frequency in this way yields an accurate depiction of the frequency modulation, and the resulting model fits the extracted data well (Figure 7.4).

The magnetic field induced by the power line variations has a different amplitude and harmonic composition in different directions [190]. From (7.1) the Larmor frequency can be expressed in terms of the components parallel (\parallel) and perpendicular (\perp) to the bias field as

$$\omega_L \approx \gamma B_{\parallel} + \frac{\gamma B_{\perp}^2}{2B_{\parallel}}. \quad (7.11)$$

Since $B_{\parallel} \gg B_{\perp}$ by choice of the bias field, the second term is heavily suppressed and the Larmor frequency is only sensitive to the component of the power line noise parallel to the bias field.

The structure of the frequency modulation is therefore expected to depend on the direction of the bias field. Indeed distinct differences in the structure of the Larmor signal are observed with the bias field oriented⁵ along x' instead of y (Figure 7.5). Notably the amplitude of the fluctuations is increased with the bias along x' , and the contribution of higher harmonics of the line frequency are evident (Table 7.2).

Bias	f_L (kHz)	A_{50} (Hz)	A_{150} (Hz)	ϕ_{50} (rad)	ϕ_{150} (rad)	ΔB_{ac} (μ G)
y	695.841(3)	131(4)	44(4)	-0.23(3)	-2.53(9)	235(11)
x'	603.567(3)	464(4)	169(4)	2.16(1)	-1.54(2)	860(11)

Table 7.2: Comparison of fit parameters of (7.10) for Figure 7.4 and Figure 7.5, quantifying power line modulation in a y -bias and an x' -bias. Note that the long interrogation time produces a precise estimate for the centre Larmor frequency f_L . ΔB_{ac} is the maximum magnetic field deviation, taking into account the phase of the components.

The fitted Larmor frequency modulation can then be used to demodulate the spectrogram and more easily visualise how the signal changes over time. The bins of the spectrogram can be shifted to line up the instantaneous Larmor frequency with the centre Larmor frequency. This works well when the frequency slew rate is low (e.g. Figure 7.4C), but when the frequency modulation is larger, distortions are observed in the demodulated spectrogram (e.g. Figure 7.5C). In this case, the frequency is changing so rapidly that it crosses multiple frequency bins of the spectrogram, causing demodulation to incorrectly convert frequency modulation to amplitude modulation.

This can be corrected using a Matched Filter analysis [191], since the frequencies are known but the amplitudes and phases are not. This would even enable demodulation of the Faraday signal for swept-field measurement (e.g. AC magnetometry). Although we are currently pursuing this form of analysis, it is outside the scope of the present work.

⁵ The axes are defined in §8.1, though power line variation is observed for any bias orientation.

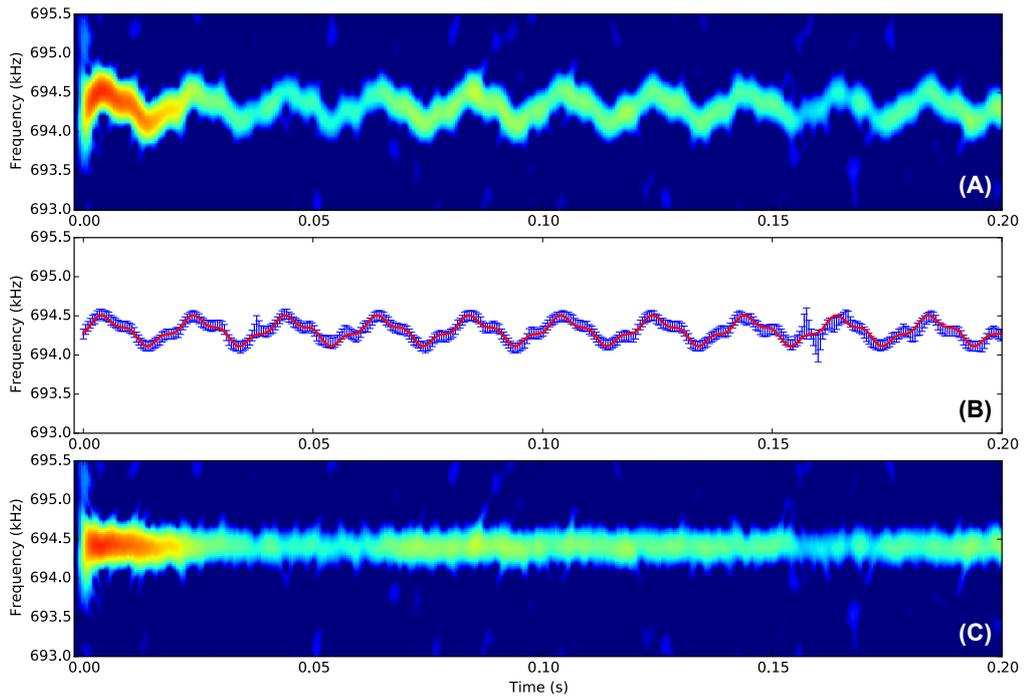


Figure 7.4: STFT of Faraday signal in a y -bias (A), showing the centre Larmor frequency changing in time. The extracted peak positions (B) are fit using (7.10) and used to demodulate the signal (C) and obtain a constant Larmor frequency.

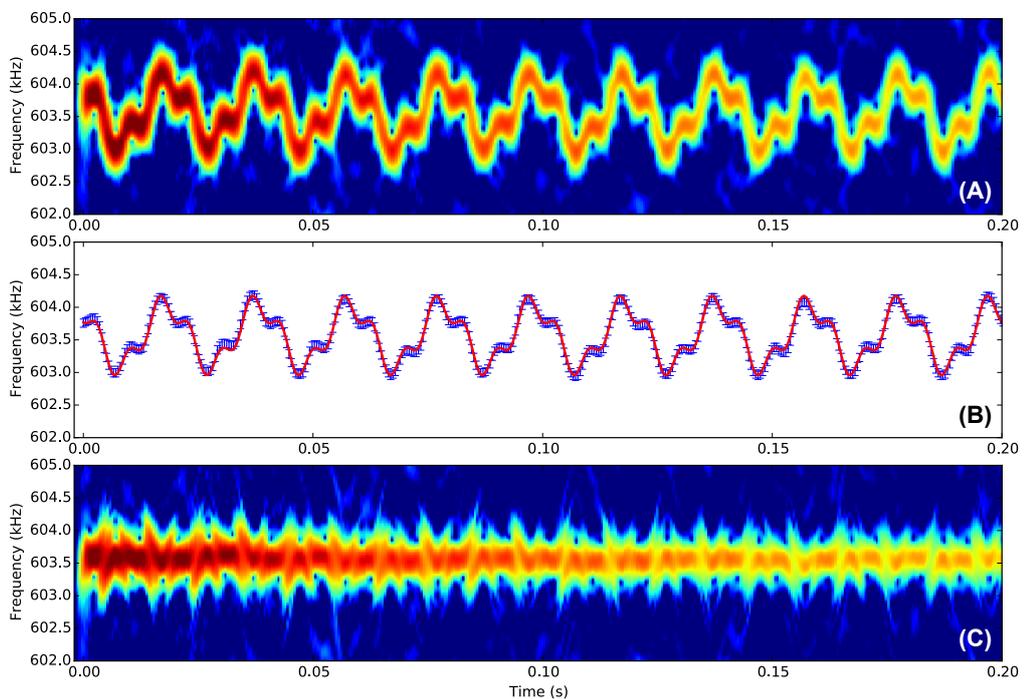


Figure 7.5: STFT showing frequency modulation of the Faraday signal in an x' -bias (A). The structure of the signal is different to a y -bias (Figure 7.4), but the centre Larmor frequency (B) is still accurately modelled by (7.10). The demodulated signal (C) shows residual 'glitches' where the frequency modulation rate of change is greater than the analysis bin size.

Note that the line noise fit is sensitive to the phase of the power line at the start of the experiment. Despite using a power line synchronisation system to synchronise the experiment to the power line cycle [138], the fit is sensitive to residual jitter in the power line harmonic composition and shot-to-shot fluctuations in the fit phase of 0.2 radians were common.

7.3: Blurring from power line modulation

The magnetic field fluctuations induced by the power line causes modulation of the Larmor frequencies of the BEC. Since frequency is linked to position in the Fourier imaging analysis (§3.2), this modulation is interpreted as motion of the atoms if unaccounted for, causing an effective blurring of the image.

Despite containing higher harmonics, the power line modulation can be approximately modelled as sinusoidal at the line frequency $f_m = 50$ Hz with a total frequency deviation of f_Δ . The Larmor frequency $f_L(t)$ is then

$$f_L(t) = f_{L0} + f_\Delta \sin(2\pi f_m t). \quad (7.12)$$

In the worst-case scenario, the measurement will coincide with a zero-crossing of the sine term. This is where the gradient is steepest, causing the biggest frequency sweep across the measurement interval, and hence producing the most pronounced blurring.

For measurement intervals τ much shorter than the power line period ($\tau \ll 20$ ms) the frequency deviation Δf_{ac} can be linearly approximated as

$$\Delta f_{ac} \leq 2\pi f_m f_\Delta \tau. \quad (7.13)$$

Conversely, if the measurement time is comparable to the period of the power line cycle or longer ($\tau \gtrsim 20$ ms), an entire cycle of the modulation is sampled and the deviation is bounded above by

$$\Delta f_{ac} \leq f_\Delta. \quad (7.14)$$

Using (3.15) to relate frequency to position, an upper bound on the blurring is

$$\Delta z_{ac} = \frac{2\pi \Delta f_{ac}}{\gamma b} \quad \text{where} \quad \Delta f_{ac} = \begin{cases} 2\pi f_m f_\Delta \tau & \text{for } \tau \ll 20 \text{ ms} \\ f_\Delta & \text{otherwise.} \end{cases} \quad (7.15)$$

Note that $\tau = \tau_f$ for a Faraday image reconstruction, and $\tau = t_{\text{window}}$ for an STFT analysis so different levels of blurring are expected for different analyses of the same data.

Based on these results, blurring can be reduced by selecting large gradients and small measurement times, which is the same conclusion reached in §3.9 to minimise Stern-Gerlach blurring. However, a specific timescale of $\tau \ll 20$ ms is imposed here, which

requires $b \gg 0.1 \text{ G/cm}$ by (3.36) to prevent Stern-Gerlach blurring dominating the resolution of the image. Furthermore, the strength f_{Δ} of the power line noise was seen to be dependent on the direction of the bias field, and so too is the resulting degree of blurring (Figure 7.6).

The power line noise therefore introduces a non-negligible degree of blurring to the Faraday imaging process. For rapid measurements, the blurring can be reduced by synchronising the experiment to expose the Faraday beam at the extrema of the power line cycle, where the frequency deviation is large but changes slowly over the duration of the measurement. Another approach is to use external coils to synthesise an opposing magnetic field to actively cancel the power line noise [192, 193]. However, the field induced by the power line has components in all directions which would require multiple sets of coils to cancel completely.

7.4: Effect of probe polarisation

The Faraday measurement itself is relatively insensitive to the polarisation state of the probe beam, as any elliptical component in the probe beam polarisation state simply decreases the contrast observed at the polarimeter. However, this elliptical component also results in a vector light shift (VLS) that creates an effective magnetic field affecting the evolution of the atomic state being observed (§2.6) and should be minimised wherever possible.

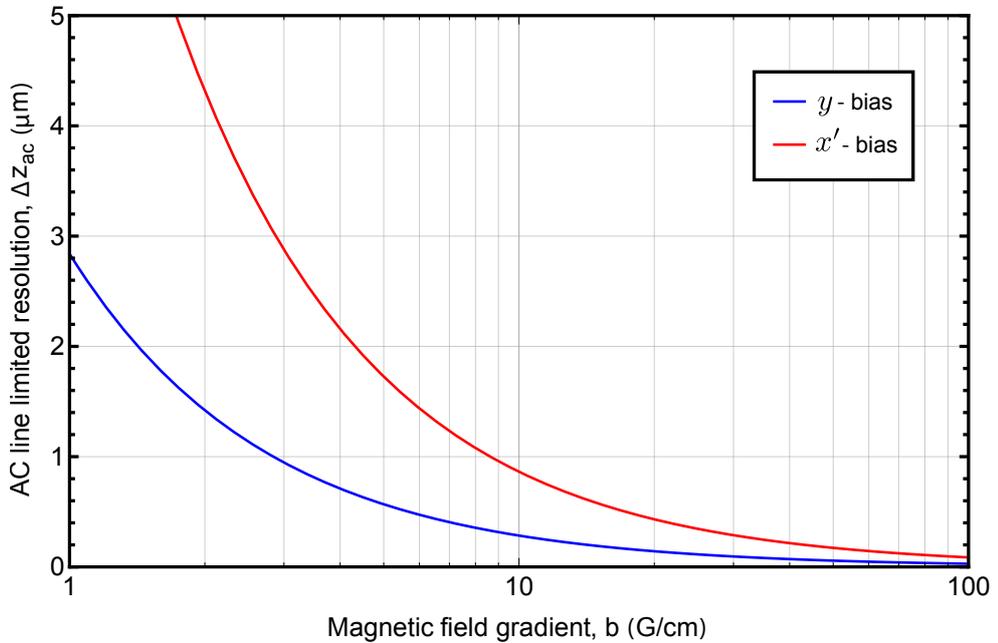


Figure 7.6: Worst-case power line blurring ($\Delta f_{ac} = f_{\Delta}$) as computed with (7.15) for varying magnetic field gradient strength in two orthogonal bias directions. Other blurring factors (Fourier/Stern-Gerlach) not considered.

Although it is straightforward to create a beam with pure linear polarisation in free-space using high-extinction polarising optics, the windows of the glass cell science chamber exhibit birefringence. A linear polarisation before the cell is therefore turned into an elliptical polarisation at the position of the atoms, resulting in unintended evolution of the condensate spin (see §2.6). It is therefore necessary to apply an opposite birefringence in advance, to counteract the ellipticity induced by the cell.

Unfortunately, cancelling the birefringence of the cell window is non-trivial as it is not possible to directly measure the polarisation state inside the vacuum chamber, and the beam experiences further birefringence as it exits the cell. The only way to reliably determine the polarisation state of the probe beam inside the chamber is to interrogate the atoms themselves.

Placing a quarter-waveplate before the cell (§5.2) allows control of the ellipticity of the beam by adjusting its angle of rotation θ . If the light is linearly polarised within the cell when the waveplate is rotated to angle θ_0 , then the ellipticity of the beam can be expressed through the Stokes parameter \hat{S}_z as

$$\langle \hat{S}_z \rangle \propto \sin(2(\theta - \theta_0)), \quad (7.16)$$

where $\langle \hat{S}_z \rangle = 0$ represents linearly polarised light. The vector light-shift acts as an effective magnetic field in the direction of propagation of the beam (§2.6), so the net magnetic field along the propagation direction is

$$B_z = B_{z0} + B_{\text{VLS}} \sin(2(\theta - \theta_0)), \quad (7.17)$$

where B_{z0} is the background magnetic field in the z -direction, and B_{VLS} is the magnitude of the maximum induced light shift (for purely circular polarisation).

Capturing the Faraday signal for various rotation angles (Figure 7.7) demonstrates how the polarisation state of the probe affects the atoms. Note that the both centre Larmor frequency and the structure of the signal change with θ .

This is because the VLS not only changes the net magnetic field strength, but the transverse intensity profile of the beam results in an effective magnetic field *gradient*,

$$\frac{\partial B_{\text{VLS}}}{\partial r} \propto \frac{\partial I(\mathbf{r})}{\partial r} \sin(2(\theta - \theta_0)) \neq 0. \quad (7.18)$$

Collapse and revival of the Faraday signal occurs on a timescale set by the *net* magnetic field gradient, with contributions from the background field gradient, applied gradient, and VLS gradient.⁶ There is potential to generate a strong effective magnetic field gradient using VLS, which could be useful for improved imaging resolution (§3.9).

⁶ The Faraday probe beam comes to a focus at the BEC so there is also an effective *longitudinal* magnetic field gradient $\frac{\partial B_{\text{VLS}}}{\partial z} \neq 0$ that contributes to dephasing. The strength of this gradient depends on the Rayleigh range of the probe beam, and is a much smaller contribution than the radial gradient.

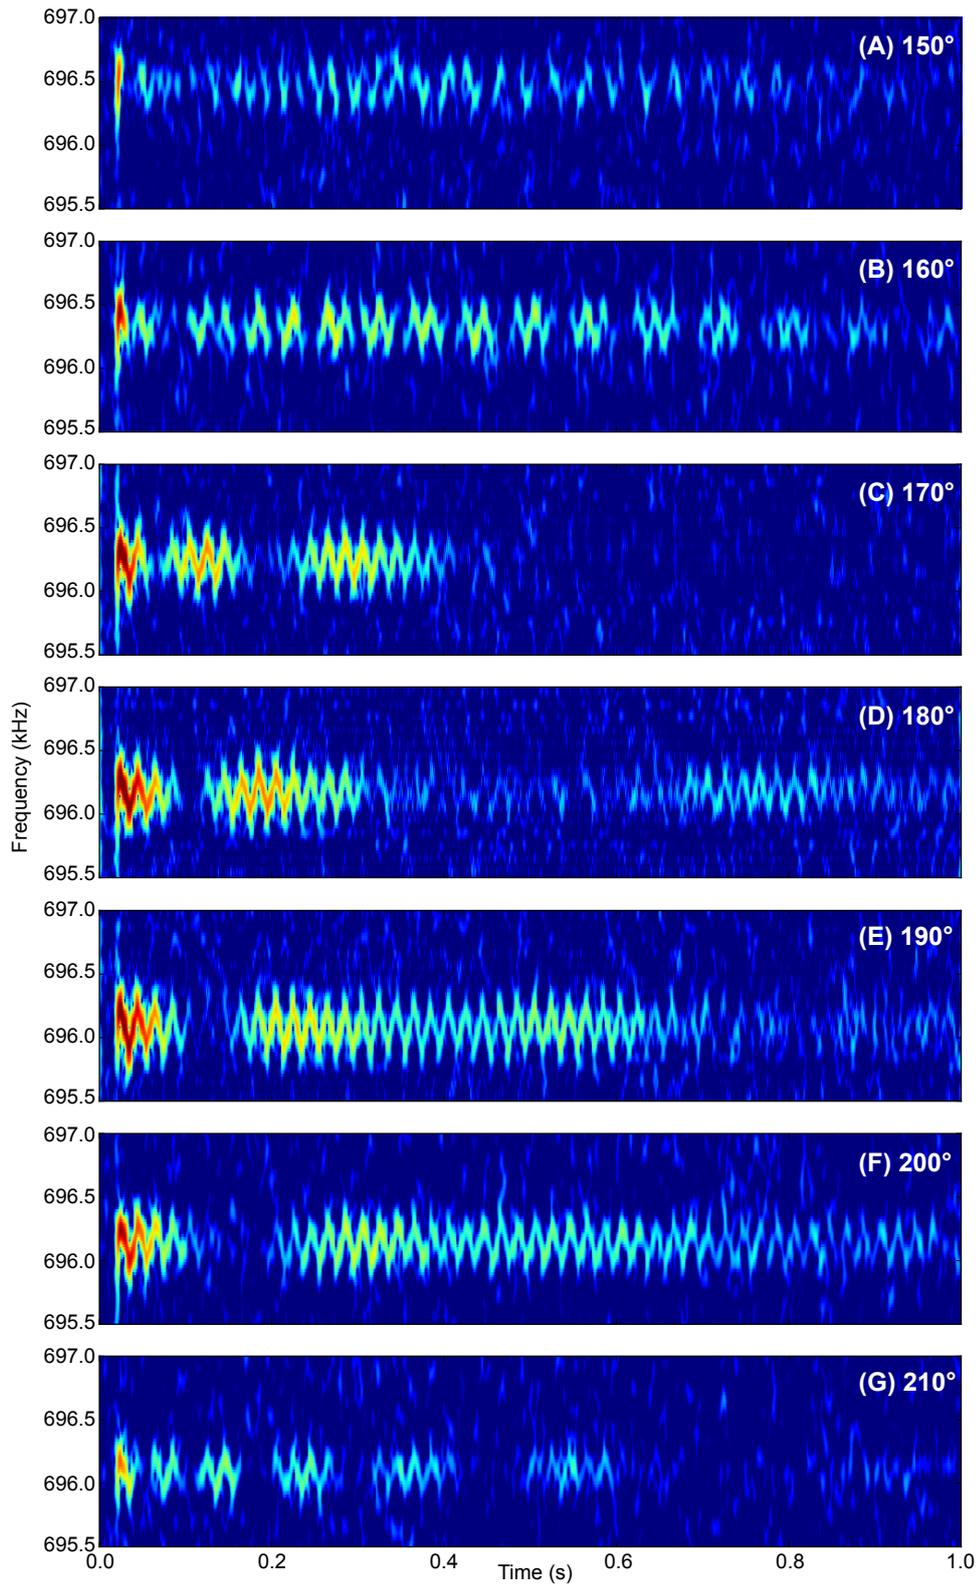


Figure 7.7: Rotating a quarter-waveplate angle (angle θ) before the cell alters the degree of ellipticity of the probe beam, affecting the structure and longevity of the Faraday signal. Note that the longest-lived signal does not correspond to minimum VLS but rather where the VLS partially cancels the background gradient.

However, the generated gradient contributes to the $\frac{\partial B_z}{\partial z}$ which is perpendicular to the bias field and does not assist Faraday imaging (see §8.2). A separate ‘VLS beam’⁷ propagating along the bias field axis (y) could be used to produce the ‘off-diagonal’ gradient ($\frac{\partial B_y}{\partial z}$) for imaging along the probe beam (z).

For a circularly polarised probe beam, no Faraday signal is observed at all as the Faraday interaction vanishes (§2.5). When the polarisation state is close to circular, the signal shows rapid decay and revival cycles, corresponding to a large effective field gradient. When it is closer to linear, longer decay times are observed as the VLS contribution to the net magnetic field gradient is reduced and the resulting spin-dependent spatial dynamics occur on a slower timescale.

The nematic interaction also causes collapse and revival of the Faraday signal (§2.7), even when the probe beam has perfectly linear polarisation. The strength of this interaction is dependent on the polarisation angle of the probe beam with respect to the bias field, and in principle a $\frac{\lambda}{2}$ -waveplate could be included before the cell to rotate the polarisation axis to the ‘magic angle’ and cancel out its contribution. In practice, however, the probe beam hits the cell at an angle of incidence near 45° , so the birefringence induced by the cell window also depends on the probe’s polarisation axis. Rotating the incident beam’s polarisation axis therefore produces an elliptical polarisation at the atoms, reintroducing a VLS which dominates the nematic contribution. As shown in §2.7, the nematic contribution is negligible in $F = 1$ ^{87}Rb atoms, so the linear polarisation axis will be taken as arbitrary.⁸

VLS-induced dephasing prevents long interrogation times from being used, and induces complex spatial dynamics in the evolution of the condensate. It should be noted however that the longest lasting signal does not correspond to perfect linear polarisation at the atoms, as the VLS can oppose the z -component of the background field gradient and reduce the net field gradient below its background value. While this *appears* to produce an improved measurement, the cancellation is sensitive to the probe beam power and fluctuations in the background magnetic field, reducing day-to-day repeatability. The VLS should therefore be nulled completely when using the Faraday effect to probe a BEC.

7.5: Eliminating the probe vector light-shift

As observed above, the probe VLS affects both the structure and longevity of the Faraday signal. Although the effect on the signal is visually pronounced, quantifying the lifetime of the Faraday signal over long interrogation times is non-trivial and sus-

⁷ Note that the Faraday probe beam must be linearly polarised to maximise the Faraday effect, whereas the VLS beam should be circularly polarised to create the largest effective magnetic field (§2.6).

⁸ Furthermore, uncertainty in determining and physically setting the optimal $\frac{\lambda}{4}$ -waveplate angle likely induces a residual VLS that is greater than the nematic contribution.

ceptible to shot-to-shot variations. Qualitative optimisation of the waveplate angle can be achieved in this way, resulting in an polarisation state at the atoms that is approximately linear. However, eliminating the VLS entirely requires a more precise approach.

The probe VLS contributes to the net magnetic field at the position of the atoms, so it can be quantified by measuring the centre frequency of the Faraday signal for a range of waveplate angles θ . As the polarisation is close to linear at the atoms, the signal exhibits a long lifetime, allowing the centre Larmor frequency to be precisely measured.

Taking the bias field to be predominately along the y -direction and using (7.17), the Larmor frequency of the atoms in the presence of vector light-shift is

$$\omega_L = \gamma|\mathbf{B}| = \gamma\sqrt{B_x^2 + B_y^2 + B_z^2} = \gamma\sqrt{B_{\perp z}^2 + [B_{z0} + B_{\text{vls}} \sin(2(\theta - \theta_0))]^2}, \quad (7.19)$$

where $B_{\perp z} = \sqrt{B_x^2 + B_y^2} \approx B_y$ is the field perpendicular to the probe's propagation direction. Since B_y dominates the other magnetic field components, the Larmor frequency becomes

$$\omega_L \approx \gamma B_y + \frac{\gamma[B_{z0} + B_{\text{vls}} \sin(2(\theta - \theta_0))]^2}{2B_y}. \quad (7.20)$$

If the VLS is small, $|B_{\text{vls}}| < |B_{z0}|$, the bracketed term above doesn't change sign as θ is varied and the Larmor frequency should appear sinusoidal in θ . However, the measured Larmor frequencies (Figure 7.8) do not vary simply sinusoidally with waveplate angle, implying that $B_{\text{vls}} > B_{z0}$. Hence the VLS dominates the residual z -component of the field when the probe beam has a (nearly) circular polarisation state at the atoms.

Performing a least-squares fit of (7.20) to the measured Larmor frequencies (Figure 7.8) results in the parameter estimates given in Table 7.3. Not only does this give an independent estimate of the bias field strength, the component of the background field along the propagation direction, and the strength of the VLS; it also determines the waveplate angle that eliminates the probe beam VLS. The waveplate was set to this angle for all future measurements to eliminate the contribution of the vector light-shift to the dynamics of the system.

Parameter	Value
$B_{\perp z} \approx B_y$	993.27(2) mG
B_{z0}	19.6(8) mG
B_{vls}	43(1) mG
θ_0	78.9(4)°

Table 7.3: Fitted parameter values for the vector light-shift measurement.

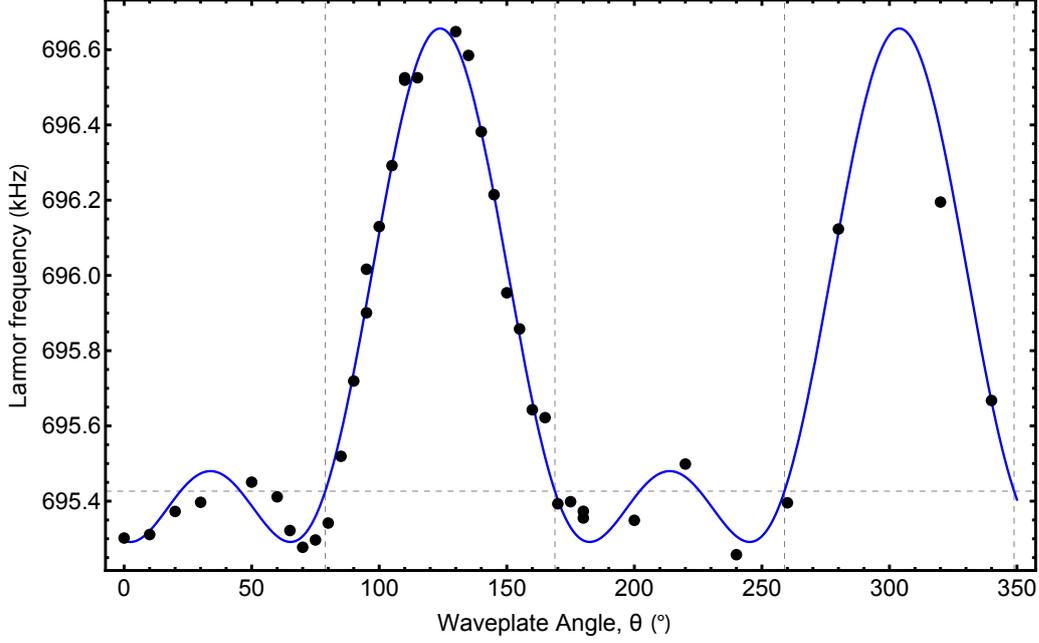


Figure 7.8: The measured Larmor frequency includes a contribution from the probe beam vector light shift (VLS), which depends on the ellipticity of the probe beam. Adjusting a $\lambda/4$ -plate (angle θ) before the science cell allows the VLS to be mapped out, which can be fitted using (7.20) to determine when the VLS is nulled (dashed lines).

7.6: Quadratic Zeeman effect

Atoms in an external magnetic field are subject to the Zeeman effect, which lifts the degeneracy between the different $|m_F\rangle$ states. Typically the Zeeman effect is taken as linear, so that the $|m_F = \pm 1\rangle$ states are shifted in opposite directions by the same amount with respect to an unshifted $|m_F = 0\rangle$ state. However, there are higher-order contributions that must be taken into account when the field is strong or the measurement occurs over a long interval.

For the ground state hyperfine levels of an alkali atom with nuclear spin I , the energy levels are described by the Breit-Rabi equation [111, 194],

$$E_m = -\frac{\Delta E_{\text{hfs}}}{2(2I+1)} + g_I \mu_B m B \pm \frac{\Delta E_{\text{hfs}}}{2} \sqrt{1 + \frac{4mx}{2I+1} + x^2}, \quad (7.21)$$

where ΔE_{hfs} is the ground-state hyperfine splitting, m the magnetic quantum number ($m \approx m_F$ for weak magnetic fields, with $m_F \in [-F, \dots, F]$ and $F = I \pm \frac{1}{2}$), g_I and g_J are the nuclear and fine structure Landé factors, and the parameter x is given by

$$x = \frac{(g_J - g_I) \mu_B B}{\Delta E_{\text{hfs}}}. \quad (7.22)$$

For the $|F = 1\rangle$ manifold of ^{87}Rb being studied, $I = \frac{3}{2}$ and the relevant sign in the expression for E_m in (7.21) is -1 .

Typically the Breit-Rabi expression is evaluated as a power series expansion in B , which to first order yields

$$E_z = -\hbar\gamma mB = -\frac{(g_J - 2(I+1)g_I)\mu_B}{2I+1}mB. \quad (7.23)$$

The next level correction is the quadratic Zeeman shift, which is obtained by taking the B^2 term in the power series expansion of the Breit-Rabi result,

$$E_{qz} = -\frac{(g_J - g_I)^2\mu_B^2}{\Delta E_{\text{hfs}}}\left(\frac{1}{4} - \left(\frac{m}{2I+1}\right)^2\right)B^2. \quad (7.24)$$

Evaluating using known constants [121] gives the following expected quadratic Zeeman shifts for ^{87}Rb in the $|F=1\rangle$ manifold,

$$E_z = -hmB \times 702.369 \text{ kHz/G}, \quad (7.25)$$

$$E_{qz} = hB^2 \times \begin{cases} -287.57 \text{ Hz/G}^2 & \text{for } |m_F=0\rangle \\ -215.68 \text{ Hz/G}^2 & \text{for } |m_F=\pm 1\rangle. \end{cases} \quad (7.26)$$

The quadratic Zeeman effect shifts the $|m_F=\pm 1\rangle$ levels by the same amount, so it effectively shifts energy level of the $|m_F=0\rangle$ state relative to the $|m_F=\pm 1\rangle$ states (Figure 7.9) by an amount hq_z , where

$$q_z = \frac{E_{+1} + E_{-1} - 2E_0}{2h} \quad (7.27)$$

$$= \frac{(g_J - g_I)^2\mu_B^2}{h(2I+1)^2\Delta E_{\text{hfs}}}B^2 + O(B^3) \quad (7.28)$$

$$= B^2 \times 71.89 \text{ Hz/G}^2, \quad (7.29)$$

and the higher order terms are taken as negligible.

In the single-mode approximation, each component of the spin-1 BEC order parameter can be written in terms of a fractional population $\rho_i(t)$ and phase $\Theta_i(t)$ as

$$\xi(t) = \left(\sqrt{\rho_-(t)}e^{i\Theta_-(t)}, \sqrt{\rho_0(t)}e^{i\Theta_0(t)}, \sqrt{\rho_+(t)}e^{i\Theta_+(t)}\right)^T, \quad (7.30)$$

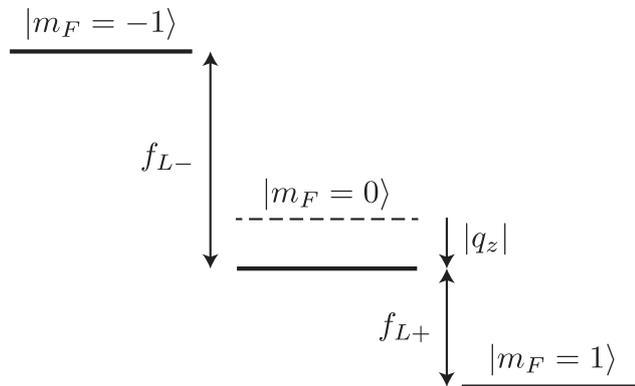


Figure 7.9: The quadratic Zeeman shift effectively shifts the energy level of the $|m_F=0\rangle$ state by q_z relative to the $|m_F=\pm 1\rangle$ states. The Faraday signal of the resulting system is interpreted as containing two separate Larmor tones $f_{L\pm} = f_L \pm q_z$.

where the components correspond to the $|m_F = -1\rangle$, $|m_F = 0\rangle$ and $|m_F = +1\rangle$ states respectively. For a condensate initially prepared in a zero magnetisation state (such as that produced by an rf $\frac{\pi}{2}$ -pulse), and taking the components as initially in phase, $\Theta_i(0) = 0$, the order parameter can be simplified to

$$\xi(t=0) = \left(\sqrt{(1-\rho_0)/2}, \sqrt{\rho_0}, \sqrt{(1-\rho_0)/2} \right)^T. \quad (7.31)$$

Temporarily neglecting the spin-dependent interaction, the time-dependent expectation value of the spin projection along the probe axis⁹ $\langle F_z \rangle$ can then be evaluated as [195]

$$\langle F_z \rangle = 2\sqrt{\rho_0(1-\rho_0)} \cos(2\pi q_z t) \cos(2\pi f_L t),$$

where $f_L \equiv \frac{E_{+1} - E_{-1}}{2h}$ and $q_z \equiv \frac{E_{+1} + E_{-1} - 2E_0}{2h}$, (7.32)

Hence the induced Faraday signal is expected to contain a single tone in the spectrogram at the centre Larmor frequency f_L , amplitude-modulated at frequency q_z . This not only provides a physical origin for (7.27), but predicts periodic collapse and revival of the Faraday signal (Figure 7.10).

Alternatively, this modulation can be written as

$$\cos(2\pi f_L t) \cos(2\pi q_z t) = \frac{1}{2}(\cos(2\pi f_{L+} t) + \cos(2\pi f_{L-} t)), \quad (7.33)$$

which is interpreted as the Faraday signal containing two frequency components, $f_{L\pm}$, where

$$f_{L\pm} = f_L \pm q_z. \quad (7.34)$$

These two interpretations are reconciled by considering the role of the spectrogram's frequency resolution. If the resolution is fine enough to resolve the splitting of the two tones, it will display two bright lines corresponding to the two sidebands at $f_{L\pm}$. However, typically the spectral resolution of the STFT ($\Delta f = 200$ Hz) is greater than the quadratic splitting at the bias field of interest (7.27). In this case, the two sidebands

⁹ The notational difference with Ref. [195] arises from the orientation of the bias field, which they take as along the z -axis but here is taken along the y -axis.

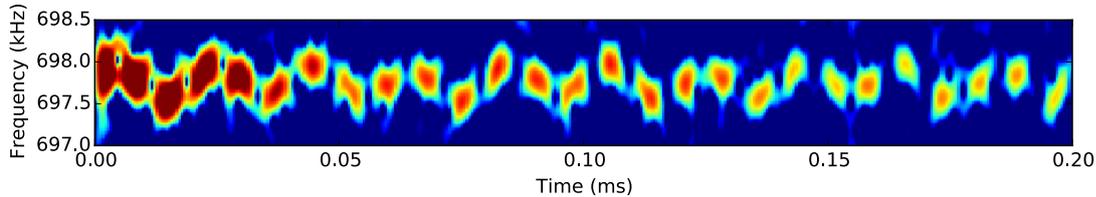


Figure 7.10: Faraday signal showing amplitude modulation arising from the quadratic Zeeman effect. The period of the nulls is 7.4(1) ms, corresponding to $q_z = 68(1)$ Hz at this bias field ($B = 0.992(1)$ G).

fall within the same spectrogram ‘frequency bin’ and interfere, resulting in amplitude modulation (a ‘beatnote’).

However, the spin-dependent interaction also affects the evolution of the condensate phase. Generalising (7.32) to consider interactions, the expectation value of the spin projection is [195]

$$\langle F_z \rangle = 2\sqrt{\rho_0(1-\rho_0)} \cos(\Theta/2) \cos(2\pi f_L t), \quad (7.35)$$

where the global condensate phase is $\Theta = \Theta_+ + \Theta_- - 2\Theta_0$. Transforming into a frame that is co-rotating with Larmor precession, the time-dependence of the system is governed by [195]

$$\begin{aligned} \frac{\partial \rho_0}{\partial t} &= \frac{2c}{\hbar} \sqrt{(1-\rho_0)^2 - m^2} \sin \Theta \\ \frac{\partial \Theta}{\partial t} &= -4\pi q_z + \frac{2c}{\hbar} (1-2\rho_0) + \left(\frac{2c}{\hbar} \right) \frac{(1-\rho_0)(1-2\rho_0) - m^2}{\sqrt{(1-\rho_0)^2 - m^2}} \cos \Theta, \end{aligned} \quad (7.36)$$

where the net magnetisation is $m = \rho_+ - \rho_-$ and c is a measure of the spin-dependent interaction strength.

Depending on the relative contributions of the quadratic shift q and the spin-dependent interaction c in (7.36), different modes of evolution can be observed. If the quadratic shift is strong then $\frac{\partial \Theta}{\partial t} \approx -4\pi q$, and the phase winds rapidly. The $\sin \Theta$ term therefore oscillates quickly, and time-averages to zero, resulting in $\frac{\partial \rho_0}{\partial t} = 0$. Spin mixing due to c is therefore ‘frozen out’. Typically the quadratic shift dominates for $B \gtrsim 500$ mG, but when microwaves are applied to cancel the quadratic shift (see §7.7), spin mixing is expected to dominate.

Finally, although the quadratic Zeeman effect causes amplitude modulation, it does not shift the centre Larmor frequency f_L since by (7.32),

$$f_L \equiv \frac{E_{+1} - E_{-1}}{2h} = \left(\frac{\gamma}{2\pi} \right) B + O(B^3), \quad (7.37)$$

which to third order is simply the linear Zeeman effect (7.23), as the quadratic terms (7.24) cancel exactly. Hence the centre Larmor frequency is taken to be directly proportional to the magnetic field strength, simplifying analysis of the Faraday signal.

7.7: Microwave dressing

The quadratic Zeeman shift resulted in an effective level shift of the $|m_F = 0\rangle$ substate, resulting in modulation of the Faraday signal. This level shift can be eliminated by applying a level shift in the opposite direction by dressing the $|m_F = 0\rangle$ state with microwave radiation (Figure 7.11). Applying microwaves that couple the Zeeman states $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$ results in an AC Stark shift of $|F = 1, m_F = 0\rangle$

in the dressed state picture, producing an effective quadratic shift. For microwave detunings Δ_{mw} much larger than the microwave Rabi frequency Ω_{mw} , the microwave-induced quadratic shift is [196]

$$q_{\text{mw}} = -\frac{\Omega_{\text{mw}}^2}{4\Delta_{\text{mw}}}. \quad (7.38)$$

The quadratic Zeeman shift q_z can then be cancelled by setting $q_{\text{mw}} = -q_z$ to get

$$\Delta_{\text{mw}} = \frac{\Omega_{\text{mw}}^2}{4q_z}. \quad (7.39)$$

However, the system does not comprise only two levels, and at the bias field strengths of interest (~ 1 G), the linear Zeeman splitting (~ 700 kHz) is larger than the typical microwave detuning (~ 300 kHz) although comparable in magnitude. The microwaves therefore result in some cross-coupling between the different $|m_F\rangle$ sub-states,¹⁰ which results in a deviation of the optimal microwave detuning from the above estimate (7.39).

The clock transition $|F = 1, m_F = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$ experiences its own quadratic shift, which is given by

$$\frac{E_{|F=2, m_F=0\rangle} - E_{|F=1, m_F=0\rangle}}{h} = \frac{(2I + 1)^2}{2} q_z = 8q_z, \quad (7.40)$$

¹⁰ Furthermore, limitations of the microwave antenna and its orientation mean the microwaves are not perfectly linearly polarised, so it is also possible to drive σ^\pm -transitions and increase cross-coupling.

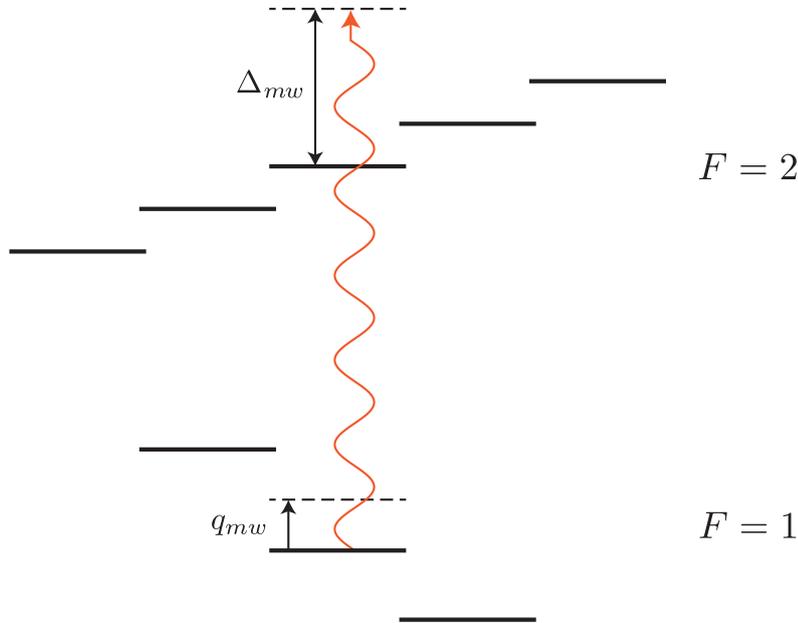


Figure 7.11: Microwave dressing the $|F = 1, m_F = 0\rangle$ to $|F = 2, m_F = 0\rangle$ transition effectively shifts the energy of the $|F = 1, m_F = 0\rangle$ state by an amount q_{mw} dependent on the detuning Δ_{mw} , and can be used to counter the quadratic Zeeman shift.

where the energy levels $E_{|F,m_F\rangle}$ are computed with¹¹ (7.24), q_z is given by (7.28) and $I = \frac{3}{2}$ for ^{87}Rb . The detuning of the microwaves¹² is therefore given by

$$\Delta \equiv f_{\text{mw}} - f_{\text{clock}} \quad \text{and} \quad f_{\text{clock}} = f_{\text{hfs}} + 8q_z \quad (7.41)$$

where $f_{\text{hfs}} = 6,834,682,610.904$ Hz is the ground-state hyperfine splitting [121].

The microwaves are generated by a PHASEMATRIX FSW-0010 DDS, which features programmatic control of microwave frequencies of up to 10 GHz with 1 mHz resolution [197]. The detuning can then be easily scanned by setting

$$f_{\text{mw}} = f_{\text{hfs}} + 8q_z + \Delta_{\text{mw}} \quad (7.42)$$

to observe the effect on the structure of the Faraday signal (Figure 7.12). The optimum detuning is expected to be at the detuning given by (7.39), though imperfect measurement of the Rabi frequency and deviations from ideal two-level behaviour lead to a deviation.

Depending on the sign and magnitude of the microwave detuning, the microwave dressing field can either enhance or counteract the quadratic Zeeman shift. Enhancing the shift leads to a larger quadratic Zeeman splitting and makes the two tones resolvable (Figure 7.12B–C). Note that when the two tones are resolved, the amplitude of each tone decays exponentially without the amplitude modulation, implying it indeed originates from the quadratic Zeeman shift and not another effect. However, when the frequency resolution of the spectrogram is too low to resolve the tones (Figure 7.12D–F), the Faraday signal is amplitude-modulated at the rate $q_{\text{net}} = q_z + q_{\text{mw}}$.

Cancelling the net quadratic shift q_{net} by choosing the appropriate microwave detuning Δ_{mw} (Figure 7.12G) produces a single Larmor tone that demonstrates no periodic amplitude modulation. The exponential decay of the signal is due to coherent dephasing by the background magnetic field gradient (§3.5), which must be independently cancelled by applying an external magnetic field gradient, or reduced through selection of the bias field direction (see §8.3), to produce a longer lasting signal.

7.8: Measuring the microwave Rabi frequency

When the microwave detuning Δ_{mw} is reduced, the microwave dressing dominates the quadratic Zeeman shift ($q_{\text{mw}} \gg q_z$) and the two Larmor sidebands f_{\pm} are split sufficiently to be resolved as separate tones in the Faraday spectrum. Because the signal persists for ~ 100 ms, the centre Larmor frequencies of the two tones can be

¹¹ The Landé g -factor for the $F = 2$ manifold has the opposite sign of $F = 1$ in (7.21), so the two quadratic shifts do not cancel.

¹² The detuning is expressed here in hertz, although rad/s is the more common convention.

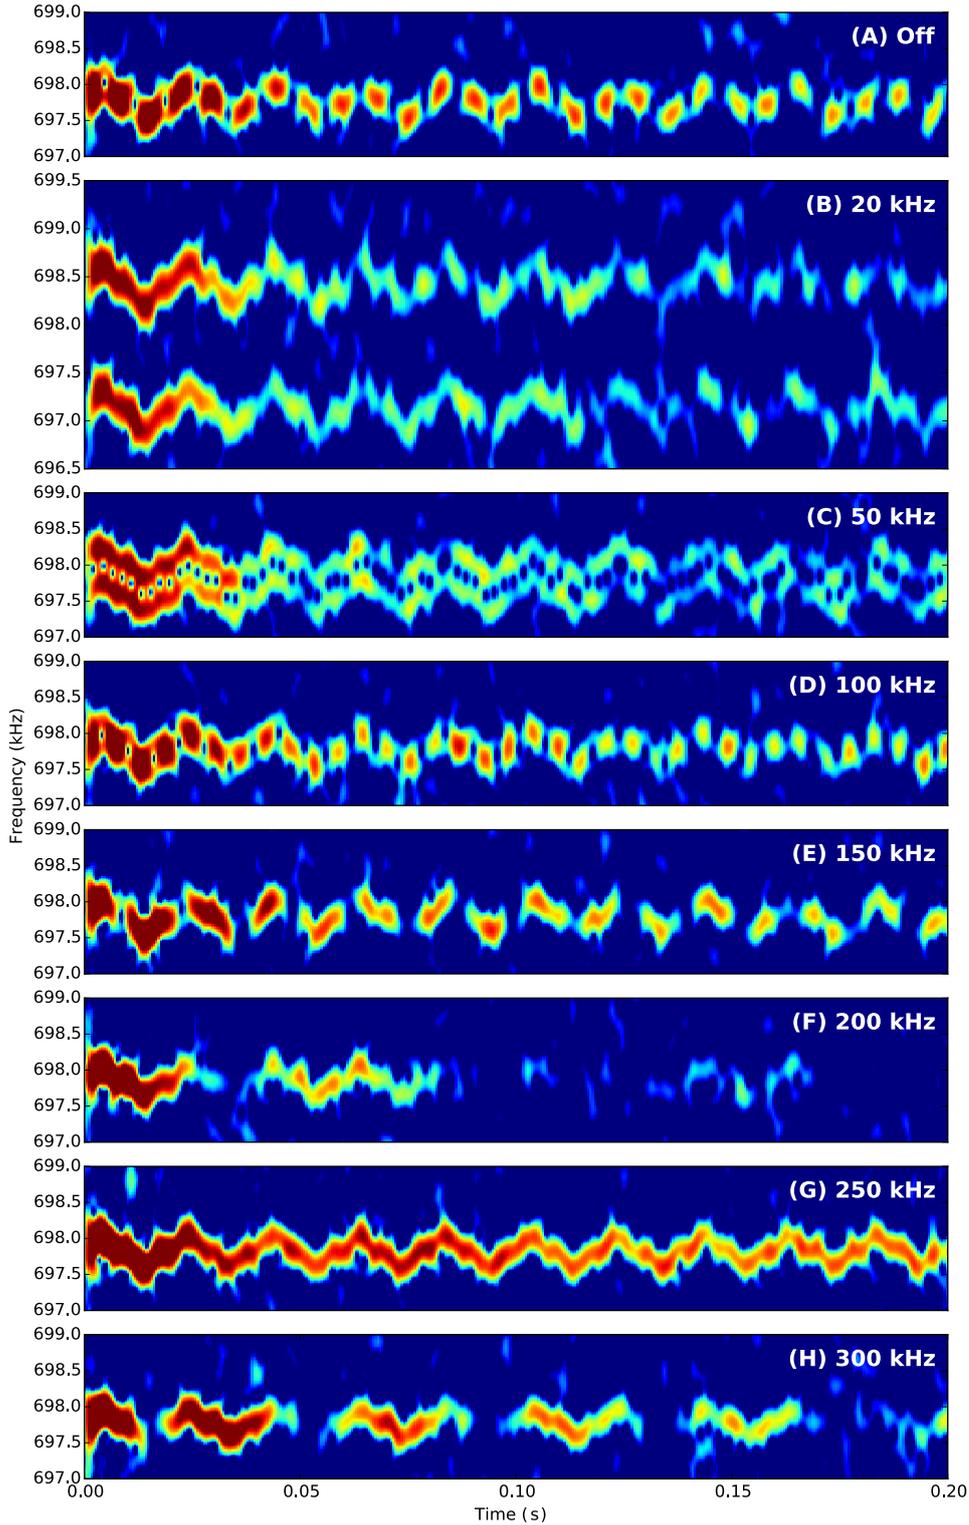


Figure 7.12: Comparison of the Faraday signal without microwaves (A) to microwaves with Rabi frequency $\Omega_{mw} = 8.49(2)$ kHz and varying detuning (B)–(H). Varying the detuning from the $|F = 1, m_F = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$ clock transition enables the quadratic shift q_{net} to be increased to resolve the individual Larmor sidebands (B)–(C), inverted (D) or suppressed (E)–(H). Cancelling the quadratic Zeeman shift (G) eliminates the periodic collapse and revival of the signal. Spectrogram window was 5 ms.

precisely measured in spite of the power line variation. This measurement provides a sensitive probe for splitting induced by the net quadratic shift,

$$q_{\text{net}} = q_z + q_{\text{mw}}, \quad (7.43)$$

and since the detunings are known, this provides a way to calibrate the microwave Rabi frequency Ω_{mw} via (7.38).

The Larmor frequencies of the two tones are $f_{L\pm} = f_L \pm q$, so the deviation from the original Larmor frequency is

$$\Delta f_{\text{mw}} \equiv \left| \frac{f_{L+} - f_{L-}}{2} \right| = |q_{\text{net}}| \equiv |q_z + q_{\text{mw}}| = \left| q_z - \frac{\Omega_{\text{mw}}^2}{4\Delta_{\text{mw}}} \right|. \quad (7.44)$$

Changing the detuning Δ_{mw} induces a different microwave shift, so measuring the splitting for different detunings allows the microwave Rabi frequency to be measured (Figure 7.13). Note that detuning closer to resonance ($\Delta_{\text{mw}} \rightarrow 0$) causes a significant transfer of population into the $|F = 2\rangle$ manifold, resulting in rapid decay of the Faraday signal and preventing accurate measurement of the Larmor frequencies. Detuning further from resonance reduced the splitting of the two tones such that their spectrogram patterns overlapped, resulting in amplitude-modulation in the spectrogram, preventing their centre frequencies from being distinguished.¹³

¹³ The restriction against using large detuning also prevents cross-coupling between the $|m_F\rangle$ sub-states, validating the two-level approximation (7.38) and enabling the Rabi frequency to be inferred.

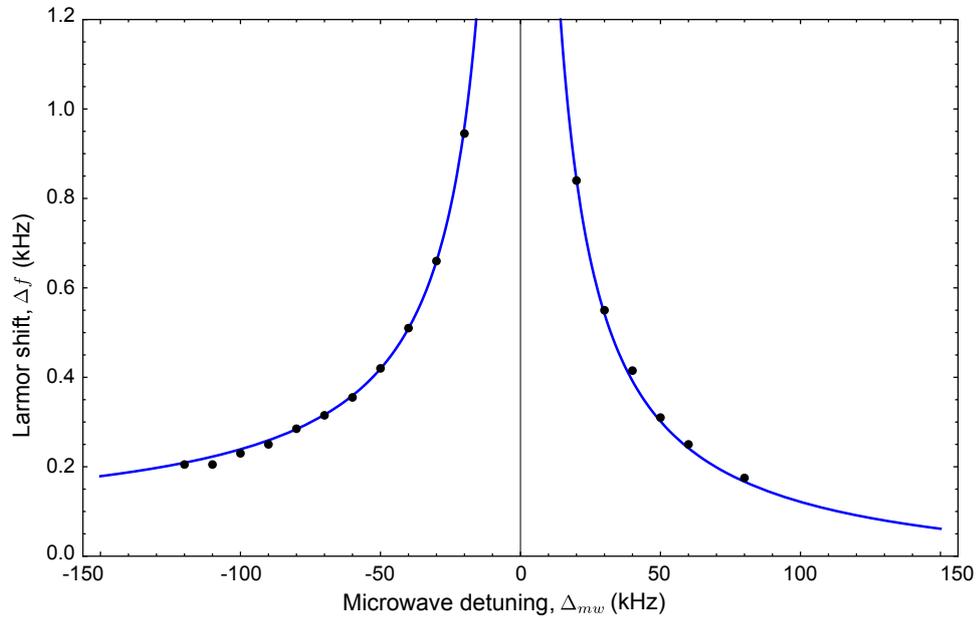


Figure 7.13: Separation Δf of the two Larmor tones caused by the quadratic Zeeman shift at different microwave detunings. Fitting (7.44) to the data (blue) gives the microwave Rabi frequency as $\Omega_{\text{mw}} = 8.49(2)$ kHz. Note that smaller detunings caused significant outcoupling into the $|F = 2\rangle$ manifold, and larger detunings did not split the levels sufficiently to distinguish the two tones.

This measurement of the microwave Rabi frequency was consistency-checked against observation of Rabi cycling between the $|F = 1\rangle$ and $|F = 2\rangle$ levels. The BEC is prepared in the $|F = 1, m_F = -1\rangle$ state then rf $\frac{\pi}{2}$ -pulsed into a $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ -state. Resonant microwaves were applied for a fixed interval t_{mw} to couple the $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$ states, after which the population in the $|F = 2\rangle$ state was measured with absorption imaging. Varying t_{mw} enabled observation of Rabi cycling between the two states, from which the Rabi frequency can be fitted. The resulting¹⁴ Rabi frequency is $\Omega_{\text{mw}} = 8.6(2)$ kHz, in agreement with the value obtained by measuring the Larmor shift, $\Omega_{\text{mw}} = 8.49(2)$ kHz.

7.9: BEC component separation

The Faraday measurement permits long interrogation times, during which the trapped BEC may exhibit spatial dynamics that are negligible in shorter timescale experiments. In particular, the magnetic field gradient induces a Stern-Gerlach force that acts to separate the $|m_F = \pm 1\rangle$ components, which causes the magnetisation of the BEC to become spatially dependent. Since the local magnetisation is given by the spin projection along the bias field, a nonzero magnetisation means the spin vectors have deviated out of the transverse plane, which reduces the magnitude of the transverse spin component and hence the amplitude of the Faraday signal (see [Figure 2.3B](#)).

However, since the BEC is being held in the dipole trap, the components do not separate indefinitely. The Stern-Gerlach force accelerates the $|m_F = \pm 1\rangle$ states up opposite sides of the harmonic trapping potential, which acts to push them back towards the centre. The resulting motion behaves like a driven harmonic oscillator whereby the Zeeman substates are separated into ‘lobes’ that oscillate spatially. [Figure 7.14A](#) shows a typical example of this in a y -bias, where the initially symmetric cloud is seen to deform within the first 50 ms of hold time. The local magnetisation becomes positive in the bottom-left area of the cloud, and negative in the top-right.

The corresponding Faraday signal measurement ([Figure 7.15B](#)) shows collapse and revival, where the minima correspond to the greatest separation between the $|m_F = \pm 1\rangle$ components, and the maxima to where overlap between the components is observed. The background gradient drives the $|m_F = \pm 1\rangle$ components up the trapping potential like a driven harmonic oscillator, causing perturbations (‘sloshing’) in the size and shape of the lobes in [Figure 7.14A](#). Where the overlap is stronger, the local magnetisation is closer to zero, Larmor precession has a larger amplitude and the Faraday signal demonstrates revivals.

¹⁴ This technique can produce precise estimates of the Rabi frequency [138]. Because it was only used as a consistency check, few data points were collected which resulted in a large uncertainty of Ω_{mw} .

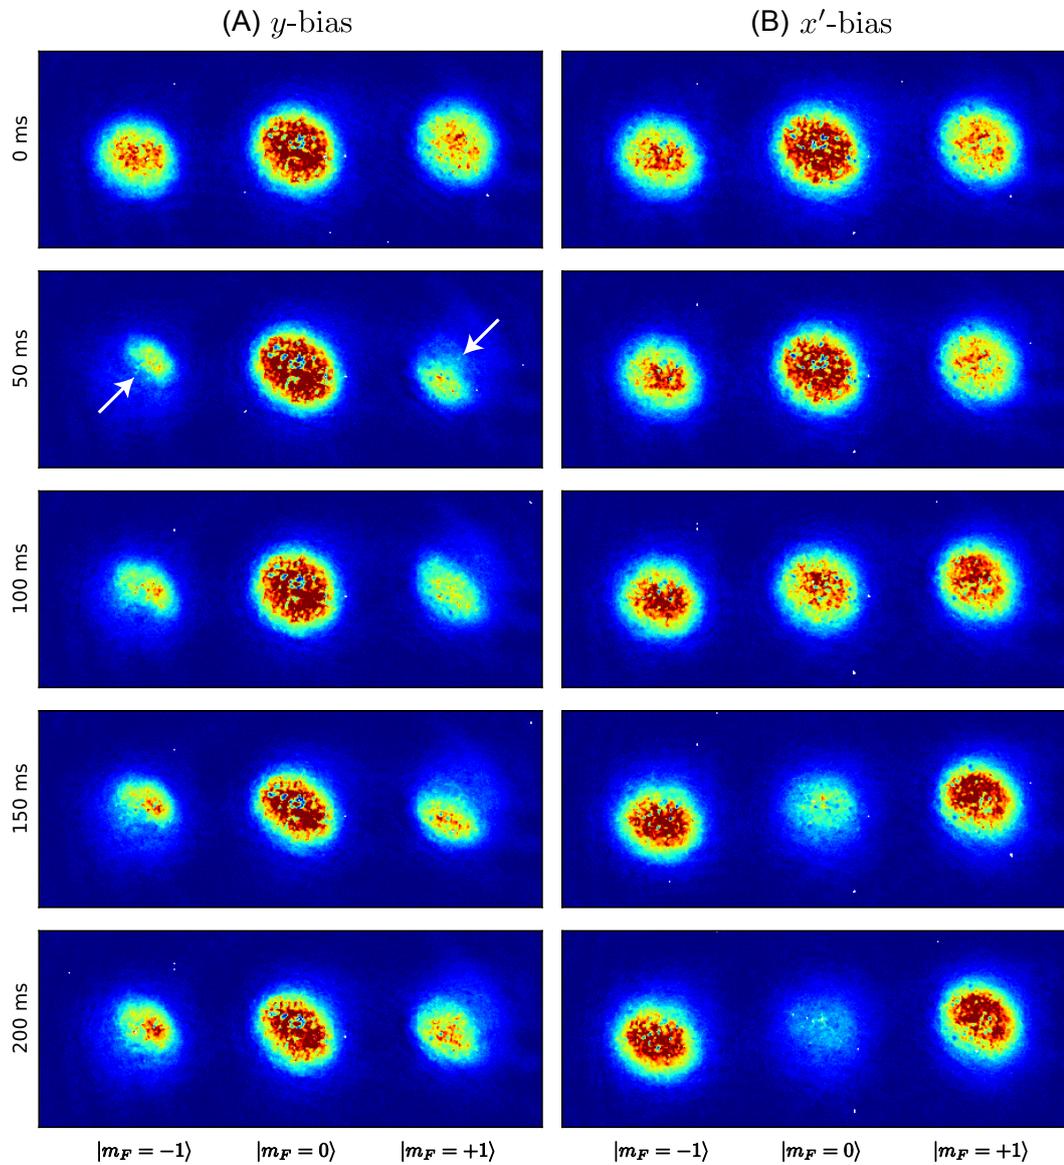


Figure 7.14: Stern-Gerlach-separated absorption images captured by the ‘side’ camera after various hold-times followed by 23 ms of TOF expansion. In a y -bias (A) the $|m_F = \pm 1\rangle$ states spatially separate (white arrows) whereas in an x' -bias (B) the three components maintain the same shape and demonstrate coherent spin mixing.

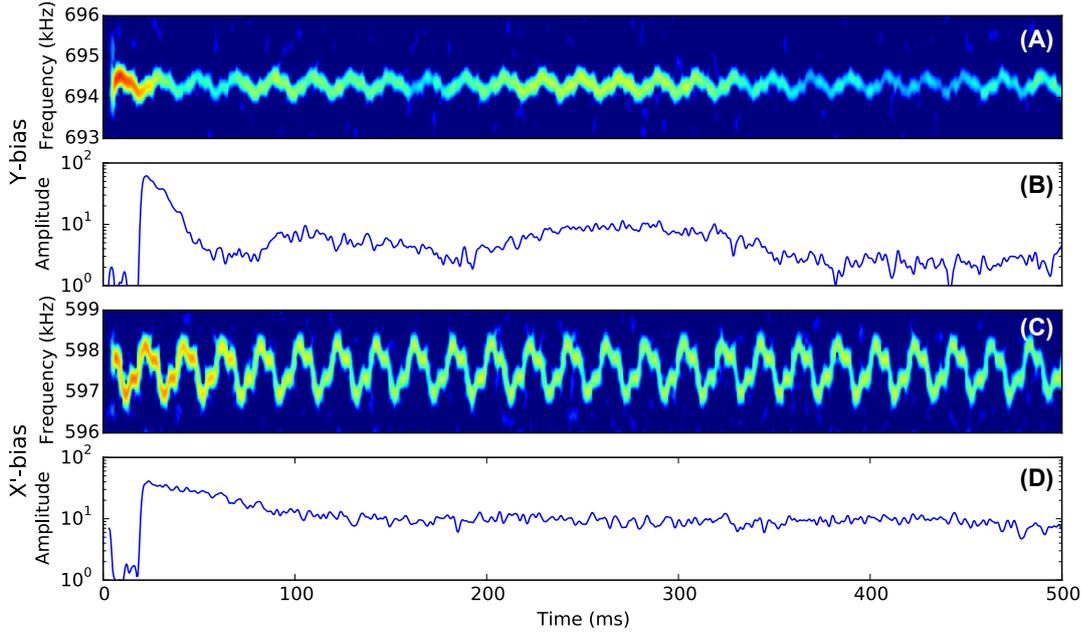


Figure 7.15: Spectrograms comparing microwave-nulled Faraday measurement with a y -bias (A) to an x' -bias (C). The demodulated signal shows decay and revival from spatial dynamics in the BEC in y -bias (B), which are not observed in an x' -bias (D).

By contrast, in an x' -bias (Figure 7.14B) the three components are observed to maintain the same overall shape, remaining in the single-mode approximation. Population transfer due to spin-exchanging collisions is observed, with the population of the initially higher density $|m_F = 0\rangle$ decreasing as the $|m_F = \pm 1\rangle$ populations increase. The corresponding Faraday measurement (Figure 7.15D) shows much more gradual decay in amplitude over time corresponding to atom number loss in the BEC. The bias field in each case was 1 G, and spin mixing was expected to be observed in both cases. However, in the y -bias case, spatial separation occurred on a faster timescale than the spin mixing.

At much longer hold times ($t \sim 200$ ms), some spatial separation can be observed in the cloud structure in an x' -bias. The direction of the separation is different to in a y -bias as the background gradient points in a different direction (see §8.3). The x' -bias was independently measured to have a significantly smaller background field gradient than the y -bias [188], so the Stern-Gerlach separation force is significantly weaker. The degree of separation is therefore much smaller, and occurs on a much longer timescale.

The Stern-Gerlach effect was previously analysed as a blurring of the Faraday image (§3.8), which limited the individual Faraday measurements to be short. However, these results demonstrate that even the background magnetic field gradient is strong enough to induce spatial separation over the timescales of interest.

In principle the component-separation can be reversed by inverting the magnetic field gradient to change the sign of the force, or applying a π -pulse to invert the spin vector which exchanges the $|m_F = \pm 1\rangle$ states. Since the background gradient cannot be reversed, either a bias orientation with low background gradient must be used to slow the separation rate, or rapid π -pulses used to periodically reverse the evolution and prevent spatial separation of the components.

The distinction should be emphasised between gradient-induced dephasing, which occurs when the spin vectors spread out in the Larmor plane, and component separation, which causes the spin vectors leave the Larmor plane. Both are produced by magnetic field gradients, and both result in a reduction of the amplitude of the Faraday signal. Although component-separation induces periodic sloshing which revives the signal, the dynamics of component separation cannot be simply cancelled out by reversing the separation and recombining the cloud, as the spatial variation of the magnetisation during separation irreversibly affects the spin-dependent interaction.

7.10: Summary

In this chapter I analysed the Faraday signal in more detail, considering the dependence of the Larmor frequency on the applied magnetic field, the structure of the power line variation, the role of probe polarisation, and the quadratic Zeeman effect. Measurements of the Larmor frequency were used to calibrate the bias coils, without requiring that the ‘nulling’ point for each coil be found first.

Magnetic field fluctuations synchronised to the 50 Hz power line cycle were observed, with the Faraday signal being frequency modulated at odd harmonics of the line frequency. The structure of the line noise was shown to depend on the bias field direction, and a method to fit the frequency modulation was demonstrated, enabling the spectrogram to be demodulated.

Birefringence induced by the glass window of the science chamber was seen to result in an elliptical polarisation at the position of the atoms, resulting in a vector light-shift. The tight focus of the probe beam caused longitudinal variation of the probe intensity, resulting in an effective magnetic field gradient, causing periodic collapse and revival of the Faraday signal. Introducing a $\frac{\lambda}{4}$ -waveplate allowed the ellipticity to be adjusted, and a method was devised that used the Faraday signal to determine when the beam was linearly polarised and eliminate the vector light-shift.

The quadratic Zeeman shift was found to be non-negligible at the bias fields of interest, resulting in two Larmor contributions to the Faraday signal, effectively resulting in rapid amplitude modulation. Microwave dressing was used to radiatively-shift the $|F = 1, m_F = 0\rangle$ state to enhance the quadratic shift and visibly split the two tones, or cancel the quadratic shift to eliminate the modulation and produce a long-lived signal.

The splitting between the two tones for small microwave detunings was applied to measure the microwave Rabi frequency in agreement with the value obtained by direct observation of Rabi oscillations on the clock transition.

Characterisation of the Faraday signal is now complete, and long-timescale Faraday measurements of the BEC are possible having eliminated the vector light-shift and the quadratic Zeeman shift. I shall next discuss applying a magnetic field gradient to achieve magnetic resonance imaging of the condensate.

Magnetic resonance imaging of BEC

In this chapter I discuss how to produce the specific magnetic field gradients required for MRI, without the need for dedicated gradient coils. I demonstrate using these applied gradients to cancel the background field gradient, and perform radiative spin echo to rephase the signal. I apply rf π -pulses to obtain multiple spin echoes, and use a long pulse sequence in close analogy with NMR to measure the decoherence time.

With control over magnetic field gradients and spin echo obtained, I achieve the first MRI images of a BEC. I reconstruct images of a partially evaporated BEC, showing a bimodal profile, and of twin BECs in a split dipole trap, clearly resolving the individual BECs. I further show that rephasing with spin echo enables time-resolved multi-shot imaging.

8.1: Coordinate systems

The role of the *strength* of the magnetic field bias and gradient have been considered, but not how they relate to the orientation of components in the laboratory reference frame (Figure 8.1). The following considerations dictate the orientations of the magnetic field and its gradient for Faraday MRI:

1. The propagation axis of the Faraday beam is the BEC spin projection component probed by the Faraday effect (see §2.5). However, it is primarily determined by optical access, collection numerical aperture and beam pointing stability (see §5.4 and §5.6).
2. The orientation of the bias field relative to the probe propagation axis determines the strength of the Faraday interaction, as it defines the plane of Larmor precession. It should be perpendicular to the probe's propagation axis where possible.
3. The bias field direction also sets which terms of the background field gradient tensor contribute to $\nabla|\mathbf{B}|$ and thus dephasing (discussed in §8.3), sensitivity to ambient field fluctuations (observed in §7.2), and the Rabi frequency of the microwave antenna used for cancelling the quadratic Zeeman shift (see §7.7).
4. The direction of the net magnetic field gradient $\nabla|\mathbf{B}|$ sets the imaging axis, i.e. the axis of the reconstructed density profile (discussed in §8.2).

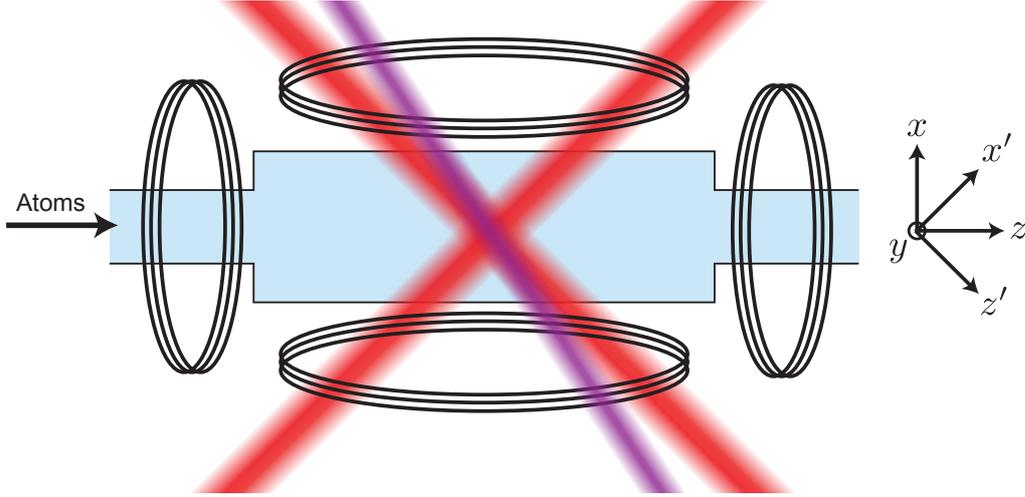


Figure 8.1: Schematic of the science chamber showing the apparatus coordinate systems. Unprimed coordinates correspond to the axes of the vacuum system and magnetic field bias coils, while primed coordinates correspond to the orientation of the trapping beams (red). The y and y' axes are both oriented out of the page, and the Faraday beam (purple) propagates at $\theta_F \approx 3^\circ$ to the trap beams.

In particular, the type and orientation of available magnetic field coils limits which magnetic field gradients can be easily generated, and what the strength of the applied gradient is. The gradient strength limits the imaging resolution by (3.15) and must be chosen to limit Stern-Gerlach blurring (see §3.8 and Table 3.1). However, the probe beam does not propagate along the axes of symmetry of the coil pairs (Figure 8.1) so it is important to clarify the coordinate system.

The discussion presented thus far has considered Faraday measurement in a geometry where the propagation direction of the probe beam is $+z$, the magnetic field is oriented along y and the gradient is $\frac{\partial B_y}{\partial z}$. For consistency with our other publications, the $+z$ direction is henceforth taken as the primary axis of the atom beam (i.e. the long axis of the vacuum system in Figure 4.1), and a primed coordinate system is defined corresponding to the orientation of the trapping beams. The primed axes are at 45° to the unprimed axes, though both take gravity to be in the $-y$ -direction. Under this designation, the Faraday beam propagates at a small angle θ_F to the $+z'$ -direction.

The distinction is significant, as our apparatus is only capable of generating magnetic field gradients along the axes of the bias coils, whereas a gradient along the dipole trap (z') would be required to image a quasi-1D condensate. Ideally the coils would be centred around the trapping beams for this reason, but the geometric limitations of the science cell would require significantly larger coils at a greater distance from the cell driven by a significantly larger current. It is still possible to generate the required $\frac{\partial B}{\partial z}$ gradient by generating gradients with both the x - and z -coils, but limitations on the driving currents (see §8.5) prevent large gradients from being produced.

8.2: Imaging axis

Faraday MRI uses a magnetic field gradient to induce a spatially varying Larmor frequency along the imaging axis. This encodes spatial information about the atom cloud in frequency space and allows an image to be reconstructed (see §2.5). However, in that discussion it was necessary to generate $\frac{\partial B_x}{\partial z}$ as the dominant gradient, which is an off-diagonal gradient that cannot be produced without specialised coil configurations.¹ I shall therefore discuss how the image is formed in the magnetic field gradients that are physically achievable using the existing (near-) Helmholtz coils present in our apparatus.

Remaining in the linear Zeeman regime but allowing the magnetic field to vary in all directions, the expression for the local Larmor frequency (3.6) can be linearly approximated² as

$$\omega_L(\mathbf{r}) = \gamma|\mathbf{B}(\mathbf{r})| \approx \omega_{L0} + (\nabla|\mathbf{B}|) \cdot \mathbf{r}. \quad (8.1)$$

The measured Faraday signal (3.9) is therefore

$$\Delta P \propto \iiint_{-\infty}^{\infty} \cos(\omega_L(\mathbf{r})t) \rho(\mathbf{r}) d^3\mathbf{r}. \quad (8.2)$$

Consider a new coordinate system \mathbf{r}'' where z'' lies in the direction $\nabla|\mathbf{B}|$. Then $\rho(\mathbf{r})$ is the only term in the integral that depends on x'' and y'' , so the integral over those directions can be performed and the signal becomes

$$\Delta P \propto \int_{-\infty}^{\infty} \cos(\omega_{L0} + \gamma b z'') \bar{\rho}(z'') dz'', \quad (8.3)$$

$$\text{where } b = \nabla|\mathbf{B}| \cdot \mathbf{r}'' = \left| \nabla|\mathbf{B}| \right| = \frac{\partial|\mathbf{B}|}{\partial z''} \quad (8.4)$$

$$\text{and } \bar{\rho}(z'') = \iint_{-\infty}^{\infty} \rho(x'', y'', z'') dx'' dy''. \quad (8.5)$$

Now the line density $\bar{\rho}$ corresponds to integrating the density of the cloud over the axes perpendicular to z'' , with each frequency component of the Faraday signal corresponding to a different position along z'' . Hence $\nabla|\mathbf{B}|$ is the imaging axis for 1D Faraday imaging.

Since the imaging axis is usually specified with respect to the orientation of the probe beam (z'), the derivative can be evaluated in that coordinate system such that the imaging axis is defined by

$$\nabla|\mathbf{B}| = \left(\frac{\partial|\mathbf{B}|}{\partial x'}, \frac{\partial|\mathbf{B}|}{\partial y'}, \frac{\partial|\mathbf{B}|}{\partial z'} \right), \quad (8.6)$$

¹ A dedicated set of gradient coils were designed and constructed [173], but are not yet installed.

² A magnetic field with non-trivial field curvature means the Larmor frequency cannot be easily inverted for position and is therefore of limited use for imaging.

where each of the spatial derivatives can be calculated in the primed or unprimed coordinates as

$$\frac{\partial|\mathbf{B}|}{\partial x_j} = \frac{\mathbf{B}}{|\mathbf{B}|} \cdot \frac{\partial\mathbf{B}}{\partial x_j} = \frac{B_{x'}}{|\mathbf{B}|} \frac{\partial B_{x'}}{\partial x_j} + \frac{B_y}{|\mathbf{B}|} \frac{\partial B_y}{\partial x_j} + \frac{B_{z'}}{|\mathbf{B}|} \frac{\partial B_{z'}}{\partial x_j}. \quad (8.7)$$

Hence it is possible to control both the imaging direction and resolution through the gradient terms $\frac{\partial B_i}{\partial x_j}$. As discussed in §3.8 and §7.3, the gradient should be as large as possible to permit rapid measurements and minimise the effect of blurring. However, only certain gradient terms can be produced by the available coil configurations, which restricts the imaging process.

8.3: Magnetic field gradient synthesis

In our apparatus the bias field is generated by pairs of coils in a near-Helmholtz configuration (see §4.1). It is possible to create a small magnetic field gradient by driving unequal current through the coils in a pair, enabling generation of non-zero $\frac{\partial B_x}{\partial x}$, $\frac{\partial B_y}{\partial y}$ or $\frac{\partial B_z}{\partial z}$. To quite a good approximation,³ the BEC is at the magnetic centre of the bias coils and the off-diagonal gradients generated by the coils vanish [138]. However, background field gradients exist in all directions so the net gradient terms in (8.7) are

$$\frac{\partial|\mathbf{B}|}{\partial x_j} \approx \frac{B_j}{|\mathbf{B}|} \frac{\partial B_{\text{coils},j}}{\partial x_j} + \frac{\mathbf{B}}{|\mathbf{B}|} \cdot \frac{\partial \mathbf{B}_{\text{bg}}}{\partial x_j}, \quad (8.8)$$

The background gradient of our apparatus has been previously characterised by tensor gradiometry using Ramsey interferometry, and the tensor in unprimed coordinates is [188]

$$\left[\frac{\partial B_i}{\partial x_j} \right]_{\text{bg}} = \begin{pmatrix} -5.71(7) & -6.92(4) & 14.70(7) \\ -6.92(4) & 15.18(8) & 2.66(4) \\ 14.95(3) & 2.66(4) & -9.47(3) \end{pmatrix} \text{mG/cm}. \quad (8.9)$$

These field gradients are primarily due to the unshielded permanent magnets of an ion pump located near the science chamber.

To prevent the background gradient from dominating the dephasing rate and interfering with image formation, the applied gradients should dominate the background. This again encourages strong applied gradients and short measurement times when performing Faraday imaging. The result (8.8) also demonstrates that the choice of bias field direction is important, since the dominant background gradient varies with the chosen orientation. In particular, the net magnetic field gradient is weaker with the bias aligned along x' than along y .⁴

³ The BEC is within 1 mm of the magnetic centre, compared to the separations (110, 157, 210) mm between the coils.

⁴ The exact orientation of the x' -bias was empirically optimised to minimise the contribution of background gradients, i.e. to minimise $|\nabla|\mathbf{B}_{\text{bg}}||$ with the background gradient tensor in (8.9).

First consider a bias field direction perpendicular to the propagation direction, such as along the y -direction. The resulting imaging axis is well approximated by

$$\nabla|B| = \left(\frac{\partial B_{\text{bg},y}}{\partial x'}, \frac{\partial B_{\text{coils},y}}{\partial y} + \frac{\partial B_{\text{bg},y}}{\partial y}, \frac{\partial B_{\text{bg},y}}{\partial z'} \right), \quad (8.10)$$

which resolves along the y -axis provided the applied gradient $\frac{\partial B_y}{\partial y}$ is much stronger than the background gradient.

It is also possible to choose the applied gradient so that it cancels the background gradient term $\frac{\partial B_{\text{bg},y}}{\partial y}$. This results in prolonging the free-induction decay of the Faraday signal, allowing for longer interrogation times as limited by the remaining background terms. However, this orients the imaging axis in the x - z plane in a way that is sensitive to the background gradient. Such a configuration would only be beneficial for diagnostic purposes as it produces the longest-lived signal, as the net gradient is very weak and the imaging axis cannot be controlled.

It is also possible to image directly along the x' -direction by applying a bias field in that direction as this is also perpendicular to the direction of optical propagation. As the primed coordinates (x', z') are oriented at approximately 45° to the unprimed (x, z) , the gradient of the magnetic field strength in this scenario is

$$\nabla|B| = \left(\frac{1}{\sqrt{2}} \left(\frac{\partial B_{\text{coils},x}}{\partial x} - \frac{\partial B_{\text{coils},z}}{\partial z} \right) + \frac{\partial B_{\text{bg},x'}}{\partial x'}, \frac{\partial B_{\text{bg},x'}}{\partial y}, \frac{\partial B_{\text{bg},x'}}{\partial z'} \right), \quad (8.11)$$

where the relevant diagonal gradient terms $\frac{\partial B_x}{\partial x}$ and $\frac{\partial B_z}{\partial z}$ can be produced by either x - or z -coils respectively, or some combination of the two.

To image along the probe propagation direction (z') without the use of off-diagonal gradients, the dominant gradient must be

$$\frac{\partial|B|}{\partial z'} \gg \frac{\partial|B|}{\partial x'}, \frac{\partial|B|}{\partial y}, \quad (8.12)$$

which implies that $B_{z'}$ should be made as large as possible.

However, this tips the Larmor precession axis towards z' , which reduces projection of the Larmor precessing spins onto the probe propagation axis. This introduces a gradient at the expense of reducing the amplitude of the Faraday signal (see §3.2).

The bias field should therefore be oriented primarily perpendicular to the propagation direction, although with some component along z' to produce a gradient. Taking the bias to be inclined at an angle ϕ to the z' -axis which, without loss of generality, can be taken to be in the y - z' plane,

$$B_{x'} = 0, \quad B_y = |B| \sin \phi, \quad B_{z'} = |B| \cos \phi. \quad (8.13)$$

This enables the synthesis of a linear gradient along the z' direction (Figure 8.2). Note that without the bias component $B_{z'}$, $\frac{\partial B_{z'}}{\partial z'}$ would still contribute to a gradient of $|\mathbf{B}|$ along the z' direction, but it would be quadratic with position,

$$|\mathbf{B}(z')| \approx |\mathbf{B}| + \frac{1}{2|\mathbf{B}|} \left(\frac{\partial B_{z'}}{\partial z'} \right)^2 z'^2, \quad (8.14)$$

which is of limited use for imaging applications as it is weak and cannot be uniquely inverted for position.

The components of this field gradient in the presence of a bias are then

$$\begin{aligned} \frac{\partial |\mathbf{B}|}{\partial x'} &\approx \left(\frac{\partial B_x}{\partial x} - \frac{\partial B_z}{\partial z} \right) \frac{\cos \phi}{\sqrt{2}} + \frac{\partial B_{\text{bg},y}}{\partial x'} \sin \phi + \frac{\partial B_{\text{bg},z'}}{\partial x'} \cos \phi \\ \frac{\partial |\mathbf{B}|}{\partial y} &\approx \frac{\partial B_y}{\partial y} \sin \phi + \frac{\partial B_{\text{bg},y}}{\partial y} \sin \phi + \frac{\partial B_{\text{bg},z'}}{\partial y} \cos \phi \\ \frac{\partial |\mathbf{B}|}{\partial z'} &\approx \left(\frac{\partial B_x}{\partial x} + \frac{\partial B_z}{\partial z} \right) \frac{\cos \phi}{\sqrt{2}} + \frac{\partial B_{\text{bg},y}}{\partial z'} \sin \phi + \frac{\partial B_{\text{bg},z'}}{\partial z'} \cos \phi. \end{aligned} \quad (8.15)$$

Hence to image along the z' -axis an unequal current is run through both the x - and z -coil pairs to achieve $\frac{\partial B_x}{\partial x} = \frac{\partial B_z}{\partial z} = b$, a bias is applied along the y -axis ($B_y \neq 0$) but without a gradient ($\frac{\partial B_y}{\partial y} = 0$). Applying a bias component in the z' direction ($B_{z'} \neq 0$) then gives

$$\nabla |\mathbf{B}| \approx (0, 0, \sqrt{2}b \sin \phi). \quad (8.16)$$

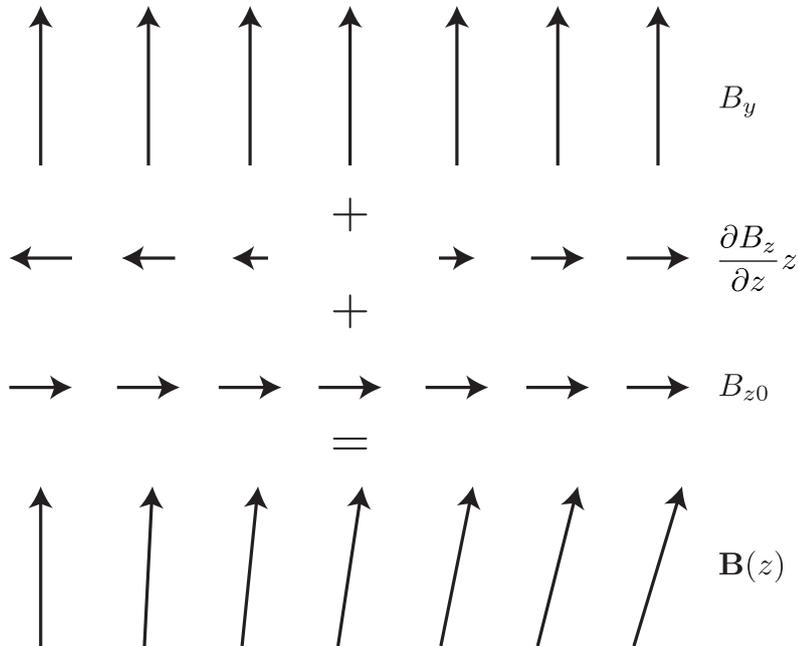


Figure 8.2: Generation of a linear magnetic field gradient along the z -axis. Note that the bias field component B_{z0} weights the contribution from $\frac{\partial B_z}{\partial z}$, so must be nonzero.

With an applied gradient chosen, and taking the background fields as small, the local Larmor frequency (8.1) varies spatially in a y -bias as

$$\omega_L(\mathbf{r}) \approx \omega_{L0} + \gamma \frac{\partial B_y}{\partial y} y, \quad (8.17)$$

and in a yz' -bias (8.13) as

$$\omega_L(\mathbf{r}) \approx \omega_{L0} + \gamma \frac{\sin \phi}{\sqrt{2}} \left(\frac{\partial B_x}{\partial x} - \frac{\partial B_z}{\partial z} \right) x' + \gamma \frac{\sin \phi}{\sqrt{2}} \left(\frac{\partial B_x}{\partial x} + \frac{\partial B_z}{\partial z} \right) z'. \quad (8.18)$$

Choosing $\frac{\partial B_x}{\partial x} = \frac{\partial B_z}{\partial z} = b$ to image along the z' -axis, this simplifies to

$$\omega_L(\mathbf{r}) \approx \omega_{L0} + (\sqrt{2}\gamma b \sin \phi) z'. \quad (8.19)$$

Therefore despite lacking special purpose gradient coils,⁵ our apparatus is capable of generating magnetic field gradients along different axes.

8.4: Faraday gradiometry

The Faraday measurement allows the centre Larmor frequency to be measured to high precision using STFT analysis, even when the free-induction decay of the signal in a gradient is rapid. This enables gradiometry to be performed by translating the cloud and looking at how the Larmor frequency changes with position. Such translation can be achieved in the x' - and z' -directions by shifting the rf frequency that drives one of the AOMs that deflect the trapping beams [138].

Measuring gradients along the z' direction in this way is most convenient, because the cloud does not leave the Faraday probe beam, whereas in the x' and y directions the cloud can only be translated by $\sim 10 \mu\text{m}$ before the image of the cloud is clipped by the iris on the Faraday imaging stage. The iris can be opened further to enable translation over greater distances, but this admits more shot noise into the measurement, decreasing the signal-to-noise ratio (see §3.3). However, the diameter of the probe beam is an upper limit on the distance the BEC can be translated and a Faraday signal obtained.

Furthermore, one of the dipole trapping beams can be split by driving the AOM with a combination of two frequencies, each of which produces a diffracted order. This results in two proximate crossed-beam dipole traps (Figure 8.3). Forced evaporation in these split traps results in two BECs,⁶ separated by $28.3(1) \mu\text{m}/\text{MHz}$ (when splitting along the z -axis). The resulting Faraday signal will have contributions from the Larmor frequencies of each cloud, enabling the magnetic field to be measured at both points simultaneously for common-mode rejection of changes in the magnetic field strength.

⁵ Typically Golay coils, a specific arrangement of saddle coils [198, 199], are used to generate the strong transverse magnetic field gradients required for medical MRI.

⁶ The frequency difference must be at least 3 MHz to ensure the traps are completely separated, as otherwise a single malformed BEC is produced.

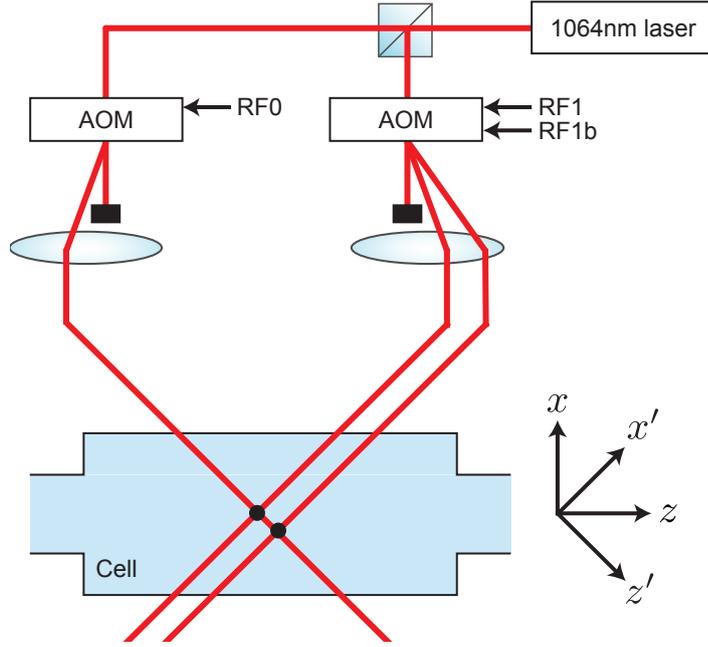


Figure 8.3: The dipole trap beams are generated as the 1st-diffracted order of AOMs, enabling them to be laterally translated inside the cell by adjusting the frequency of the driving rf. Mixing a second frequency (RF1b) into one of the AOMs results in two diffracted beams, creating a double-well trap in the cell and producing two BECs (black dots).

However, the overall amplitude of the Faraday signal is decreased as each split BEC typically contains significantly fewer atoms than a single BEC. Improving the atom number by including a ‘transfer’ stage in the experiment staging during which the BECs are formed in a small separation split trap, then moved to their final baseline by slewing the AOM drive frequency, could enable high-precision magnetometry to be performed [138].

This gradiometry method is applied in §8.6 to calibrate the magnetic field gradient generated by our quadrupole coils, and the split crossed-beam dipole trap technique is used in §8.10 to capture magnetic resonance images of twin BECs.

8.5: Bias coil gradients

As discussed in §8.3, the bias coils can be used to produce a small magnetic field gradient by running a differential current through the coil pairs. The current in each coil is controlled by an individual control voltage (e.g. the x -coils have separate V_{x+} and V_{x-} lines), which are normally set to be equal to produce a uniform field at the position of the atoms (at the ‘bias value’, $V_{x\pm} = V_x$). However, a gradient proportional to the control voltage difference $\delta V_x = V_{x+} - V_{x-}$ can be introduced by setting

$$V_{x\pm} = V_x \pm \frac{1}{2}\delta V_x. \quad (8.20)$$

This acts as a combination of a Helmholtz coil (V_x) that creates a uniform bias field and an anti-Helmholtz coil ($\frac{1}{2}\delta V_x$) that produces a gradient. Using the expression for the field generated by a loop of wire, the magnetic field gradient can be shown to be

$$\frac{\partial B_x}{\partial x} = \frac{3\alpha_x}{5\Delta x}\delta V_x, \quad (8.21)$$

where α_x is the calibration factor (§7.1) and Δx the spacing of the coils.

The coil drivers cannot change the sign of the current, so $V_{i\pm} \geq 0$. The strongest gradient that can be produced by this arrangement is therefore $\delta V_i = 2V_i$. This introduces significant limitations to the possible gradients that can be produced, as typically only one component has a large control voltage value, enabling a strong gradient to only be produced along that axis (Table 8.1). It should be noted that in our apparatus the strength of the x' - and yz' -bias fields cannot be further increased without inverting the current in one of the coil pairs,^{7,8} which is presently not implemented. However, the y -bias can be increased by simply increasing the control voltage V_y . The bias coils are therefore only able of creating large $\frac{\partial B_y}{\partial y}$ gradients in y - and yz' -biases, and $\frac{\partial B_x}{\partial x}$ gradients in an x' -bias.

Bias	V_x (V)	V_y (V)	V_z (V)	f_L (kHz)	$ B $ (G)	$t_{\pi/2}$ (μ s)
y	0.1958	0.6283	0.8744	698	0.994	14.2
x'	0.4210	0.1559	0.0211	597	0.850	16.7
yz'	0	0.6283	0.2788	843	1.200	19.0

Table 8.1: Typical bias field control voltages, measured Larmor frequencies, and $\frac{\pi}{2}$ -pulse durations for different bias field directions. These settings were used in the data presented in the remainder of this thesis. The Larmor frequency for a given set of control voltages was observed to slowly drift by 7 kHz over the course of three months, corresponding to drift in the background field.

Because of this limitation on the gradient strength, the gradient induced is typically the same order of magnitude as the background field gradient. Measuring the coherent dephasing rate of the Faraday signal gives a measure of the gradient strength (see §3.5), and it can be seen that the signal lifetime is maximised for a particular gradient where the applied gradient cancels the background gradient (Figure 8.4). The reduced dephasing at this point results in a much longer-lasting signal (Figure 8.5), with the remaining dephasing due the other (uncancelled) components of the background gradient.

⁷ Specifically, V_z in an x' -bias and V_x in a yz' -bias cannot be reduced without inverting the current.

⁸ Note that the coils are wired such that $V_i > 0$ opposes the background magnetic field. Hence reducing the coil current for $V_i < V_{i,\text{null}}$ corresponds to increasing the *net* magnetic field strength (see §7.1 and Table 7.1). In principle the same magnetic field strength could be obtained with $V_i > V_{i,\text{null}}$, but this would require much higher coil currents. The coils are not actively cooled so this would cause heating and likely affect the coil calibration.

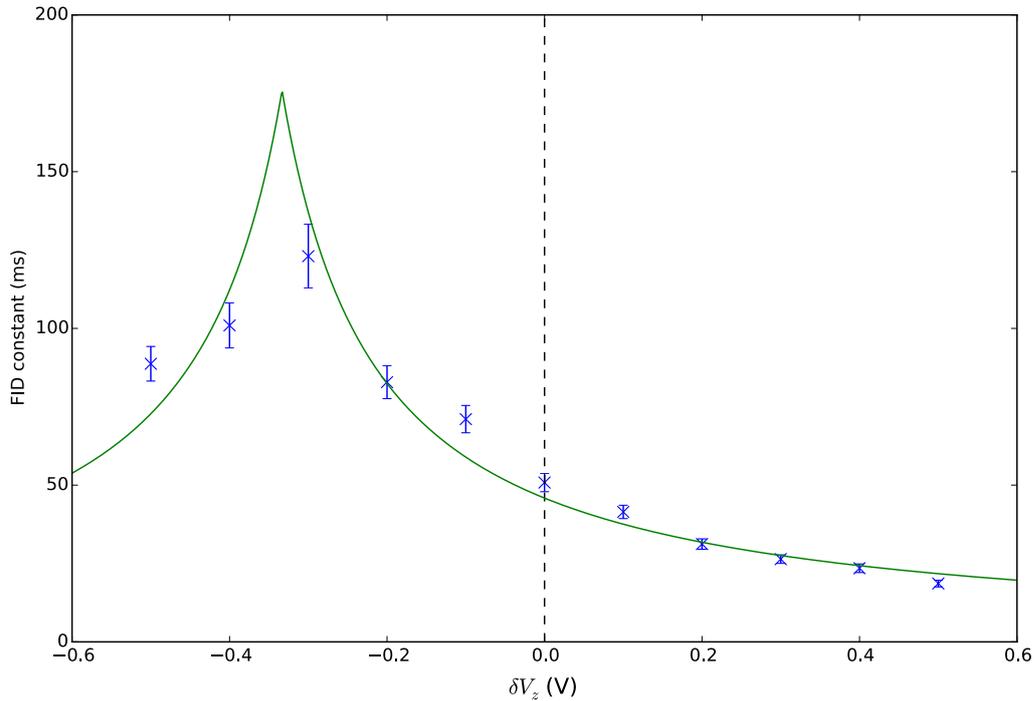


Figure 8.4: Measured dependence of the Faraday signal dephasing time-constant as a function of the applied gradient δV_z in a y -bias, indicating that the background gradient is suppressed when $\delta V_z = -0.35$ V. The dephasing rate contains contributions from all components of the background gradient, each of which could be independently suppressed for a long-lived Faraday signal.

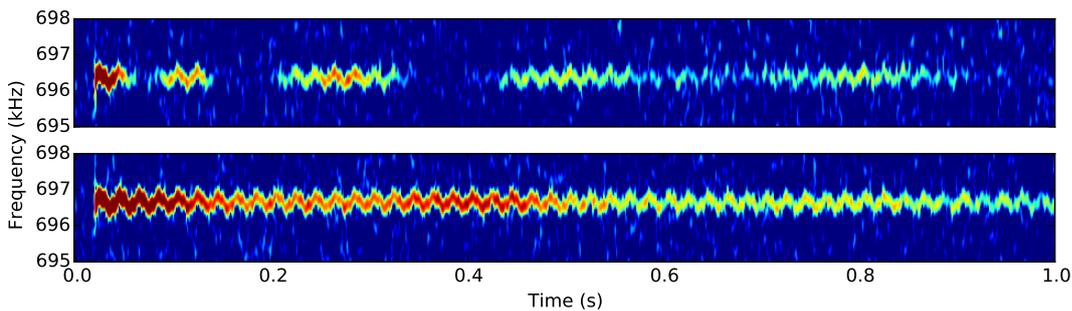


Figure 8.5: Measured Faraday signal without an applied gradient (top) and with gradient cancelling the background (bottom), resulting in reduced dephasing to achieve a long-lived signal (up to $t = 1.5$ s). y -bias, $\delta V_z = -1.0$ V. Note that the VLS (see §7.5) was not cancelled for this measurement so the applied field cancels both the background field and VLS contributions.

8.6: Quadrupole coil gradients

It is also possible to synthesise a magnetic field gradient using the quadrupole coils that form the basis of our MOT and magnetic traps. This results in diagonal gradient terms which can be made significantly stronger than those generated by the bias coils.

The field generated by the quadrupole coils is

$$\mathbf{B} = \frac{B_q}{2}(-\mathbf{x} + 2\mathbf{y} - \mathbf{z}), \quad (8.22)$$

which produces diagonal gradients

$$\left[\frac{\partial B_i}{\partial x_j} \right]_{\text{quads}} = \frac{B_q}{2} \begin{pmatrix} -1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \quad (8.23)$$

The quadrupole coils are driven by a custom IGBT-based voltage-controlled current supply,⁹ which takes a control voltage and drives 20 A/V through the coils, producing a quadrupole gradient of 1.89 G/cm/A.

It should be noted that although the quadrupole driver is capable of generating gradients up to 300 G/cm, applying gradients larger than 5 G/cm to a dipole-trapped BEC resulted in no atoms remaining at the end of the hold time. This is because the dipole trap holds the atoms against gravity in the y -direction, and applying strong $\frac{\partial B_y}{\partial y}$ further lowers the trap depth allowing atoms to fall out.

Unlike the bias coils, since the quadrupole centre is (deliberately) not coincident with the BEC position, applying a gradient through the quadrupole coils changes the net magnetic field at the position of the atoms and hence the Larmor frequency (Figure 8.6). This calibration is necessary to set the frequency of the rf pulses for a given gradient strength to prevent pulse from going off-resonance, which would significantly reduce the efficacy of the rf pulse.

The quadrupole zero is known to be located directly above the BEC, resulting in a magnetic field gradient that is primarily $\frac{\partial B_y}{\partial y}$ in a y -bias. This enables the distance Δy between the quadrupole zero and the BEC to be estimated through $\Delta\omega_L = \gamma B_q \Delta y$, resulting in $\Delta y = 85.3(7) \mu\text{m}$.

The quadrupole gradient is observed to enhance the gradient-induced signal dephasing (Figure 8.7), with the free-induction decay rate being linear in the applied gradient strength as expected (see §3.5). The dephasing time constant was measured by fitting a half-Gaussian to the averaged envelope of the Fourier-filtered signal (§6.7). This is necessary because resolving the timescale of the decay in the spectrogram would require the window size of the STFT to be reduced substantially (see §6.10).

⁹ The quadrupole driver exhibits a long switch-on transient (~ 100 ms), so it is important to apply the gradient and allow it to settle before the initial $\frac{\pi}{2}$ -pulse that tips the spins to begin Larmor precession. The settling time is longer for smaller B_q , becoming prohibitively long for $B_q \lesssim 0.1$ G/cm.

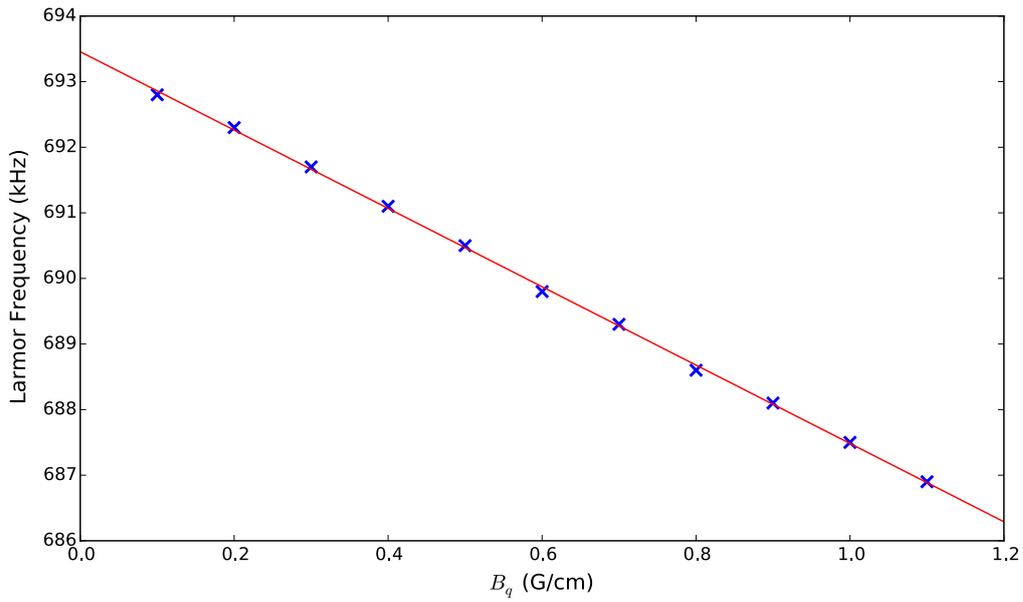


Figure 8.6: Increasing the quadrupole gradient B_q in a y -bias field changes the net magnetic field strength at the position of the BEC, and hence its Larmor frequency. The fit gives $\frac{\partial f_L}{\partial B_q} = -5.97(5)$ kHz/(G/cm).

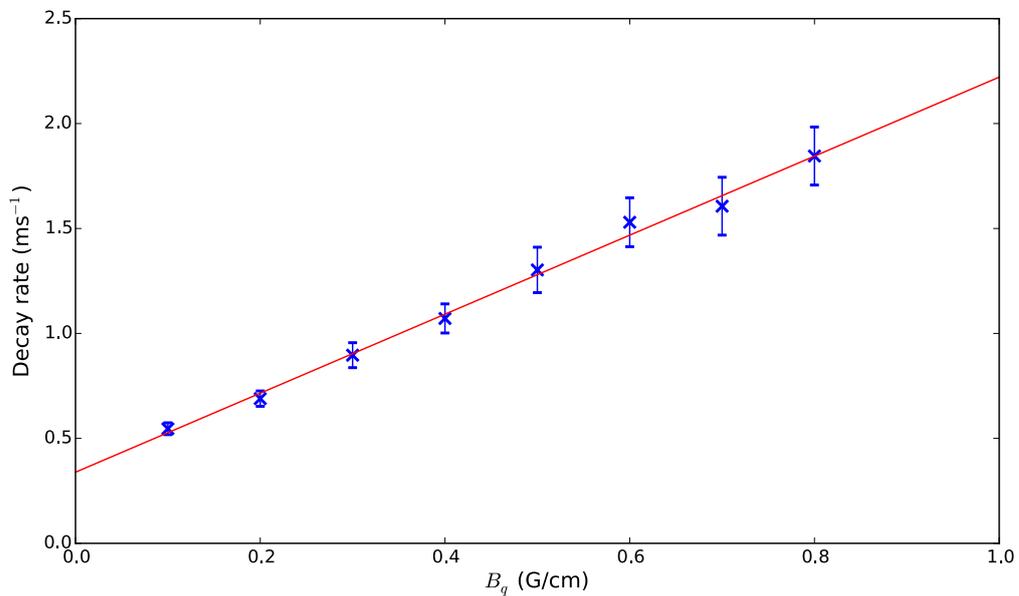


Figure 8.7: Applying a quadrupole gradient increases the gradient-induced dephasing rate, resulting in more rapid free-induction decay (FID) of the signal. The decay rate is linear in B_q when it dominates the background gradient contributions. The decay rate at higher values of B_q was too rapid to measure.

The strength of the induced gradient is measured using Faraday gradiometry (§8.4) to measure the Larmor frequency of the BEC as its location is translated along the z' -axis using the dipole trapping laser (Figure 8.8).

8.7: Radiative spin echo

Spin echo forms an integral part of the Faraday imaging procedure as it rephases the Faraday signal and enables multiple sequential images to be taken of the dynamics occurring within the BEC. Performing a spin echo also provides strong evidence that the decay of the signal is the result of coherent dephasing, and not the result of another process interfering with the measurement.¹⁰

Because the ‘bias’ coils cannot generate a gradient strong enough to reverse the background gradient, and our quadrupole coil driver is not capable of inverting the sign of B_q in (8.22), the gradient cannot be reversed in our apparatus and gradient recall echo (see §3.6) cannot be performed. However, radiative spin echo should be possible by applying rf π -pulses to the cloud.

The π -pulse is an rf pulse at the Larmor frequency lasting twice as long as a $\frac{\pi}{2}$ -pulse, which results in inversion of the spin state. Pulse fidelity affects the amplitude of the

¹⁰ Examples of other phenomena that result in signal loss are atom loss mechanisms (§5.8), quadratic Zeeman shift (§7.6) and spatial separation (§7.9).

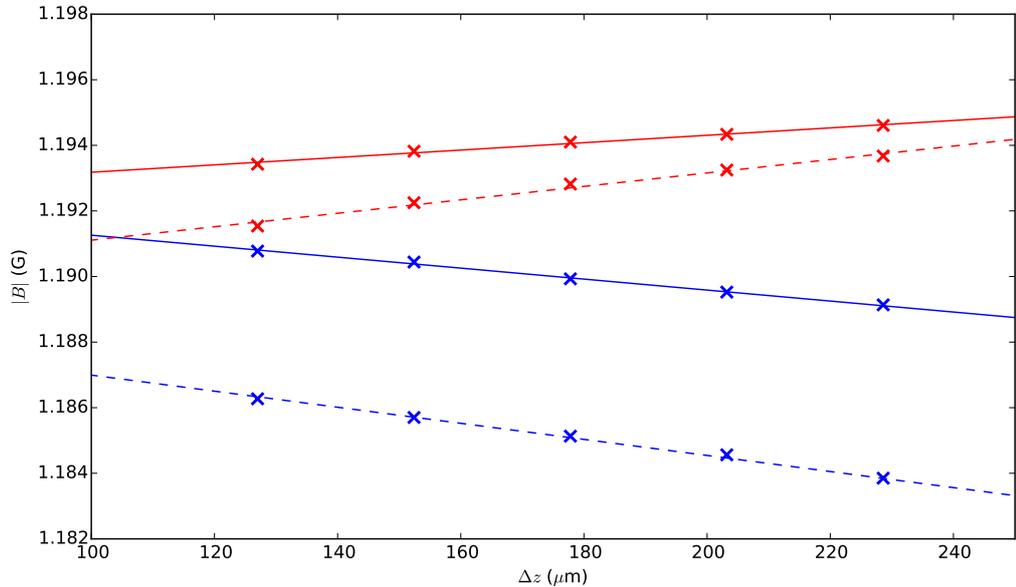


Figure 8.8: Gradiometry of the quadrupole coil (yz' -bias) performed by measuring the change in Larmor frequency with position along the $+z'$ (red) and $-z'$ (blue) directions for $B_q = 0.2$ G/cm (solid) and $B_q = 0.5$ G/cm (dashed).

echo, and can be improved by tuning the pulse duration in a given gradient. However, spin echo can be achieved using the calibrated magnetic field strength without further optimisation. Applying a π -pulse at time t_0 is seen to reverse the dephasing and produce a revival of the Faraday signal at time $2t_0$ (Figure 8.9). The echo is strong enough to be observed in a single-shot measurement but as the echo is consistent and repeatable, envelope averaging (§6.7) can be used to clarify the structure.

Further echoes can be produced with subsequent π -pulses. The shape of the signal is preserved with each echo, albeit with reduced amplitude. In accordance with (3.27), the decay envelope is modelled as approximately Gaussian, so the expected Faraday signal for a train containing N π -pulses can be written as

$$S(t) = \sum_{n=0}^N A_n \exp\left(-\frac{(t - nt_e - t_0)^2}{2t_D^2}\right), \quad (8.24)$$

where A_n is the amplitude of the n^{th} echo, t_e is the echo time,¹¹ t_0 is the time of the initial $\frac{\pi}{2}$ -pulse, and t_D is the decay constant common to all echoes. This enables the echo envelope from a train of π -pulses to be fit simultaneously for an accurate estimation of the decay constant t_D (Figure 8.10).

The amplitude A_n of the n^{th} echo decreases with n , corresponding to imperfect rephasing of the spin which results in complicated spin dynamics. Since each echo corresponds to an observation of the cloud, it is natural to ask how many sequential observations can be made before the signal decays into the noise.

¹¹ The π -pulses are applied at $t = (n - 1/2)t_e$ with corresponding echoes appearing at $t = nt_e$.

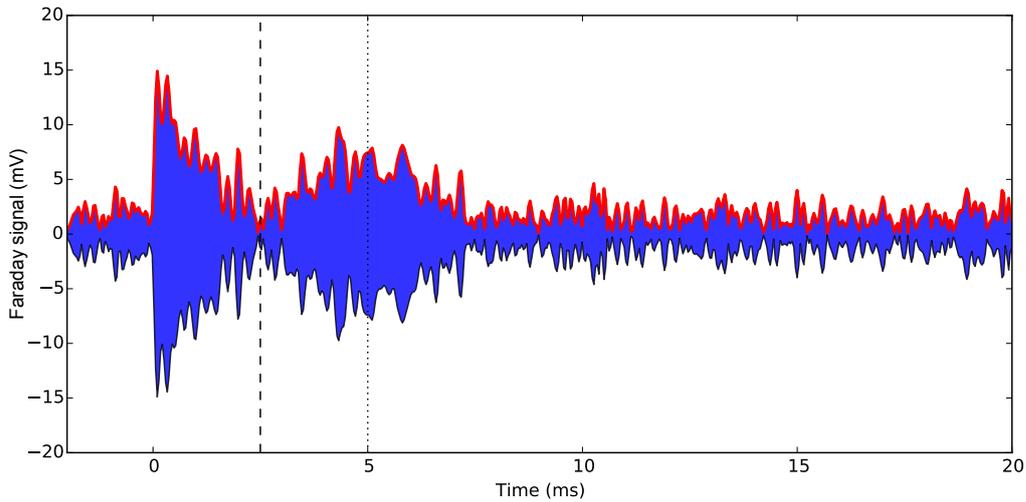


Figure 8.9: Single-shot observation of a spin echo in the filtered Faraday signal at $t_E = 5.0$ ms (dotted line) by applying an rf π -pulse at time $t_0 = 2.5$ ms (dashed line). Fourier-filtered with a 10 kHz band-pass filter, signal envelope shown in red. Captured in a y -bias with $B_q = 0.2$ G/cm.

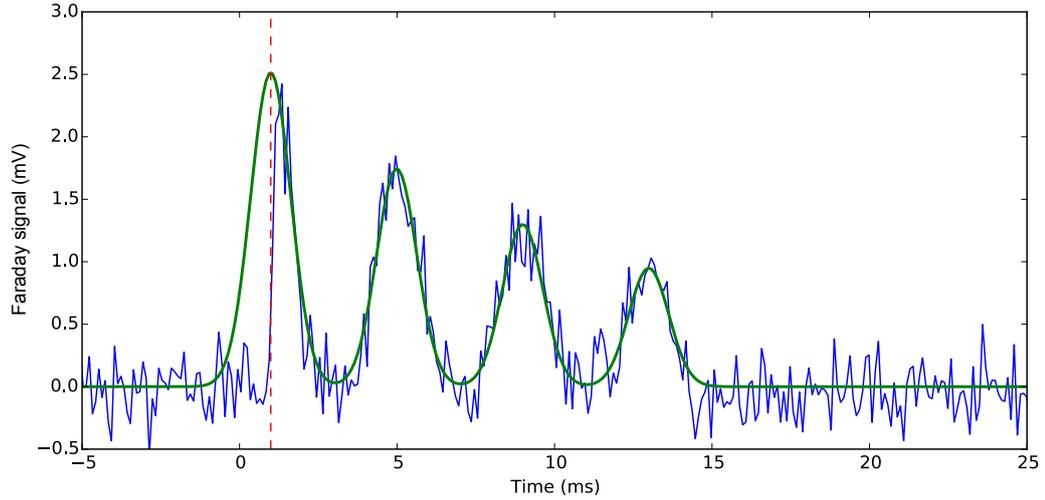


Figure 8.10: Observation of multiple spin echoes with echo time $t_e = 3$ ms, which can be simultaneously fit to accurately determine the coherent dephasing time $t_D = 0.90(3)$ ms. Dashed line corresponds to the initial $\frac{\pi}{2}$ -pulse. Envelope averaged $10\times$, captured in a yz' -bias with $B_q = 0.5$ G/cm.

Preparing a pulse train of 50 π -pulses (Figure 8.11) demonstrates that echoes can be resolved above the background noise level for the entire sequence. Fitting the height of the echoes as previously described allows the decoherence to be quantified (Figure 8.12), and is observed to decay exponentially. The characteristic timescale of this decoherence is traditionally referred to as the T_2 time, which for the captured pulse sequence was found to be $33(3)$ ms.

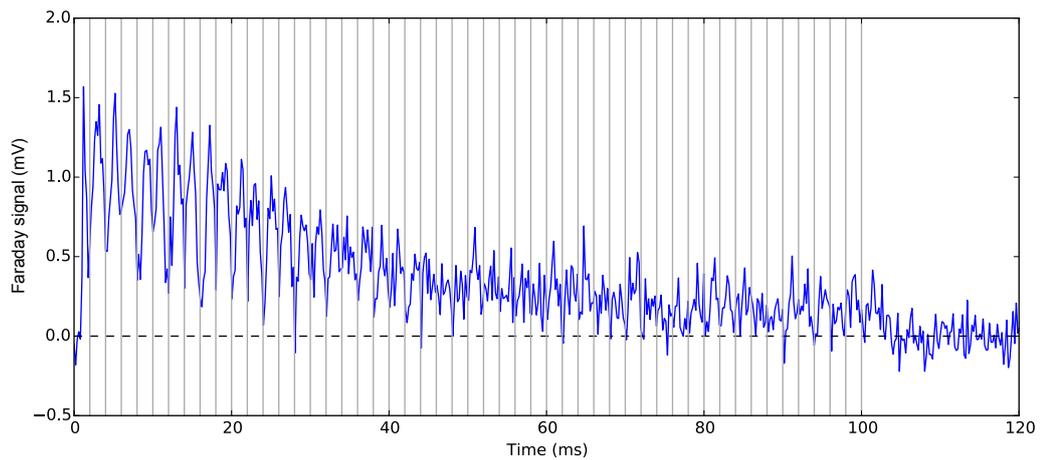


Figure 8.11: Applying a train of 50 π -pulses (grey vertical lines) with $t_e = 2$ ms results in echoes that can be clearly resolved for the first ~ 20 pulses, and distinguished above the background noise until the end of the train, demonstrating that some coherence has been retained through multiple rephasings. Envelope averaged $10\times$, captured in a yz' -bias with $B_q = 0.5$ G/cm.

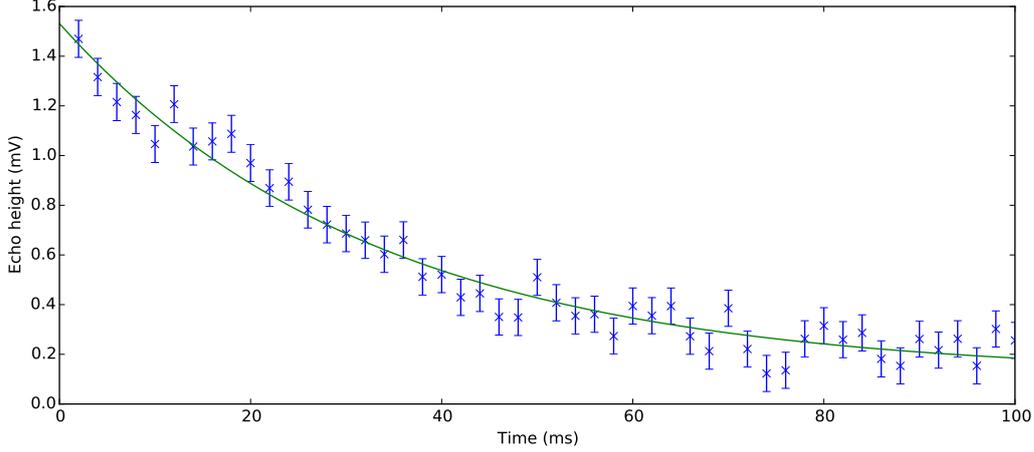


Figure 8.12: Fitting the amplitudes of the echoes in the 50 π -pulse sequence shows decay over time, with a characteristic timescale (T_2) of 33(3) ms.

The decoherence time T_2 with an applied gradient is much shorter than the observed T_2^* time of the Faraday signal when the background gradient is cancelled (observed to be ~ 1 s in Figure 8.5), supporting the hypothesis that this decoherence is limited by the fidelity of the π -pulses.¹² This train of π -pulses corresponds to a Carr-Purcell (CP) sequence, in which each pulse is identical and any error induced by the imperfection of each π -pulse is cumulative. The error can be improved by using a Carr-Purcell-Meiboom-Gill (CPMG) sequence, where the phase of consecutive π pulses is reversed, so the error introduced by one pulse can be ‘reversed’ by the subsequent pulse. This shall be investigated in future work.

The MRI community has developed a wide variety of pulse sequences [200] for coherent spin control, many of which can be immediately adapted to improve the fidelity of rephrasing pulse trains applied to a BEC. These techniques will be indispensable in developing Faraday imaging for further time-resolved studies of BEC dynamics.

8.8: Considerations for MRI of BEC

The Faraday measurements presented thus far have considered an approximately spherical BEC held in a crossed-beam dipole trap. The resulting Faraday measurement is uninteresting, as the signal envelope follows a Gaussian decay at short evolution times, while the ‘interesting’ structure corresponding to the BEC’s Thomas-Fermi profile appearing at longer times (§3.5) is potentially being buried in the measurement noise. The success of spin echo observations (§8.7) provides evidence that the decay of the Faraday signal is due to the expected gradient-induced coherent dephasing, but

¹² Although the rf pulse is unshaped, it is short and therefore spectrally broad. In a weak gradient, the pulse remains effective for the entire cloud, but shaped pulses (see §3.7) are required in a strong gradient.

the reconstructed Gaussian profile is inconclusive as other decay mechanisms could result in the same signal envelope.

Scenarios that present a more interesting profile to image will therefore be considered. Ideally, a 1D elongated BEC could be produced with a feature such as a dark soliton to be imaged. However, not only is our trap geometry currently incompatible with the production of elongated BECs, but the currently achievable magnetic field gradients mean the resolution of the image produced (see §3.9) by the current generation apparatus is severely limited and cannot presently resolve features at the healing length scale (see §3.9).

In the following sections I present two configurations that generate more interesting profiles, with structure that can be resolved through magnetic resonance imaging using currently achievable gradients. The first is a partially-evaporated BEC, where the forced rf evaporation procedure is interrupted early, resulting in a cold atom cloud that contains both significant condensed and thermal fractions. The second is to use a split dipole trap to produce two spatially-separated BECs. The Faraday beam passes through both BECs, and a gradient applied along the separation axis enables them to be resolved.

8.9: MRI of a partially evaporated BEC

Bose-Einstein condensates are produced in our apparatus by performing forced evaporation in the crossed-beam dipole trap. This is achieved by ramping down the power of the dipole trapping laser to selectively outcouple the ‘hot’ atoms from the edge of the trap, reducing the overall temperature of the cloud at the expense of reducing the number of trapped atoms (see [138] for details on our implementation).

When the cloud becomes cold enough, part of it condenses in the centre, forming a cloud that contains both thermal and condensed fractions. The two fractions have different momentum distributions, so releasing the cloud from the trap results in different expansion rates. Time-of-flight imaging therefore observes a bimodal distribution, corresponding to a tight Thomas-Fermi profile from the condensed fraction upon a wide Gaussian background from the thermal fraction. This clear distinction between thermal and condensed fractions has long been considered the experimental ‘hallmark’ of condensation.

The evaporation process can be interrupted by truncating the evaporation ramp to observe the onset of condensation. The truncation is controlled by a parameter $\xi \in [0, 1]$, where $\xi = 0$ (full truncation) results in a fully thermal cloud and $\xi = 1$ (no truncation) corresponds to complete condensation with no thermal fraction remaining. Varying the truncation parameter changes the balance of thermal and condensed fractions, which can then be quantified by capturing an MRI of the partially evaporated cloud.

This presents an *in situ* characterisation of the evaporation process, paving the way towards performing continuous MRI of a cloud as it undergoes evaporation as a minimally destructive experimental diagnostic.¹³

The distinction between condensed and thermal fractions is clear in *momentum* space, but Gross-Pitaevskii simulations [202, 203] have shown that the *spatial* profile of a partially evaporated cloud is also expected to be bimodal. The distinction between the two fractions is not as clear as in momentum space, but the thermal fraction is nevertheless expected to occupy a larger spatial extent than the condensed fraction, so the Faraday signal from the thermal fraction should dephase more rapidly than the condensate fraction. The Faraday signal in the time domain is therefore expected to be the bimodal sum of two Gaussians¹⁴ that have different time constants.

Large magnetic field gradients are not required to resolve the overall shape of the cloud, permitting the weaker bias coils to be used. This results in slow gradient-induced dephasing, allowing the Faraday signal to be captured over an extended measurement time with dense sampling of the envelope. This is necessary to correctly resolve the bimodality of the distribution in the time domain, and permits accurate fits to be made.

To reduce the influence of noise, envelope averaging (§6.7) was used to combine multiple shots together. In a more refined experiment, a brighter probe beam would reduce or eliminate the need for averaging. The probe beam brightness could be increased by up to $30\times$ without significant loss of SNR due to photon scattering, while a $3\times$ increase in brightness would give the same SNR as $10\times$ averaging.

The captured Faraday signals (Figure 8.13) are well approximated as the bimodal sum of two Gaussians, except in the case of complete evaporation to BEC ($\xi = 1$) where the signal is accurately modelled by a single component. Inverting the time-domain signal using §3.2 reconstructs the spatial profile of the cloud along the z' -axis. As the truncation parameter is increased, atoms are lost from the trap resulting in lower integrated line density $\bar{\rho}(z')$, and the overall width of the profile shrinks (Figure 8.14).

Comparing the profiles obtained by reconstructing the fitted time-domain signal demonstrates how the profile changes as the evaporation advances. This demonstrates that the MRI technique is indeed capable of distinguishing between the thermal and condensed fractions, and could be used to diagnose the evaporative process in real time.

¹³ The Faraday effect has already been applied as a diagnostic technique to measure how the atom number changes during evaporation [201]. This enables post-selection of shots based on the number of atoms at the end of evaporation, reducing the statistical error introduced by BEC atom number fluctuations.

¹⁴ The profiles of the individual fractions are only approximately Gaussian, but the corresponding Faraday signal appears Gaussian but with structure at longer evolution times (see §3.5).

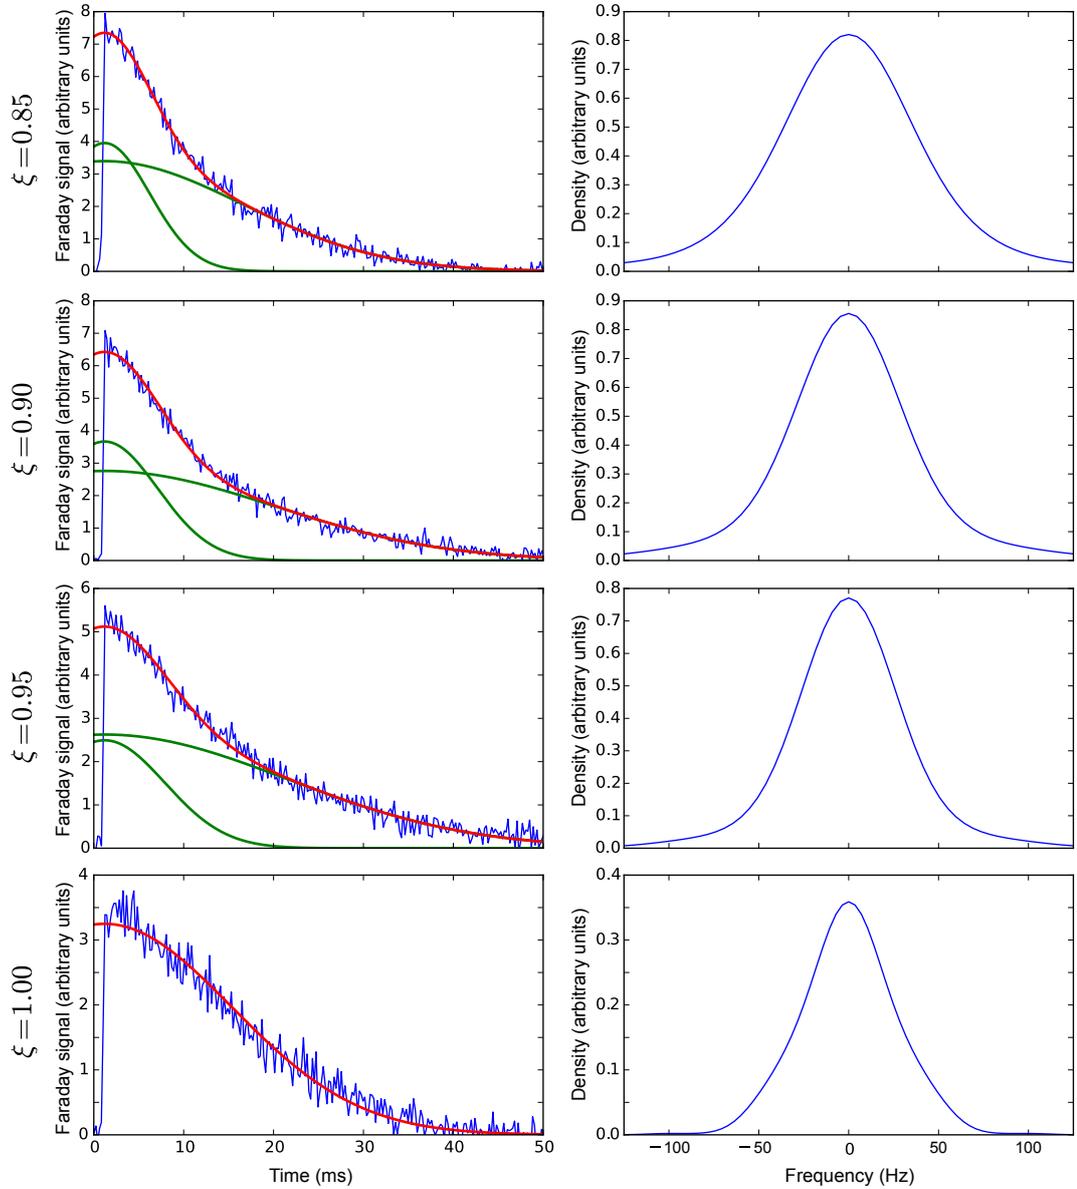


Figure 8.13: The Faraday signal of a partially evaporated BEC with truncation parameter ξ is the bimodal sum of two Gaussians (left), corresponding to the thermal and condensed parts of the cloud. The reconstructed profiles (right) show a decreased spread as the truncation is increased towards total condensation ($\xi = 1$). The frequency axis of the reconstruction is related to position by the effective gradient (8.16), which was not calibrated. Captured in a yz' -bias with $\delta V_{Bz} = 0.5$ V and averaged $\times 10$.

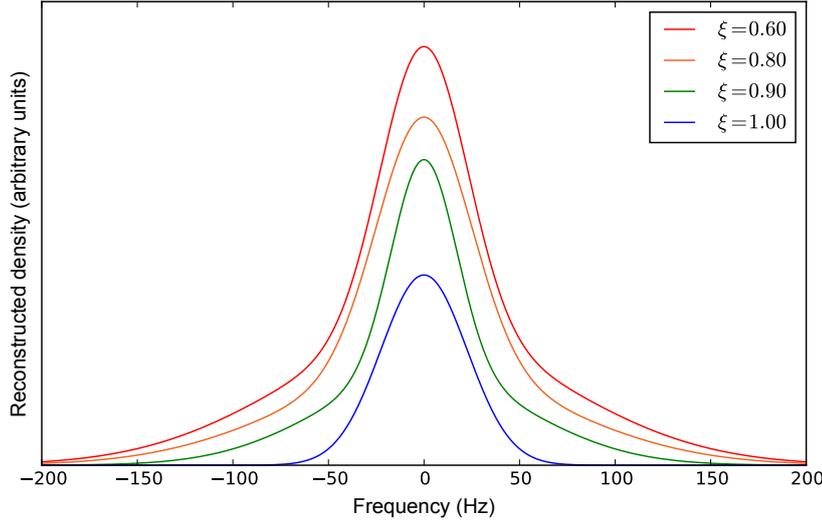


Figure 8.14: Comparison of magnetic resonance images reconstructed for partially evaporated clouds with different truncation parameters ξ using bimodal fits.

8.10: MRI of twin BECs

As described in §8.4 the dipole trapping beams can be split in two by adding a second rf source into the corresponding AOM that controls the beam position and amplitude. This forms two separate traps, enabling the creation of ‘twin BECs’ by transferring the magnetically trapped cloud directly into the separate dipole traps (provided the two beams are sufficiently separated). When splitting the beam along the z' -axis, the profile of the twin BECs along the probe beam is that of two distinct Thomas-Fermi profiles. This is in direct analogy to the first demonstration of nuclear MRI¹⁵ using two cylinders of distilled water immersed in heavy water by Lauterbur in 1973 [97].

In order to resolve the two BECs individually in the spectrogram,¹⁶ the difference between their centre Larmor frequencies must be much greater than the spectrogram resolution, requiring that the gradient be

$$b \gg \frac{2\pi \Delta f}{\gamma \Delta z'}, \quad (8.25)$$

where Δf is the spectrogram resolution (typically 200 Hz) and $\Delta z'$ is the separation (‘splitting distance’) between the trap centres. The traps are loaded most efficiently for small splitting distances (see §8.4), with the minimum splitting given by the dipole trap beam waist ($\sim 70 \mu\text{m}$). The magnetic field gradient must therefore be

$$b \gg 40 \text{ mG/cm}. \quad (8.26)$$

¹⁵ Specifically, this was the first image that could distinguish between molecules containing the same elements but with different isotopic composition (hydrogen and deuterium).

¹⁶ As discussed in §7.6, if the Faraday signals from the two clouds fall within the same ‘bin’, they form a beatnote and the resulting spectrogram appears to be amplitude-modulated.

This does not appear to be a demanding gradient, but as the clouds are split along the z' -axis, the imaging axis is also the probe propagation axis. As discussed in §8.2 this requires a yz' -bias field, which by (8.16) reduces the efficacy of the applied gradient. Keeping the bias field tip angle (ϕ in (8.13)) small gives a strong Faraday signal¹⁷ but requires a strong applied gradient to resolve the cloud separation.

The apparatus is also capable of preparing twin BECs separated along the x' -axis, which removes these requirements on the bias field direction and gradient strength. However, the tight focal waist of the probe laser means that the twin BECs produced are no longer contained within the Faraday beam when split along this axis, preventing reconstruction of an image.

Therefore an MRI was captured using a yz' -bias and the quadrupole coils to generate the gradient necessary to resolve the BECs (Figure 8.15). In this spectrogram, the $\frac{\pi}{2}$ -pulse occurred at $t = 2$ ms and the reconstructed profile of the twin BECs can be seen as a line-out along the frequency axis, with two ‘bumps’ corresponding to the individual BECs clearly seen. From the known splitting distance $\Delta z' = 283(1) \mu\text{m}$ and measured centre Larmor frequencies, by (8.16) the *effective* gradient in this bias field is

$$b \equiv \frac{\partial |\mathbf{B}|}{\partial z'} = B_q \sin \phi = 118(5) \text{ mG/cm}. \quad (8.27)$$

¹⁷ The split trap loads less efficiently than a single crossed-beam trap, so the Faraday signal from each BEC is already reduced compared to an unsplit trap because it contains fewer atoms in total.

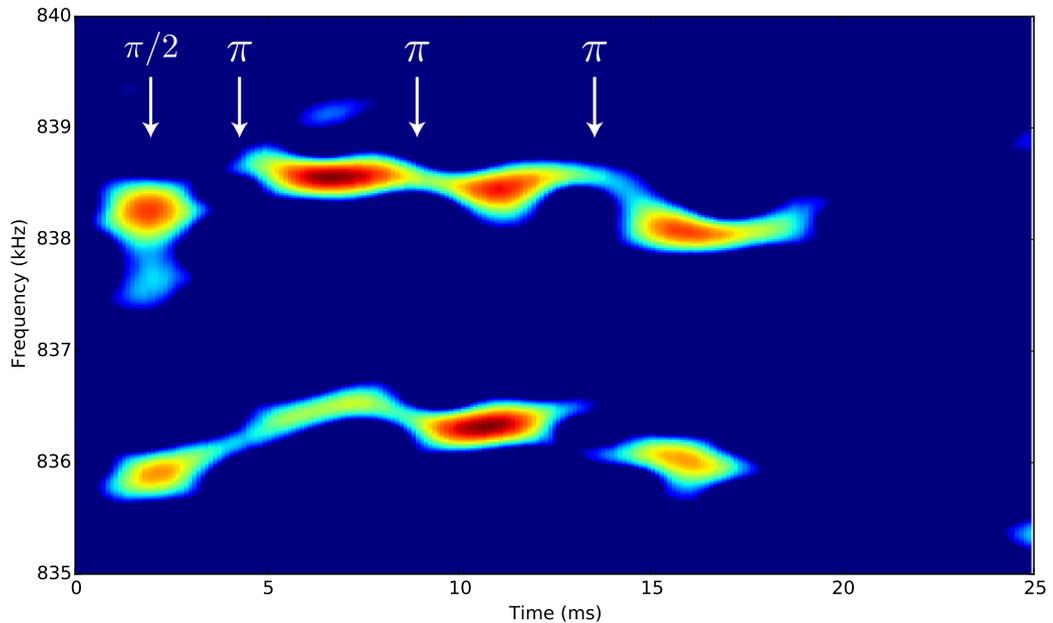


Figure 8.15: Spectrogram of twin BECs in a split dipole trap, clearly resolving the two BECs in the frequency domain, and successfully rephased by three subsequent echoes in the time domain. Acquired in a yz' -bias with $B_q = 0.2 \text{ G/cm}$ and no averaging.

This is a demonstration of the simultaneous common-mode gradiometry method suggested in §8.4, although the rapid dephasing prevents more accurate determination of the centre Larmor frequency.

Time-resolved imaging is demonstrated by periodically rephasing the spins with multiple π -pulses. Magnetic field fluctuations originating from power-line noise can be seen in the time-evolution of the spectrogram, with the centre Larmor frequency of each BEC clearly showing frequency modulation (as discussed in §7.2). The important observation is that the frequency modulation is *common mode*, and it could be eliminated from the measurement with matched filter signal processing.

The analysis presented here is not the conventional inverse Fourier transform analysis of (3.14), but is equivalent as every column of the spectrogram itself is a Fourier transform. In a conventional analysis the signal would be divided into windows corresponding to each echo in the measurement, whereas the overlapped spectrogram provides a visualisation that shows how these snapshots change over time. In particular, the frequency modulation from power line fluctuations would produce significant blurring in a conventional image that can be readily identified in the spectrogram.

However, the present inability to generate strong magnetic field gradients in this bias configuration severely limits the resolution of the reconstructed image. This proves that magnetic resonance imaging with the Faraday effect can resolve structure *in situ*, and solutions to generating stronger magnetic field gradients are outlined in §9.2.

8.11: Summary

In this chapter I discussed applying magnetic field gradients to the BEC, and presented the first magnetic resonance images of Bose-Einstein condensates.

I derived the relationship between the applied magnetic gradient and the imaging axis which the profile is reconstructed over. I outlined the generation of magnetic field gradients and how imaging along the probe propagation axis can be achieved without specialised coils to create the ‘off-diagonal’ magnetic field gradients. I discussed the gradients that can be generated with our current coil configurations, and the limitations. I presented coil calibrations and observed how the free-induction decay rate was affected by the applied gradient – and that cancelling the background gradient resulted in a significantly reduced dephasing rate.

I applied rf π -pulses to produce radiative spin echoes, proving dephasing is coherent and allowing precise determination of coherent dephasing time. A train of pulses was seen to rephase the spin and produce echoes out to a large time, and the decay in the amplitude of the echoes was used to quantify the decoherence rate (T_2 time). This decoherence was attributed to the imperfect fidelity of the π -pulses, and the CPMG sequence was proposed as a way to circumvent imperfect pulses.

I captured magnetic resonance images of a partially evaporated BEC containing a thermal and condensed fraction for varying truncation parameters to investigate the change in cloud shape approaching the condensation transition. The time-domain signal was shown to be bimodal, with the signal corresponding to the condensed fraction persisting to longer times. Reconstructing the cloud profiles showed a reduction in the spread as the cloud was brought closer to the condensation point. Although not as clear as time-of-flight absorption images, the thermal and condensed fractions could be clearly distinguished.

I then captured magnetic resonance images of twin BECs formed in a split dipole trap. The two BECs could be clearly resolved in the spectrogram, and were successfully rephased with multiple spin echo pulses. This demonstrates both the minimally-destructive nature of the measurement, and the ability to observe time-resolved spatial dynamics.

At present the resolution of the magnetic resonance images is limited by the magnetic field gradients that can currently be synthesised. Quadrupole gradients were applied in a yz' -bias to image along the z' axis, but dedicated gradient coils could vastly improve the resolution.

Conclusions and outlook

9.1: Conclusions

This thesis has presented a theoretical and experimental study of a novel magnetic resonance imaging technique for *in situ* imaging of Bose-Einstein condensates (BECs) using the Faraday effect. The minimally destructive nature of the measurement enabled long interrogation times multi-shot imaging to be performed, paving the way for time-resolved studies of *in situ* dynamics.

In the first chapter, I discussed the physics of Bose-Einstein condensates at the healing-length scale, and the dynamics which cannot be resolved *in situ* by existing imaging techniques. The Faraday effect was introduced as a technique to probe the condensate state *in situ* using off-resonant light, and I described how a magnetic field gradient could be used to produce a 1D profile of the BEC. I described how techniques from magnetic resonance imaging (MRI) using magnetic field gradients along multiple axes can be applied to build up 2D and even 3D images of the BEC. The resolution of such an imaging technique would not be subject to the diffraction limit, raising the prospect of multi-shot imaging of features that cannot otherwise be resolved without destructively expanding the BEC. This opens up a range of new possibilities in the use of condensates to simulate complicated systems where the same initial condition cannot be replicated, such as quantum turbulence.

Building up a 3D image requires many ‘shots’ to be taken in sequence, necessitating long exposure times. It is therefore desirable for the interaction to be minimally destructive and not cause any extra dynamics that perturb the BEC being imaged. In the second chapter I provided a theoretical analysis of the atom-light interaction in the far off-resonant limit, where the dynamics are governed by a tensor polarisability that couples atomic spin to the polarisation of the probe beam. The scalar polarisability demonstrated that the trapping potential vanished at specific ‘magic’ wavelengths, from which $\lambda = 790\text{ nm}$ was chosen as the probe wavelength. The Faraday effect was shown to arise from the vector polarisability, which also produced an unintended effective magnetic field if the probe beam had an elliptical polarisation. The tensor polarisability, which would induce complicated dynamics in the BEC, was shown to be negligible in the regime of interest.

Using this characterisation of the Faraday effect, in the third chapter I developed the Faraday interaction from a simple spin probe to a method for extracting a 1D profile of the cold atom cloud by using a magnetic field gradient. The signal-to-noise ratio of this measurement was shown to depend on the ratio of measurement duration to scattering lifetime. I considered coherent spin dephasing in terms of the time-domain Faraday signal, and saw that small features generated revivals at longer measurement times. I quantified the resolution of the reconstructed image in terms of the applied magnetic field gradient, with large magnetic field gradients and short measurement times encouraged to prevent Stern-Gerlach blurring. I showed that sub-micron resolution was shown possible, although the necessary gradients are beyond the capabilities of our current apparatus.

From this theoretical base, I began to describe the experimental implementation in the fourth chapter. At the commencement of this project, there were no cold atom experiments at our institution, so our group designed and constructed an apparatus for the production and study of spinor condensates from the ground up. The design was centred around a large glass science chamber, with atoms sourced from a Zeeman slower and condensed in a crossed-beam dipole trap. As this was a collaborative effort spanning several projects, my in-depth discussion covers only my primary contributions, which have not been described elsewhere. This included a digital beatnote lock that enabled the Zeeman repump laser to be reliably locked to an arbitrary detuning and maximise the load rate of the magneto-optical trap. A microcontroller-based interlock was created to control and monitor the oven, ensuring the oven remains in a safe state at all times. The modular imaging application, BIAS, which was created to interactively control the scientific cameras used for absorption imaging, and its role in the wider control system was detailed.

With reliable production of spinor BECs with large atom number, in the fifth chapter I discussed the experimental considerations for generating the Faraday beam to probe them. The Faraday measurements rely on polarisation stability of the probe beam, and the polarisation-maintaining fiber was found to induce large polarisation fluctuations when the polarisation axis did not match the fiber. Methods for aligning the polarisation axis, and for centering the small probe waist on the BEC were described. The lifetime of the BEC in the trap was measured, and the amplified spontaneous emission of the diode laser generating the Faraday beam was found to contain a component resonant with the atoms causing rapid scattering. An interference filter blocked the resonant components, and the resulting scattering rate agreed with the theoretical predictions. The tightly focused probe beam therefore does not perturb the trap or cause resonant scattering, enabling the BEC to be probed for long timescales.

In the sixth chapter I described polarimetry of the probe beam after passing through the BEC, and constructed a shot-noise-limited photodetector which was calibrated using an 'electro-optic BEC'. Faraday measurement of a BEC was performed, and the

structure of the resulting signal analysed. Short-time Fourier transforms were used to generate spectrograms to facilitate interpretation of the signal, which was seen to be both frequency- and amplitude-modulated on the same timescale.

The structure of the Faraday signal was investigated further in the seventh chapter. The frequency modulation was seen to result from magnetic field fluctuations induced by electrical components near the science cell as the modulation was formed by harmonics of the power line frequency. The windows of the science cell induced birefringence on the probe beam, resulting in the atoms experiencing an effective magnetic field because of the vector light-shift. This was quantified by precision measurement of the Larmor frequency, and cancelled using a quarter-waveplate before the cell. Rapid amplitude modulation of the Faraday signal was shown to result from the quadratic Zeeman effect, which produced two Larmor tones that formed a beat-note. Microwave dressing was applied to cancel the quadratic shift and eliminate the amplitude modulation.

Having eliminated the undesirable dynamics induced by the probe beam, in the eighth chapter I performed Faraday imaging experimentally. Our apparatus is not capable of producing 'off-diagonal' gradients, so imaging with the weak 'diagonal' gradients that can be produced was considered. Faraday measurement of the Larmor frequency was used to calibrate the gradients, and applying a gradient to cancel the background gradient was observed to increase the signal lifetime. Radiative π -pulses were then observed to rephase the Faraday signal and generate a spin. This proves that the dephasing process is coherent, which is crucial for building up profiles to produce higher-dimensional images or enable time-resolved studies to be performed.

Magnetic resonance images were then captured of two configurations with non-trivial structure. A partially evaporated atom cloud was seen to produce a bimodal Faraday signal with components from both the thermal and condensed fractions, enabling the two to be distinguished. Varying the truncation parameter showed the thermal fraction disappears, leaving only the condensed component – displaying the potential for use as an experimental diagnostic during evaporation to BEC. A split dipole trap was used to produce twin spatially separated BECs, which were clearly resolved by Faraday imaging. The twin BECs were rephased by up to four spin echoes, demonstrating the spatial and temporal resolution of the technique.

9.2: Outlook

This work demonstrated that magnetic resonance imaging with the Faraday effect can be used for minimally destructive observation of a spinor BEC *in situ*. There are several immediate extensions that can be applied to the work presented, and multiple directions that could be pursued in the near-future. Some modifications are conceptually straightforward but technically challenging to implement, such as the introduction of a cavity *in vacuo* to suppress spontaneous emission, and some involve the application of more advanced MRI techniques to BEC.

The resolution of the magnetic resonance images presented in this work was limited by the strength of the available magnetic field gradients. Without specialised gradient coils, the existing bias and quadrupole coils were used instead. The bias coils could only generate a weak gradient, and while the quadrupole coils were capable of generating a much stronger gradient, this weakened the trap against gravity resulting in the condensate ‘falling out’ of the trap. Control of the magnetic field gradient direction and strength is currently the greatest weakness of our apparatus, which we plan to address in the near future.

Alternatively, changing the trap geometry to be quasi-1D would allow the condensate to expand along the long-axis and relaxing the constraints on the strength of the applied gradient since the large spatial extent results in a wider spread of Larmor frequencies. In principle this can be achieved by weakening one of the dipole trapping beams in the crossed-dipole trap, but in practice this results in significant atom number loss. Creating a dedicated 1D-trap would require reconstruction of the dipole trap optics, which the group intends to pursue.

Installation of dedicated gradient coils – such as current bars or Golay coils – could create strong gradients, but installing them around the vacuum system is a technical challenge. In particular, to image along the z' -axis (which would be the long-axis of the quasi-1D trap), they would need to be placed above and below the fragile science cell in the few millimetres of clearance between the cell window and the quadrupole coils. Geometric constraints introduced by the optics and components around the cell restrict the size and shape of the coils, and the only truly feasible design is to align the coils along the z -axis. Although this does not produce the ideal gradient terms, imaging along the required axis can still be achieved by choice of the bias field direction.

Coils capable of generating magnetic field gradients that would enable MRI at sub-micron resolution with experimentally feasible drive currents have been designed and constructed for our apparatus [173], but are yet to be installed. These coils require a drive current of > 50 A to achieve the necessary gradient for the desired resolution ($b > 33$ G/cm), and will therefore require a pulsed current source to prevent significant Joule heating. However, the coils form an inductive load, and pulsing such high

currents into them requires prohibitively large voltages. An insightful solution was devised by P. Pakkiam [173] to apply sinusoidal instead of pulsed gradients, and recover the density profile by an integral transform.

It may also be possible to generate a large *effective* magnetic field gradient using a circularly-polarised beam propagating along the x' -axis. The intensity profile of the beam will result in a non-zero gradient term $\frac{\partial B_{x'}}{\partial z'}$, which is ideal to image along the z' -axis if the bias field points along x' . As discussed in §8.3, this is not possible using diagonal gradients without tipping the bias field axis and reducing the amplitude of the Faraday signal, making an effective gradient an attractive alternative. However, this laser must also be at the magic wavelength to avoid perturbing the trap (§2.4), which would require a new high-power laser, and the condensate lifetime will be reduced by induced scattering (§5.9). It should also be noted that the intensity profile along the y -axis will potentially result in an effective magnetic field gradient that weakens the trap against gravity, as the quadrupole coils did (see §8.6).

Another limitation is that the intensity profile is Gaussian, so B_{vls} does not vary monotonically across the cloud. This means two points on opposite sides of the cloud have the same Larmor frequency, so frequency does not uniquely map to position, and the inversion (3.11) cannot be performed. This could be solved positioning the centre of the beam to the side of the BEC, so that the condensate is contained in one half of the beam, ensuring that B_{vls} does vary monotonically across the cloud. An image can be reconstructed, but the non-linear profile of the beam means the ‘bin’/‘pixel’ size will not be uniform, making interpretation difficult.

However, the strength of the magnetic field gradient can not be increased arbitrarily, as it induces Stern-Gerlach separation and separates the spin components of the cloud. While this evolution can be reversed by inverting the gradient or applying a radiative π -pulse, the Faraday measurement must be completed before the separation induces significant blurring of the reconstruction. Technical limitations such as the experimentally feasible bias field strength, the photodetection bandwidth and acquisition rate were shown to place an upper bound on the potential gradient strength (see §3.9), limiting the gradient to 290 G/cm in the present apparatus. Furthermore, while the gradient switching time doesn’t affect the imaging resolution, it must be fast enough to prevent spatial separation of the BEC (as observed in §7.9) from occurring between measurements. Other more fundamental limitations are likely to arise from a strongly varying magnetic field across the BEC, such as a spatially-dependent quadratic Zeeman shift (see §7.6) that cannot simply be cancelled by applying an opposite microwave dressing.

The real test of the Faraday imaging technique will require a feature at the healing length scale to be produced. While it would be possible to stir the condensate to produce a vortex lattice, in the first instance a quasi-1D feature such as a dark soliton

could be imaged to enable characterisation of the system before extending to higher dimensional measurements with spin sequencing. The group has a new theoretical proposal to generate black solitons using magnetic resonance control (MRC) [135, 136], which presents an attractive target for Faraday imaging as a demonstration of minimally-destructive time-resolved imaging. However, the operating principle of the MRC protocol is spatially-selective population transfer, which also requires strong magnetic field gradients, and has only been studied as a proof-of-principle in a quasi-1D condensate. The requirement to generate magnetic field gradients outlined above for imaging therefore need also be addressed for magnetic resonance control.

More techniques can potentially be adapted from medical MRI, such the application of diffusion-weighted or flow-encoding gradients to directly observe persistent flow within the condensate, and would be beneficial in applications studying quantum turbulence. The spin sequencing considered in this thesis are only the simplest possible combinations, and there is a wide range of literature (see, e.g. [200]) on combinations of gradients and pulse sequences to achieve the highest sensitivity, and compensate for experimental imperfections. As observed in §8.7, fidelity of the π -pulses was observed to limit the number of spin echoes that could be observed, and techniques such as CPMG could be implemented to overcome this. Resolving this decoherence is critical to developing Faraday MRI as a higher-dimensional imaging technique, as many sequential measurements are required to build up the reconstruction.

Although time-resolved imaging was demonstrated with twin BECs, the decoherence time observed with the train of π -pulses limits the timescale over which the time-evolution of the density can be observed. One solution (proposed in §3.6) is to create a condensate in a single Zeeman substate with an ‘interesting’ density profile, then tip the spins with a $\frac{\pi}{2}$ -pulse, capture a magnetic resonance image, apply a π -pulse to rephase the spins, then apply another $\frac{\pi}{2}$ -pulse coincident with the spin echo that transfers the population back into the original state. Evolving the system for some time and repeating the process prevents decoherence and residual dephasing from degrading the measured Faraday signal over long timescales, which is critical for time-resolved studies.

Another benefit of Faraday measurement over camera-based imaging techniques is that there is no readout time. The Faraday signal could potentially be processed electronically in real-time by a field-programmable gate array (FPGA) implementing a matched filter. This processing could generate feedback for the control system during a running experiment, which has great potential for quantum control and wavefunction engineering. The possibility exists of creating a closed-loop feedback system between MRC and MRI to enable the engineering of highly complex quantum systems.

However, this kind of magnetic resonance imaging uses the spin degree of freedom solely as a quantum-photon interface to infer the density profile of the BEC, which

ignores the richness of structure presented by the spinor. Spin textures could be studied directly through observation of the local spin projection, and Faraday imaging has already been used to observe the formation of spin domains [42]. Probing from multiple directions enables measurement of different projections of the local spin, and has been shown to produce squeezed states [204, 205].

These applications demonstrate how the fundamentally quantum nature of Faraday measurement results in a wide range of potential applications in creation and observation of complex quantum systems. This closes the circle of reciprocity between fundamental science and medicine, applying a diagnostic imaging technique borne out of nuclear physics 40 years ago to cutting-edge atomic physics.

References

- [1] E. A. Cornell and C. E. Wieman. *Nobel lecture: Bose-Einstein condensation in a dilute gas, the first 70 years and some recent experiments*. Rev. Mod. Phys., 74:875–893, 2002.
- [2] W. Ketterle. *Bose-Einstein Condensates: The Coldest Matter in the Universe*. MIT World Talk, 2001. Available at <http://mit.tv/zWLrk4>.
- [3] A. Einstein. *Quantentheorie des einatomigen idealen Gases. 2. Abhandlung*. Sitz. Ber. Preuss. Akad. Wiss., 23:3–14, 1925. Available at http://www.lorentz.leidenuniv.nl/history/Einstein_archive/Einstein_1925_publication.
- [4] A. Einstein. *Quantum Theory of the Ideal Monoatomic Gas*. Available at <http://wavewatching.net/lost-papers/>. Translation by H. Dekant and P. Terlunen.
- [5] F. London. *The λ -Phenomenon of Liquid Helium and the Bose-Einstein Degeneracy*. Nature, 141:643–644, 1938.
- [6] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. *Observation of Bose-Einstein condensation in a dilute atomic vapor*. Science, 269(5221):198–201, 1995.
- [7] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. *Bose-Einstein Condensation in a Gas of Sodium Atoms*. Phys. Rev. Lett., 75:3969–3973, 1995.
- [8] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman, and E. A. Cornell. *Vortices in a Bose-Einstein Condensate*. Phys. Rev. Lett., 83:2498–2501, 1999.
- [9] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. *Vortex Formation in a Stirred Bose-Einstein Condensate*. Phys. Rev. Lett., 84:806–809, 2000.
- [10] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari. *Theory of Bose-Einstein condensation in trapped gases*. Rev. Mod. Phys., 71:463–512, 1999.
- [11] A. L. Fetter. *Rotating trapped Bose-Einstein condensates*. Rev. Mod. Phys., 81:647–691, 2009.
- [12] S. Burger, K. Bongs, S. Dettmer, W. Ertmer, K. Sengstock, A. Sanpera, G. V. Shlyapnikov, and M. Lewenstein. *Dark Solitons in Bose-Einstein Condensates*. Phys. Rev. Lett., 83:5198–5201, 1999.

- [13] J. Denschlag, J. Simsarian, D. Feder, C. W. Clark, L. Collins, J. Cubizolles, L. Deng, E. Hagley, K. Helmerson, W. Reinhardt, et al. *Generating solitons by phase engineering of a Bose-Einstein condensate*. *Science*, 287(5450):97–101, 2000.
- [14] B. P. Anderson, P. C. Haljan, C. A. Regal, D. L. Feder, L. A. Collins, C. W. Clark, and E. A. Cornell. *Watching Dark Solitons Decay into Vortex Rings in a Bose-Einstein Condensate*. *Phys. Rev. Lett.*, 86:2926–2929, 2001.
- [15] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet. *Formation and propagation of matter-wave soliton trains*. *Nature*, 417(6885):150–153, 2002.
- [16] L. Khaykovich, F. Schreck, G. Ferrari, T. Bourdel, J. Cubizolles, L. D. Carr, Y. Castin, and C. Salomon. *Formation of a matter-wave bright soliton*. *Science*, 296(5571):1290–1293, 2002.
- [17] S. Inouye, M. Andrews, J. Stenger, H.-J. Miesner, D. Stamper-Kurn, and W. Ketterle. *Observation of Feshbach resonances in a Bose-Einstein condensate*. *Nature*, 392(6672):151–154, 1998.
- [18] F. K. Fatemi, K. M. Jones, and P. D. Lett. *Observation of Optically Induced Feshbach Resonances in Collisions of Cold Atoms*. *Phys. Rev. Lett.*, 85:4462–4465, 2000.
- [19] D. J. Papoular, G. V. Shlyapnikov, and J. Dalibard. *Microwave-induced Fano-Feshbach resonances*. *Phys. Rev. A*, 81:041603, 2010.
- [20] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga. *Feshbach resonances in ultracold gases*. *Rev. Mod. Phys.*, 82:1225–1286, 2010.
- [21] W. Ketterle, D. Durfee, and D. Stamper-Kurn. *Making, probing and understanding Bose-Einstein condensates*. In *Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics "Enrico Fermi"*, pages 67–176, 1999. arXiv:cond-mat/9904034 v2.
- [22] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard. *Optics and interferometry with atoms and molecules*. *Rev. Mod. Phys.*, 81:1051–1129, 2009.
- [23] Y. Le Coq, J. Retter, S. Richard, A. Aspect, and P. Bouyer. *Coherent matter wave inertial sensors for precision measurements in space*. *Appl. Phys. B*, 84(4):627–632, 2006.
- [24] M. J. Snadden, J. M. McGuirk, P. Bouyer, K. G. Haritos, and M. A. Kasevich. *Measurement of the Earth's Gravity Gradient with an Atom Interferometer-Based Gravity Gradiometer*. *Phys. Rev. Lett.*, 81:971–974, 1998.
- [25] J. M. McGuirk, G. T. Foster, J. B. Fixler, M. J. Snadden, and M. A. Kasevich. *Sensitive absolute-gravity gradiometry using atom interferometry*. *Phys. Rev. A*, 65:033608, 2002.

- [26] 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*. *Phys. Rev. A*, 84:033610, 2011.
- [27] S. Wildermuth, S. Hofferberth, I. Lesanovsky, E. Haller, L. M. Andersson, S. Groth, I. Bar-Joseph, P. Krüger, and J. Schmiedmayer. *Bose-Einstein condensates: microscopic magnetic-field imaging*. *Nature*, 435(7041):440–440, 2005.
- [28] 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*. *Phys. Rev. Lett.*, 98:200801, 2007.
- [29] I. Bloch, J. Dalibard, and S. Nascimbène. *Quantum simulations with ultracold quantum gases*. *Nature Phys.*, 8(4):267–276, 2012.
- [30] M. Kobayashi and M. Tsubota. *Quantum turbulence in a trapped Bose-Einstein condensate*. *Phys. Rev. A*, 76:045603, 2007.
- [31] E. A. L. Henn, J. A. Seman, G. Roati, K. M. F. Magalhães, and V. S. Bagnato. *Emergence of Turbulence in an Oscillating Bose-Einstein Condensate*. *Phys. Rev. Lett.*, 103:045301, 2009.
- [32] A. J. Allen, N. G. Parker, N. P. Proukakis, and C. F. Barenghi. *Quantum turbulence in atomic Bose-Einstein condensates*. In *Knotted, Linked and Tangled Flux in Quantum and Classical Systems*, 2013. arXiv:1302.7176 [cond-mat.quant-gas].
- [33] Q. Chen, J. Stajic, S. Tan, and K. Levin. *BCS–BEC crossover: From high temperature superconductors to ultracold superfluids*. *Physics Reports*, 412(1):1–88, 2005.
- [34] 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*. *Phys. Rev. Lett.*, 80:2027–2030, 1998.
- [35] Y. Kawaguchi and M. Ueda. *Spinor Bose-Einstein condensates*. *Phys. Rep.*, 520(5):253–381, 2012.
- [36] D. M. Stamper-Kurn and M. Ueda. *Spinor Bose gases: Symmetries, magnetism, and quantum dynamics*. *Rev. Mod. Phys.*, 85(3):1191, 2013.
- [37] D. M. Stamper-Kurn. *Seeing spin dynamics in atomic gases*. arXiv:1501.00057 [cond-mat.quant-gas], 2015.
- [38] J. Stenger, S. Inouye, D. Stamper-Kurn, H.-J. Miesner, A. Chikkatur, and W. Ketterle. *Spin domains in ground-state Bose-Einstein condensates*. *Nature*, 396(6709):345–348, 1998.

- [39] 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*. Phys. Rev. Lett., 92:140403, 2004.
- [40] C. D. Hamley, E. M. Bookjans, G. Behin-Aein, P. Ahmadi, and M. S. Chapman. *Photoassociation spectroscopy of a spin-1 Bose-Einstein condensate*. Phys. Rev. A, 79:023401, 2009.
- [41] H. Jing, Y. Jiang, W. Zhang, and P. Meystre. *Laser-catalyzed spin-exchange process in a Bose-Einstein condensate*. Phys. Rev. A, 81:031603, 2010.
- [42] L. Sadler, J. Higbie, S. Leslie, M. Vengalattore, and D. Stamper-Kurn. *Spontaneous symmetry breaking in a quenched ferromagnetic spinor Bose-Einstein condensate*. Nature, 443(7109):312–315, 2006.
- [43] M. Vengalattore, J. Guzman, S. R. Leslie, F. Serwane, and D. M. Stamper-Kurn. *Periodic spin textures in a degenerate $F = 1$ ^{87}Rb spinor Bose gas*. Phys. Rev. A, 81:053612, 2010.
- [44] 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*. Phys. Rev. Lett., 95(5):050401, 2005.
- [45] C. J. Pethick and H. Smith. *Bose-Einstein condensation in dilute gases*. Cambridge University Press, 2002.
- [46] E. G. M. van Kempen, S. J. J. M. F. Kokkelmans, D. J. Heinzen, and B. J. Verhaar. *Interisotope Determination of Ultracold Rubidium Interactions from Three High-Precision Experiments*. Phys. Rev. Lett., 88:093201, 2002.
- [47] M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, 1999.
- [48] Y. Kagan, E. L. Surkov, and G. V. Shlyapnikov. *Evolution of a Bose-condensed gas under variations of the confining potential*. Phys. Rev. A, 54:R1753–R1756, 1996.
- [49] M. Holland and J. Cooper. *Expansion of a Bose-Einstein condensate in a harmonic potential*. Phys. Rev. A, 53:R1954–R1957, 1996.
- [50] Y. Kagan, E. L. Surkov, and G. V. Shlyapnikov. *Evolution of a Bose gas in anisotropic time-dependent traps*. Phys. Rev. A, 55:R18–R21, 1997.
- [51] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag New York, 1999.
- [52] E. Lundh, C. J. Pethick, and H. Smith. *Vortices in Bose-Einstein-condensed atomic clouds*. Phys. Rev. A, 58:4816–4823, 1998.

- [53] F. Dalfovo and M. Modugno. *Free expansion of Bose-Einstein condensates with quantized vortices*. Phys. Rev. A, 61:023605, 2000.
- [54] L. D. Turner. *Holographic imaging of cold atoms*. PhD thesis, School of Physics, University of Melbourne, 2004.
- [55] G. S. Settles. *Schlieren and Shadowgraph Techniques*. Springer-Verlag Berlin Heidelberg, 2001.
- [56] E. Hecht. *Optics*. Addison-Wesley, San Francisco, 4th edition, 2002.
- [57] M. Andrews, M.-O. Mewes, N. Van Druten, D. Durfee, D. Stamper-Kurn, and W. Ketterle. *Direct, nondestructive observation of a Bose condensate*. Science, 273(5271):84–87, 1996.
- [58] K. E. Wilson, Z. L. Newman, J. D. Lowney, and B. P. Anderson. *In situ imaging of vortices in Bose-Einstein condensates*. Phys. Rev. A, 91(2):023621, 2015.
- [59] M. Andrews, D. Stamper-Kurn, H.-J. Miesner, D. Durfee, C. Townsend, S. Inouye, and W. Ketterle. *Propagation of sound in a Bose-Einstein condensate*. Phys. Rev. Lett., 79(4):553, 1997.
- [60] S. Kadlecik, J. Sebby, R. Newell, and T. Walker. *Nondestructive spatial heterodyne imaging of cold atoms*. Opt. Lett., 26(3):137–139, 2001.
- [61] L. D. Turner, K. F. E. M. Domen, and R. E. Scholten. *Diffraction-contrast imaging of cold atoms*. Phys. Rev. A, 72:031403, 2005.
- [62] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels, and P. van der Straten. *Thermodynamics of Bose-Einstein-condensed clouds using phase-contrast imaging*. Phys. Rev. A, 81:053632, 2010.
- [63] D. Freilich, D. Bianchi, A. Kaufman, T. Langin, and D. Hall. *Real-time dynamics of single vortex lines and vortex dipoles in a Bose-Einstein condensate*. Science, 329(5996):1182–1185, 2010.
- [64] A. Ramanathan, S. R. Muniz, K. C. Wright, R. P. Anderson, W. D. Phillips, K. Helmerson, and G. K. Campbell. *Partial-transfer absorption imaging: A versatile technique for optimal imaging of ultracold gases*. Rev. Sci. Instrum., 83(8):083119, 2012.
- [65] M.-O. Mewes, M. R. Andrews, D. M. Stamper-Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. *Output Coupler for Bose-Einstein Condensed Atoms*. Phys. Rev. Lett., 78:582–585, 1997.
- [66] H. Steck, M. Naraschewski, and H. Wallis. *Output of a Pulsed Atom Laser*. Phys. Rev. Lett., 80:1–5, 1998.

- [67] J. E. Debs, D. Döring, P. A. Altin, C. Figl, J. Dugué, M. Jeppesen, J. T. Schultz, N. P. Robins, and J. D. Close. *Experimental comparison of Raman and rf outcouplers for high-flux atom lasers*. Phys. Rev. A, 81:013618, 2010.
- [68] S.-w. Chiow, T. Kovachy, H.-C. Chien, and M. A. Kasevich. *102ħk Large Area Atom Interferometers*. Phys. Rev. Lett., 107:130403, 2011.
- [69] M. Faraday. *Experimental Researches in Electricity. Nineteenth Series*. Phil. Trans. Roy. Soc., 136(1):1–20, 1846.
- [70] M. Faraday. *Experimental Researches in Electricity, Vol III*, chapter 1, pages 1–11. Richard Taylor and William Francis, 1855.
- [71] D. A. V. Baak. *Resonant Faraday rotation as a probe of atomic dispersion*. Am. J. Phys., 64(6):724–735, 1996.
- [72] M. J. Freiser. *A Survey of Magneto-optic Effects*. IEEE Trans. Mag., 4(2):152–161, 1968.
- [73] D. Budker, W. Gawlik, D. F. Kimball, S. M. Rochester, V. V. Yashchuk, and A. Weis. *Resonant nonlinear magneto-optical effects in atoms*. Rev. Mod. Phys., 74(4):1153–1201, 2002.
- [74] A. Kuzmich, N. Bigelow, and L. Mandel. *Atomic quantum non-demolition measurements and squeezing*. Europhys. Lett., 42(5):481, 1998.
- [75] A. Kuzmich, L. Mandel, J. Janis, Y. E. Young, R. Eijnisman, and N. P. Bigelow. *Quantum nondemolition measurements of collective atomic spin*. Phys. Rev. A, 60:2346–2350, 1999.
- [76] Y. Takahashi, K. Honda, N. Tanaka, K. Toyoda, K. Ishikawa, and T. Yabuzaki. *Quantum nondemolition measurement of spin via the paramagnetic Faraday rotation*. Phys. Rev. A, 60:4974–4979, 1999.
- [77] J. F. Sherson and K. Mølmer. *Polarization Squeezing by Optical Faraday Rotation*. Phys. Rev. Lett., 97:143602, 2006.
- [78] S. R. de Echaniz, M. W. Mitchell, M. Kubasik, M. Koschorreck, H. Crepaz, J. Eschner, and E. S. Polzik. *Conditions for spin squeezing in a cold 87 Rb ensemble*. J. Opt. B, 7(12):S548, 2005.
- [79] M. Takeuchi, S. Ichihara, T. Takano, M. Kumakura, T. Yabuzaki, and Y. Takahashi. *Spin Squeezing via One-Axis Twisting with Coherent Light*. Phys. Rev. Lett., 94:023003, 2005.
- [80] B. Julsgaard, A. Kozhekin, and E. S. Polzik. *Experimental long-lived entanglement of two macroscopic objects*. Nature, 413(6854):400–403, 2001.

- [81] L. Brillouin. *A Theorem of Larmor and Its Importance for Electrons in Magnetic Fields*. Phys. Rev., 67(7-8):260–266, 1945.
- [82] H. G. Dehmelt. *Modulation of a Light Beam by Precessing Absorbing Atoms*. Phys. Rev., 105:1924–1925, 1957.
- [83] T. Isayama, Y. Takahashi, N. Tanaka, K. Toyoda, K. Ishikawa, and T. Yabuzaki. *Observation of Larmor spin precession of laser-cooled Rb atoms via paramagnetic Faraday rotation*. Phys. Rev. A, 59(6):4836–4839, 1999.
- [84] G. A. Smith, S. Chaudhury, and P. S. Jessen. *Faraday spectroscopy in an optical lattice: a continuous probe of atom dynamics*. J. Opt. B, 5(4):323, 2003.
- [85] M. L. Terraciano, M. Bashkansky, and F. K. Fatemi. *Faraday spectroscopy of atoms confined in a dark optical trap*. Phys. Rev. A, 77(6):063417, 2008.
- [86] M. Koschorreck, M. Napolitano, B. Dubost, and M. W. Mitchell. *Sub-Projection-Noise Sensitivity in Broadband Atomic Magnetometry*. Phys. Rev. Lett., 104(9):093602, 2010.
- [87] F. Wolfgramm, A. Cerè, F. A. Beduini, A. Predojević, M. Koschorreck, and M. W. Mitchell. *Squeezed-Light Optical Magnetometry*. Phys. Rev. Lett., 105(5):053601, 2010.
- [88] D. Budker and M. Romalis. *Optical magnetometry*. Nature Phys., 3(4):227–234, 2007.
- [89] C. C. Bradley, C. A. Sackett, and R. G. Hulet. *Bose-Einstein Condensation of Lithium: Observation of Limited Condensate Number*. Phys. Rev. Lett., 78:985–989, 1997.
- [90] M. Gajdacz, P. L. Pedersen, T. Mørch, A. J. Hilliard, J. Arlt, and J. F. Sherson. *Non-destructive Faraday imaging of dynamically controlled ultracold atoms*. Rev. Sci. Instrum., 84(8):083105, 2013.
- [91] F. Kaminski, N. S. Kampel, M. P. Steenstrup, A. Griesmaier, E. S. Polzik, and J. H. Müller. *In-situ dual-port polarization contrast imaging of Faraday rotation in a high optical depth ultracold 87Rb atomic ensemble*. Eur. Phys. J. D, 66(9):1–8, 2012.
- [92] *Nobel Prize for Physics, 1952*. Nature, 170:911–912, 1952.
- [93] J. Alger. *The 1991 Nobel Prize in chemistry awarded to an MRI investigator*. J. Comp. Assist. Tomogr., 16(1):1–2, 1992.
- [94] Nobel Media AB. *The Nobel Prize in Chemistry 1991, 1991*. Available at http://www.nobelprize.org/nobel_prizes/chemistry/laureates/1991.

- [95] R. Damadian. *Tumor Detection by Nuclear Magnetic Resonance*. *Science*, 171(3976):1151–1153, 1971.
- [96] I. D. Weisman, L. H. Bennett, L. R. Maxwell, M. W. Woods, and D. Burk. *Recognition of Cancer in vivo by Nuclear Magnetic Resonance*. *Science*, 178(4067):1288–1290, 1972.
- [97] P. C. Lauterbur. *Image formation by induced local interactions: examples employing nuclear magnetic resonance*. *Nature*, 242(5394):190–191, 1973.
- [98] P. Mansfield. *Multi-planar image formation using NMR spin echoes*. *J. Phys. C*, 10(3):L55, 1977.
- [99] E. L. Hahn. *Spin Echoes*. *Phys. Rev.*, 80(4):580–594, 1950.
- [100] H. Y. Carr and E. M. Purcell. *Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments*. *Phys. Rev.*, 94(3):630–638, 1954.
- [101] P. Mansfield. *Nobel Lecture: Snap-Shot MRI*, 2003.
- [102] Nobel Media AB. *The Nobel Prize in Physiology or Medicine 2003*, 2003. Available at http://www.nobelprize.org/nobel_prizes/medicine/laureates/2003.
- [103] J. Skalla, G. Wäckerle, M. Mehring, and A. Pines. *Optical magnetic resonance imaging of Rb vapor in low magnetic fields*. *Phys. Lett. A*, 226(1-2):69–74, 1997.
- [104] J. Skalla, G. Wäckerle, and M. Mehring. *Optical magnetic resonance imaging of atomic diffusion and laser beam spatial profiles*. *Opt. Comm.*, 143(4-6):209–213, 1997.
- [105] A. R. Young, S. Appelt, A. B.-A. Baranga, C. Erickson, and W. Happer. *Three-dimensional imaging of spin polarization of alkali-metal vapor in optical pumping cells*. *Appl. Phys. Lett.*, 70(23):3081–3083, 1997.
- [106] K. Toyoda, Y. Takahashi, and T. Yabuzaki. *Magnetic resonance imaging of Bose-Einstein condensates*. *Appl. Phys. B*, 74:115–120, 2002.
- [107] 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*. *Phys. Rev. Lett.*, 102(12):125301, 2009.
- [108] M. Yasunaga and M. Tsubota. *Spin Echo in Spinor Dipolar Bose-Einstein Condensates*. *Phys. Rev. Lett.*, 101(22):220401, 2008.
- [109] 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*. *Phys. Rev. A*, 84:021605, 2011.
- [110] J. R. Ackerhalt and P. W. Milonni. *Interaction Hamiltonian of quantum optics*. *J. Opt. Soc. Am. B*, 1(1):116–120, 1984.
- [111] D. A. Steck. *Quantum and Atom Optics*, 2014. Revision 0.10.0, available at <http://atomoptics.uoregon.edu/~dsteck/teaching/quantum-optics/>.

- [112] F. Bloch and A. Siegert. *Magnetic Resonance for Nonrotating Fields*. Physical Review, 57(6):522–527, 1940.
- [113] W. Happer and B. S. Mathur. *Effective Operator Formalism in Optical Pumping*. Phys. Rev., 163:12–25, 1967.
- [114] A. R. Edmonds. *Angular momentum in quantum mechanics*. Princeton University Press, 1996.
- [115] R. Littlejohn. *Irreducible Tensor Operators and the Wigner-Eckart Theorem*, 2010. UC Berkeley Physics 221A lecture notes, available at <http://bohr.physics.berkeley.edu/classes/221/1011/notes/wigeck.pdf>.
- [116] E. Neil. *Spherical tensors*, 2014. JILA Physics 5250 lecture notes, available at http://www.colorado.edu/physics/phys5250/phys5250_fa14.
- [117] C. Eckart. *The Application of Group theory to the Quantum Dynamics of Monatomic Systems*. Rev. Mod. Phys., 2:305–380, 1930.
- [118] E. P. Wigner. *Group theory and its application to the quantum mechanics of atomic spectra*. Academic Press, 1959.
- [119] J. M. Geremia, J. K. Stockton, and H. Mabuchi. *Tensor polarizability and dispersive quantum measurement of multilevel atoms*. Phys. Rev. A, 73(4):042112, 2006.
- [120] I. H. Deutsch and P. S. Jessen. *Quantum control and measurement of atomic spins in polarization spectroscopy*. Opt. Comm., 283(5):681–694, 2010.
- [121] D. A. Steck. *Rubidium 87 D Line Data*, 2010. Revision 2.1.4, available at <http://steck.us/alkalidata/>.
- [122] L. J. LeBlanc and J. H. Thywissen. *Species-specific optical lattices*. Phys. Rev. A, 75:053612, 2007.
- [123] B. Arora, M. Safronova, and C. W. Clark. *Tune-out wavelengths of alkali-metal atoms and their applications*. Phys. Rev. A, 84(4):043401, 2011.
- [124] R. H. Leonard, A. J. Fallon, C. A. Sackett, and M. S. Safronova. *High precision measurement of the ^{87}Rb D-line tune-out wavelength*. arXiv:1507.07898 [physics.atom-ph], 2015.
- [125] J. P. Barrat and C. Cohen-Tannoudji. *Optical pumping and the density matrix formalism*. J. Phys. Le Rad., 22(6):329–336, 1961. Étude du pompage optique dans le formalisme de la matrice densité.
- [126] G. A. Smith, S. Chaudhury, A. Silberfarb, I. H. Deutsch, and P. S. Jessen. *Continuous Weak Measurement and Nonlinear Dynamics in a Cold Spin Ensemble*. Phys. Rev. Lett., 93(16):163602, 2004.
- [127] D. A. Steck. *Caesium D Line Data*, 2010. Revision 2.1.4, available at <http://steck.us/alkalidata/>.

- [128] P. C. D. Hobbs. *Ultrasensitive laser measurements without tears*. *Appl. Opt.*, 36(4):903–920, 1997.
- [129] M. Lewenstein and L. You. *Quantum Phase Diffusion of a Bose-Einstein Condensate*. *Phys. Rev. Lett.*, 77:3489–3493, 1996.
- [130] A. N. Garroway, P. K. Grannell, and P. Mansfield. *Image formation in NMR by a selective irradiative process*. *J. Phys. C*, 7(24):L457, 1974.
- [131] M. Andrews, C. Townsend, H.-J. Miesner, D. Durfee, D. Kurn, and W. Ketterle. *Observation of interference between two Bose condensates*. *Science*, 275(5300):637–641, 1997.
- [132] J. Abo-Shaeer, C. Raman, J. Vogels, and W. Ketterle. *Observation of vortex lattices in Bose-Einstein condensates*. *Science*, 292(5516):476–479, 2001.
- [133] F. T. Hioe. *Solution of Bloch equations involving amplitude and frequency modulations*. *Phys. Rev. A*, 30:2100–2103, 1984.
- [134] M. S. Silver, R. I. Joseph, and D. I. Hoult. *Selective spin inversion in nuclear magnetic resonance and coherent optics through an exact solution of the Bloch-Riccati equation*. *Phys. Rev. A*, 31:2753–2755, 1985.
- [135] L. M. Bennie, P. B. Wigley, S. S. Szigeti, M. Jasperse, J. J. Hope, L. D. Turner, and R. P. Anderson. *Precise wavefunction engineering with magnetic resonance*. arXiv:1412.6854 [quant-ph], revised 2015.
- [136] L. M. Bennie. *The fine art of wavefunction engineering: Precise engineering of the Bose-Einstein condensate wavefunction using magnetic resonance control*. PhD thesis, School of Physics and Astronomy, Monash University, (in prep.).
- [137] C. J. Foot. *Atomic physics*. Oxford University Press, New York, 2005.
- [138] A. A. Wood. *Spinor Bose-Einstein condensates in magnetic field gradients*. PhD thesis, School of Physics and Astronomy, Monash University, (in prep.).
- [139] Y.-J. Lin, A. Perry, R. Compton, I. Spielman, and J. Porto. *Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential*. *Phys. Rev. A*, 79:063631, 2009.
- [140] A. A. Wood. *A rubidium cold atom beam source*. Honours thesis, School of Physics, Monash University, 2010.
- [141] L. M. Bennie. *Design of a Single-Layer Zeeman Slower for ^{85}Rb and ^{87}Rb* . PHS3360 research project, School of Physics, Monash University, 2009.
- [142] S. Bell, M. Junker, M. Jasperse, L. Turner, Y.-J. Lin, I. Spielman, and R. Scholten. *A slow atom source using a collimated effusive oven and a single-layer variable pitch coil Zeeman slower*. *Rev. Sci. Instrum.*, 81(1):013105, 2010.
- [143] L. M. Bennie. *Design and Construction of a Zeeman Slower for ^{85}Rb and ^{87}Rb* . PHS3350 research project, School of Physics, Monash University, 2010.

- [144] L. M. Bennie. *A large atom number magneto-optical trap for BEC production*. Honours thesis, School of Physics, Monash University, 2010.
- [145] A. S. Arnold, J. S. Wilson, and M. G. Boshier. *A simple extended-cavity diode laser*. *Rev. Sci. Instrum.*, 69(3):1236–1239, 1998.
- [146] P. Starkey. *Construction and Locking of Two External Cavity Diode Lasers (ECDL)*. PHS3360 research project, School of Physics, Monash University, 2009.
- [147] V. Negnevitsky. *FPGA-based laser stabilisation using modulation transfer spectroscopy*. Honours thesis, Department of Electrical and Computer Systems Engineering, Monash University, 2010.
- [148] D. J. McCarron, S. A. King, and S. L. Cornish. *Modulation transfer spectroscopy in atomic rubidium*. *Meas. Sci. Tech.*, 19(10):105601, 2008.
- [149] N. F. Chilton. *Design and Construction of a Laser Amplifier*. PHS3350 research project, School of Physics, Monash University, 2010.
- [150] D. A. Smith and I. G. Hughes. *The role of hyperfine pumping in multilevel systems exhibiting saturated absorption*. *Am. J. Phys.*, 72(5):631–637, 2004.
- [151] P. Starkey. *An optical, experiment control and data acquisition system for a BEC lab*. Honours thesis, School of Physics, Monash University, 2010.
- [152] V. Negnevitsky. *Modulation transfer spectroscopy for fast, accurate laser stabilisation*. Honours thesis, School of Physics, Monash University, 2009.
- [153] J. Appel, A. MacRae, and A. I. Lvovsky. *A versatile digital GHz phase lock for external cavity diode lasers*. *Meas. Sci. Tech.*, 20(5):055302, 2009.
- [154] Hamamatsu Photonics. *Ultrafast MSM Photodetectors G4176 Series*, 2013. Available at <https://www.hamamatsu.com/resources/pdf/lsr/G4176E.pdf>.
- [155] Analog Devices. *ADF4108: PLL Frequency Synthesizer*, 2007. Available at http://www.analog.com/static/imported-files/data_sheets/ADF4108.pdf.
- [156] Analog Devices. *Ask the Applications Engineer-30: PLL Synthesizers*, 2002. Available at <http://www.analog.com/media/en/technical-documentation/application-notes/AN-30.pdf>.
- [157] K. L. Corwin, Z.-T. Lu, C. F. Hand, R. J. Epstein, and C. E. Wieman. *Frequency-stabilized diode laser with the Zeeman shift in an atomic vapor*. *Appl. Opt.*, 37(15):3295–3298, 1998.
- [158] P. Starkey, C. Billington, S. Johnstone, M. Jasperse, K. Helmersen, L. Turner, and R. Anderson. *A scripted control system for autonomous hardware-timed experiments*. *Rev. Sci. Instrum.*, 84(8):085111, 2013.
- [159] The HDF Group. *Hierarchical Data Format, version 5*, 1997–2015. <http://www.hdfgroup.org/HDF5/>.
- [160] M. Jasperse. *h5labview: HDF5 bindings for LabVIEW*, 2012–2015. Available at <http://h5labview.sf.net/>.

- [161] M. Jasperse. *PyByWire*, 2012–2015. Available at <http://pybywire.sf.net>.
- [162] iMatix Corporation. *ZeroMQ: Code connected*, 2007–2014. Available at <http://zeromq.org/>.
- [163] S. Furuhashi. *MessagePack*, 2008–2013. Available at <http://msgpack.org/>.
- [164] M. Jasperse. *labview-zmq: ZeroMQ bindings for LabVIEW*, 2012–2015. Available at <http://labview-zmq.sf.net/>.
- [165] M. Jasperse. *MessagePack bindings for LabVIEW*, 2012–2015. Available at <http://msgpack-labview.sf.net/>.
- [166] C. Hawthorn, K. Weber, and R. Scholten. *Littrow configuration tunable external cavity diode laser with fixed direction output beam*. *Rev. Sci. Instrum.*, 72(12):4477–4479, 2001.
- [167] MOG Laboratories. *External Cavity Diode Laser Controller, Rev 6*. Available at <http://www.moglabs.com>.
- [168] L. Maguire, S. Szilagyi, and R. Scholten. *High performance laser shutter using a hard disk drive voice-coil actuator*. *Rev. Sci. Instrum.*, 75(9):3077–3079, 2004.
- [169] L. M. Bennie, P. T. Starkey, M. Jasperse, C. J. Billington, R. P. Anderson, and L. D. Turner. *A versatile high resolution objective for imaging quantum gases*. *Opt. Expr.*, 21(7):9011–9016, 2013.
- [170] B. E. A. Saleh and M. C. Teich. *Fundamentals of photonics*. John Wiley & Sons, New York, 1991.
- [171] Eagleyard Photonics. *Data sheet for EYP-RWE-0810-03010-1300*, 2012. Available at <http://www.eagleyard.com/en/products/>.
- [172] Semrock Corporation. *808 nm MaxLine laser clean-up filter*. Available at <http://www.semrock.com/FilterDetails.aspx?id=LL01-808-12.5>.
- [173] P. Pakkiam. *Towards Magnetic Resonance Imaging of Bose-Einstein Condensates*. Honours thesis, School of Physics, Monash University, 2014.
- [174] Hamamatsu Photonics. *High-speed photodiodes (S5973 series)*, 2012. Available at http://www.hamamatsu.com/resources/pdf/ssd/s5971_etc_kpin1025e.pdf.
- [175] P. C. D. Hobbs. *Building Electro-Optical Systems: Making It All Work*. John Wiley & Sons, Inc., 2000.
- [176] Analog Devices. *MT-047: Op Amp Noise*, 2008. Available at <http://www.analog.com/media/en/training-seminars/tutorials/MT-047.pdf>.
- [177] Analog Devices. *OP470: Very Low Noise Quad Operational Amplifier*, 2002. Available at <http://www.analog.com/media/en/technical-documentation/data-sheets/OP470.pdf>.
- [178] C. M. Caves. *Quantum-mechanical noise in an interferometer*. *Phys. Rev. D*, 23(8):1693–1708, 1981.

- [179] C. M. Caves. *Quantum-mechanical radiation-pressure fluctuations in an interferometer*. Phys. Rev. Lett., 45(2):75–79, 1980.
- [180] R. Loudon. *The Quantum Theory of Light*. Oxford University Press, New York, third edition, 2001.
- [181] M. Jasperse. *High resolution magnetic resonance imaging of Bose-Einstein condensates*. Mid-candidature review, School of Physics, Monash University, 2012.
- [182] W. Kester. *Undersampling applications*. Available at <http://www.analog.com/media/en/training-seminars/design-handbooks/3689418379346Section5.pdf>.
- [183] W. Kester. *Fundamentals of Sampled Data Systems: Sampling theory*. Available at <http://www.analog.com/library/analogDialogue/archives/39-06/Chapter2SampledDataSystemsF.pdf>.
- [184] National Instruments Corporation. *Break Nyquist Criterion Using Undersampling Technique*, 2007. Available at <http://www.ni.com/newsletter/50078/en/>.
- [185] National Instruments. *Device Specifications: NI 6363*, 2015. Available at <http://www.ni.com/pdf/manuals/374651a.pdf>.
- [186] A. H. Nuttall. *Some windows with very good sidelobe behavior*. IEEE Trans. Acoust., Speech & Sig. Proc., 29(1):84–91, 1981.
- [187] D. Gabor. *Theory of communication. Part 1: The analysis of information*. J. Inst. Elec. Eng. III, 93(26):429–441, 1946.
- [188] 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*. arXiv:1408.0944 [quant-ph], revised 2015.
- [189] J. Arrillaga and N. Watson. *Power system harmonics*. John Wiley & Sons, 2003.
- [190] N. Behbood, F. M. Ciurana, G. Colangelo, M. Napolitano, M. W. Mitchell, and R. J. Sewell. *Real-time vector field tracking with a cold-atom magnetometer*. Appl. Phys. Lett., 102(17):173504, 2013.
- [191] W. Humphreys. *Match Filtering BEC Signals*. PHS3350 research project, School of Physics and Astronomy, Monash University, 2015.
- [192] Y. Eto, H. Ikeda, H. Suzuki, S. Hasegawa, Y. Tomiyama, S. Sekine, M. Sadgrove, and T. Hirano. *Spin-echo-based magnetometry with spinor Bose-Einstein condensates*. Phys. Rev. A, 88:031602, 2013.
- [193] D. K. Skjellerup-Wakefield. *AC Magnetic Field Noise Cancellation for Cold-Atom Experiments*. PHS3350 research project, School of Physics, Monash University, 2013.
- [194] G. Breit and I. I. Rabi. *Measurement of Nuclear Spin*. Phys. Rev., 38:2082, 1931.
- [195] W. Zhang, D. L. Zhou, M.-S. Chang, M. S. Chapman, and L. You. *Coherent spin mixing dynamics in a spin-1 atomic condensate*. Phys. Rev. A, 72:013602, 2005.

- [196] F. Gerbier, A. Widera, S. Fölling, O. Mandel, and I. Bloch. *Resonant control of spin dynamics in ultracold quantum gases by microwave dressing*. Phys. Rev. A, 73:041602, 2006.
- [197] National Instruments Corporation. *QuickSyn Microwave Frequency Synthesizers*, 2014. Available at http://ni-microwavecomponents.com/datasheets/DS_FSW-0010_0020.pdf.
- [198] M. J. Golay. *Field homogenizing coils for nuclear spin resonance instrumentation*. Rev. Sci. Instrum., 29(4):313–315, 1958.
- [199] D. I. Hoult and R. E. Richards. *Critical Factors in the Design of Sensitive High Resolution Nuclear Magnetic Resonance Spectrometers*. Proc. Royal Soc. London A, 344(1638):311–340, 1975.
- [200] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI Pulse Sequences*. Elsevier, 2004.
- [201] B. J. Sawyer, A. B. Deb, T. McKellar, and N. Kjærgaard. *Reducing number fluctuations of ultracold atomic gases via dispersive interrogation*. Phys. Rev. A, 86:065401, 2012.
- [202] K. Góral, M. Gajda, and K. Rzażewski. *Thermodynamics of an interacting trapped Bose-Einstein gas in the classical field approximation*. Phys. Rev. A, 66:051602, 2002.
- [203] P. B. Blakie and M. J. Davis. *Projected Gross-Pitaevskii equation for harmonically confined Bose gases at finite temperature*. Phys. Rev. A, 72:063608, 2005.
- [204] R. Sewell, M. Koschorreck, M. Napolitano, B. Dubost, N. Behbood, and M. Mitchell. *Magnetic sensitivity beyond the projection noise limit by spin squeezing*. Phys. Rev. Lett., 109(25):253605, 2012.
- [205] R. Sewell, M. Napolitano, N. Behbood, G. Colangelo, and M. Mitchell. *Certified quantum non-demolition measurement of a macroscopic material system*. Nature Phot., 7(7):517–520, 2013.
- [206] M. Jasperse. *GalilLogger*, 2011. Available at <http://becwiki.physics.monash.edu.au/GalilLogger>.
- [207] J. Söding, D. Guéry-Odelin, P. Desbiolles, F. Chevy, H. Inamori, and J. Dalibard. *Three-body decay of a rubidium Bose-Einstein condensate*. Appl. Phys. B, 69(4):257–261, 1999.
- [208] E. A. Burt, R. W. Ghrist, C. J. Myatt, M. J. Holland, E. A. Cornell, and C. E. Wieman. *Coherence, Correlations, and Collisions: What One Learns about Bose-Einstein Condensates from Their Decay*. Phys. Rev. Lett., 79:337–340, 1997.
- [209] P. S. Julienne, F. H. Mies, E. Tiesinga, and C. J. Williams. *Collisional Stability of Double Bose Condensates*. Phys. Rev. Lett., 78:1880–1883, 1997.
- [210] D. Döring. *Interferometry and precision measurements with Bose-condensed atoms*. PhD thesis, Australian National University, 2011.
- [211] M. J. Martin. *Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock*. PhD thesis, University of Colorado, 2013.

- [212] B. H. Bransden and C. J. Joachain. *Physics of Atoms and Molecules*. Prentice-Hall, Harlow, 2nd edition, 2003.
- [213] J. K. Stockton. *Continuous Quantum Measurement of Cold Alkali-Atom Spins*. PhD thesis, California Institute of Technology, 2007.
- [214] R. A. Cline, J. D. Miller, M. R. Matthews, and D. J. Heinzen. *Spin relaxation of optically trapped atoms by light scattering*. *Opt. Lett.*, 19(3):207–209, 1994.
- [215] T. Bergeman, C. Tsai, D. Heinzen, and C. Amiot. Private communication, 2015.



Oven controller implementation

This appendix contains details of the implementation of the oven controller and interlock described in §4.3. The controller processes analog inputs from multiple gauges, and makes decisions based on its state machine. It produces analog outputs that regulate the oven temperature, and digital outputs that control the atom beam. Failures disable the relays and the oven, ensuring the system ‘fails safe’.

The state machine is updated every 150 ms, and the sensors are read every iteration. An ‘error counter’ is used to debounce error situations, as spurious noise or rf-pickup can cause sensors to momentarily return an invalid reading. Once the error state has been entered into, manual intervention is required to return to normal operation, which can be done using a physical switch on the unit, or remotely using TELNET if measurement logs indicate the situation is safe.

A.1: Inputs and outputs

The analog inputs/outputs are 0–5 V, with inputs taken as single-ended with common ground. The digital inputs are optocoupled with HI voltages between 5 and 24V. The voltage of the digital output lines are externally set to +24 V using the OP0A/OP1A pins.

Line	Purpose
DO0	Enable Peltiers
DO4	Open gate valve
DO13	Alarm buzzer
DO14	Backlight for ‘start’ button
DO15	Backlight for ‘stop’ button
DI0	‘Peltiers override’ switch
DI3	Peltiers flow meter
DI4	‘Start’ button
DI5	‘Stop’ button
DI6	‘Manual override’ switch

Line	Purpose
AO0	Bellows heater (PID)
AO1	Collimation tube heater (PID)
AI0	Bellows thermistor
AI1	Collimation tube thermistor
AI2	Cold-cup thermistor (<i>in vacuo</i>)
AI3	Water supply thermistor
AI4	UHV pressure (ion gauge)
AI5	LV can pressure (Pirani gauge)
AI6	Peltier thermocouple (<i>in vacuo</i>)

The oven is heated by two 1m lengths of heater tape (insulated nichrome wire), which are wrapped around the bellows and collimation tube. They are fed by MANSON 3402 HCS power supplies (max 20 A/30 V) operated in externally controlled constant-voltage mode (0 – 5 V FSD). Constant-voltage mode had improved dynamic range over constant-current mode, and the PID control accounts for any change of heater tape resistance with temperature. Details on the PID loop and the tuning of its gain constants are presented in [140].

A HTC GVB-SS-CF35-P pneumatic gate valve between the cold-cup and Zeeman slower isolates the oven from the rest of the vacuum system. It is actively forced open by compressed-air regulated by a spring-return NUMATICS L01SA459G00040 solenoid, which ensures the gate valve is closed in the event of power loss.

The Peltiers are formed by two MELCOR MS2-11 TECs run at 35 V/6 A to achieve a steady-state cold-side temperature of -32°C . The Peltier's power supply is activated by a relay switch on its 'remote control' port, which is disabled unless permitted by the PLC. Active water-cooling of the 'hot' side is required to prevent runaway heating of the unit, which is provided by a closed-loop water chiller (JULABO HE).

Flow is measured by the onboard pulse counter using a GEMS SENSORS FT-210 flow sensor which produces 22,000 pulses/L. Note that the GALIL pulse counter input is only guaranteed to 300 Hz max (max flow 0.8 L/min), although some units went higher. A divide-by-4 counter was implemented (using a 74C74) to increase the max flow rate to 2.4 L/min. Later interlock designs used an RS 257-149 flow sensor (4,600 pulses/L) and an LM2907N-8 to perform frequency-voltage conversion.

The type-E thermocouples were read using OCEAN CONTROLS TCAMPV2 thermocouple amplifiers (based on the AD594). An RC-filter (time constant 1 s) was added to the output to smooth out high-frequency noise on the output voltage (reduced from 20 mV_{RMS} to 2 mV_{RMS}). The Peltiers thermocouple amplifier is configured for negative polarity to provide a positive voltage for temperatures below 0°C .

Thermistors are NTC type EC95 (F-material) with nominal resistance 3 k Ω at 25°C . The non-linear relationship between resistance and temperature over the operating range 25 – 235°C was phenomenologically modelled as

$$R(T) = \frac{1}{aT^4 + bT^2 + c} \quad \Rightarrow \quad T(R) = \sqrt{\frac{-b}{2a} + \frac{1}{2}\sqrt{\left(\frac{b}{a}\right)^2 - \frac{4}{a}\left(c - \frac{1}{R}\right)}} \\ = 53.85\sqrt{\sqrt{29500/R + 1} - 1.376} \quad (\text{calibrated}),$$

for T in $^{\circ}\text{C}$ and R in Ohms. The thermistor resistances were measured using voltage dividers, taking the analog input across a known resistance R_0 ,

$$R(T) = \left(\frac{V_0}{V} - 1\right) R_0,$$

where $V_0 = 5\text{ V}$ is the reference voltage, and R_0 is chosen for each thermistor to match the thermistor resistance at the expected operating temperature.

The analog outputs of the SRS IGC100 are $0 - 12\text{ V}$, requiring voltage dividers to be read by the $+5\text{ V}_{\text{max}}$ analog inputs. The voltage V is logarithmic in pressure, with

$$P(\text{Torr}) = \begin{cases} 10^{V-12} & \text{for an ion gauge} \\ 10^{V-5} & \text{for a Pirani gauge.} \end{cases} \quad (\text{A.1})$$

The GALIL unit has no exponential function, but the following approximation can be used to compute the pressure in Torr,

$$10^x \approx \tan[((ax + b)x + c)x + d] \quad (\text{A.2})$$

where $a = 11.0$, $b = -49.7$, $c = 78.9$, $d = 44.0$, and the argument of $\tan[x]$ is specified in degrees.

The current state of the machine is displayed on a CRYSTALFONTZ CFA634 LCD display connected to the serial port of the GALIL (handle 'P1'). The unit maintains a TCP connection to a logging server (handle 'EC'), and every 55 s the latest measurements are sent to it using custom protocol [206]. A message is also sent whenever the state machine changes state. TCP logging was chosen to ensure the messages are resent if a collision occurs, and that both the unit and server know when the connection is severed and intervention is required.

A.2: Code listing

```

1 | REM =====[ GALIL OVEN INTERLOCK/CONTROLLER ]=====
2 | #AUTO; 'resume from here on power-cycle
3 | ZC1; 'reset state variable
4 | IQ65535; 'ensure logic is the correct way around
5 | CB13; 'buzzers off
6 | CBO; 'peltiers off at start
7 | A00,0; A01,0; 'heaters off
8 | SM255,255,255,128; 'permit only lab subnet
9 | JS#RECONN; 'enable logging
10 | TMSG=TIME
11 |
12 | MG{P1}{^9},{^6}{N}; 'boot screen
13 | WT1000
14 | MG{P1}{^4},{^22},{^27}{N}; 'config screen
15 |
16 | DMT[6]; 'allocate temperature array
17 |
18 | REM ***** PID LOOP SETTINGS *****
19 | ' loop A = oven, loop B = collim tube
20 | AZ-1,-1; 'ensure disabled
21 | CL500; 'set update interval to 500ms
22 | AF0,1; 'analog input chans
23 | KP350,400; KI 2.0,1.5; KD5,5; 'P,I,D gains
24 | DB0,0; 'no control deadband
25 | A00,0; A01,0; 'zero control signal
26 |
27 | TPTR=-20.0; 'peltier temperature (negative)
28 | PTRCONV=0; 'peltiers have converged
29 | TOVN1=80; TOVN2=120; 'oven temperatures
30 | PS2.25,3.25; 'oven set-point voltages
31 | NCONV=10; 'number of loops before convergence
32 |
33 | REM ***** BEGIN CODE *****
34 | FAIL=0; 'nothing is wrong yet
35 | SHOWTPTR=0; SHOWTOVN=0; 'display settings
36 | COOLFAIL=0; '# loops with coolant failure
37 |
38 | JS#RUNCHK; 'we all good to start?
39 |
40 | REM ***** INITIALISATION STATE *****
41 | #WAITING
42 | ZC1
43 | SHOWTPTR=0; SHOWTOVN=0
44 |
45 | MG"NOTICE,standby"
46 | SB14;CB15;'buttons
47 |
48 | MG{P1}{^4},{^12}
49 | #WAITLP
50 | JS#RUNCHK
51 | MG{P1}{^1},"STANDBY"
52 | MG{P1}"<start> peltiers"

```

```

53 JP#COOLPTR,NEXT>0
54 WT150
55 JP#WAITLP
56
57
58 REM ***** ACTIVATE PELTIERS *****
59 #COOLPTR; 'COOL PELTIERS
60 CB13;'turn off the siren
61 SB0;'start peltiers
62 SHOWTPTR=1;'now track temperature
63
64 ZC2
65 MG"NOTICE,cool peltiers"
66 CB14;SB15;'buttons
67
68 PTRCONV=0;'no convergence
69 NSET=0
70 #CLPTRLP
71 JS#RUNCHK
72 MG{P1}{^1},"Peltiers cooling"
73 JP#WARMPTR,PREV>0
74 IFT[4]<TPTR
75 NSET=NSET+1
76 ELSE
77 NSET=0
78 ENDIF
79 WT150
80 JP#CLPTRLP,NSET<NCONV
81 PTRCONV=1
82
83 MG"NOTICE,status=cold",'
84 "Tptra=",T[4]{F2.3}
85 JP#COLD
86
87 REM ***** PELTIERS COLD *****
88 #COLD
89 SHOWTOVN=0
90 ZC4
91 MG"NOTICE,waiting"
92 SB14;SB15;'buttons
93
94 #COLDLP
95 JS#RUNCHK
96 MG{P1}{^1},"Peltiers cold"
97 MG{P1}"<start> oven"
98 JP#HEATOVN,NEXT>0
99 JP#WARMPTR,PREV>0
100 WT150
101 JP#COLDLP
102
103
104 REM ***** ACTIVATE OVEN *****
105 #HEATOVN
106 AZ0,1;'activate control loop outputs
107 ILO.3,0.1;'set integrator limits
108 SHOWTOVN=1
109
110 ZC5
111 MG"NOTICE,heating"
112 CB14;SB15;'buttons
113
114 IF(TOVN1<0);'make sure it's positive
115 TOVN1=-TOVN1
116 TOVN2=-TOVN2
117 ENDIF
118 'wait for convergence
119 NSET=0
120 #HTOVNLP
121 JS#RUNCHK
122 MG{P1}{^1},"Heating oven"
123 JP#COOLOVN,PREV>0
124 IF(@ABS[T[0]-TOVN1]<5)&'
125 (@ABS[T[1]-TOVN2]<5)
126 NSET=NSET+1
127 ELSE
128 NSET=0
129 ENDIF
130 WT150
131 JP#HTOVNLP,NSET<NCONV
132 TOVN1=-TOVN1; TOVN2=-TOVN2;'convergence=negative
133
134 MG"INFO,status=hot",'
135 "Tovn1=",T[0]{F2.3}
136 JP#READY
137
138 REM ***** SYSTEM READY/WAITING *****
139 #READY
140 ZC7
141 MG"NOTICE,ready"
142 SB14;SB15;'buttons
143 #READYLP
144 JS#RUNCHK
145 MG{P1}{^1},"System ready"
146 MG{P1}"<start> open g.v."
147 JP#ACTIVE,NEXT>0
148 JP#COOLOVN,PREV>0
149 WT150
150 JP#READYLP
151
152 REM ***** ACTIVE BEAM *****
153 #ACTIVE;'ANY JUMP AWAY MUST CALL #KILLGV
154 MG"NOTICE,active"
155 ZC8
156 CB14;SB15;'buttons
157
158 JS#RUNCHK;'paranoia
159
160 SB4; ' *** OPEN GATE VALVE ***
161 MG"NOTICE, gate=open"
162
163 #ACTLP
164 JS#RUNCHK
165 MG{P1}{^1},"ACTIVE BEAM"
166 MG{P1}"<stop> close g.v."
167 IF PREV>0
168 JS#KILLGV
169 JP#READY
170 ENDIF
171 WT150
172 JP#ACTLP
173
174
175 REM ***** DEACTIVATE OVEN *****
176 #COOLOVN
177 AZ-1,-1;'kill PID loops
178 A00,0;A01,0;'slam outputs
179 'reset trackers
180 IF(TOVN1<0)
181 TOVN1=-TOVN1;TOVN2=-TOVN2
182 ENDIF
183
184 JP#WARMPTR,FAIL>0; 'jump on
185
186 ZC6
187 MG"NOTICE,cooling oven"
188 SB14;SB15;'buttons
189
190 NSET=0

```

```

191 #CLOVNLP
192 JS#RUNCHK
193 JP#HEATOVN,NEXT>0
194 JP#COLD,PREV>0
195 MG{P1}{^1},"Cooling oven"
196 IF (T[0]<50)&(T[1]<50)
197     NSET=NSET+1
198 ELSE
199     NSET=0
200 ENDIF
201 JP#COLD,NSET>NCONV
202 WT150
203 JP#CLOVNLP
204
205 REM ***** DEACTIVATE PELTIERS *****
206 #WARMPTR
207 CBO;'disable peltier
208 PTRCONV=0;'not converged
209
210 ZC3
211 MG"NOTICE, stopping peltiers"
212 SB14;SB15;'buttons
213
214 NSET=0
215 #WMOVNLP
216 JS#RUNCHK
217 JP#COOLPTR,NEXT>0
218 JP#WAITING,PREV>0
219 MG{P1}{^1},"Stopping peltiers"
220 IF T[0]<2
221     NSET=NSET+1
222 ELSE
223     NSET=0
224 ENDIF
225 JP#WAITING,NSET>NCONV
226 WT150
227 JP#WMOVNLP
228
229
230 REM ***** COMPUTE MEASUREMENTS *****
231 #COMPUTE
232 ' check vacuum pressures (measured constants)
233 P1=@AN[4]*2.4481-12; 'log10(P) - ION gauge
234 P2=@AN[5]*2.4423-5; 'log10(P) - PIRANI gauge
235 ' check water flow
236 PC-1; 'activate pulse counter
237 WT100
238 'flow rate in L/min (110/3=36.666)
239 FLOW=_PC*6/55.0;
240 PC0; 'deactivate counter
241 ZD 1000*FLOW
242 ' compute thermistor readings
243 ' resistances chosen for max sensitivity
244 R=(5.0/(@AN[0]+0.01)-1)*2095.0;'
245 T[0]=53.85*@SQR[@SQR[29500/R+1]-1.376]
246 R=(5.0/(@AN[1]+0.01)-1)*1397.0;'
247 T[1]=53.85*@SQR[@SQR[29500/R+1]-1.376]
248 R=(5.0/(@AN[2]+0.01)-1)*3300000;'
249 T[2]=-13.25*(@AN[2])-24.38
250 R=(5.0/(@AN[3]+0.01)-1)*20100.0;'
251 T[3]=53.85*@SQR[@SQR[29500/R+1]-1.376]
252
253 ' compute thermocouple readings
254 T[4]=@AN[6]*-124.28+52.545
255 T[5]=@AN[7]*94.6-3.8
256 MG{P1}{^12},{^10}
257 IF (SHOWTOVN)
258     MG{P1}"TOVEN=",T[0]{Z3.1},"", '
259     T[1]{Z3.1}
260 ELSE
261     MG{P1}{^10}{N}
262 ENDIF
263 IF (SHOWTPTR)
264     MG{P1}"COLD CUP=",T[4]{Z2.1}{N}
265 ENDIF
266
267 ' external input
268 IF (FAIL=0)
269     IF (@IN[4]=0)|(@IN[5]=0)
270         NEXT=@IN[4]*(2*(NEXT==0)-1);
271         '>1 on first loop, <=0 after
272         PREV=@IN[5]*(2*(PREV==0)-1);
273         '>1 on first loop, <=0 after
274     ENDIF
275 ELSE
276     NEXT=0
277     PREV=0
278 ENDIF
279 EN
280
281 REM ***** STATUS UPDATE *****
282 #STATUS
283 JS#COMPUTE; 'compute things
284
285 'check we're connected to logger
286 IF(_IHC2=0)&(_XQ1=-1)
287     XQ#RECONN,1
288 ENDIF
289 'periodically send message
290 IF (TIME>=TMSG)
291     CW2;CFC;'reset output controls to be safe
292     MG"INFO,state=",_ZC{Z2.0}, '
293     ",Tovn1=",T[0]{Z3.1},"",Tovn2="{N}
294     MG T[1]{Z3.1},"",Tptr=",T[4]{Z3.1}, '
295     ",flow=",FLOW{Z3.1}
296     TMSG=TIME+55000;'slightly less than 60s
297 ENDIF
298
299 EN;#STATUS
300
301
302 REM ***** FAILSAFE CHECK *****
303 #RUNCHK; 'Run through checks
304 JS#STATUS
305
306 'are we TRANSITIONING TO manual override?
307 IF(_ZC>-3)&((@IN[0]>0)|(@IN[6]=0))
308     JP#MANUAL
309 ENDIF
310 MG{P1}{^1}
311
312 REM ***** RUN CRITICAL CHECKS FIRST *****
313 'check thermocouples
314 REM digital check is susceptible to transients,
315 REM -- check analog voltages instead
316 IF(@AN[6]<0.01)|(@AN[7]<0.01)
317     MG{P1}"Thermocouple error"
318     IF (FAIL=0)
319         MG"CRIT,TC=",@AN[6]{Z1.3}, '
320         @AN[7]{Z1.3}
321         JP#FAIL
322     ENDIF
323 EN
324 ENDIF
325 ' check water flow at least 1.5L/min
326 ' require failure for 5 loops before failsafe
327 IF (FLOW<1.5)
328     IF (COOLFAIL<5)

```

```

329     COOLFAIL=COOLFAIL+1
330     ENDIF
331 ELSE
332     IF(COOLFAIL>0)
333         COOLFAIL=COOLFAIL-0.5
334     ENDIF
335 ENDIF
336 ' error: failed to be cool
337 IF(COOLFAIL>=5)
338     MG{P1}"Flow rate error"
339     IF(FAIL=0)
340         MG"CRIT,flow,rate=",FLOW{Z1.2}
341         JP#FAIL
342     ENDIF
343 EN
344 ENDIF
345 ' check the peltiers
346 IF(T[4]>30)|((PTRCONV>0)&(T[4]>-12))
347     ' CRITICAL TEMPERATURE ERROR
348     MG{P1}"Critical peltier"
349     IF(FAIL=0)
350         MG"CRIT,peltier,T=",T[4]{F2.3}
351         JP#FAIL
352     ENDIF
353 EN
354 ENDIF
355 IF(PTRCONV>0)&(T[4]>(TPTR+2));
356     ' peltiers too hot, respond by cooling oven
357     MG{P1}"Peltiers too hot"
358     IF(FAIL=0)
359         MG"WARN,peltier,Tptr=",T[4]{F2.3}
360         JP#SOFTERR
361     ENDIF
362 EN
363 ENDIF
364
365 ' check the oven
366 IF((TOVN1<0)&((@ABS[T[0]+TOVN1]>10)|‘
367     (@ABS[T[1]+TOVN2]>10)))
368     MG{P1}"Oven control fail"
369     IF(FAIL=0)
370         MG"WARN,",‘
371         "heater,Tovn1=",T[0]{F2.3}{N}
372         MG",Tovn2=",T[1]{F2.3}
373         JP#SOFTERR
374     ENDIF
375 EN
376 ENDIF
377
378 'check thermistors
379 IF(@AN[0]<0.05)|(@AN[1]<0.05)|‘
380     (@AN[2]<0.02)|(@AN[3]<0.05)
381     MG{P1}"Thermistors discon."
382     IF(FAIL=0)
383         MG"WARN, thermistors"
384         JP#SOFTERR
385     ENDIF
386 EN
387 ENDIF
388
389 'check vacuum gauges
390 IF(@AN[4]<0.1)|(@AN[4]>11)|‘
391     (@AN[5]<0.1)|(@AN[5]>11)
392     MG{P1}"Vacuum gauge discon."
393     IF(FAIL=0)
394         MG"WARN, vacuum gauge"
395         JP#SOFTERR
396     ENDIF
397 EN
398
399
400 ' check vacuum pressure
401 IF(P1>-6)|(P2>2)
402 IF(FAIL=0)
403     MG{P1}"Vacuum press. fail"
404     x=@FRAC[P1]
405     x=@TAN[((11*x-50)*x+79)*x+44]
406     MG"WARN, iongauge=",x{F1.2},‘
407     "*10~",@INT[P1]{F2.0}{N}
408     x=@FRAC[P2]
409     x=@TAN[((11*x-50)*x+79)*x+44]
410     MG", pirani=",x{F1.2},"*10~",‘
411     @INT[P2]{F2.0}
412     MG"WARN, vacuum"
413     JP#SOFTERR
414 ENDIF
415 EN
416 ENDIF
417
418 FAIL=0; 'nothing was wrong
419 EN; 'end #RUNCHK
420
421
422 REM ***** FAILSAFE ACTIVATIONS *****
423 #KILLGV
424 IF@OUT[4]
425     MG"NOTICE, gate=shut"
426     CB4; 'slam gate valve
427 ENDIF
428 EN
429
430 #DISABLE
431 JS#KILLGV; 'shut the gate valve
432 IF(FAIL<>2); 'disable peltiers
433     CBO
434     PTRCONV=0
435 ENDIF
436 AZ-1,-1; 'shutdown oven
437 A00,0; A01,0
438 IF(TOVN1<0)
439     TOVN1=-TOVN1; TOVN2=-TOVN2
440 ENDIF
441 CB14;CB15; 'no buttons
442 SHOWTPTR=1;SHOWTOVN=1;'debug info
443 NSET=0
444 EN
445
446 REM ***** CRITICAL FAILURE *****
447 #FAIL
448 ZS0;'forget the stack
449 FAIL=1
450 JS#DISABLE
451 SB13; 'make noise
452 ZC-1
453 MG"EMERG,failsafe"
454 #FAILED
455 JS#RUNCHK
456 MG{P1}{~1},"FAILSAFE ACTIVATED"
457 WT250
458 JP#FAILED
459
460 REM ***** RECOVERABLE FAILURE *****
461 #SOFTERR
462 ZS0;'forget the stack
463 FAIL=2
464 JS#DISABLE
465 JS#BUZZ; 'make noise
466 ZC-2

```

```

467 #SOFTLP
468 JS#RUNCHK
469 MG{P1}{^1},"ERROR"
470 WT250
471 JP#SOFTLP,FAIL>0
472 JP#COLD,PTRCONV>0
473 JP#WAITING
474
475 #BUZZ
476 SB13
477 WT150
478 CB13
479 EN
480
481 REM ***** MANUAL OVERRIDE *****
482 #MANUAL
483 ZS0;'forget the stack
484 FAIL=-1
485 ZC-3
486 JS#DISABLE
487 MG"NOTICE, manual"
488 #OVERLP
489 JS#RUNCHK
490 IF @IN[0]>0
491 MG{P1}{^1},"* PELTIER OVERRIDE *"
492 ELSE
493 MG{P1}{^1},"* MANUAL OVERRIDE *"
494 ENDIF
495 OB13,@IN[6]&(@IN[0]>0);'siren
496 WT250
497 JP#OVERLP,(@IN[0]>0)|(@IN[6]=0)
498 CB13;'siren off
499 FAIL=0
500 JP#WAITING

```

```

501 EN
502
503 REM ***** ETHERNET RECONNECTION *****
504 #RECONN
505 'do not call me from the main thread
506 IHC=>-3
507 IHC=130,194,171,188 <519>2; 'syslog server (TCP)
508 WT250
509 JP#RECONN,_IHC2<>-2; 'keep trying
510 CW2;CFC
511 EN
512
513 REM ***** RUN-TIME ERRORS *****
514 #CMDERR
515 CB4;'CLOSE GATE VALVE
516 ZS0;'empty the stack
517 FAIL=1;'required for #DISABLE
518 JS#DISABLE;'system in unknown state- make safe
519 MG{P1}{^12},"CODE ERROR",_TC{Z3.0}
520 MG{P1}"Line",_ED{Z3.0}
521 CFC
522 MG"CRIT,error,line=",
523 _ED{Z3.0},"",descr="{N}
524 TC1;'outputs to CF
525 SB13;'activate noise maker
526 #ERRLP
527 JS#STATUS;'continue to send status messages
528 WT30000
529 JP#ERRLP
530 EN;'end the program (halt)
531
532 #TCPERR
533 '_IA4 contains broken handle
534 RE;'return from error state

```


B

Loss mechanisms

Following Refs. [207, 208], we model the density-dependent number loss in a BEC with spatially-varying number density $n(\mathbf{r})$ as

$$\frac{dn(\mathbf{r})}{dt} = -Ln(\mathbf{r})^3 - Gn(\mathbf{r})^2 - \gamma n(\mathbf{r}), \quad (\text{B.1})$$

where L , G and γ are the three-, two- and one-body loss coefficients respectively.

One-body losses arise from collisions of condensed atoms with the background gas (caused by an imperfect vacuum), as well as scattering from the off-resonant dipole trap and Faraday probe beams. The background collision rate is quantified by the vacuum lifetime τ_0 , which depends on the quality of the vacuum apparatus, and the off-resonant scattering rate γ_s for a particular beam depends on both the power and detuning (see §B.2). The total one-body scattering rate is then $\gamma = \gamma_s + \tau_0^{-1}$.

Potential sources of two-body loss are spin-exchange, dipolar relaxation, and photoassociation. However, there is no antitrapped state in an optical dipole trap so spin exchange does not result in atom number loss, and the dipole relaxation rate in an $F = 1$ BEC is very low [207, 209]. Potentially the high intensity of the Faraday probe beam combined with long interrogation times could induce losses through photoassociation. However, there are no known photoassociation lines near the probe wavelength (see §B.3) and the dependence of the atom number decay rate (§5.9) is consistent with a negligible photoassociation rate.

Finally, three-body losses occur by recombination, when two atoms collide and form a dimer. The liberated molecular binding energy is carried away as kinetic energy by a third atom, expelling it from the trap. Typically this is the dominant loss process in a BEC [207, 208].

The total number of atoms in the BEC and its associated rate of change is therefore

$$N \equiv \int n(\mathbf{r}) d^3\mathbf{r} \quad (\text{B.2})$$

$$\begin{aligned} \Rightarrow \frac{dN}{dt} &= \frac{d}{dt} \int n(\mathbf{r}) d^3\mathbf{r} = \int \frac{dn(\mathbf{r})}{dt} d^3\mathbf{r} \\ &= -L \int n(\mathbf{r})^3 d^3\mathbf{r} - G \int n(\mathbf{r})^2 d^3\mathbf{r} - \gamma \int n(\mathbf{r}) d^3\mathbf{r} \\ &= -LN \langle n(\mathbf{r})^2 \rangle - GN \langle n(\mathbf{r}) \rangle - (\gamma_s + \tau_0^{-1})N, \end{aligned} \quad (\text{B.3})$$

where the density-weighted average of a function $f(\mathbf{r})$ is defined as

$$\langle f(\mathbf{r}) \rangle \equiv \frac{1}{N} \int f(\mathbf{r}) n(\mathbf{r}) d^3\mathbf{r}. \quad (\text{B.4})$$

Note that the weighted averages in (B.3) indicate the form of the loss rate equation depends on the shape of the density profile. For a given profile shape, the integrals can be evaluated to obtain an expression in terms of the total atom number N , which affects the form of solutions to the differential equation (B.3). Thermal and condensed clouds therefore exhibit different loss rate curves for the same total atom number [208].

B.1: Thomas-Fermi profile

The peak density of a harmonically trapped BEC is closely described by the Thomas-Fermi model, with corrections near the edges arising from the nonzero kinetic energy of the condensate. The tails of the distribution where the approximation becomes poorer have low density and therefore do not contribute significantly to (B.4), reinforcing the approximation.

A BEC with Thomas-Fermi radii $(x_{\text{TF}}, y_{\text{TF}}, z_{\text{TF}})$ and chemical potential μ has a spatial profile given by [10]

$$n(\mathbf{r}) = \frac{\mu}{g} \max \left(1 - \frac{x^2}{x_{\text{TF}}^2} - \frac{y^2}{y_{\text{TF}}^2} - \frac{z^2}{z_{\text{TF}}^2}, 0 \right). \quad (\text{B.5})$$

Rescaling and transforming to spherical coordinates, the number averages become

$$\langle n(\mathbf{r})^k \rangle = \frac{4\pi}{N} \left(\frac{2\mu}{m\bar{\omega}^2} \right)^{3/2} \left(\frac{\mu}{g} \right)^{k+1} \int_0^1 r^2 (1-r^2)^{k+1} dr, \quad (\text{B.6})$$

where the Thomas-Fermi radii and trapping frequencies are related by $\mu = \frac{1}{2} m \omega_x^2 x_{\text{TF}}^2$, and $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean trapping frequency.

Normalising the integral by setting $\langle 1 \rangle \equiv 1$ for $k = 0$, the chemical potential follows

$$\frac{\mu}{g} = c_1 N^{2/5} \quad \text{where} \quad c_1 = \frac{15^{2/5}}{8\pi a^3} \left(\frac{a}{\bar{a}} \right)^{12/5}, \quad (\text{B.7})$$

where $a = mg/4\pi\hbar^2$ is the interatomic scattering length and $\bar{a} = \sqrt{\hbar/m\bar{\omega}}$ is the characteristic harmonic oscillator length.

Evaluating the averages in (B.3) for the density profile (B.5) gives

$$\langle n(\mathbf{r}) \rangle = \frac{4}{7} c_1 N^{2/5}, \quad \text{and} \quad \langle n(\mathbf{r})^2 \rangle = \frac{8}{21} c_1^2 N^{4/5}. \quad (\text{B.8})$$

Note that $\langle n(\mathbf{r})^2 \rangle = \frac{7}{6} \langle n(\mathbf{r}) \rangle^2$, in agreement with Ref. [207]. Hence the total number of atoms remaining in a Thomas-Fermi BEC follows the model

$$\frac{1}{N} \frac{dN}{dt} = -\frac{8c_1^2}{21} L N^{4/5} - \frac{4c_1}{7} G N^{2/5} - \gamma_s - \frac{1}{\tau_0}. \quad (\text{B.9})$$

This model is applied in §5.7–5.9 to model the decay in BEC atom number over time.

B.2: Photon scattering rate

In the far-detuned limit, the probability of direct excitation of an atom into an excited state through absorption of a photon becomes negligible. However, adiabatic elimination of the excited state (§2.1) predicts the virtual transitions between ground states.

Even though this interaction is ‘second order’, it still results in scattering of photons and reduces the lifetime of the BEC.¹ The scattering rate can be seen as a result of the Kramers-Kronig relation, which states that a real polarisability must have an accompanying imaginary component to satisfy continuity [111, §1.4.1.3, §14.1.4.2]. The imaginary component describes absorption, traditionally expressed in terms of a scattering cross-section relation called the Kramers-Heisenberg expression. Equivalently, Fermi’s (second) golden rule predicts the scattering rate as arising directly from the interaction Hamiltonian (2.4) [180, §8.7].

The brief discussion here follows Refs. [180, 210, 211], presented in notations and conventions consistent with the rest of this thesis. Consider an atom initially in state $|a\rangle$. Through the process of absorption of a photon (frequency ω and polarisation ϵ) and emission of a scattered photon (frequency ω_{sc} and polarisation ϵ_{sc}), the atom ends up in state $|b\rangle$. The scattering rate for emitting the photon into a solid angle element $d\Omega$ is given by the coherent sum of transition amplitudes through any accessible intermediate state $|j\rangle$,² giving [211]

$$\frac{d\gamma_{a\rightarrow b}}{d\Omega} = \frac{I_0\omega_{\text{sc}}^3}{(4\pi\epsilon_0)^2\hbar^3c^4} \left| \sum_{|j\rangle} \frac{\langle b|\epsilon_{\text{sc}} \cdot \hat{\mathbf{d}}|j\rangle \langle j|\epsilon \cdot \hat{\mathbf{d}}|a\rangle}{\omega_{ja} - \omega} + \frac{\langle b|\epsilon \cdot \hat{\mathbf{d}}|j\rangle \langle j|\epsilon_{\text{sc}} \cdot \hat{\mathbf{d}}|a\rangle}{\omega_{ja} + \omega_{\text{sc}}} \right|^2, \quad (\text{B.10})$$

where $\hbar\omega_{ja}$ is the energy difference between $|a\rangle$ and $|j\rangle$, and I_0 is the intensity of the probe. The first term in the sum corresponds to absorption followed by emission, whereas the second term is emission followed by absorption. Since we consider a BEC in the electronic ground state, all intermediate states have $\omega_{ja} > 0$ and the second term can be disregarded.

Following [211], note that the scattered photon polarisation ϵ_{sc} is simply a Cartesian unit vector and can be factored out to give

$$\frac{d\gamma_{a\rightarrow b}}{d\Omega} = \frac{I_0\omega_{\text{sc}}^3}{(4\pi\epsilon_0)^2\hbar^3c^4} \left| \epsilon_{\text{sc}} \cdot \sum_{|j\rangle} \frac{\langle b|\hat{\mathbf{d}}|j\rangle \langle j|\epsilon \cdot \hat{\mathbf{d}}|a\rangle}{\omega_{ja} - \omega} \right|^2 \equiv \frac{I_0\omega_{\text{sc}}^3}{(4\pi\epsilon_0)^2\hbar^3c^4} |\epsilon_{\text{sc}} \cdot \mathbf{D}|^2. \quad (\text{B.11})$$

where \mathbf{D} represents an ‘effective’ dipole moment induced by the laser.

¹ The momentum kick from absorption expels the atom from the BEC, so the photon scattering rate is equal to the atom loss rate due to scattering.

² The kets $|a\rangle$, $|j\rangle$ and $|b\rangle$ are shorthand for a full set of quantum numbers describing the state, each corresponding to some $|n, L, J, F, m_F\rangle$. Therefore the sum over $|j\rangle$ corresponds to summing over *all* its associated quantum numbers.

The net scattering rate for the $|a\rangle \rightarrow |b\rangle$ transition is obtained by integrating (B.11) over all emission solid angles ($d\Omega$) and summing over the scattered photon polarisations (denoted $\epsilon_{\text{sc},\parallel}$ and $\epsilon_{\text{sc},\perp}$),

$$\begin{aligned}\gamma_{a\rightarrow b} &= \int \frac{d\gamma_{a\rightarrow b}}{d\Omega}(\epsilon_{\text{sc},\parallel}) + \frac{d\gamma_{a\rightarrow b}}{d\Omega}(\epsilon_{\text{sc},\perp}) d\Omega \\ &= \frac{I_0\omega_{\text{sc}}^3}{(4\pi\epsilon_0)^2\hbar^3c^4} \int |\epsilon_{\text{sc},\parallel} \cdot \mathbf{D}|^2 + |\epsilon_{\text{sc},\perp} \cdot \mathbf{D}|^2 d\Omega.\end{aligned}\quad (\text{B.12})$$

The scattered wave is locally planar in the far field, so only two polarisations need to be considered. Both polarisations are perpendicular to the propagation vector \mathbf{k}_{sc} , which by definition is normal to the surface element $d\Omega$. The orientation of the two polarisations is arbitrary (provided $\epsilon_{\text{sc},\perp} \times \epsilon_{\text{sc},\parallel} = \hat{\mathbf{k}}_{\text{sc}}$), allowing $\epsilon_{\text{sc},\perp}$ to be chosen perpendicular to \mathbf{D} and eliminating the second term in (B.12).

Evaluating the integral is then a textbook problem [212, §4], yielding

$$\gamma_{a\rightarrow b} = \frac{I_0\omega_{\text{sc}}^3}{(4\pi\epsilon_0)^2\hbar^3c^4} \left(\frac{8\pi}{3} |\mathbf{D}|^2 \right).\quad (\text{B.13})$$

The square magnitude of the induced dipole is calculated in the spherical-basis,

$$|\mathbf{D}|^2 = \sum_q |D_q|^2 \quad \text{where} \quad D_q = \epsilon_q \cdot \mathbf{D} \quad \text{and} \quad \mathbf{A} \cdot \mathbf{B} = \sum_q (-1)^q A_q B_{-q} \quad (\text{B.14})$$

is the spherical vector dot-product [111, §7.3.3]. Applying conservation of energy requires the frequency of the scattered photon to be $\omega_{\text{sc}} = \omega - \omega_{ba}$, giving

$$\gamma_{a\rightarrow b} = \frac{I_0(\omega - \omega_{ba})^3}{6\pi\epsilon_0^2\hbar^3c^4} \sum_q \left| \sum_{|j\rangle} \frac{\langle b|\hat{d}_q|j\rangle \langle j|\epsilon \cdot \hat{\mathbf{d}}|a\rangle}{\omega_{ja} - \omega} \right|^2,\quad (\text{B.15})$$

where q now represents the polarisation of the scattered photon.

The final states are distinguishable, so the scattering rates into each possible final state $|b\rangle$ are summed to obtain the total scattering rate for the initial state $|a\rangle$,

$$\gamma_a \equiv \sum_{|b\rangle} \gamma_{a\rightarrow b} = \frac{I_0}{6\pi\epsilon_0^2\hbar^3c^4} \sum_{|b\rangle} (\omega - \omega_{ba})^3 \sum_q \left| \sum_{|j\rangle} \frac{\langle b|\hat{d}_q|j\rangle \langle j|\epsilon \cdot \hat{\mathbf{d}}|a\rangle}{\omega_{ja} - \omega} \right|^2.\quad (\text{B.16})$$

We are interested in a linearly polarised probe beam, and choose the quantisation axis along the polarisation axis, giving $\epsilon \cdot \hat{\mathbf{d}} = \hat{d}_0$. The transitions of interest are of the form

$$\underbrace{|n J F m_F\rangle}_{|a\rangle} \rightarrow \underbrace{|n J' F' m_{F'}\rangle}_{|j\rangle} \rightarrow \underbrace{|n J F'' m_{F''}\rangle}_{|b\rangle},$$

with the initial state $J = 1/2, F = 1$.

Explicitly including the quantum numbers in the sum, and taking the Zeeman splitting within each hyperfine level as negligible compared to the detuning, gives³

$$\gamma_F = \frac{I_0}{6\pi\epsilon_0^2\hbar^3c^4} \sum_{F''m_{F''}q} (\omega - \omega_{F''F})^3 \left| \sum_{J'F'm_{F'}} \frac{\langle JF''m_{F''}|\hat{d}_q|J'F'm_{F'}\rangle \langle J'F'm_{F'}|\hat{d}_0|JFm_F\rangle}{\omega_{F''F} - \omega} \right|^2. \quad (\text{B.17})$$

The dipole matrix elements can be computed using the Wigner-Eckart theorem. As in §2.1, in the convention of Ref. [121] this gives⁴

$$\langle JFm_F|\hat{d}_q|J'F'm_{F'}\rangle = (-1)^{F'-1+m_F} \sqrt{2F+1} \begin{pmatrix} F' & 1 & F \\ m_{F'} & q & -m_F \end{pmatrix} \langle JF||\hat{d}||J'F'\rangle. \quad (\text{B.18})$$

The hyperfine reduced matrix-element is related to the fine-structure reduced matrix element through the Wigner-6j, including the nuclear isospin I_s as

$$\langle JF||\hat{d}||J'F'\rangle = (-1)^{F'+J+I+1} \sqrt{(2F'+1)(2J+1)} \left\{ \begin{matrix} J & J' & 1 \\ F' & F & I_s \end{matrix} \right\} \langle J||\hat{d}||J'\rangle. \quad (\text{B.19})$$

This reduced element is itself related to the fine-structure transition linewidth $\Gamma_{J'}$ as⁵

$$\Gamma_{J'} = \frac{8\pi^2}{3\epsilon_0\hbar\lambda_{J'}^3} \frac{2J+1}{2J'+1} |\langle J||\hat{d}||J'\rangle|^2. \quad (\text{B.20})$$

Hence using measured values of the linewidth [121], the reduced dipole element can be computed and the scattering rate obtained for a given detuning.

In comparing different fine structure transitions, it is instructive to further decompose the reduced dipole element in terms of the orbital angular momentum L using the Wigner-6j expression (B.19) and $\hat{J} = \hat{L} + \hat{S}$.

$$\langle J||\hat{d}||J'\rangle = (-1)^{J'+L+S+1} \sqrt{(2J'+1)(2L+1)} \left\{ \begin{matrix} L & L' & 1 \\ J' & J & S \end{matrix} \right\} \langle L||\hat{d}||L'\rangle. \quad (\text{B.21})$$

Hence the dependence on J' in (B.20) can be written as

$$\lambda_{J'}^3 \Gamma_{J'} = \frac{8\pi^2}{3\epsilon_0\hbar} \frac{2J+1}{2L+1} \left\{ \begin{matrix} L & L' & 1 \\ J' & J & S \end{matrix} \right\}^2 |\langle L||\hat{d}||L'\rangle|^2. \quad (\text{B.22})$$

The D1 and D2 transitions are both $S \rightarrow P$ transitions, so both have $L = 0$ and $L' = 1$. The L -reduced matrix element is therefore the same and evaluating the Wigner-6j symbols gives

$$\lambda_{D1}^3 \Gamma_{D1} = \lambda_{D2}^3 \Gamma_{D2}. \quad (\text{B.23})$$

³ Although the expression involves m_F , the net scattering rate is independent of the chosen value.

⁴ For details on the decomposition, see for example [111, §7.3.4.1, §7.3.7.1], [213, §6.1.2], [211, §B].

⁵ This expression explicitly refers to spontaneous decay $J' \rightarrow J$, so ordering of J and J' matters.

It is important to note that the reduced matrix elements are not symmetric, and should not be thought of as expectation values. The symmetry of the reduced element is related to the symmetry of Clebsch-Gordan coefficients, giving [111, §7.3.5.1]

$$\langle J' || \hat{d} || J \rangle = (-1)^{J'-J} \sqrt{\frac{2J+1}{2J'+1}} \langle J || \hat{d} || J' \rangle^*. \quad (\text{B.24})$$

Using (B.23), this leads to the somewhat counter-intuitive result that

$$\frac{\langle J = \frac{1}{2} || \hat{d} || J' = \frac{3}{2} \rangle}{\langle J = \frac{1}{2} || \hat{d} || J' = \frac{1}{2} \rangle} = \sqrt{2} \quad \text{but} \quad \frac{\langle J' = \frac{3}{2} || \hat{d} || J = \frac{1}{2} \rangle}{\langle J' = \frac{1}{2} || \hat{d} || J = \frac{1}{2} \rangle} = -1. \quad (\text{B.25})$$

B.2.1 Far-detuned limit

For the large detunings of interest, the hyperfine splitting of both the ground and excited states can be taken as negligible (see §2.2). Recalling that $|J'F'\rangle$ is the intermediate (excited) state and $|JF''\rangle$ is the final (ground) state,

$$\omega_{F''F} \ll \omega \quad \Rightarrow \quad (\omega - \omega_{F''F})^3 \approx \omega^3, \quad (\text{B.26})$$

$$\omega_{F'F} \approx \omega_{J'J} \quad \Rightarrow \quad \omega_{ja} - \omega \approx \omega_{J'J} - \omega = -\Delta_{J'}, \quad (\text{B.27})$$

where $\Delta_{J'}$ is the detuning from the associated line centre. The sum over intermediate states then separates into fine structure contributions,

$$\gamma_a \approx \frac{I_0 \omega^3}{6\pi \epsilon_0^2 \hbar^3 c^4} \sum_{F'' m_{F''} q} \left| \frac{A_{D1}}{\Delta_{D1}} + \frac{A_{D2}}{\Delta_{D2}} \right|^2, \quad (\text{B.28})$$

$$\text{with } A_{J'} = \sum_{F' m_{F'}} \langle JF'' m_{F''} | \hat{d}_q | J'F' m_{F'} \rangle \langle J'F' m_{F'} | \hat{d}_0 | JF m_F \rangle. \quad (\text{B.29})$$

The only potentially nonzero terms in the sum have $m_{F'} = m_F$ and $q = m_{F'} - m_{F''}$, so the sums over $m_{F'}$ and q vanish.⁶ There are two kinds of scattering processes, corresponding to either $q = 0$ or $q = \pm 1$. These describe Rayleigh scattering and Raman scattering respectively, as $q = 0$ gives leaves the atomic state unchanged⁷ and $q \neq 0$ corresponds to transitions between different states. Evaluating the coefficients for the transitions of interest shows [210]

$$A_{D2} = \begin{cases} 2A_{D1} & \text{for Rayleigh scattering } (q = 0) \\ -A_{D1} & \text{for Raman scattering } (q = \pm 1). \end{cases} \quad (\text{B.30})$$

⁶ Sometimes this is written with the sum over q put *inside* the inner sum (e.g. [210, 214]). Strictly this is incorrect as the photon is observable and the scattering *rates* sum together, not the *amplitudes*. However, since $m_{F'}$ is summed over in the internal summation, and non-zero terms have $m_{F'} = m_{F''} + q$, these are equivalent if notationally confusing formulations.

⁷ While it is obvious that $q = 0$ requires $m_{F''} = m_F$, the scattering rate $\gamma_{|a\rangle \rightarrow |b\rangle} = 0$ for $F'' \neq F$.

Hence the contributions $A_{J'}/\Delta_{J'}$ to (B.28) can suppress each other in specific regimes. For detunings far from both lines (resulting in $\Delta_{D1} \approx \Delta_{D2}$), Raman scattering is suppressed. At the magic wavelength (with $\Delta_{D2} = -2\Delta_{D1}$, see §2.4), Rayleigh scattering vanishes.

This is consistent with the Kramers-Kronig interpretation of the spherical tensor decomposition of the interaction Hamiltonian (2.7). Rayleigh scattering corresponds to the scalar term $\hat{\mathcal{H}}^{(0)}$ whereas Raman scattering arises from the vector term $\hat{\mathcal{H}}^{(1)}$. The scattering rates vanish when the Hamiltonian terms become small (see (2.30) and (2.46) respectively).

Introducing the coupling constant

$$\gamma_0 = \frac{\omega^3}{72\pi\epsilon_0^2\hbar^3c^4} |\langle J = \frac{1}{2} || \hat{d} || J' = \frac{3}{2} \rangle|^4 = \frac{\omega^3}{18\pi\epsilon_0^2\hbar^3c^4} \alpha_0^2 = \frac{\lambda_{J'}^6 \Gamma_{J'}^2}{8\pi\hbar c \lambda^3}, \quad (\text{B.31})$$

the decay rate corresponding to Rayleigh scattering is⁸

$$\gamma^{(q=0)} = 3I_0\gamma_0 \left(\frac{1}{3\Delta_1} + \frac{2}{3\Delta_2} \right)^2 = \frac{I_0\gamma_0}{3} \left(\sum_{J'F'} \frac{\alpha_{J'F'}^{(0)}}{\alpha_0\Delta_{J'}} \right)^2, \quad (\text{B.32})$$

whereas considering Raman scattering instead gives

$$\gamma^{(q=1)} + \gamma^{(q=-1)} = 6I_0\gamma_0 \left(\frac{1}{3\Delta_1} - \frac{1}{3\Delta_2} \right)^2 = 6I_0\gamma_0 \left(\sum_{J'F'} \frac{\alpha_{J'F'}^{(1)}}{\alpha_0\Delta_{J'}} \right)^2. \quad (\text{B.33})$$

The total scattering rate is the sum of the two, which gives rise to a cancellation of cross-terms between the D1 and D2 amplitudes, leaving only the ‘incoherent’ sum of fine-structure contributions,

$$\begin{aligned} \gamma &= \gamma^{(q=-1)} + \gamma^{(q=0)} + \gamma^{(q=1)} \\ &= I_0\gamma_0 \left(\frac{1}{\Delta_1^2} + \frac{2}{\Delta_2^2} \right) = I_0\gamma_0 \sum_{J'} \left(\sum_{F'} \frac{\alpha_{J'F'}^{(0)}}{\alpha_0} \right)^2 \frac{1}{\Delta_{J'}^2}. \end{aligned} \quad (\text{B.34})$$

At the magic wavelength, given by (2.30), the far-detuned limit §2.2 applies and the scattering rate is

$$\gamma = \frac{(\Gamma_{J'}\lambda_{J'}^3)^2}{128c^3\pi^4\hbar} \frac{(\lambda_{D1} + 2\lambda_{D2})^3}{(\lambda_{D1} - \lambda_{D2})^2\lambda_{D1}\lambda_{D2}} I_0. \quad (\text{B.35})$$

It is worth restating that the net scattering rate is independent of m_F , and therefore applies to superpositions of m_F states as well, such as the $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ population distribution studied in this thesis. Also note that this result only holds for linearly polarised (π) incident light because of the simplification in (B.29) that only one component of the dipole operator contributes to absorption.

⁸ The expression of the scattering rate in terms of $\alpha^{(i)}$ is true for *this* system and is conjectured (but not proven here) to be true in general.

B.3: Photoassociation resonances

Photoassociation occurs when a strong optical field induces two atoms to form a molecular dimer, with the released binding energy promoting the dimer to an excited rovibrational state. Spectroscopically, photoassociation appears as a resonance with a particular lineshape depending on the properties of the inter-atomic potential.

Although photoassociation near the probe wavelength of interest ($\lambda = 790.0$ nm) has not yet been studied, studies have been carried out near the D2 line at 780.241 nm (Figure B.1). Extrapolating the known resonance line centres enables estimation of any resonances near the probe wavelength.

Phenomenologically, we model the wavenumber associated with photoassociation resonance as a quartic in the vibrational quantum number v as

$$\frac{1}{\lambda_v} = a_0 - \sum_{n=2}^4 a_n^n (v - v_0)^n. \quad (\text{B.36})$$

This is a generic quartic model with 5 parameters, expressed in a way that reduces the covariances between the fit parameters and allows more independent estimation of their values and uncertainties (Table B.1).

The model predicts the closest resonance occurs with a vibrational quantum number of $v = 94$ at a detuning of 56 GHz. The line spacing of the nearest resonances was estimated as 134 GHz, indicating it is unlikely that another resonance is nearby. The measured resonance linewidths were 1.8(4) GHz, which is small compared to the extrapolated detuning, even with the expected power-broadening from the intense Faraday probe beam. Estimating an uncertainty on the detuning is difficult as the resonance is extrapolated by 1.1 THz and covariance between the fit parameters overestimates the uncertainties.

Furthermore, as the magic wavelength is detuned between the D1 and D2 lines, it is possible that a photoassociation resonance arising from the D1 line could result in another loss channel. No data detuned sufficiently far from the D1 line for extrapolation was available for fitting, preventing an estimation from being made.

Including two-body loss in the model would complicate the fitting procedure and reduce the numerical stability (see §5.7), which already required approximation using (5.4) to achieve reliable convergence. If photoassociation were a prevalent loss mechanism, neglecting it from the model would result in an inferred one-body loss rate that would be non-linear in probe power. However, the gradient of the loss rate curve (Figure 5.10) agrees with the scattering rate prediction (3.19), indicating that losses due to off-resonant scattering dominated any photoassociation processes.

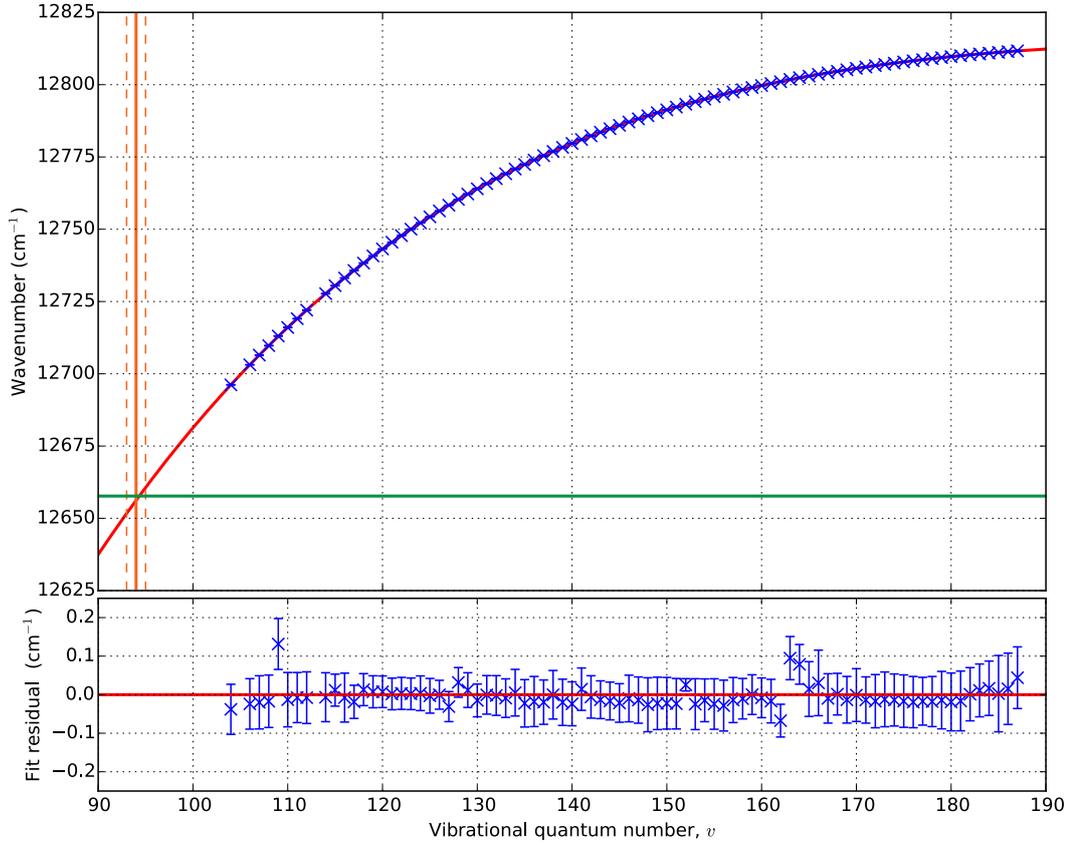


Figure B.1: Known $1g$ -photoassociation resonances for ^{87}Rb between 780.5 nm and 787.6 nm (top). The phenomenological model (B.36) fits well, with residuals (bottom) mostly falling within the measured linewidth (error bars) of each resonance. Extrapolating the fit to the magic wavelength $\lambda_{\text{magic}} = 790.0$ nm (green) shows nearby photoassociation resonances (orange). The $v = 94$ resonance (solid orange) is estimated to be detuned by 56 GHz, and the adjacent lines (dashed orange) give the line spacing as 134 GHz. Data provided by [215].

Parameter	v_0	a_0	a_2	a_3	a_4
Fit value	210(3)	12814.3(4)	0.069(4)	0.0223(9)	0.0279(3)

Table B.1: Fit parameters for the line centres of photoassociation resonances in ^{87}Rb .