An Ultra-Cold Lithium Source For Investigating Cold Dipolar Gases

Jacob Oliver Stack

Department of Physics

Imperial College London

University of London

This thesis is submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy of the University of London and the Diploma of Membership of Imperial College.

September, 2010

Declaration of Originality

I declare that all the work presented in this thesis is my own unless otherwise referenced. The only exception is the results for tests 6 through 9 presented in table 3.1. These results were collected by Anna Koveleva and James Millen under my partial supervision during their MSci final project.

Abstract

This thesis describes progress towards the formation of an ultra-cold dipolar gas for the study of strongly-correlated quantum systems. To this end, two lines of research have been followed: the direct formation of an electrically polarised atomic gas, and the use of these ultra-cold atoms to sympathetically cool a molecular gas to low temperatures. For use in these experiments a Zeeman slower and magneto-optical trap (MOT) have been designed, constructed and demonstrated. Up to 2×10^{10} ^{7}Li atoms are loaded into the MOT with a minimum temperature of 0.85 mK. To transport the atoms from the MOT, for further study, a pair of overlapping magnetic quadrupole traps are constructed. By carefully adjusting the currents in the trap, the quadrupole potential can be moved from the one trap to the next. Atoms are loaded into the first magnetic trap where the trap lifetime exceeds 1 second. They are then transported 27.5 mm across the chamber to the second trap, held there for a period of time and finally returned to the first trap to be imaged. The efficiency of the transport procedure is measured and compared to the results predicted by simulations. Measurements of the number of atoms remaining in the trap after transport indicate a transport efficiency of 70% per trip for transport times of both 50 and 100 ms. A lifetime of 111$\pm$19 ms is measured in the second trap. To electrically polarise the atomic gas such that dipole-dipole interactions become significant, an electric field of the order of 1 MV/cm is needed. The feasibility of producing such a large field is investigated, and a field of 0.81 MV/cm is realised between two stainless steel, optically-polished electrodes in ultra-high vacuum. Finally, the possibility of using indium-tin-oxide (ITO) coated glass plates is investigated for this application.

Acknowledgments

There are many people I would like to thank including family, friends and colleagues, without whom, this work would not have been possible. There isn't room to list everyone here but there are a select few I would like to acknowledge especially.

First and foremost is Mike, my supervisor, for his constant guidance as well as challenges, which have helped me mature into the semi-respectable scientist I am now. Without his help this work would not have been possible. I have found through my time in the cold matter group a huge ensemble of friends and colleagues who are always willing to help at any inconvenience and were often the source of many an interesting discussion. A special thank you is included to all those in the group who helped to wind the dozens of coils for the Zeeman slower and magnetic traps!

A very special thanks is reserved for Jon Dyne, who, time and time again has gone beyond the call of duty in the workshop to make sure the experiment has everything it needs and remains in good working order. He has provided both invaluable consultations on the design of various components as well as a valued friendship and squash partner. I would also like to thank David Williams and Steve Maine for their advice and help in the workshop, particular Dave who spent many an hour machining the formers for the Zeeman Slower!

I have had support from many friends during thesis, both old and new, the list for which would require a thesis of its own. However, I want to give special thanks to a few, so thank you to Sofia Olsson for your love and support while I wrote this thesis, despite my grumbling(!); to Andy Smith for your constantly uplifting charisma and generally keeping me entertained; and to Noé Kains and Ben Beiny for your friendship and many good times.

Lastly, but by no means least, I have to thank my family; my sister, Torie; my Mum; my grandparents, Julian and Vivien; and my Dad who sadly passed away in June 2007. They have always supported me unconditionally in all my past and future endeavours and I always appreciate them.

Oh, and hello and thanks to my examiners, Danny Segal and Donatella Cassettari for reading this thesis!

"With equal passion I have sought knowledge. I have wished to understand the hearts of men. I have wished to know why the stars shine. And I have tried to apprehend the Pythagorean power by which number holds sway above the flux"

Bertrand Russell, Autobiography

"Be excellent to each other"

Bill and Ted

Contents

1 Introduction **14**

1.1 Background and Motivation for this Work 14

1.2 Aims 20

1.3 Overview 21

2 Making a Li MOT **25**

2.1 A Brief Introduction to Laser Cooling 25

2.1.1 Radiation Pressure 26

2.1.2 Laser Cooling and Trapping 27

2.1.3 Sub-Doppler Cooling 31

2.2 Lithium 33

2.3 The Laser System 34

2.3.1 Frequency Stabilisation 36

2.3.2 Optical Set-up 45

2.4 The Lithium Beam 48

2.4.1 Effusive Beams 49

2.4.2 Beam Properties 50

2.4.3 The Oven 51

2.4.4 Absolute Beam Flux 54

2.5 The Zeeman Slower 57

2.5.1 Theory 58

2.5.2 Length Calculations 67

2.5.3 Vacuum Conductance 70

2.5.4 Magnetic Field Optimisation 72

2.5.5 Construction 78

2.5.6 Performance and Optimisation 81

2.6 The Magneto-Optical Trap 85

2.6.1 The Chamber Design 85

2.6.2 MOT coils 88

2.7 Atom Number 89
2.7.1 Fluorescence measurements 89
2.7.2 Absorption Imaging 92
2.7.3 Characterisation 96
2.8 Cloud Temperature 101
2.8.1 Release-recapture 102
2.8.2 Ballistic expansion 104
2.9 Summary 107

3 Production of Very Large Electric Fields 109
3.1 Theoretical Overview 110
3.1.1 Field Emission 110
3.1.2 Breakdown Mechanisms 112
3.2 Experimental Set-up 114
3.2.1 Outline 114
3.2.2 Vacuum System 115
3.2.3 Measurement Circuitry 117
3.2.4 The Electrodes 118
3.3 Polishing Techniques 119
3.3.1 Hand-polishing 119
3.3.2 Machine-polishing 120
3.3.3 Electro-polishing 121
3.3.4 Optical-polishing 123
3.4 Results 124
3.5 Atoms Under Electric Fields 127
3.5.1 ITO coated electrodes 130
3.6 Polarising the Atoms 136

4 Transporting the Atoms 142
4.1 Magnetic Trapping of Neutral Atoms 144
4.2 Designing the Magnetic Traps 145
4.2.1 First Coil Pair 146
4.2.2 Second Coil Pair 148
4.3 Realising the Magnetic Trap 152
4.4 Lowering the Cloud Temperature 156
4.5 Trap Lifetime Measurements 160
4.6 A Deeper Trap 162
4.7 Transferring Between the Traps 164

5 Conclusion and Outlook **171**

5.1 Final Conclusions 171

5.2 Outlook 173

A Installation and Assembly of the Magnetic Traps **176**

List of Figures

1.1 Dipole-dipole interactions 15
1.2 A schematic diagram giving an overview of the main experiment. . . 22

2.1 The force on an atom sitting in a 1D optical molasses 28
2.2 The restoring forces acting on an atom sitting in a MOT 30
2.3 An illustration depicting the mechanisms behind Sisyphus cooling . . 32
2.4 The energy level diagram for ^{7}Li, highlighting the cooling and repump transition frequencies required for laser cooling. 34
2.5 A schematic of the optical set-up used for saturation spectroscopy . 37
2.6 Saturation spectroscopy data showing the ground state hyperfine features including the cross-over for the D2 transition in ^{7}Li 39
2.7 The optical set up used to perform polarisation spectroscopy of lithium 40
2.8 Polarisation spectroscopy data showing the ground state hyperfine features including the cross-over for the D2 transition in ^{7}Li. 44
2.9 The energy level diagram showing the required frequencies for the experiment 46
2.10 The full optical set-up used to supply the 5 required frequencies to the experiment 47
2.11 The optical set-up used to deliver the trapping light, the slowing light and repump beams to the experiment. 48
2.12 Vapour pressure as a function of temperature for lithium 49
2.13 The Maxwell-Boltzmann curves showing the atomic speed distribution of lithium and rubidium as a function of velocity 50
2.14 The lithium oven 52
2.15 The new lithium oven 53
2.16 Lithium, modelled as a three-level atom 55
2.17 A comparison of experimental and theoretical data for an effusive beam exiting the oven 56
2.18 The ideal magnetic field profile required to decelerate ^{7}Li atoms from 900 ms^{-1} to 50 ms^{-1} 61

2.19 The full energy level diagram for the $^2S_{1/2}$ and $^2P_{3/2}$ lithium states as a function of magnetic field 63
2.20 The transition frequencies for (a) the $|1,-1\rangle \rightarrow |2,-2\rangle$ transition and (b) the $|1,1\rangle \rightarrow |2,2\rangle$ transition as a function of magnetic field. . . . 64
2.21 A plot of the transition rates as a function of magnetic field for the two decay channels pumping atoms into the $^2S_{1/2}\,|F=1,\ M_F=1\rangle$ state 65
2.22 The probability of pumping an atom into the $^2S_{1/2}\,|F=1,\ M_F=1\rangle$ state within a given section of the Zeeman slower 66
2.23 The normalised atom flux leaving the slower as a function of Zeeman slower length with and without the effects of transverse cooling from a converging laser cooling beam. 69
2.24 The normalised atom flux exiting the slower as a function of slower length with and without a differential pumping tube 72
2.25 The magnetic field residues obtained by subtracting the ideal field from the field of a slower built from 8 and 13 coils. 75
2.26 The total magnetic field profile for the Zeeman slower for a set of 13 coils run from 4 power supplies 76
2.27 Numerical simulations showing the number of atoms as a function of velocity exiting the Zeeman slower 78
2.28 A photograph of the Zeeman slower components prior to construction 79
2.29 A technical drawing for a typical coil former used in the slower . . . 80
2.30 A schematic of the circuit diagram used to measure and adjust the current in each coil of a given coil set 81
2.31 Comparison of the theoretical and experimental magnetic field data for each of the 13 coils used to construct the Zeeman slower 82
2.32 The magnetic field on the axis of each of the 13 coils used to construct the Zeeman slower 83
2.33 The experimentally optimised full magnetic field profile of the slower compared to the expected ideal case 84
2.34 A true-colour photograph of a cloud of $\sim 1 \times 10^{10}$ neutral lithium atoms trapped in our MOT. 85
2.35 A schematic of the MOT chamber, illustrating the trapping beams, the position of the MOT cloud and the trapping coils 87
2.36 The magnetic field profile in the transitional period between the end of the slower and the centre of the trapping volume 89
2.37 Experimental data depicting the dependence of the number of trapped atoms on the temperature of the oven 97

2.38 Experimental data for the dependence of atom number on laser power 98
2.39 Experimental data for the dependence of atom number on the laser detunings 99
2.40 Experimental data for the dependence of atom number on the magnetic field gradient of the MOT 101
2.41 Release-recapture data for atoms in the magneto-optical trap 104
2.42 A series of absorption images of the atom cloud taken at successively longer times after release, illustrating ballistic expansion. 104
2.43 The digital triggering schedule delivered by the computer to the experimental apparatus 105
2.44 Ballistic expansion temperature data for the MOT 106

3.1 A simplified representation of the potential barrier at the surface of a metal at $x = 0$ as seen by an electron with and without an electric field present 110
3.2 The field emission from a metal surface as a function of electric field strength 111
3.3 The modification of an electric field between two plates due to the presence of a microscopic protrusion. 113
3.4 A photograph of the experimental apparatus used to measure the effects of surface quality on the electric field threshold of the electrodes 115
3.5 A scaled diagram of the electrode test rig used to support the stainless steel electrodes at a separation of 0.5 mm 116
3.6 A flow chart highlighting the basic operations of the program implemented in LabView to record the current and voltage data from the experiment. 118
3.7 Photographs of unpolished electrodes viewed through an optical microscope at various magnifications 119
3.8 Photographs of hand polished electrodes viewed through an optical microscope at various magnifications 120
3.9 Photographs of machine polished electrodes viewed through an optical microscope at various magnifications 121
3.10 Photographs of electro-polished electrodes viewed through an optical microscope at various magnifications 122
3.11 Photographs of electro-polished electrodes polished using an improved recipe, viewed through an optical microscope at various magnifications 123
3.12 Photographs of optically polished electrodes viewed through an optical microscope at various magnifications 123

3.13 A graph presenting the current and voltage data collected by the DAQ board for a typical experimental run 124
3.14 The maximum acceleration along the axis normal to the plate surface experienced by a lithium atom as a function of distance from the plate edge 128
3.15 A photograph of the full assembly of the ITO coated electrode plates and their support rig 131
3.16 A basic schematic illustrating how the plates are secured to the support rig 132
3.17 Detailed drawings of the adjustable capacitor assembly design for the ITO coated electrodes. 133
3.18 A schematic illustrating how a relative curvature between the electrode plates will lead to interference fringes, separated by a distance x. 134
3.19 An illustration of the optical set-up used to measure the separation of the electrodes 134
3.20 The intensity profile of the reflected beams from the surface of parallel glass plates 135
3.21 A schematic depicting the model used to calculate the effects of an applied electric field between two plates 139

4.1 Diagrams illustrating the magnetic traps and their position in the MOT chamber 147
4.2 Diagrams illustrating the shape and dimensions of the coil former used for the second trap. 148
4.3 A photograph of the full trap assembly showing both magnetic traps and their supporting structure 151
4.4 A graph depicting the number of atoms remaining in the magnetic trap once released from the MOT as a function of the time since release 155
4.5 A graph presenting the number of atoms remaining in the magnetic trap, 10 ms since release from the MOT for various lengths of cooling pulse applied prior to release. 158
4.6 A graph depicting the decay curves for atoms in a magnetic trap as a function of time since release for different lengths of cooling pulse applied prior to release. 158
4.7 Ballistic expansion data used to determine the temperature of the MOT. The graph shows the Gaussian width of the cloud as a function of time since release 159

4.8 A graph presenting the number of atoms remaining in the magnetic trap as a function of time since being released from the MOT 160
4.9 A graph showing the absolute magnetic field profile along the axis of the slower, including field effects from Zeeman slower, magnetic trap 1 and the compensation coils 162
4.10 A graph depicting the increase in atom number collected by the magnetic trap for an increased trap depth 163
4.11 A magnetic trapping potential as it is transferred between the centres of two quadrupole traps. 164
4.12 A series of 1D and 2D plots of the magnetic field profile created by traps 1 and 2 during transport, calculated in the lab frame. 165
4.13 The effective potentials during transport in the frame of an atom beginning with zero velocity at the centre of the trap 167
4.14 Phase-space plots showing the results of 1-D simulations used to calculate the number of atoms expected to remain trapped after a single transport from trap 1 to trap 2 168
4.15 The fraction of atoms remaining after transport to trap 2 and back as a function of time held in trap 2 169
4.16 Three images of trap 1, before, during and after transport 169

A.1 Detailed drawings of the MOT chamber and the in-vacuum magnetic traps 178

List of Tables

1.1 The strength of dipole-dipole interactions relative to s-wave scattering for three cases. 16

2.1 Physical properties of bulk and atomic lithium 33
2.2 Properties of the ^{7}Li D2 transition 34
2.3 Data for the coils designed to produce the required magnetic field. . 76
2.4 The currents supplied to the coils of the Zeeman slower after optimisation. 85

3.1 A summary of the results for the largest obtained electric field strengths for various polishing techniques. 126

4.1 Coil data for the magnetic travellator 151

Chapter 1

Introduction

1.1 Background and Motivation for this Work

Over recent years there has been considerable interest growing in the atomic physics community in the possibility of studying the behaviour of cold, polar gases. These strongly correlated quantum gases can be used to simulate prototypes of condensed matter systems that are currently not well understood [1], including higher temperature superconductivity systems. At present, the collision physics of ultra-cold gases formed in the lab is dominated by van der Waals interactions. These collisional processes can be described by a partial-wave expansion of the angular-momentum components of the interacting wave functions. For very low collision energies the partial wave expansion is dominated by the zero angular momentum component, known as the s-wave component. For alkali atoms the energy threshold at which the s-wave component of the expansion becomes dominant equates to a temperature of approximately 1 mK. Thus, in the regime of ultra-cold gases and Bose-Einstein condensates (BECs) the particle-particle interactions are almost entirely governed by s-wave scattering and consequently are described by only a single parameter, the s-wave scattering length. These interactions are short range, isotropic and energy-independent in the low-energy regime of the cold gas. The dominating interaction of these van der Waals forces scales as C_6/r^6, where C_6 is a species-specific coefficient and r is the internuclear distance. The nature and stability of such ultra-cold clouds can be manipulated by adjusting this scattering length through the application of magnetic fields [2]. This control has allowed researchers to explore the physics of Feshbach resonances and create ultra-cold molecules in this way [3].

A wider range of many-body physics phenomena could be explored if the gas comprised of particles with more complex interactions, such as would occur in an ultra-cold *polar* gas. Such a polarised state could be realised by the application of an electric or magnetic field. In contrast to s-wave scattering interactions, multipole

interactions between particles are longer range and anisotropic, with the longest range interactions being due to dipole-dipole interactions, scaling with distance as $1/r^3$. By sufficiently polarising a cold gas these long-range interactions will dominate over the usual s-wave interaction processes providing a novel and promising testing ground for theories of quantum phase transitions. The number of predicted quantum phases achievable increases significantly if such a gas is coupled to a periodic potential, such as an optical lattice. These phases include Mott insulators, supersolids and checkerboard phases [4]. Confirmation of or deviations from the expected results will provide valuable insights into the physics of condensed matter problems. The dynamics of such an experiment provide a high level of control over the parameters determining each phase and allows one to tune across a rich variety of quantum phase transitions. Others have proposed schemes of using such a polar gas for quantum information processing [5, 6].

In a simple dipolar BEC the forces between neighbouring dipoles will either be attractive or repulsive depending on their relative position, as illustrated in figure 1.1. The relative positions of the positive and negative components of the dipole determine the direction of the Coulomb force between the neighbouring dipoles.

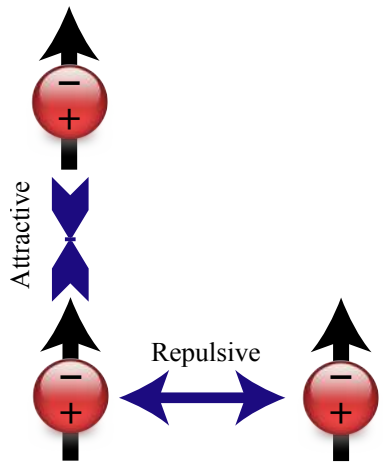

Figure 1.1: Dipole-dipole interactions

This allows the nature and stability of the BEC to be adjusted simply by adjusting the aspect ratio of the trap containing the cloud. Within an optical lattice this level of control is achieved by modifying the wavelength or intensity of one or more of the lattice beams, which in turn allows control of the occupation number of the lattice sites. It is also possible to modify the angle of an applied electric field relative to a 2D lattice plane, thereby adjusting the strength and anisotropy of the forces between neighbouring dipoles and consequently making significant modifications to the phase diagram. Through careful adjustment of these parameters one can realise and study many of the quantum phases predicted by many-body quantum theory.

To create a dipolar gas one may employ either magnetic or electric fields. The strength of the dipole interactions will depend on the particle of choice, namely atoms or molecules, and the field type. The feasibility of each case is discussed below. The dimensionless parameter that quantifies the strength of the interactions relative to s-wave scattering processes can be written as [7]

$$\epsilon_{dd} = \frac{a_{dd}}{|a_s|} = \frac{1}{3}\frac{M}{m_e}\frac{a_0}{|a_s|}d^2 \tag{1.1}$$

where a_s is the s-wave scattering length and a_{dd} is an effective dipole-dipole interaction length. M is the atomic or molecular mass, m_e is the electron mass, a_0 is the Bohr radius and d is the effective dipole strength, in atomic units. The effects of dipole-dipole interactions in a cold gas become dominant for values of ϵ_{dd} greater than 1. Both atomic and molecular gases are currently being investigated under the influence of magnetic and electric fields. Table 1.1 lists the values of the dipole-moments and the corresponding values for ϵ_{dd} for the various magnetic and electric field cases, providing examples for atoms and molecules.

Table 1.1: The strength of dipole-dipole interactions relative to s-wave scattering for three cases. The effective dipole strength in atomic units in each case is given by d, where α is the fine-structure constant, α_{at} is the polarisability of the atom, μ is magnetic dipole of the atom and μ_e is the induced electric dipole. The term $\alpha/2$ is the Bohr magneton, re-expressed in atomic units.

Case	d (atomic units)	Example	ϵ_{dd}
Magnetic	$\frac{\mu}{\mu_B}\frac{\alpha}{2}$	^{7}Li ($\mu = \mu_B,\ a_s = 26.9a_0$)	0.002
		^{52}Cr($\mu = 6\,\mu_B,\ a_s = 40a_0$)	0.36
Electric (weak-field limit)	$\frac{\alpha_{at}E}{ea_0}$	^{7}Li ($a_s = 5.1\,a_0$, $E \approx 0.8$ MV/cm and $\alpha_{at} = 40.8$ GHz/(MV/cm)^2	0.55
Electric (strong-field limit)	μ_e/ea_0	Molecule, $M = 50$ a.u., $\mu_e = 1D$, $a_s = 100a_0$	47

In a spin-polarised gas the magnetic dipole moment is independent of field strength and the dipole moment is dependent only on the spin state of the particle. The magnetic dipole-dipole interactions are unimportant for most atomic and molecular species available because $\epsilon_{dd} \ll 1$. However, one particular exception of note is ^{52}Cr which has a relatively large magnetic moment. ^{52}Cr has recently been demonstrated in a BEC state [8] and observed under an applied homogeneous magnetic field [9]. Despite a value of $\epsilon_{dd} < 1$ the effects of dipole-dipole interactions were observed. The long-range and anisotropic magnetic dipole-dipole interaction leads to an anisotropic deformation of the expanding chromium condensate upon release. This anisotropy depends on the orientation of the atomic dipole moments, which in

turn is determined by the aspect ratio of the trap.

Europium may also be a promising candidate for studying magnetic dipole interactions. It has a magnetic moment of $7\,\mu_B$, although its scattering lengths are unknown. Europium has already been buffer gas cooled using ^{3}He to a temperature of 250 mK, trapping as many as $1{\times}10^{12}$ atoms in a magnetic trap [10]. While laser cooling of Eu has not been experimentally explored, evaporative cooling alone could achieve sub-mK temperatures while retaining a substantial number of atoms.

Most molecules are unsuitable for the study of magnetic dipole interactions. Diatomic oxygen is one of the few molecules which is paramagnetic in the ground state. A supersonic beam has recently been decelerated using electromagnetic coils [11] in a analogous method to Stark deceleration. The magnetic moment for O_2, however, is only $1.8\,\mu_B$ giving a value for ϵ_{dd} of $\sim$0.03. For most other magnetically polarised molecules the value of ϵ_{dd} is even smaller. It is clear from these results that to reach values of ϵ_{dd} greater than 1 an electrically polarised gas is required.

For the electric field cases in the weak-field limit the Stark shift is proportional to the square of the electric field and thus the dipole moment grows linearly with the field strength. For atoms, the fields required to reach a value of ϵ_{dd} close to 1 still fall in the quadratic, weak-field regime. To reach a sufficient degree of polarisation atoms require extremely large applied electric fields . For ^{7}Li in the $|F = 2, M_F = 2\rangle$ ground state the scattering length is -26.9 a_0. This state is most commonly used to confine the atoms as it is a stretched state with the largest magnetic dipole moment. To obtain a significant value for ϵ_{dd} (>0.1) in this state, extremely large field strengths, greater than 1 MV/cm are required. It should also be noted that the negative sign for the scattering length means that a BEC in this state is unstable and will collapse above a critical atom number of $\sim$1400 atoms in accordance with many-body quantum theory [12, 13, 14]. ^{7}Li in the $|F = 1, M_F = -1\rangle$ ground state, however, has a scattering length of 5.1 a_0 and atoms in this state can be magnetically confined, albeit more weakly compared to the $|F = 2, M_F = 2\rangle$ state for a given field strength. To sufficiently polarise atoms in this state a field strength of only 0.8 MV/cm is required to reach a value of $\epsilon_{dd} = 0.55$. It should be noted that field strengths above 100 kV/cm are often considered unattainable in the lab, despite encouraging theoretical predictions. However, recent work presented in this thesis demonstrate that field strengths greater than 0.8 MV/cm can be achieved.

The simple analysis used here to estimate the required field strengths is confirmed by more sophisticated calculations [15] which also predict the presence of shape resonances at relatively weaker fields. For lithium in the $|F = 2, M_F = 2\rangle$ state a resonance is expected at 540 kV/cm and is many tens of kV/cm wide. Although this value is calculated at zero temperature, the results are not expected to change

significantly for atoms in the ultra-cold regime. These resonances allow us to tune the scattering length from negative to positive values, allowing one to realise a stable BEC that would otherwise collapse in the absence of these applied fields. An atomic BEC under electric fields thus provides a very direct method to study an array of interesting quantum phases in the alkali metals. To date an ultra-cold polar gas of this kind has not been realised.

For a cold atomic gas the electric fields required to reach the strong-field limit are unattainable. The dipolar nature of most heteronuclear molecular species, however, puts the strong-field limit within comfortable reach. In this limit the particles are fully polarised and the strength of the induced electric dipole, μ_e, is independent of field strength. For a typical fully polarised molecule we find a value for $\epsilon_{dd} \gg 1$ and the dipole-dipole interactions will dominate the collisional nature of the gas. The formation of ultra-cold molecular gases, however, is particularly challenging. A wide variety of experiments [16, 17] are currently aiming to create polar molecules in the cold and ultra-cold regimes, including buffer-gas cooling, Stark deceleration, and the photoassociation and Feshbach association of laser cooled atoms. However, neither buffer-gas cooling nor Stark deceleration have achieved temperatures below 100 mK, while the association methods are significantly limited by the choice of atoms that can be laser cooled. In addition, association typically creates molecules in higher ro-vibrational states that must be pumped into the ground state to avoid losses from the trap due to inelastic scattering events. In order to stitch ultracold atoms together in this way to create molecules, a magnetically- or optically-induced Feshbach resonance can be used. To date sub-mK temperatures have been achieved in the ground state for several heteronuclear molecules, including KRb, LiCs and RbCs via (optical) photoassociation [18, 19, 20]. The use of magneticlly-induced Feshbach-association has also been demonstrated [21]: If the atoms are paramagnetic the interaction potential can be tuned through the application of a magnetic field. A Feshbach resonance occurs when the energy of the unbound state (the atomic pair) and of the bound state (the diatomic molecule) become equal. By tuning the magnetic field through such a resonance, atoms can be adiabatically brought into a bound state. Typically the binding energy of a these associated molecules is very small, so methods have been developed to coherently transfer the molecules into a deeply bound state. One of the most efficient methods of transfer is to use a technique called stimulated Raman adiabatic passage (STIRAP) [22]; in simple terms two laser fields are used to couple the initial and ground states of the molecule with an intermediate state whose wavefunction strongly overlaps with both states. This technique can be used to transfer the molecules into the singlet or triplet electronic ground state of the ro-vibrational ground state with a very good efficiency

[21]. Despite such tremendous progress in the production of cold molecules in recent years however, to date no experimental realisation of such molecules under electric fields has been made.

One relatively recent avenue of research is the possibility of laser cooling molecules directly. Some diatomic molecules meet certain spectroscopic criteria that allow Doppler cooling methods to be applied in ways analogous to the successful techniques that have been employed for atoms. If a molecule can be found with highly diagonal Franck-Condon factors then relatively few laser frequencies are needed to form a closed cycling transition [23]. The application of radiative forces to molecules has recently been successfully demonstrated by DeMille *et al.* [25] by deflecting a molecular beam of strontium fluoride. Laser cooling of other molecules, including calcium fluoride, is also currently being investigated at Imperial College [26]. These molecules, amongst several others amenable to laser cooling, including CaH and BeH, are paramagnetic, opening up the possibility of realising a molecular magneto-optical trap [23, 24].

To achieve molecular temperatures similar to those achieved for the alkali metals the possibility of sympathetically cooling Stark-decelerated or buffer-gas cooled molecules using an ultra-cold atomic species as a refrigerant is being explored. By overlapping a large, dense sample of ultra-cold atoms with the trapped molecules the molecules will thermalise with the atoms through elastic collisions. It should be noted, however, that if one or both species are not in the absolute ground state inelastic collision rates may dominate and sympathetic cooling is not possible. Theoretical work on the ratio of such elastic to inelastic collisional rates for various species is a current active area of research, e.g. [27].

One additional possibility for achieving a polarised gas, not discussed above, is the use of atoms in large Rydberg states. The electrical polarisability of such atoms can be very large, scaling as n^2, where n is the occupied principal quantum number. This allows for large values of ϵ_{dd} to be obtained for relatively weak fields. Several problems, however, need to be addressed before a cold Rydberg gas is realised. Rydberg blockades limit the fraction of atoms in a gas that can be excited to a Rydberg state [28]. An irradiating laser beam resonant with the electronic transition of an atom will excite it. If the atom is driven to a high n-state the electric field generated by the induced dipole will interact significantly with neighbouring atoms, shifting them out of resonance with the laser and thus preventing their excitation. This is known as a Rydberg blockade. This process can be generalised to a many-atom ensemble, however, the picture becomes more complicated as many interacting dipoles create a band of energy levels. This may allow the Rydberg blockade to be circumvented via multi-photon excitations [29].

The initial model used to describe a Rydberg gas was based on the 'frozen-Rydberg gas' hypothesis [30] which treats the cloud as a stable ensemble of atoms in a Rydberg state. However, it was quickly recognised that such a gas is inherently unstable due to the dipole-dipole interactions between Rydberg atoms and that the gas may spontaneously form an ultra-cold plasma [31], something that itself is still of scientific interest. This inherent instability ultimately limits the study of dipole-dipole interactions in a polarised Rydberg gas.

The motivation for this thesis is to study an ultra-cold dipolar gas in order to explore the rich variety of quantum phases attainable and provide insight into some of the unsolved problems of condensed matter theory. It is of interest to the scientific community to create such a gas directly using an atomic ensemble under a very large applied electric field as well as through more indirect methods via the sympathetic cooling of molecular samples that can be studied under relatively weak fields.

The electric field strengths required to sufficiently polarise an atomic ensemble such that dipole-dipole interactions are dominant is approximately the same for all the laser-cooled alkali metals. Lithium was thus selected for this project for its suitability both for electric polarisation experiments and as a refrigerant for sympathetic cooling. In our group we have been studying the possibility of Stark decelerating lithium hydride as a molecular candidate for sympathetic cooling [32]. LiH is a natural choice for Stark deceleration because of its large dipole to mass ratio, while the relatively low number of electrons makes it more amenable to *ab initio* calculations. However, without a good knowledge of the potential energy surfaces of the atom-molecule interactions the likelihood of success of sympathetic cooling for a given set of species is almost completely unknown. The only parameter one can choose that is advantageous to the experiment is the atom-molecule mass ratio. In a gas mixture with a energy independent cross section, the average energy transferred per collision is $k_B \Delta T \xi$, where ΔT is the temperature difference and ξ is a scaling factor given by $\xi = 4m_{at}m_{mol}/(m_{at} + m_{mol})^2$ [33]. To minimise the number of collisions required for the two species to reach thermal equilibrium it is beneficial to chose a mass ratio close to unity, such that ξ is maximised. To this end, lithium, with a mass almost identical to LiH (and thus $\xi \approx 1$) was chosen as the atomic candidate most suitable to meet the aims of this thesis.

1.2 Aims

The aim of this experiment is to make significant progress towards the realisation of an ultra-cold dipolar gas, whose electric dipole-dipole interactions dominate over the usual s-wave scattering processes present in the low temperature range. To

this end we plan to trap a large cloud of neutral lithium atoms in the sub-mK regime in a magneto-optical trap for later use in two separate experiments; (1) as a refrigerant for the sympathetic cooling of light, polar molecules and (2) for experiments with electrically polarised, ultra-cold lithium. To date several lithium magneto-optical traps have been realised by other groups [34, 35, 36, 37]. The most common technique for providing a source of slow lithium atoms is to laser cool an effusive beam from an oven using Zeeman slowing techniques. This leads to substantial atom numbers in the MOT, exceeding 10^9 atoms at densities of 1×10^{11} atoms/cm^3, with temperatures below 1 mK. Further cooling techniques demonstrate that temperatures close to the Doppler limit are also obtainable. Once in the MOT the atoms can be loaded into a second trap to be transported to a new location to study their behaviour under large electric fields. The ability to efficiently transport our atoms is also an important step towards the sympathetic cooling of molecules, which requires the molecular and atomic traps to be spatially overlapped. To this end we plan to construct a mechanism for efficiently transporting a large sample of laser-cooled atoms over many centimetres while maintaining sub-mK temperatures.

The design of the experiment has been kept open to allow us to explore both the possibility of achieving an electrically polarised atomic gas directly through the application of very large electric fields or to use atoms as a refrigerant in the sympathetic cooling of cold polar molecules. A concurrent project is investigating the production of a cold source of both LiH and CH molecules with the intention of using our Li source as the refrigerant for sympathetic cooling. Much theoretical work is being carried out to calculate the potential energy surfaces for these atom-molecule interactions to assess the ratio of elastic and inelastic collisions and determine whether such cooling mechanisms are possible. This is being explored both specifically for the case of LiH and more generally to gain an understanding of the physics of atom-molecule and molecule-molecule scattering. We aim to further the work on the production of extremely large DC electric fields ($0.1 - 1$ MV/cm) already underway in the group to show that the realisation of an electrically polarised atomic or molecular gas is feasible in the near future. The experimental difficulties of placing cold atoms under such large electric fields will also be discussed.

1.3 Overview

The thesis is divided into three main chapters, each describing in detail the experiments performed to fulfill the main aims discussed above. Chapter 2 discusses the design and implementation of the lithium source, the Zeeman slower and the magneto-optical trap constructed to trap a large number of ultra-cold lithium atoms.

These were all constructed from the ground up, there being no previous experience of laser cooling lithium in the group. An overview of the apparatus that was built is shown in figure 1.2.

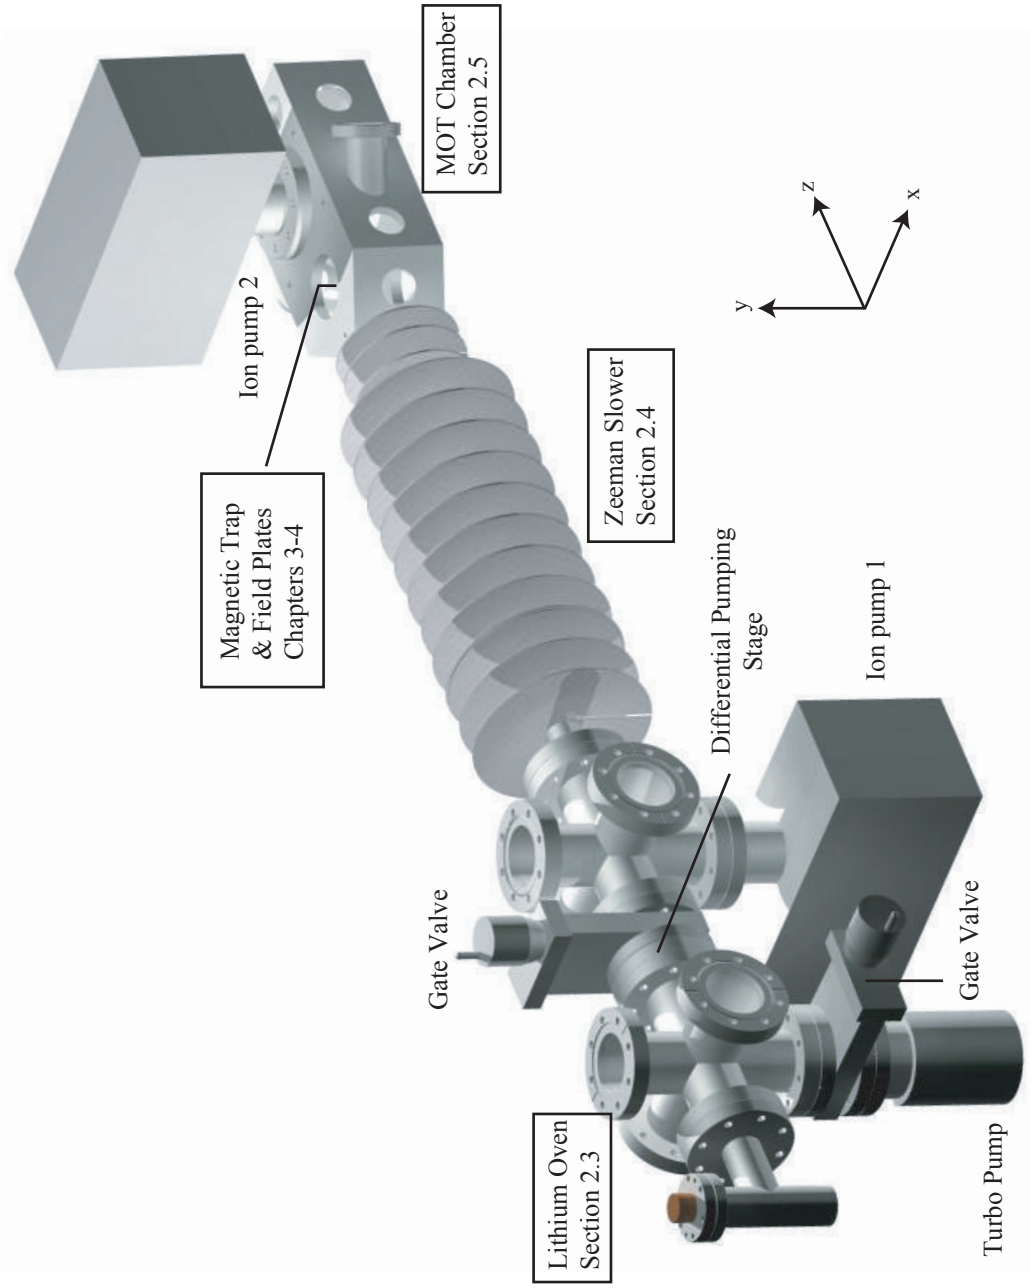

Figure 1.2: A schematic diagram giving an overview of the main experiment.

The experiment begins with a lithium oven which provides an intense beam of atomic lithium. This beam passes through a differential pumping tube and enters a Zeeman slower. The Zeeman slower decelerates a large fraction of the beam to a design velocity of approximately 55 ms^{-1}. This low velocity beam is then captured in a magneto-optical trap. All of these components are discussed in chapter 2 and

results are given to show how the atom number depends on various tunable parameters of the experiment such as the beam intensities and detunings. We also measure the temperature of the cloud. We have explored several techniques to determine the temperature and atom number and their merits are discussed here.

Chapter 3 discusses the results of a feasibility study designed to determine whether large enough electric field values can be created in order to sufficiently polarise an ultra-cold atomic cloud such that dipole-dipole interactions dominate the collisional physics of the system. The aims of these experiments were to investigate the production of large electric fields between two broad-area stainless steel electrodes and explore a number of surface treatment techniques designed to increase the strength of the electric field prior to electrical breakdown. An investigation into the effects of a number of polishing methods on the maximum field strengths attainable is performed and the results are documented in section 3.4. The chapter concludes with a section on the possibility of applying these fields across a cloud of ultra-cold lithium atoms. The chapter also discusses some of the recent work we undertook in the study of glass plates coated in an electrically conductive transparent medium, as an alternative material for the electrodes. The use of glass plates allows us to measure the degree of parallelism between the plates and their separation using laser light, thereby avoiding contact with the polished surfaces of the plates. The design for these plates is described in detail followed by the experimental results demonstrating the principle of achieving very parallel electrode plates with an accurately known separation. The final section provides an overview of the current electrode plate designs and the outlook for future designs is discussed.

Chapter 4 describes the methodology for loading the lithium atoms from the magneto-optical trap into a solely magnetic trap and transporting them to a second magnetic trap several centimetres away. The MOT is positioned directly downstream from the Zeeman slower and was positioned as close as possible the exit aperture of the slower to maximise the flux of atoms reaching the trapping region. The magnetic traps are formed by pairs of in-vacuum quadrupole coils. The first magnetic trap sits directly over the MOT to trap the maximum number of atoms. The second magnetic trap partially overlaps the first to aid the simplicity of the transportation mechanism. The field gradient of the second trap is steeper than the first to increase the compression of the cloud in preparation for placement between the narrowly-spaced electrodes. The design of the in-vacuum magnetic coils had to meet several stringent requirements due to the little available volume inside the chamber. The reasons behind the choice of the final design for both the coil formers and electrode mount are therefore discussed carefully in section 4.2. Due to a number of experimental difficulties the atom number in the MOT prior to loading

the magnetic traps is smaller than in the original MOTs created. Therefore, several methods were employed to increase the atom number loaded into the magnetic trap, the results for which are detailed through sections 4.3 to 4.6. The theory behind transportation and the trap transfer efficiency is then discussed and the experimental results are analysed in section 4.7. Some disagreement was noted between the theory and experiment and a discussion of these results can be found towards the end of this section.

The thesis concludes with a summary of the main results of this work and an overall outlook for the experiment, exploring the future for a polarised, ultra-cold atomic gas as well as the progress being made towards the realisation of sympathetically cooling a dipolar molecular species. Various improvements to the apparatus are also discussed in an effort to try and circumvent some of the problems highlighted within this thesis, particularly with regards to the transportation and polarisation of the atoms.

Chapter 2

Making a Li MOT

In the following chapter I will discuss the experimental apparatus designed and built for trapping a large cloud of neutral lithium atoms at temperatures approaching 1 mK. Section 2.1 provides a basic introduction into the theory and implementation of laser cooling. This is by no means an exhaustive review of the subject but provides an appropriate grounding for the following sections. Excellent pedagogical texts on the subject include [38] and chapters 1 to 8 of [39]. Sections 2.3 to 2.6 provide a tour of the experimental apparatus used to construct a lithium magneto-optical trap. Section 2.3 describes the laser system and optics used to provide the cooling and trapping radiation at 671 nm. Section 2.4 describes the lithium oven built to provide a continuous effusive beam of lithium to the Zeeman slower, documented in section 2.5. Finally, section 2.6 describes the machinery of the trap itself and discusses the characterisation of the lithium MOT cloud. Where appropriate, additional theoretical background has been provided to support the measurements and results communicated throughout this chapter.

2.1 A Brief Introduction to Laser Cooling

Over the last 30 years the field of laser cooling and trapping has flourished significantly. A number of sophisticated schemes underpinned by the interaction of atoms and molecules with laser light have been developed allowing physicists to understand and explore the subject of quantum mechanics and its influences in the ultra-cold regime. Experiments now routinely reach micro-Kelvin and even nano-Kelvin temperatures for a variety of atomic species leading to the formation of various quantum phases including Bose-Einstein condensates and Mott insulators. These provide a better understanding of quantum matter and have lead to the development of a number of applications including atomic clocks and progress towards quantum com-

puting.

The following section will revise the optical forces that can be applied to atoms and their applications in laser cooling and trapping. The physics will initially be described as a one-dimensional problem before being extended to three dimensions in the case of optical molasses. For a more detailed explanation of the results below, please refer to chapters 3, 5 and 7 of [39].

2.1.1 Radiation Pressure

The treatment below is a semi-classical approach applied to a two-level atom interacting with a near-resonant light field. In conformance with Ehrenfest's theorem, the force, F, on an atom is defined as the expectation value of the quantum mechanical force operator, $\mathcal{F}$, given by

$$F = \langle \mathcal{F} \rangle = \frac{\mathrm{d}}{\mathrm{d}t} \langle p \rangle \ . \tag{2.1}$$

Expressing the momentum operator in equation 2.1 as a function of the Hamiltonian describing the atom-light interaction we arrive at the result [39]

$$F = \hbar \left(\frac{\partial \Omega}{\partial z} \rho_{eg}^{*} + \frac{\partial \Omega^{*}}{\partial z} \rho_{eg} \right) , \tag{2.2}$$

where Ω is the Rabi frequency, ρ_{eg} describes the optical coherence between the ground and excited states of the atom and z is the direction of light propagation. This equation is valid for any light field applied to a two-level atom providing the electric dipole and rotating wave approximations can be made. This is a very general result that can be used to find the total scattering force for an atom in any situation for which the optical Bloch equations can be solved. To solve equation 2.2 it is useful to express $\mathrm{d}\Omega/\mathrm{d}z$ into its real and imaginary parts, equating to the gradient of the amplitude of the light field and the gradient of the phase, respectively. For the case of a travelling wave as present in typical laser cooling experiments the amplitude of the field is constant and only the term for the gradient of the phase remains, given by $ik\Omega$. The steady-state solutions of the Bloch equations for an atom at rest in a near-resonant light field then provide us with an expression for ρ_{eg} and the appropriate substitutions in equation 2.2 can be made. The formula for the scattering force applied to an atom undergoing absorption and emission is thus given by:

$$F = \frac{\hbar k s_0 \Gamma/2}{1 + s_0 + (2\delta/\gamma)^2} \ . \tag{2.3}$$

Here, s_0 is the saturation parameter given by $3I\lambda^3\tau/\pi hc$ and I is the intensity of the applied radiation. The wavevector, k, of the light field is related to the

wavelength by $k = 2\pi/\lambda$, δ is the detuning between the frequency of the light field and the atomic transition and γ is the linewidth of the transition. Equation 2.3 is a very intuitive formula; it is simply the momentum of a single photon multiplied by the scattering rate of the light by the atom. Light emitted via spontaneous emission is in a random direction and averaged over many scattering events will remove momentum from the atom in the direction of the radiation. However, as the intensity is increased the force saturates at a value of $\hbar k\gamma/2$. This is because stimulated emission, which increases with laser intensity, returns the momentum lost back to the light field.

2.1.2 Laser Cooling and Trapping

Laser cooling a two-level atom is simply a continuous application of the force described in equation 2.3. To begin with, consider an atom travelling along z, exposed to a counter-propagating, coherent radiation source, i.e. a laser. If one directs a laser at a frequency resonant with the transition of the two-level atom a quantum of its momentum, given by $\hbar k_z$, is removed from the atom along the vector, -z, upon the absorption of a photon, where k_z is the wavevector of the light. In the regime where stimulated emission can be neglected the atom will spontaneously emit a photon in a random direction and is ready to absorb another. For a typical beam the momentum lost per scattering event will be a very small fraction of the total momentum of the atom. In fact, one can show that for lithium travelling at a speed of 1 $\mathrm{km\,s^{-1}}$ over 10^4 scattering events on the D2 line are required to bring an atom to rest.

The changing Doppler shift experienced by the atom as it decelerates quickly brings the atom out of resonance with the light field, preventing us from scattering more photons. To compensate for this Doppler shift one must change either the laser frequency or the atomic transition frequency as a function of the atom's velocity. This requires either a time varying frequency or a position dependent one. The first method is known as chirp cooling [40, 41]; as a group of atoms decelerate the frequency of the laser is scanned at an appropriate rate to maintain the resonance condition. This, however, has the obvious disadvantage that it can only provide atoms in discrete packages rather than a continuous beam, reducing the loading rate of any trap. The second method changes the atomic transition frequency as a function of position. It achieves this by applying an electric, [42], or magnetic shift, [43], to the atoms, varying in magnitude along the deceleration path. In general the magnetic field sizes required to shift an atomic transition frequency by a given amount are easier to achieve in the laboratory compared to the electric field sizes required to Stark shift the energy level the same amount. Consequently Zeeman

slowing remains at the forefront in the design of most laser cooling experiments of this kind. The details of Zeeman slowing are discussed in further detail in section 2.5.

Equation 2.3 describes the force acting on an atom from a single light beam. In a 1-dimensional molasses the situation is described by two counter-propagating laser beams with equal intensity, frequency and polarisation. In the low intensity limit such that the beams are not coupled by stimulated emission, the net force is just the summation of the radiative forces with opposing k vectors. For an atom at rest this is clearly zero, however, for an atom with a velocity, v, there is an associated Doppler shift, $k.v$. For light detuned to the red of the transition this Doppler shift will bring the atom closer to resonance with the beam opposing its motion and further from the beam co-propagating with the atom. The total force is then simply the sum of the forces from each beam. Writing the total detuning to include the Doppler shift for each beam as $\delta \pm k.v$ and Taylor expanding the sum about $v = 0$ we find the net force to be

$$F_{net} \cong \frac{-8\hbar k^2 \delta s_0 v}{\gamma(1 + s_0 + (2\delta/\gamma)^2)^2} \ , \tag{2.4}$$

for which terms of order $(kv/\gamma)^3$ and higher have been neglected. As such, there is a velocity dependent force cooling the atom, as illustrated in figure 2.1.

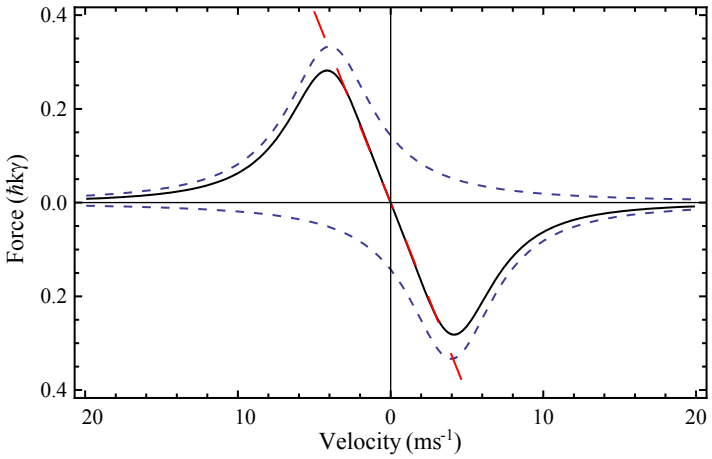

Figure 2.1: The damping force applied to an atom in a 1D molasses for a detuning of γ and a saturation parameter $s_0 = 2$. The dashed blue lines are determined from equation 2.3 and the black solid line is their complete sum. The red dashed line is given by equation 2.4

A molasses of this kind will not reduce the momentum of the atoms indefinitely due to heating contributions from the lasers. After emission the atom will receive a momentum impulse equal to the momentum carried by the photon. This leads to a steady-state temperature limit known as the Doppler temperature. The temperature can be derived by equating the heating and cooling rates for these processes

and a minimum for the steady-state kinetic energy of the atom can be found (see section 4.4) at a detuning of $\delta = -\Gamma/2$. The final result for the minimum Doppler temperature is

$$T_D = \frac{\hbar\Gamma}{2k_B}. \tag{2.5}$$

While many methods have been employed to trap neutral atoms none have been employed more frequently and successfully than the magneto-optical trap (MOT), particularly for the alkali metals. The first MOT was implemented by Raab *et al.* [44]; the basic techniques described here have changed very little since then.

The magneto-optical trap is a hybrid trap employing both optical and magnetic fields. The radiation pressure is provided by a three dimensional molasses, with each axis contributing a damping force as described by equation 2.4. This alone, however, is not sufficient to trap the atoms as there is no restoring force towards the centre of the trap. To achieve this an inhomogeneous magnetic field is applied by a pair of anti-Helmholtz coils centred about the trap. Across the trapping region this provides a constant field gradient with a zero-field point at the centre. Note that the gradient of the field along the coils' axis is twice the gradient in the radial direction. To describe the light-atom interaction in this picture I will reduce the problem again to one-dimension along z. I will also simplify the atom to a two level atom whose total angular momentum quantum numbers are $F = 0$ and $F = 1$, in the ground, $|g\rangle$, and excited state, $|e\rangle$, respectively. If an atom in state $|g\rangle$ moves from the origin along positive z the degeneracy of the upper M_F values is lifted such that the $|+1\rangle$ state is at a higher energy and the $|-1\rangle$ state at a lower. If we detune the laser frequency, ω_ℓ, to the red, relative to the $M_F = 0$ state and apply a σ_- polarisation to the counter-propagating beam and a σ_+ polarisation to the co-propagating beam the atom will couple much more strongly to the beam opposing its motion than the other.

This can be seen in figure 2.2: for an atom at positive values of z the detuning from the $M_F = -1$ level, δ_-, will be smaller than from the $M_F = +1$ level, δ_+, and the σ_- beam will provide greater light pressure. Equally, for negative z values, $\delta_+ < \delta_-$ and the σ_+ beam will restore the atoms towards $z = 0$. Since the field magnitude is position dependent the force will act as a restoring force towards the centre. The problem can be expanded to three-dimensions without further complication and applied to atoms with larger F numbers in the ground and excited states.

For multilevel atoms the situation for cooling and trapping is slightly more complex. The D2 transition for Li for example involves a doublet ground state and four unresolved hyperfine levels in the upper state, amounting to a gallimaufry of 24 M_F levels with significant mixing even at relatively small fields. The general principles,

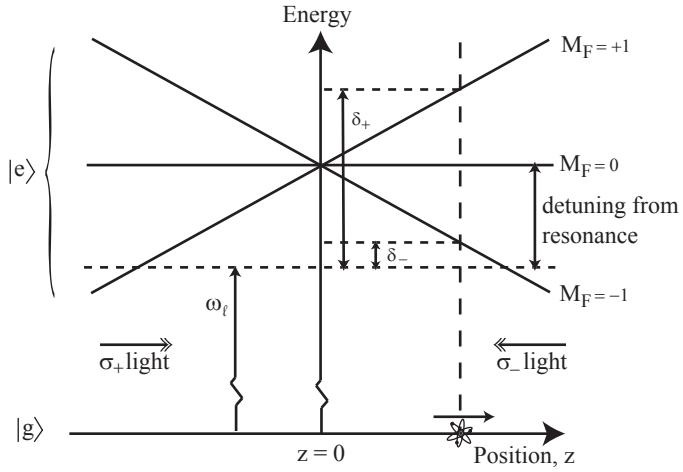

Figure 2.2: An energy level diagram portraying the mechanism through which a 2D optical molasses in a quadrupolar magnetic field can provide a position-dependent restoring force to the zero-field point.

however, may still be applied providing a closed transition can be maintained. For most alkalis, including Li, this requires a repumper laser to restore atoms that have entered the non-resonant hyperfine level back into the cooling cycle.

The capture process in a MOT is complicated and will depend on the trapping beam radius, detuning and intensity and the magnetic field gradient. The general result, however, is that for a given set of parameters there is a maximum capture velocity above which atoms will not be trapped. For most alkalis this capture velocity ranges in the region of a few tens of metres per second. If we are trapping atoms from a thermal background it is obviously beneficial to have as large a capture velocity, v_c, as possible since the number of trappable atoms scales as v_c^4 for atoms caught from the low velocity tail of the Maxwell-Boltzmann distribution. For the case of a Zeeman slower we have the freedom to choose the maximum velocity of atoms entering the trapping region and need only ensure that the capture velocity of the MOT is greater than or equal to this. The capture velocity of the MOT may be serendipitously enhanced because a traversing atom will experience a decreasing magnetic field in red detuned light just as in Zeeman deceleration. In the simplest case then, if we assume any atom decelerated to rest within the trapping volume, of diameter $\mathrm{d} \sim 10 - 20$ mm, will be captured then the capture velocity is given by $v_c = \sqrt{2ad} \approx 300\ \mathrm{ms}^{-1}$, where $a = \frac{\hbar k \Gamma}{2m}$ is the maximum deceleration. This value will typically be a *vast* over-estimate of the true capture velocity. To calculate a more realistic value, parameters such as the laser detuning, intensity, spot size and polarisation; the atom trajectory; the true magnetic field profile and the multi-level nature of the atom must all be taken into account. To avoid such a complicated numerical simulation we will make use of the experimental findings given in [37]. The experimental apparatus described here matches closely with our set-up given in

chapter 1; the authors found that for detunings of 2 linewidths the capture velocity of the MOT was measured to be approximately 35 ms^{-1} and increases to 50 ms^{-1} for a detuning of 5Γ. This provides an upper-limit for the designed exit-velocity of the Zeeman slower.

2.1.3 Sub-Doppler Cooling

Sub-Doppler cooling techniques allow us to break the temperature limit given by equation 2.5 for optical cooling. The phenomenon was originally discovered by P. Lett *et al.* [45] and later theoretically described by Dalibard and Cohen-Tannoudji [46] in the late 1980s. For lithium these techniques do not work well because of the strong coupling between the excited state and the $F = 1$ component of the ground state. To understand the reasons for this a brief description of the cooling mechanism is given here.

In a scenario where two counter-propagating beams of equal intensity and frequency are carrying orthogonal polarisation, such as along the axis of a MOT, a position-dependent polarisation gradient is set-up repeated with a period of half a wavelength. The cooling mechanisms can be categorised into two types, those for which the beams are composed of two orthogonal linearly polarised beams (lin$\perp$lin) and those for which they are circular. In the first case, as we move along the beam, the polarisation of the total field rotates from linear (45 degrees relative to the two beam polarisations), to σ_-, back to linear (orthogonal to the original polarisation) and finally to σ_+ before repeating. An atom in this field will be optically pumped into a specific magnetic sub-level depending on the polarisation of the field. In addition, each sub-state will experience a different light shift, described by the Stark effect in a light field, depending on the local polarisation of the field. For example, as shown in figure 2.3, an atom in a sub-level with a positive M_F value starting at a position of linear polarisation will climb a potential hill as it approaches a point where the field is σ_- polarised and the light shift is largest. Here it will be repumped into a sub-level with a negative M_F value whose light shift is opposite to that in the presence of σ_- light. The atom now continues to climb a second potential as it travels towards a region of σ_+ light which will again redistribute the population of the magnetic sub-levels. Consequently, an atom will find itself continually climbing a potential as it moves towards any beam providing an additional damping force to the usual Doppler cooling. This is known as Sisyphus cooling.

In the case of a MOT where the polarisation of the beams is circular and counter-propagating beams have the same helicity (relative to their direction), the polarisation of the total light field for a particular beam axis is linear everywhere, but the angle rotates through 2π during the length of one wavelength. Unlike in the

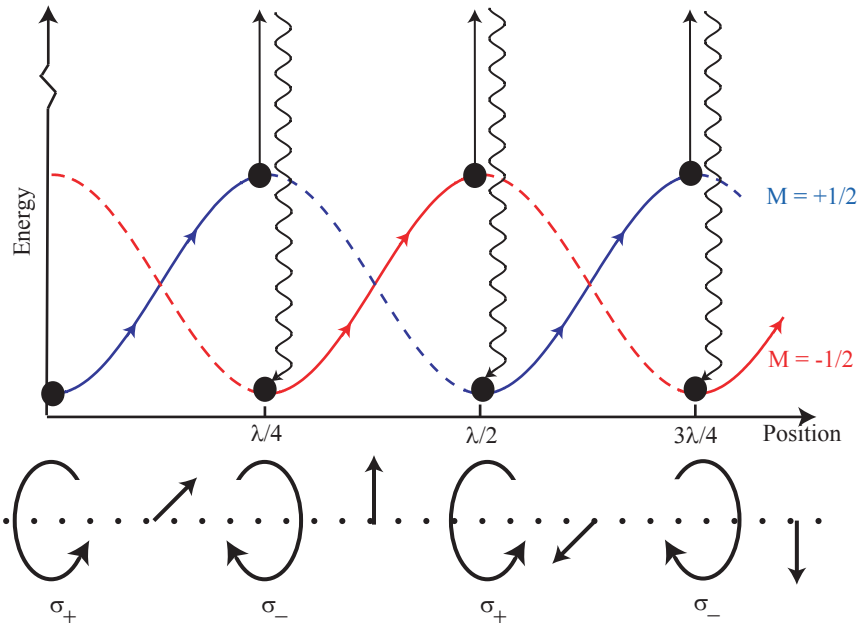

Figure 2.3: Sisyphus cooling in the lin⊥lin configuration: A ground state atom will follow the solid line, continually relinquishing its kinetic energy as it climbs each potential hill. The frequency of absorption is less than the frequency of spontaneous emission so energy is always lost. The light shift is dependent on the local polarisation of the light.

case of Sisyphus cooling, however, there is no polarisation-dependent light shift for the ground state magnetic sublevels. An atom moving through a rotating light field sees a rotating quantisation axis and must be optically pumped in order to follow it. As a consequence the population of the ground state magnetic sub-levels does not adiabatically follow the local polarisation direction but experiences a lag [46]. This non-adiabaticity can be shown to favour the population of more positive M_F values for atoms moving towards the σ_- polarised beam and contrariwise for atoms opposing the σ_+ beam. Again, the atomic motion in this scenario is damped. There are limits to sub-Doppler cooling imposed by the recoil the atom receives at the point of emission of a photon. This simple association leads to a minimum temperature dictated by the momentum of the emitted photon, i.e. $k_B T_{min} = (\hbar k)^2/2M$. Interestingly, the cooling method described above requires the atom to be localisable to within the scale of $\sim 2\pi/\lambda$. The uncertainty principle then requires that the atom has a momentum spread of at least $\hbar k$. These arguments are verified and discussed in detail in [47]. Several methods exist to cool below the recoil limit using sophisticated techniques to manipulate the atoms with laser light, including Raman cooling, but are beyond the scope of this discussion.

In a magneto-optical trap the atoms sit at the intersection of six beams each carrying circular polarisation in the presence of a magnetic quadrupole field. While there is clearly some degree of gradient in the polarisation the details are not obvious

in three-dimensions. Reports of sub-Doppler temperatures, however, are common in MOTs for other alkali species such as Na and Cs. The absence of any such observation for lithium is because the effects of the redistribution of the ground state magnetic sub-levels are much weaker due to the strong coupling to the neighbouring hyperfine state, thereby diluting the population. This has been theoretically shown for lithium [48], and also for potassium [49] whose atomic structure is similar. These theoretical results are consistent with the experimental results [37, 49]. To achieve sub-Doppler temperatures for lithium alternative techniques such as evaporative [50] and sympathetic cooling [51] must be employed.

2.2 Lithium

Lithium is the lightest of all the metals and has been selected for this project for both its suitability for the sympathetic cooling of light molecules and for the study of ultra-cold atomic gases under large electric fields. The properties for bulk and atomic ^{7}Li are presented in table 2.1. Both lithium isotopes are naturally abundant; for these experiments we have chosen to use bosonic ^{7}Li atoms to allow for the option of creating a Bose-Einstein condensate in the future.

Table 2.1: Physical properties of bulk and atomic lithium

Property	Symbol	Value	Unit
Isotope abundance	$^6\mathrm{Li}:{}^7\mathrm{Li}$	7.5 : 92.5	%
Density (300 K)	ρ	534	$\mathrm{kg\,m^{-3}}$
Melting Point	T_M	453.7	K
Boiling Point	T_B	1615	K
Enthalpy of Vaporisation	h	20845.7	$\mathrm{kJ\,kg^{-1}}$
Polarisability	α	40.8	$\mathrm{GHz/(MV/cm)^2}$
Atomic Number	Z	3	
Neutron Number	N	4	
Atomic Mass (u)	m_a	7.016	
Nuclear Spin	I	3/2	

To laser cool and trap ^{7}Li a closed transition must be maintained; this is most easily attainable on the D2 transition. The optical properties associated with this transition are given in table 2.2.

Table 2.2: Properties of the ^{7}Li D2 transition

Property	Symbol	Value	Unit
Wavelength	λ	670.961	nm
Frequency	ν	446.8094	THz
Lifetime	τ	26.87	ns
Natural Linewidth	γ	5.92	MHz
Saturation Intensity	I_s	2.56	mW cm^{-2}
Absorption Cross-section	$\sigma = 3\lambda^2/2\pi$	2.15×10^{-9}	cm^2
Doppler Temperature	T_D	142.1	μK
Doppler Velocity	v_D	41.04	cm s^{-1}
Recoil Temperature	T_r	6.06	μK
Recoil Velocity	v_r	8.47	cm s^{-1}

2.3 The Laser System

To laser cool and trap lithium atoms in a magneto-optical trap a narrow-bandwidth laser source operating at a wavelength of 671 nm with at least 100 mW of optical power is required. The most common lasers used for this purpose are dye-lasers and diode lasers. The energy level diagram for ^{7}Li is given in figure 2.4, highlighting the hyperfine states relevant to the D2 transition.

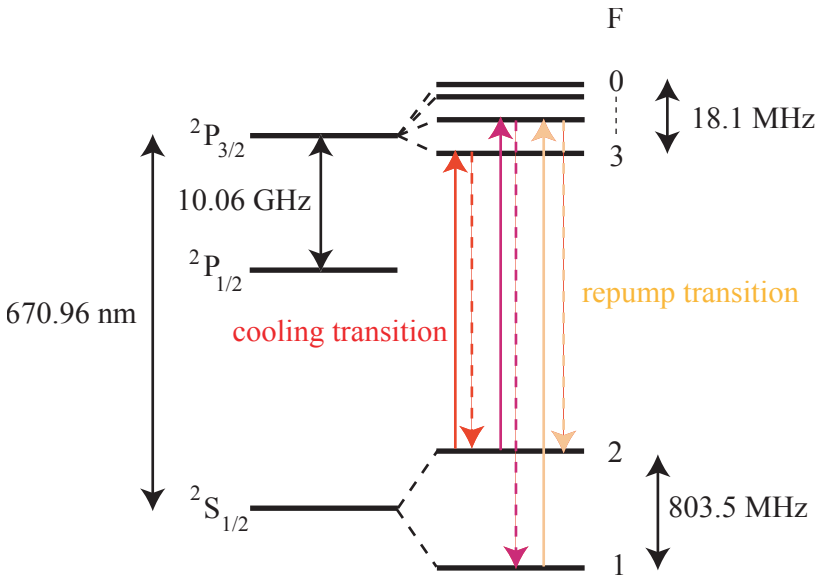

Figure 2.4: The energy level diagram for ^{7}Li, highlighting the cooling and repump transition frequencies required for laser cooling.

The $2s^2S_{1/2}(F = 2) \rightarrow 2s^2P_{3/2}(F' = 3)$ transition (red) is closed and thus most suitable for laser cooling. However, with a natural transition linewidth of 5.92 MHz

the hyperfine states of the upper level are partially unresolved and simultaneous excitations to neighbouring hyperfine states can occur (magenta). Consequently decay routes to the $^2S_{1/2}(F = 1)$ are opened, requiring a repump beam (yellow) to maintain a closed transition. In the case of lithium, the coupling from the excited state to the $F = 1$ component of the ground state is relatively strong in comparison to the $F = 2$ state. The branching ratio of the coupling strengths from the $2s^2P_{3/2}$ state to the two hyperfine ground states is calculated to be approximately $2 : 1$. As a result the repump beam requires as much as 50% of the total cooling power to close this transition. A closed transition can be maintained by applying circularly polarised light to cycle between either the $|F = 2, M_F = 2\rangle$ and $|F' = 3, M'_F = 3\rangle$ or $|F = 2, M_F = -2\rangle$ and $|F' = 3, M'_F = -3\rangle$ transition, however, in practice obtaining a pure enough sample of a given polarisation is very difficult and decay routes to the $F = 1$ ground state will always occur.

The transition frequencies required for the Zeeman slower and the MOT typically will be different and it is useful to have independent control of each set of beams. Without being able to form a closed transition two independent repump beams are also required for cooling and trapping, amounting to four frequencies necessary to operate the experiment; additional frequencies may also be required for imaging. While obtaining four low-powered lasers at 671 nm for each transition is possible [37], ensuring they are all locked to the correct frequency relative to each other can become complicated. To circumvent this we use a single laser and obtain the other frequencies through acousto-optical modulation. The hyperfine splitting in the $^2S_{1/2}$ state is 803.5 MHz which can be bridged with one or more acousto-optical modulators (AOMs). A disadvantage of using acousto-optical modulation is the conversion efficiency; for a single pass through a good quality AOM at 200 MHz, only 70% of the light will be converted to the correct frequency and often double passes are required to reach the appropriate frequencies. Given this low efficiency, and the number of AOMs required we find that a total power of several hundred milliWatts is required for collecting a large number of atoms.

For most of the experiments reported here we used a TOPTICA TA100 tapered amplifier laser system providing 500 mW at 671 nm. However, when setting up the experiment, we did not have this laser and instead a COHERENT 699 dye laser was used for testing the laser-locking system and for performing some of the initial atomic beam characterisation experiments. The dye of choice at 671 nm is DCM (4-dicyanomethylene-2-methyl-6-p-dimethylaminostyryl-4H-pyran). The dye was pumped by 5 W of laser power from a COHERENT VERDI V5 solid state laser. The dye laser provided up to 350 mW of power at 671 nm with fresh dye. However, as with most dyes, this initial power would decay falling to as low as 80 mW within

two weeks and regular dye changes were necessary.

The first and most important task for any laser system is to lock the laser to provide a stable, narrow bandwidth source. The following section outlines the locking scheme used to keep the laser frequency fixed to a precision of better than 1 MHz.

2.3.1 Frequency Stabilisation

The aims of these stabilisation experiments were to achieve a laser frequency as accurately matched to a lithium transition as possible, providing a stable lock with sub-MHz precision. This provides a suitable repeatability for our experiments, essential for laser cooling and trapping. It is also beneficial for the locking system to be very easy to use and require little to no day-to-day management, particularly when the experiment is running in full operation. For these reasons, and because there is a considerable amount of experience in the group employing this method we chose to use polarisation spectroscopy for our locking scheme. However, upon a literature review it came to light that there were no reports of any group previously implementing polarisation spectroscopy for a lithium cell, possibly due to the difficulties arising from the lack of a closed cycling transition. For this reason we chose to begin with saturation spectroscopy as a proof of principle before exploring the implementation of polarisation spectroscopy.

Saturation Spectroscopy

Saturation spectroscopy is an invaluable technique used to perform Doppler-free measurements of atomic and molecular gases. The technique has been used extensively in the field of spectroscopy and is frequently employed to provide a reference point to lock a laser to.

To perform saturation spectroscopy a pump beam of at least saturation intensity is passed through a vapour cell and scanned about the resonance of an atomic transition. A second much weaker probe beam of the same frequency is then passed through the cell, counter-propagating to the pump beam, and measured on a photodiode as a function of frequency. For comparison, a reference beam also passes through the cell, but does not overlap with the pump beam. This reference beam allows the small Doppler-free signal to be isolated from the much larger Doppler-broadened background, and provides common mode rejection of intensity noise, pressure and temperature fluctuations and mechanical vibrations. Figure 2.5 depicts the set-up used for our experiment.

As discussed later in section 2.4 the vapour pressure for lithium at room temperature is negligible and so the lithium must be heated to achieve an atomic vapour

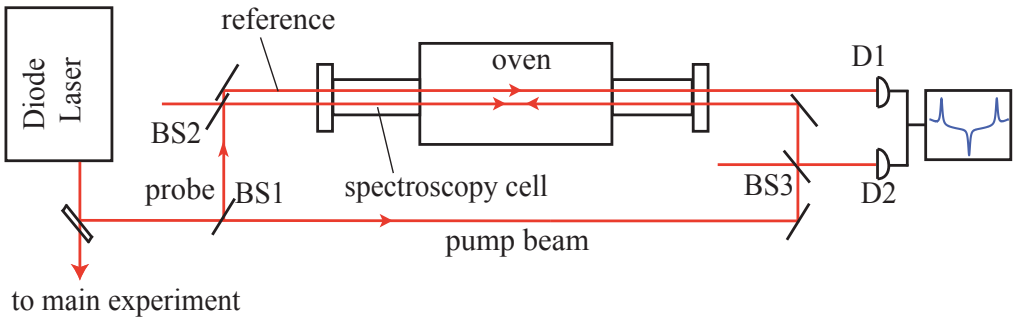

Figure 2.5: The experimental apparatus used to perform saturation spectroscopy in a lithium cell. A 10% beam splitter is used to divide a beam into counter-propagating pump and probe beams. A reference beam is passed through the cell but does not overlap with pump beam. This allows us to subtract the background signal and reduce the signal to noise ratio. The signals are recorded on a pair of photodiodes and subtracted electronically.

suitable for spectroscopic measurements. The spectroscopy cell, sketched in figure 2.5, is a cylindrical steel chamber of length 150 mm and diameter 52 mm, adjoined by two 100 mm long arms, 40 mm in diameter, one at each face of the chamber. Each arm is terminated by a DN40CF viewport for optical access. The central chamber is encased in an insulating, cylindrical brick laced medially with a current carrying wire to heat the oven. The wire is wound in a zig-zag pattern along the length of the cylinder covering the whole circumference. This pattern minimises any uncontrolled magnetic field bias to the atoms. The adjoining arms are wrapped in 6 mm diameter copper piping to cool the walls to encourage condensation of lithium on the walls before reaching the windows. If allowed to react for a sufficient length of time lithium will permanently reduce the transparency of the quartz windows. Operating temperatures in the oven are typically around 750 K providing a pressure of 5-10$\times 10^{-3}$ mbar in agreement with the expected vapour pressure for lithium. These parameters yield an absorption of approximately 50% of a 400 μW on-resonant beam. The pump beam carries a power of approximately 4 mW during operation. Vacuum pumping is initially provided by a 6 $\mathrm{l\,s^{-1}}$ rotary vane roughing pump, but this is closed off from the cell by a valve once the pressure is stable without pumping.

In the absence of the pump beam, the absorption profile of the probe beam, as measured by the photodiode is a Doppler-limited profile. At 750 K the full width at half maximum of this profile is approximately 3.3 GHz. The laser frequency is tuned to scan about the point of maximum absorption and the pump beam is introduced. With the counter-propagating pump beam present the absorption profile is unaffected as each beam interacts with a different velocity class, except, at those

frequencies corresponding to the field-free, stationary atomic transitions, i.e. with no velocity component in the direction of the beam. If the laser is resonant with the F=1 component of the D2 line then only the atoms with a Doppler shift smaller than the natural linewidth of the transition, i.e. moving almost perpendicular to the beams, will interact appreciably with both beams simultaneously. The higher intensity pump beam excites a large proportion of the atoms in this small velocity class into the excited state and thus there will be a drop in the population of the $F = 1$ state. This population is further reduced by optical pumping into the $F = 2$ state. Consequently there is a further drop in the absorption, resulting in a Doppler-free, Lamb-dip at the atomic resonance. This dip in the profile is Lorentzian in shape, whose width is approximately 18 MHz, three times the width of the natural linewidth of the transition due to the unresolved nature of the excited state hyperfine levels. A similar scenario occurs if the laser is tuned to the $F = 2$ component of the D2 line, depleting the $F = 2$ population and producing another Lamb-dip. The polarisation of the light and the exact direction of the local magnetic field are unknown so it assumed that the atoms can be excited to any hyperfine state that obeys the $\Delta F = \{0, \pm 1\}$ selection rule.

There is also an additional Doppler-free feature observed, lying midway between the $F = 1$ and $F = 2$ transition frequencies of the D2 line, known as a cross-over resonance. At this frequency some atoms in the $F = 2$ state will have the correct Doppler shift to be resonant with the D2 transition and will be optically pumped into the $F = 1$ state. These atoms in the $F = 1$ state will now be Doppler-shifted to be resonant with the counter-propagating probe beam and the increase in the $F = 1$ population will lead to an increase in the amount of absorption. In addition there will be a second velocity class, resonant with the $F = 1$ component that will be excited to the $F = 2$ state, which will also be resonant with the probe beam. This produces a Doppler free signal that is inverted [1] relative to the other two features and typically will have the largest amplitude of the three. This is because the Doppler broadened width is greater than the ground-state hyperfine splitting allowing the probe beam to interact with both the F=1 and F=2 states

[1] Those familiar with saturation spectroscopy in rubidium will note that the sign of the cross-over resonance signal is not inverted relative to the neighbouring Lamb-dips. In rubidium the ground-state hyperfine levels are separated by more than the Doppler width. Consequently, cross-over resonances due to interplay between the ground-level hyperfine states are negligible. In contrast to Li, however, the excited-state hyperfine components of Rb are separated by more than the natural line width and thus resolvable. This allows the pump laser at the cross-over resonance frequency to interact with two excited hyperfine components rather than one, thereby further depleting the population of the ground state resonant with the light and thus appears as a decrease in the absorption profile of the probe laser.

at the cross-over frequency, leading to an increase in the probe-beam absorption profile. By subtracting the absorption profile of this probe beam from that of the reference beam, a set of low-noise, power-independent, narrow linewidth signals can be produced, as recorded in figure 2.6. The pump and probe beams used to record the data below intersected at a small angle, rather than perfectly counter-propagating. Thus the linewidth of the spectroscopic features are broadened from the expected 18 MHz to approximately 40 MHz. If the pump beam is of high intensity power-broadening effects are also observed.

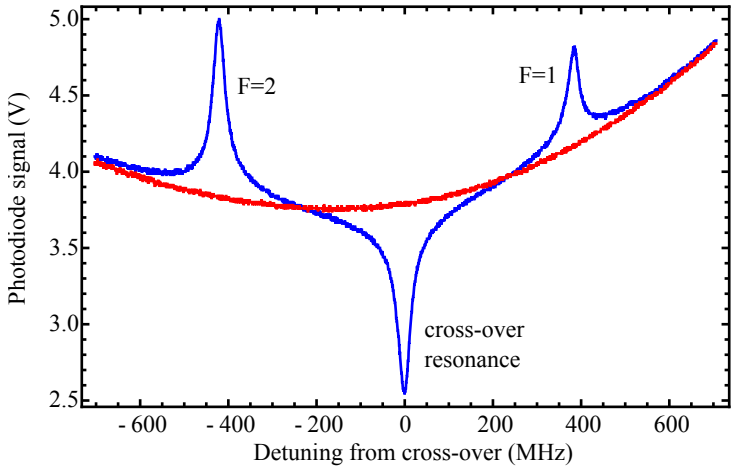

Figure 2.6: Saturation spectroscopy data showing the ground state hyperfine features including the cross-over for the D2 transition in ^{7}Li. The blue curve is measured with the pump beam present and the red curve without.

The left hand peak corresponds to the F=2 hyperfine transition red detuned from the cross-over resonance, centre, while the right hand peak is the F=1 transition. Note that the upper state hyperfine levels are not resolved. The frequency axis in figure 2.6 is not a measured parameter, but inferred from the hyperfine splitting in the ground state of 803.5 MHz, determined experimentally by previous groups' work. The height of the two hyperfine peaks is approximately in the ratio 5:3. This ratio would be expected if one assumes that the population of the atoms is evenly distributed amongst the 8 M_F levels in the ground multiplet. The size of the dip at the cross-over is larger than the heights of the two peaks, as expected since two velocity classes are involved in the formation of this spectroscopic feature. The red curve is the signal in the absence of the pump beam. This background Gaussian profile is a result of the background Doppler broadening which has not been subtracted here.

To lock the laser at this frequency the laser is dithered about the peak of the Lamb-dip and the resulting signal is differentiated using standard electronics. The derivative is zero at the centre of the Lamb-dip. An electronic feedback circuit is

then used to maintain the frequency about which the laser is dithered at the zero crossing and thus the atomic transition frequency. The amplitude and frequency of the dither determine the precision and stability of the lock, depending on the nature of any background noise. The frequency is typically accurate to within 1 MHz of the desired frequency and is stable to within 1-2 MHz. Once the saturation spectroscopy was working well the feasibility of using polarisation spectroscopy was investigated.

Polarisation Spectroscopy

Polarisation spectroscopy was first proposed and demonstrated by Wieman and Hänsch as an advancement of saturation absorption spectroscopy [52]. The technique is based on light-induced birefringence and dichroism of an absorbing gas, rotating the polarisation of an on-resonant probe beam. It is related to saturation spectroscopy and is very similar in set-up, but has a number of advantages. In general it is reported to offer a considerably better signal-to-noise ratio and sub-MHz precision without the need to dither the laser frequency. For these reasons we were keen to implement the technique for lithium. The set-up used for our experiment is given below in figure 2.7.

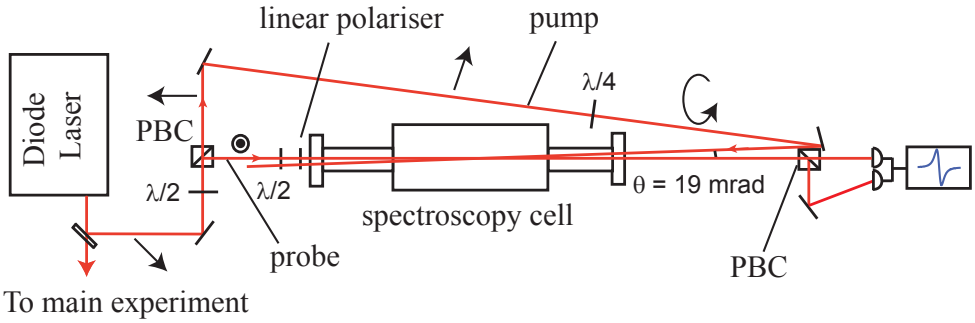

Figure 2.7: The experimental apparatus set up to perform polarisation spectroscopy of lithium. A rotation in the polarisation of the probe beam is translated into an intensity difference at a polarising beam cube (PBC) and measured by a pair of photodiodes.

As with saturation spectroscopy the output from the laser is split into a weak probe beam and a stronger pump beam using a half waveplate ($\lambda/2$) and a polarising beam splitting cube (PBC) to provide greater control over the power ratios. The output beams from the PBC are linearly polarised and orthogonal to one another. The Doppler-free nature of the spectroscopy is understood by the same phenomena operating in saturation spectroscopy. The linearly polarised probe beam is passed through a $\lambda/2$ waveplate, the spectroscopy cell and finally split at a second PBC. The polarisation of the probe is rotated by the waveplate such that equal intensities are measured at the output of the two arms from the second PBC. The two output

beams are focussed onto a pair of photodiodes whose signals are subtracted from one another and the difference is amplified. This set-up is similar to that used by M. L. Harris *et al* [53]. Note that this method provides twice the sensitivity than the original set-up described in [52] in which the beam splitting cube is replaced by a polariser rotated perpendicular to the probe beam polarisation and only one photodiode is employed.

Without the presence of the pump beam the resulting signal voltage will be zero, independent of frequency or laser intensity. Although the signal is intensity and frequency independent only very small changes in the polarisation are required to inflict significant noise on the measurement. The polarisation of the laser source is very stable, but small vibrations from the table couple to the $\lambda/2$ plate and modulate the polarisation. To try and counteract this a linear polariser was placed between the plate and the cell such that any polarisation changes are almost completely converted into intensity changes. This significantly increases our signal-to-noise ratio.

The pump beam is passed through a $\lambda/4$ waveplate to obtain circularly polarised light and passes through the cell almost completely counter-propagating to the probe beam to maximise the region of overlap between the two beams. The pump beam is expanded to twice the diameter of the probe to further increase the overlapping volume. At the resonance frequency the atoms are excited on the transition $|F, M_F\rangle \rightarrow |F', M_F'\rangle$ by the pump beam. The quantum number, M_F, that describes the projection of $\boldsymbol{F}$ onto the direction of light propagation follows the selection rule $\Delta M_F = +1$ for transitions induced by σ_+ polarised light. Over several cycles the population of the degenerate M_F levels of the ground state become imbalanced, depleting the lowest M_F values and populating the highest. This anisotropic distribution changes the polarisation state of the incident linearly polarised probe beam. This can be understood by expanding the linear polarised probe beam into a superposition of σ_+ and σ_- polarised light. When resonant with such an ensemble of atoms whose M_F state populations are favoured towards more positive M_F states, the σ_+ light is absorbed less than the σ_- light. Consequently, the plane of polarisation of the exiting beam is slightly rotated creating an imbalance in the intensities landing on the pair of photodiodes.

The degree of imbalance across the M_F states depends on the pump intensity, the absorption cross-section, the transit time of atoms in the laser beam and any possible relaxation processes that may repopulate lower F or M_F states. These relaxation processes are particularly strong in lithium due to the unresolved transition lines in the excited state of the D2 line. To fully understand the line profiles of the polarisation signals a more in-depth analysis of the processes involved is required. An overview is given below and a detailed analysis can be found in [53, 54].

The polarisation of the linear polarised probe beam is rotated by a half-wave plate and enters the cell at an angle, ϕ_0 relative to the horizontal axis of the beam splitter on the far side of the cell. The electric field of the beam is thus given by

$$\underline{E} = E_0 \cos\phi_0 \ \underline{\hat{x}} + E_0 \sin\phi_0 \ \underline{\hat{y}} \ . \tag{2.6}$$

However, it is easier to understand the interaction of the linear polarised beam with the lithium by writing the electric field as a superposition of σ_+ and σ_- circularly polarised waves. The basis sets for the left and right circular components can be written as $\underline{\hat{\epsilon}}_\pm = \underline{\hat{x}} \pm i\underline{\hat{y}}$. The two components experience different absorption coefficients, α_+ and α_-, and different refractive indicies, n_+ and n_- due to the unbalanced level populations induced by the pump beam. After an interaction length, L, along the z axis the electric field components after the cell can be expressed as

$$\begin{aligned} \underline{E}_+ &= \frac{1}{2} E_0 \ e^{i\phi_0} \ e^{i[\omega t - k^+ L + i(\alpha^+/2)L]} \ \underline{\hat{\epsilon}}_+ \ , \\ \underline{E}_- &= \frac{1}{2} E_0 \ e^{-i\phi_0} \ e^{i[\omega t - k^- L + i(\alpha^-/2)L]} \ \underline{\hat{\epsilon}}_- \ , \end{aligned} \tag{2.7}$$

where $k^\pm = n^\pm \omega/c$. The total transmitted amplitude, $\underline{E}_+ + \underline{E}_-$, reaching the polarising beam splitter after the cell is subsequently divided into its $\underline{\hat{x}}$ and $\underline{\hat{y}}$ components

$$\begin{aligned} \underline{E}_x &= E_0 e^{i\omega t} e^{-i[\frac{\omega L n}{c} - \frac{i\alpha L}{2}]} \cos(\Delta + \phi_0) \ \underline{\hat{x}} \ , \\ \underline{E}_y &= E_0 e^{i\omega t} e^{-i[\frac{\omega L n}{c} - \frac{i\alpha L}{2}]} \sin(\Delta + \phi_0) \ \underline{\hat{y}} \ , \end{aligned} \tag{2.8}$$

where $\Delta = \frac{wL\Delta n}{c} - \frac{i\Delta\alpha L}{2} = \Delta_R + i\Delta_I$. Here $n = \frac{1}{2}(n^+ + n^-)$, $\Delta n = \frac{1}{2}(n^+ - n^-)$ and likewise for α and $\Delta\alpha$. For these experiments $\phi_0 = \pi/4$, in order to ensure the probe beams reaching the photodiodes are balanced in the absence of the pump beam. The wavevector, k, has been replaced by the equivalent expression $\frac{\omega n}{c}$. In practice, the differences $\Delta\alpha L$ and Δn induced by the pump are typically very small such that $\Delta \ll 1$. The change in absorption, $\Delta\alpha$, is caused by those atoms within the velocity interval, $\Delta v_z = \pm\Gamma/k$, that simultaneously interact with both beams, where Γ is the natural linewidth of the transition. The line profile of $\Delta\alpha(\omega)$ is therefore a Lorentzian one, given by

$$\Delta\alpha(\omega) = \frac{\Delta\alpha_0}{1+x^2}, \quad \text{with} \quad x = \frac{\omega_0 - \omega}{\Gamma/2}, \quad \text{and} \quad \alpha_0 = \alpha(\omega_0). \tag{2.9}$$

$\Delta\alpha_0$ is the difference between the absorption coefficients for left and right-circularly polarised light on resonance, $\alpha_{0+} - \alpha_{0-}$.The refractive index, $n(\omega)$, is related to $\alpha(\omega)$ through the Kramers-Kronig dispersion relation providing the dispersion profile of the spectroscopy signal;

$$\Delta n(\omega) = \frac{c}{\omega_0}\frac{\Delta\alpha_0 x}{1+x^2}. \tag{2.10}$$

The signal reaching each photodiode is proportional to the square of the amplitude of the electric field. The error signal is simply the difference of these two intensities, given by

$$\begin{aligned}\Delta S(\omega) &= \underline{E}_x\underline{E}_x^* - \underline{E}_y\underline{E}_y^* \\ &= E_0^2\,\cos\left(2\phi_0 + 2\Delta_R\right)e^{-\alpha L} \\ &= -I_0\,e^{-\alpha L}\,\frac{2\Delta n L}{c} \quad \text{for} \quad \phi_0 = \frac{\pi}{4},\ \Delta_R \ll 1\ .\end{aligned} \tag{2.11}$$

The profile for $\Delta S(\omega)$ only depends on the difference in the refractive indicies but not on the the difference in absorption. This is what provides the dispersive properties of the cell; if $\Delta S(\omega)$ is plotted over the transition frequencies of the atom, the familiar dispersion signals found in figure 2.8 are reproduced.

The arguments above overlook contributions from the windows of the cell to the line profile of the signal. The windows provide a small amount of absorption and the pressure difference across the quartz will induce a degree of birefringence, rotating the polarisation of the light. In general the magnitude of these effects are very small; they are calculated in detail in [54].

To calculate the degree of anisotropy in the medium and thus the expected value of Δ requires an in-depth analysis of the rate equations between the various M_F states of the upper and lower levels for the various times-of-flight for atoms traversing the pump beam. The required calculations to determine the expected size of Δn for a given atomic species are detailed in [53]. Their theoretical findings for rubidium and ceasium match well with their experimental results. An equivalent analysis for lithium has not yet been performed.

The polarisation spectroscopy data for our locking scheme is given on the following page. As with figure 2.6 the left and right dispersion profiles correspond to the $F = 2$ and $F = 1$ transitions respectively and the central signal is the dispersion signal corresponding to the cross-over resonance. The relative sizes of the signals and the sign of their gradient depend upon the interplay of a large number of details and cannot be determined through intuitive analysis alone. Instead, detailed calculations including all M_F states and their respective transitions are required. One detail that does seem to be consistent amongst polarisation spectroscopy features, including the data presented here, is that the amplitude of the signals corresponding to the two hyperfine components of the ground state scale inversely with the number of M_F states for that level; i.e. it is expected that the signal for the $F = 1$ component will be larger than the signal for the $F = 2$ component [53]. This is

opposite to the case for saturation spectroscopy, where the state with the largest F number produces the largest signal. This is because cycling atoms towards the largest M_F value is statistically less likely to be successful in the $F = 2$ transition due to losses from the cooling cycle to the $F = 1$ component, while cycling in the $F = 1$ hyperfine state is more robust.

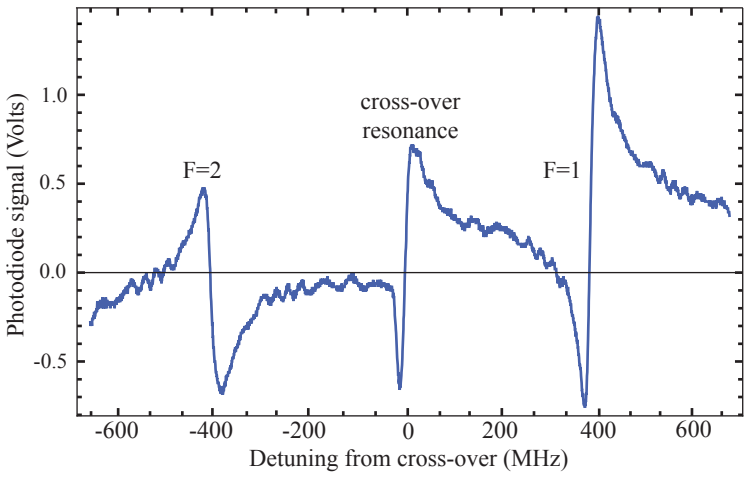

Figure 2.8: Polarisation spectroscopy data showing the ground state hyperfine features including the cross-over for the D2 transition in ^{7}Li.

Under ideal circumstances the background against which the signal sits should be flat and lie along zero. In practice, however, several factors can distort the signal. Some of these are frequency independent, such as the effects from the windows mentioned above. These can be accounted for by balancing the outputs of the beam splitter cube in the absence of the pump beam. Other effects are more complicated.

The signal is very sensitive to any background magnetic field due to the Faraday effect. The presence of a magnetic field induces a birefringence in the atomic gas causing the σ^+ and σ^- components of the probe beam to propagate at different velocities through the medium. The induced phase difference is tantamount to a rotation in the polarisation of the linearly polarised beam. In the absence of a pump beam, the effect will interact with any velocity group resonant with the laser as the frequency is scanned. The result is a Doppler broadened background superimposed onto the signal recorded, offset from zero. The offset may be tuned by adjusting the strength of the field. When the pump beam is added, Doppler free features appear, whose sign and magnitude may be controlled by the strength and direction of the applied magnetic field vector in a manner equivalent to how the Doppler limited profiles are manipulated.

Originally, a coil was wound around the spectroscopy oven to create a weak magnetic field to counteract the Earth's field as recommended by [53]. However, experimentally we have found that the height of the dispersion signals can be doubled

with the application of a 2 gauss field parallel to the probe beam, significantly larger than the measured Earth field. Decreasing the field will decrease the signal height even for negative fields, i.e. if the field is reversed. For large enough reversed fields the signal becomes heavily distorted and finally return with the gradient of the slope reversed. This signal is typically very small and far offset from zero. Obviously, any magnetic field lifts the background offset, however, this can be compensated for by tweaking the angle of polarisation entering the cell. For fields parallel to the probe beam, larger than 2 gauss, the signal continues to increase, however, it becomes difficult to compensate for the background without distorting the flatness of the background significantly. Applying such a magnetic field, however, reduces the accuracy of the lock to a given spectroscopy feature. Because we are probing on the most positive M_F sub-level for each of the F states the Zeeman shift displaces the spectroscopic features. For a 2 gauss field the cross-over frequency is shifted by approximately 3 MHz and has been experimentally verified by direct comparisons with saturation spectroscopy measurements made in the absence of any magnetic field. When characterising the MOT it is important to include this shift in calculations in order to provide the correct atom number and measure the dependence on the laser detunings.

The lock is provided by a PI circuit whose gain is maximised without causing op-amp oscillations and whose feedback is tuned to respond most strongly to frequencies below 100 kHz, as these are expected to dominate the power-spectrum in the lab. The locking electronics attempt to maintain the frequency corresponding to the zero-crossing of the dispersion profile. If the signal voltage drifts from zero, the integrator changes the frequency of the laser such that it moves towards the desired zero-crossing point. As can be seen in figure 2.8 this creates a broad re-capture range, only failing if the voltage reaches a new zero-crossing, for example at the $F = 2$ resonance. This range can extend to several hundred MHz, significantly broader than achieved in saturation spectroscopy, providing a very robust locking mechanism. When locking to the cross-over feature, the measured noise in the error signal corresponds to frequency deviations below 500 kHz and the laser remains locked for many hours. Very slow drifts in the orientation of the $\lambda/2$ plate need to be corrected for once a day.

In the following section I discuss the rest of the optical set-up and describe how the necessary frequencies for laser cooling and trapping are generated.

2.3.2 Optical Set-up

The Zeeman slower requires both a coolig laser beam and a repumper beam, separated in frequency by the hyperfine splitting of the ground state. These frequencies,

however, are shifted from the MOT and repumper frequency of the MOT due to the presence of the magnetic fields inside the slower. The design of the Zeeman slower requires the associated light to be shifted 803.5 MHz to the red of the crossover resonance (as described in section 2.5), placing the Zeeman repump light exactly at the crossover frequency, to which the laser is locked. This arrangement minimises the number of AOMs required while optimising the use of the available laser power. The shifts required for each frequency are outlined schematically in figure 2.9. The figure also includes an additional frequency required for an absorption beam used to image the MOT cloud and characterise the atom number and temperature.

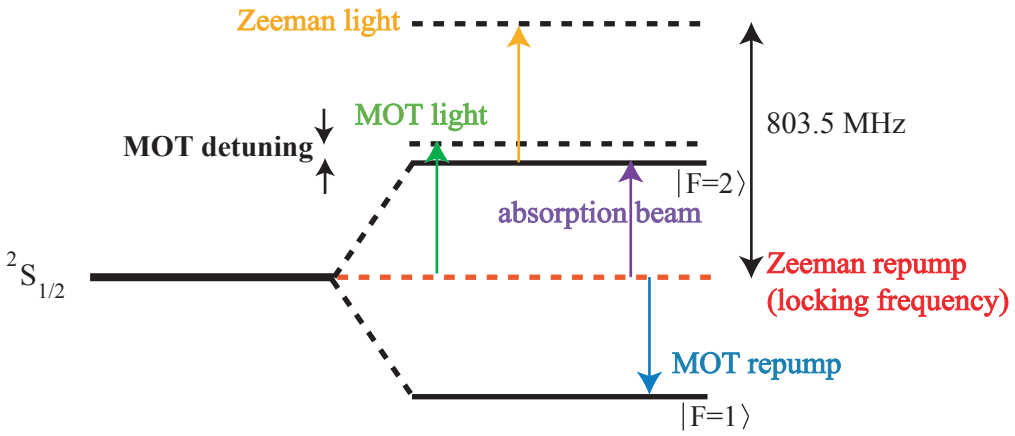

Figure 2.9: The energy level diagram showing the required frequencies for the Zeeman slower, the MOT and the absorption imaging beam. Each frequency is reached by one or more AOMs from the locking frequency at the cross-over feature.

The complete optical set-up for the experiment is given in figure 2.10. To obtain the five frequencies required, four AOMs are employed. For each beam delivered to an AOM a half waveplate and a polarising beam splitting cube are used to pick off the required power from the original 500 mW beam and send it to the appropriate AOM. Each AOM applies a shift of approximately 200 MHz and the first-order harmonic output beam is double passed back through the AOM with a total efficiency of 55-65% and a total shift of 400 MHz. Two 300 mm plano-convex lenses are used to focus the beam at the AOM to a spot size of 400 μm and recollimate it to the original diameter on the far side. By introducing a quarter waveplate into the beam path the polarisation is rotated so that when the returning beam passes back through the polarising beam cube, the polarisation is orthogonal to the original beam and consequently separated from the input beam. In two cases, the zeroth-order beam from the AOM is picked off to recycle some of the power that would otherwise be lost. This is either reintroduced into the original beam path to be used later or in the case of the absorption beam, sent directly to the experiment. The Zeeman cooling beam and the Zeeman repump beam are combined onto a 75:25 beam splitter.

While this leads to 30 to 40 mW of power loss the ratio of the beam splitter was chosen specifically to minimise this loss. By choosing high quality optics and AOMs enough beam power is obtained from the 500 mW, amplified diode laser to operate the Zeeman slower and MOT. The Zeeman slowing light has a typical power of 25-30 mW with a repump power of 2-3 mW. The beams were expanded to a diameter of 20 mm and focussed towards the oven. The focal point of the beam was chosen to maximise the overlap between the light and the diverging atomic beam, while in addition, providing a transverse cooling component to the atoms. The total power delivered to the MOT beams was approximately 50 mW with a total repump power of 7 mW.

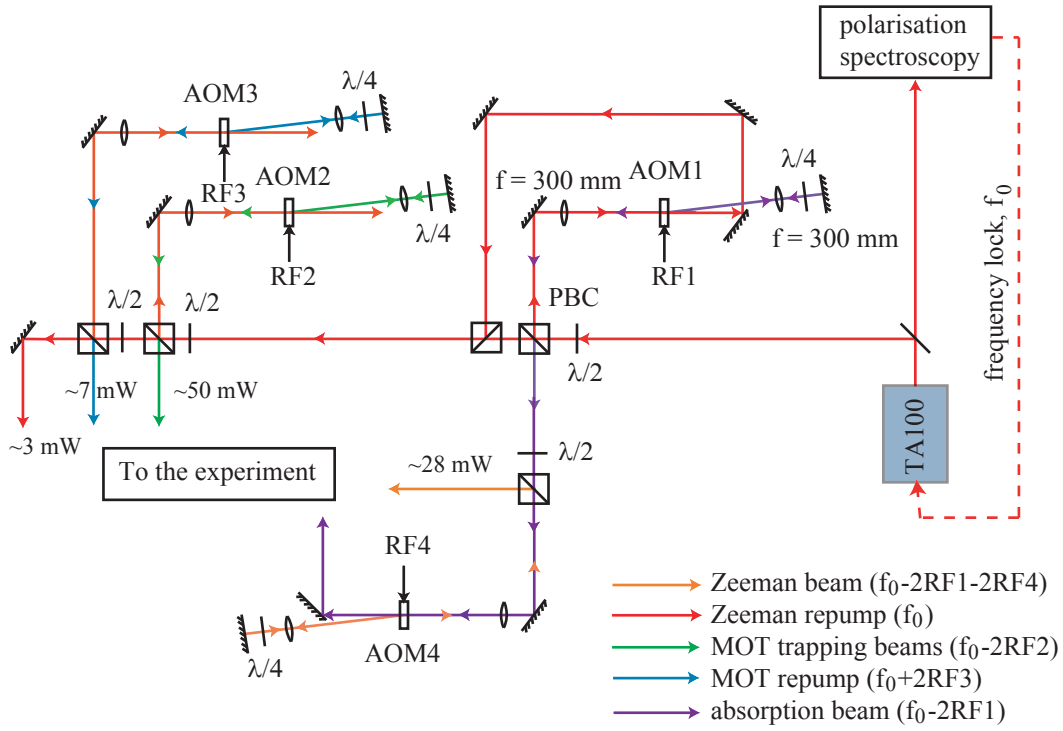

Figure 2.10: The full optics set-up used to deliver the 5 frequencies from the TA100 via four 200 MHz AOMs. The colour coding for the frequencies is used throughout the chapter.

The MOT beams were divided into 3 equally intense, 25 mm diameter beams and retroreflected along the three axes of the MOT, providing a trapping volume of ~15.6 cm^3. Two quarter waveplates were introduced into each beam before and after the MOT to maintain the correct helicity for the trap. The repump beam is combined with the MOT beams at the 33:67 ratio beam splitter that picks off the first MOT beam as shown in figure 2.11. Part of the repump is then further split in two equal parts along with the trapping beams. This gives a repump beam power ratio of 4:1:1 across the three axes of the trap; such anisotropy in the beam powers should not effect the success of trapping an atom providing the powers of

the retro-reflected beams closely match their original beams.

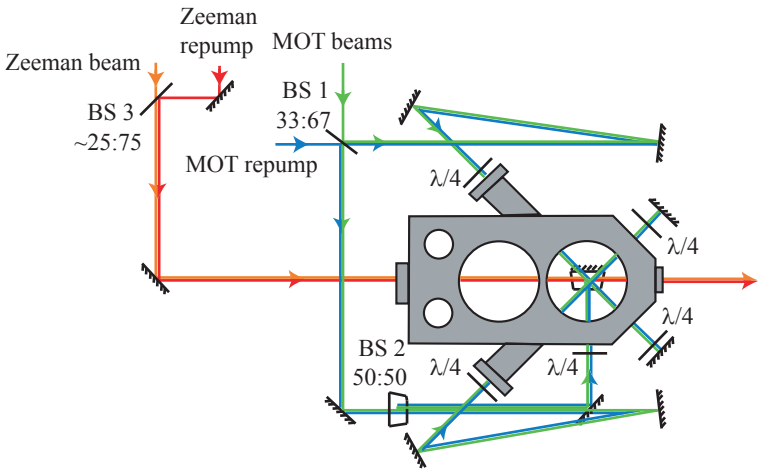

Figure 2.11: The optical set-up used to deliver the trapping light, the slowing light and repump beams to the experiment.

The dependence of the number of atoms trapped in the MOT on beam power and detuning will be discussed further in section 2.7.

2.4 The Lithium Beam

In order to load a magneto-optical trap it is, of course, vital to have some atoms with which to load it. A number of methods exist to achieve the vapour pressures in the low velocity regime, required to load a trap, while maintaining large atom numbers and long trap lifetimes. The first experiments to successfully trap atoms in a magneto-optical trap employed a chirped laser to decelerate a beam of sodium atoms into the trap [44]. This method of laser cooling, however, provides a pulsed beam. The Zeeman slower by contrast provides a continuous beam which will load a trap much faster [43], while also providing a differential pumping tube to lower the background pressure in the trapping chamber. As people sought to improve and simplify these experiments other methods came to be employed. For example, rubidium and caesium have a high enough vapour pressure at room temperature that a MOT can be loaded directly, using heated dispensers inside the trapping chamber [55]. For these atoms, a reasonable fraction of the thermal distribution can be captured, thereby avoiding the complications of building an effusive beam source and decelerator. While in principle this method can be implemented for lithium, the temperature required to reach the necessary vapour pressure is quite a bit higher and background pressures limit the trappable atom number. Inside an evacuated chamber the vapour pressure as a function of temperature is well described by the Clausius-Clapeyron equation, given by:

$$P(T) = 760\, e^{\frac{h}{8.31}\frac{m}{}(\frac{1}{T_{boil}} - \frac{1}{T})}, \tag{2.12}$$

where h is the enthalpy of vaporisation, in kJkg^{-1}, m is the molecular weight for the species of interest and T_{boil} is the boiling point at 1 atmosphere in Kelvin. Figure 2.12 compares the vapour pressures for lithium (blue) and rubidium (red) as a function of temperature.

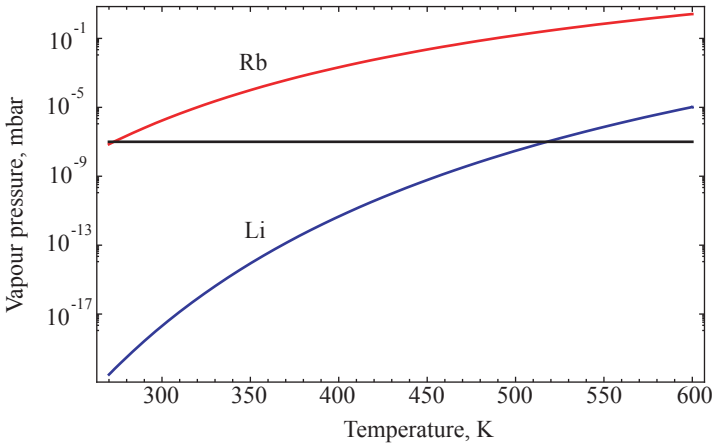

Figure 2.12: Vapour pressure as a function of temperature for lithium (blue), and rubidium (red). The black line indicates a pressure of 10^{-7} mbar.

At room temperature the vapour pressure for rubidium is already more than adequate and often large pumps are required to maintain a background pressure low enough to sustain long trapping lifetimes. For Li on the other hand, to reach a vapour pressure of 10^{-7} mbar (black line) temperatures of over 500 Kelvin are required. While these temperatures are easily acquirable with dispensers, the large differences in temperature and mass for the two species lead to significantly different Maxwell-Boltzmann distributions, meaning a much smaller fraction of lithium atoms will fall within the capture velocity of the trap in comparison to rubidium. This is depicted in figure 2.13.

Other methods such as light induced atom desorption [56] have been tested for rubidium successfully, however, how well the results translate for other alkali species has not been well explored and it would be difficult to estimate what sort of atom numbers can be trapped. As a large density is desirable for this experiment we have chosen to load our trap from a laser-cooled effusive beam source.

2.4.1 Effusive Beams

The first effusive beams of neutral particles moving in straight lines with thermal velocities were produced almost 100 years ago in laboratory experiments by Dunoyer

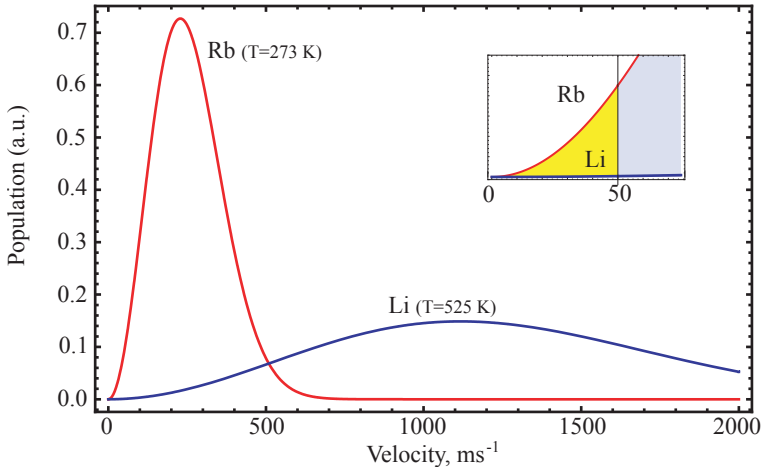

Figure 2.13: The Maxwell-Boltzmann curves showing the atomic speed distribution of lithium (blue) and rubidium (red) as a function of velocity for a vapour pressure of 10^{-7} bar. The temperature required for Li to reach this vapour pressure is 525 K while for Rb a temperature of 273 K is required. The inset highlights the velocity region up to 50 ms^{-1}; there are significantly more rubidium atoms than lithium in this velocity class, given the same vapour pressure.

[57, 58, 59]. The basic principles he used to create such beams have not changed since then and are applicable to both atomic and molecular species. The beams are formed by allowing a vapour to effuse from a closed chamber through a small aperture into a chamber evacuated to pressures several orders of magnitude lower. The vapour pressure is kept low enough to maintain the conditions of molecular flow, meaning the particles may move through the aperture and within the beam without undergoing collisions. The simplicity of such a set-up means that the properties of the beam can all be accurately predicted from classical gas kinetics. The details of these calculations may require some modification depending on the dimensions of the aperture [60], but in the simplest case may be described by effusion from a thin-walled orifice. Because the area of the aperture is typically much smaller than the area of the walls of the source chamber the thermal equilibrium of the vapour is not disturbed by the effusion of particles. Thus, in the regime of molecular flow, the flux, angular distribution and velocity profile of the beam can be predicted without any further assumptions.

2.4.2 Beam Properties

The number of molecules, dN, leaving an aperture area element, $d\sigma$, during a time, dt, with a velocity between v and $v+dv$ into a solid angle element, $d\omega$, at an angle θ relative to the normal of σ can be shown through simple kinetic gas theory [61] to be

$$dN = nf(v)\frac{d\omega}{4\pi}\, v\, \cos\theta\, d\sigma\, dt\, dv, \tag{2.13}$$

where n is the number density and *f(v)*dv is the normalised function describing the velocity distribution of the atoms inside the oven. The form of the function, $f(v)$, is the Maxwell-Boltzmann distribution for a gas at thermal equilibrium in a cell. The additional factor of v in equation 2.13 is included to account for the faster particles striking the hole more often than the slower ones. The normalised function, $f_{beam}(v)$, is

$$f_{beam}(v) = \frac{4}{\sqrt{\pi}}\frac{v^2}{\alpha^3}\, e^{-(\frac{v}{\alpha})^2}, \tag{2.14}$$

where $\alpha = \sqrt{2k_BT/m}$, k_B is the Boltzmannn constant, T is the temperature of the gas and m is the mass of the particle.

In the case where we no longer have a thin-walled orifice, but instead a long channel of length, l, and radius, r, the total number of particles leaving the source chamber is reduced by a factor, $1/\kappa$. For long cylindrical tubes, such that $l \gg r$, $1/\kappa$ is given by $\frac{8}{3}\frac{r}{l}$ [60]. The angular distribution for such a channel is modified from the expected $\cos\theta$ distribution to a more forward focussed one, as calculated by Claussing [62]. This reduces the total number of atoms leaving the aperture without diminishing the flux at angles close to $\theta = 0$. This improvement of forward beam intensity per quantity of source material is of great value in many beam experiments.

2.4.3 The Oven

The oven in use throughout the experiments discussed in this chapter is a very simple cylindrical tube of length 85 mm and diameter, 40 mm, welded to a DN63CF flange. The aperture is a cylindrical channel of length 2 mm and radius 0.35 mm and is recessed 28.5 mm from the face of the large flange. Figure 2.14 depicts a schematic drawing of the oven.

To heat the oven, a 2 metre long heater wire is wrapped around the front half of the cylinder and a second, similar wire around the back half, allowing for a more tunable control of the temperature gradient along the length of the oven pipe. This is important as the walls of the aperture must be kept warmer than the rest of the oven to prevent lithium condensing in the channel and clogging the exit. The currents for the heater wires are supplied by a pair of variacs, capable of providing up to 120 V each. The typical resistance of our heater wires is ∼26 Ω and draw 2.5 A and 1 A for the front and back heater, respectively. The temperature is monitored

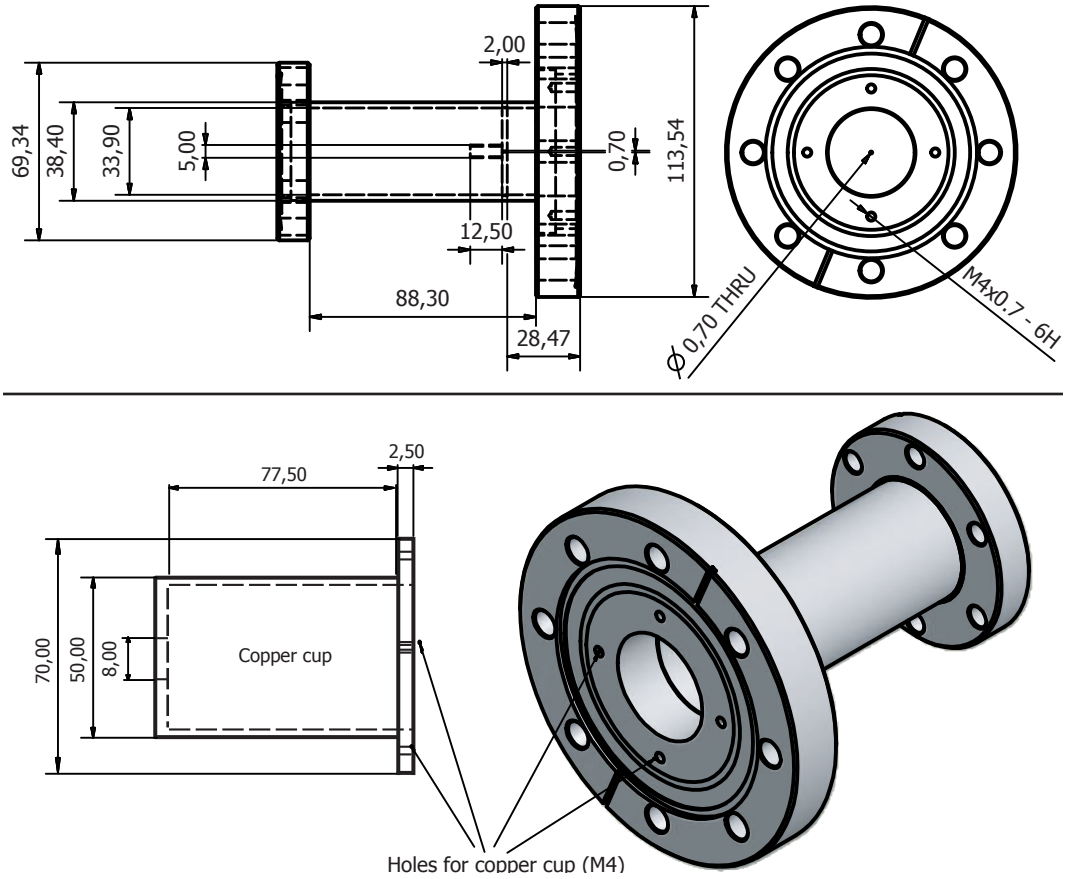

Figure 2.14: The lithium oven. An effusive Li beam exits the oven through a 0.7 mm diameter hole in the front wall.

by a pair of k-type thermocouples mounted to the outside of the chamber, one above the aperture and a second by the rear of the oven. During typical operation a difference of $\sim$30-40 °C is maintained between the front and back thermocouples. The oven is insulated from the external environment by layers of fibre-frac and encased with aluminium foil. Because of the reactive properties of lithium with copper the DN40CF flange is sealed using a pure nickel gasket, rather than a copper one.

This oven design was plagued by problems due to lithium condensing around the seal of the small flange of the oven whenever the oven was cooled down. Several efforts were made to keep the face of the flange warmest during cooling, but to no avail. This issue was later addressed by renewing the design and replacing the oven with a new one. This oven was used during the experiments outlined in chapter 4. A schematic for the new design is given in figure 2.15.

The oven is formed from a single T-piece; the trunk of the T is mounted with a DN63CF flange and mounted to the experiment. The upper junction of the T accommodates the DN40CF flange and is used to load the lithium into the lower

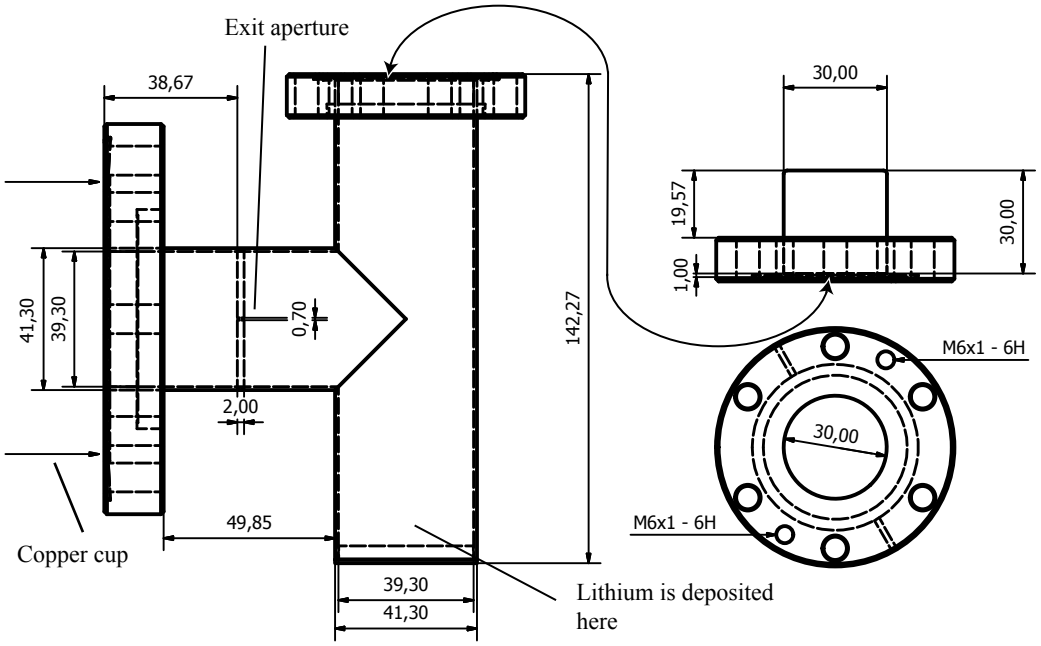

Figure 2.15: The new lithium oven. An effusive Li beam exits the oven through a 0.7 mm diameter hole in the front wall.

arm of the oven. It was hoped then, that any lithium condensing on the roof of this flange would drip down under gravity into the lower arm to be recycled. Despite these improvements lithium persisted in condensing around the seal between the flange and the aperture wall. The reason for this is due to poor heat conduction through the thick flange to the inner walls of the oven. To compensate for this a hole from the DN40CF flange was bored out in the centre, leaving a 1 mm thick wall between the inside and outside of the flange. A tight-fitting copper bobbin was then silver soldered into the steel flange, protruding 20 mm above the flange face (right-hand side of figure 2.15). To heat the oven the same heater wires were used to wrap around the chamber. One wire was wrapped around the lower arm and used to heat the lithium. The second wire was wrapped around the trunk of the T piece and the upper arm, ensuring significant heat was delivered to the protruding copper bobbin. This allows precise control of the temperature of both the 700 μm aperture and the upper flange; during cooling these regions can be kept warmest, encouraging lithium to condense in the lower arm. The results of these modifications were very successful and to date there have been no further issues.

Lithium exits the oven into a solid angle of almost 2π, meaning much of it never reaches the Zeeman slower and must be pumped away by the vacuum pumps. Large quantities of lithium in the pump will significantly reduce its operational lifetime and should be avoided if possible. To collect a large proportion of the unused lithium a 80 mm deep copper cup was bolted to the outside face of the DN63CF flange of

the oven, figure 2.15. The cup has an 8 mm diameter aperture at one end to allow the useful lithium to exit.

A diagram of the experimental set-up used in this project is given in figure 1.2, including the most recent oven design, described here. The oven is attached to one port of a DN63CF 6-way cross. The port opposite the oven attaches to a second 6-way cross, separated by an all metal gate-valve. The crosses are used to mount pumps, pressure gauges and viewports to allow optical access for characterisation of the effusive beam. An aluminium flag (not shown) is suspended from the roof of the first chamber. The flag can be lowered and raised by a hydraulic translational stage to block or unblock the lithium beam as necessary. By closing the gate valve between the two 6-way crosses the oven can be opened without letting the Zeeman slower and MOT chamber up to air. There is a differential pumping tube between the two 6-way crosses to allow for an additional pumping stage to lower the background pressure in the MOT chamber. The dimensions of the tube are discussed in section 2.5. Chamber 1 is pumped by a LEYBOLD TW70H 60 ls^{-1} turbo pump backed by a two-stage diaphragm pump. Chamber 2 is pumped by a VARIAN NOBLE DIODE 55 ls^{-1} ion pump. There is a cold cathode pressure gauge, capable of reading pressures as low as 1×10^{-10} mbar, attached to each of the two chambers.

2.4.4 Absolute Beam Flux

To measure the flux of Li atoms leaving the oven as a function of temperature a 1 mm $1/e$-diameter gaussian probe beam of power 20 μW was aligned through the first chamber approximately 110 mm downstream of the oven exit aperture, orientated perpendicular to the atomic beam. The probe beam was scanned 1.5 GHz about the D2 transition to measure the population of the $^2\mathrm{S}_{1/2}$ doublet. To measure the degree of absorption the laser intensity was measured using a 1 MHz bandwidth amplified photodiode and recorded as a function of the laser frequency.

To convert these readings into a measurement of the absolute flux of atoms we need to know the mean number of photons scattered by an atom traversing the probe beam. The spacing of the hyperfine levels of the $^2\mathrm{P}_{3/2}$ state is on the the order of the linewidth of the transition and cannot be resolved. This means that any excitation from the ground state by the probe beam has a probability of decaying into the other hyperfine state and thus are lost from the scattering cycle. Consequently, Li must be modeled as a 3-level system and the appropriate rate equations must be solved; figure 2.16.

All the atoms are assumed to start in level 1 and are driven to level 2 by a laser with detuning, δ and intensity, I. Level 1 is assumed to be stable and the probability of spontaneous decay from level 2 to levels 1 or 3 is given by the branching ratio, r:

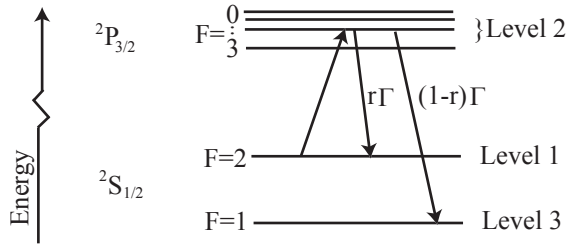

Figure 2.16: Lithium, modelled as a three-level atom. Levels 1 & 3 represent the ground-state hyperfine components, $F = 1$ and $F = 2$, and level 3 indicates the amalgamation of the four, unresolved, hyperfine components of the $P_{3/2}$ state.

Level 2 decays with a rate $r\Gamma$ to level 1 and a rate $(1-r)\Gamma$ to level 3.

The atom excitation rate is R and is directly proportional to the intensity of the laser beam, I. Under the conditions of molecular flow, such that the probability of collisions can be neglected, it can be assumed that the coherence between the atoms and the light field quickly reaches steady-state and the optical Bloch equations can be solved to give the excitation rate

$$R = \frac{I/I_{sat}}{1 + (\frac{2\delta}{\Gamma})^2} \frac{\Gamma}{2}, \tag{2.15}$$

where Γ is the spontaneous decay rate from level 2 to ground, $I_{sat} = \pi hc\Gamma/3\lambda^3$ is the saturation intensity and λ is the wavelength of the transition connecting the ground and excited states. The detuning of the radiation will depend on the laser frequency and the Doppler shift, $v_z \sin(\theta)/\lambda$, experienced by atoms in the diverging beam with velocity components along the axis of the probe beam, where v_z is the velocity parallel to the axis of the atomic beam and θ is the angle subtended by the projection of the atom and the z-axis. To calculate the number of photons scattered we must solve the rate equations to find the population of level 2 as a function of time, $N_2(t)$; it is assumed that once an atom reaches level 3 it remains there. The rate equations for the 3 level system can be solved exactly to give the population of level 2 as a function of time. By integrating $\Gamma N_2(t)$ over the atom-laser interaction time, τ, we find the number of fluorescent photons emitted per atom to be [63]

$$N_p = \frac{R\Gamma}{R_+ - R_-} \left(\frac{e^{-R_+\tau} - 1}{R_+} - \frac{e^{-R_-\tau} - 1}{R_-} \right), \tag{2.16}$$

where

$$R_\pm = R + \Gamma/2 \pm \sqrt{R^2 + rR\Gamma + \Gamma^2/4}\,. \tag{2.17}$$

The exact value for R depends on the detuning, which in turn depends upon the velocity of the atom and the solid angle, $\mathrm{d}\omega = \sin(\theta)\mathrm{d}\theta\,\mathrm{d}\phi$. To determine a more

realistic value for the number of photons scattered per atom equation 2.16 must be integrated over the velocity distribution and the solid angle of the atom beam that intersects the laser. In practice the beam from the oven exits into the full solid angle available (2π), however, this calculation may be reduced by integrating over the angle, dϕ, subtended by the height of the probe beam and dθ, subtended by the aperture of the copper cup.

The branching ratio for lithium from the $^2P_{3/2}$ state to the two hyperfine ground states is estimated to be $r = 2/3$. For a 1 mm diameter laser beam, resonant with the $F = 2$ component of the D2 line and an intensity of 20 μW ($s = 1$), a typical atom travelling perpendicular to the light beam will scatter 1.7 photons on average. The number of atoms exiting the oven per second is determined by equation 2.13, however, a simple estimate can be made by using conservation of mass flow, i.e. the number of atoms entering chamber 1 is equal to the number of atoms exiting chamber 1. Given the vapour pressure of lithium inside the oven, $P(T)$ (equation 2.12) and the conductance of the exit aperture, C (equation 2.30), the number of atoms leaving the oven per second as a function of oven temperature is given by $N_{atoms} = CP/k_BT$. For our experimental values the number of atoms exiting per second is approximately 3×10^{15} for a temperature of 500 °C. The total number of photons scattered per second is thus given by the product $N_p N_{atoms}$, integrated over the velocity distribution and a suitable solid angle. The total number of scattered photons is then compared to the number of photons in the beam to give a value for the expected degree of absorption as a function of temperature. Figure 2.17 shows the data measured in the lab compared with theoretical expectations.

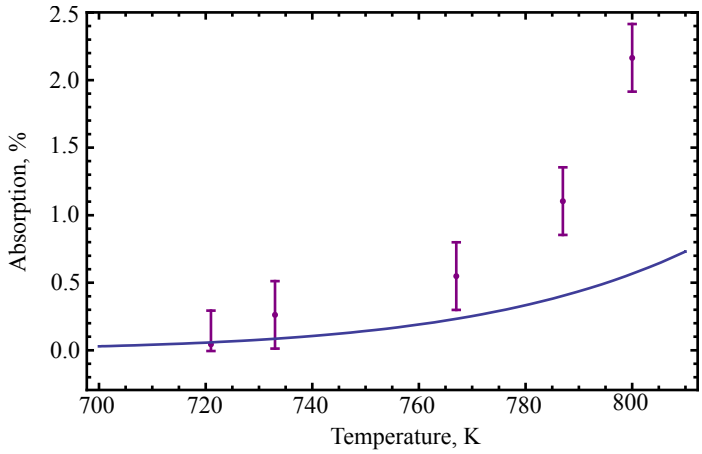

Figure 2.17: A comparison of theoretical predictions (blue curve) and absorption data for an effusive beam exiting the oven, measuring the $^2S_{1/2}$ (F=2) to $^2P_{3/2}$ transition.

The discrepancy between the data and the theoretical expectations in figure 2.17 may be due to an underestimate of the temperature. There is a significant

uncertainty in the temperature of the atom ensemble inside the oven, upon which the amount of absorption is particularly sensitive. If we assign a temperature to the data 40 °C warmer than measured by the thermocouples then there is a good agreement between data and theory. Other parameters such as the branching ratio and the solid angle of the beam subtended by the laser beam are all estimated, however, these parameters have relatively little influence on the theoretical curve within their expected uncertainties.

To estimate the flux of atoms required to load a large MOT we need to decide on some acceptable loss rates from the trap due to background collisions. The loading rate for atoms into the trap is determined by a balance between the capture rate and the loss rate,

$$\frac{\mathrm{d}N}{\mathrm{d}t} = R_c - \frac{N}{\tau}, \tag{2.18}$$

where N is the number of atoms in the trap, R_c is the capture rate and τ is the lifetime determined by collisions with background atoms. The solution to equation 2.18 tends to a steady-state for large t given by $N_{tot} = R_c\tau$, where N_{tot} is the number of trapped atoms as t tends to infinity. For our experiment we would like to load of the order of $10^{\,9}$ atoms into the trap with a loading constant of around a second. This requires a loading rate, R_c, of $1 \times 10^{\,9}$ atoms per second. Using the calculations described above, taking into account the reduced solid angle, it is estimated that for an oven temperature of 500 °C approximately 10^{11} atoms per second exit the Zeeman slower in the absence of any laser cooling. An efficiency of 1% for the number of atoms decelerated to within the capture velocity of the MOT is comfortably achievable.

To maximise the loading rate for the experiment we need to look in detail at the design of the Zeeman slower and how one can optimise its efficiency. The following section will discuss the design, construction and characterisation of our slower.

2.5 The Zeeman Slower

The Zeeman slower is an extremely useful and efficient tool and by far the most common method employed for laser cooling an atomic species. The success of laser cooling rests on the requirement that the frequency of the radiation is resonant with a closed transition in the atomic species for a sufficient length of time such that the atom can undergo multiple scattering events and dissipate its momentum. For an atom of given velocity, v, moving through a light field there is an associated Doppler shift in the wavelength experienced by the atom relative to that when at rest given

by

$$\delta_D = v\lambda/c \, . \tag{2.19}$$

If we were to begin laser cooling our atomic beam with radiation resonant with atoms at a particular velocity, the atoms would very quickly decelerate and this resonance condition would no longer be held. By adjusting the frequency of the atomic transition, so as to maintain resonance with the light as the atom decelerates, the effects of radiation pressure can be applied continuously as the atom decelerates. The simplest and most common approach to manipulate the frequency of the atomic transition is to use a spatially varying magnetic field to apply a Zeeman shift to the atoms. By carefully tuning the magnitude of the field along the axis of the atomic beamline, it is possible for the atoms to remain resonant with the radiation throughout the length of the slower. For most alkalis this equates to a length of only tens of centimetres in order to decelerate a thermal atomic beam to rest.

2.5.1 Theory

Equations of Motion

To decelerate the largest number of atoms possible it is necessary to apply the largest possible stopping force and maintain it along the complete length of the slower. The stopping force on an atom is given by equation 2.3. In the limit of high laser intensity, such that $s \gg 1$, the force is maximised and the equation is reduced to

$$F_{max} = \frac{\hbar k \Gamma}{2} \, . \tag{2.20}$$

For a particle of mass, m, the maximum deceleration applied by this force is simply $a_{max} = F_{max}/m$. The equations of motion for a particle under constant deceleration give

$$z(t) = v_0 t - \frac{1}{2} a t^2 \quad \text{and} \quad v_z(t) = v_0 - at,$$

where v_0 is the initial velocity, in this case the capture velocity of the Zeeman slower. In general it is not possible to apply this maximum deceleration to a given atom due to imperfections in the magnetic field shape and insufficient light intensities, such as in the wings of the beam profile. Our slower is constructed from a series of solenoids, which when combined will reproduce the desired field shape. However, due to the nature of its construction there will be ripples in the field profile, the amplitudes of which will depend on the number of solenoids used to construct the field. Atoms in a slower designed to decelerate at the maximum deceleration are

very intolerant to the size of these ripples and will be lost. It is therefore useful to introduce a reduction factor into the acceleration, ϵ, whose value lies between 0 and 1. The smaller the value of ϵ the more tolerant the atoms are to ripples in the field, however, a longer slower is required. For these experiments a value of $\epsilon = 0.5$ was chosen as a suitable compromise. The equations of motion can subsequently be combined to give the particle's velocity as a function of length

$$v_z(z) = v_0 \sqrt{1 - \frac{\epsilon\, z}{z_0}}\,, \tag{2.21}$$

where $z_0 = v_0^2/2a_{max}$ and $z_0/\epsilon = L$, the required length of the slower. Atoms with velocities greater than $v_0 = \sqrt{2\epsilon a_{max} L}$ are not captured in the Zeeman slower. As will be shown in subsequent sections, for a given atomic species, an optimum length can be calculated to maximise the atomic flux decelerated by a Zeeman slower.

Magnetic Field Gradient

For an atom with velocity, v, the Doppler shift, in angular frequency units, is $\Delta\omega_D = k.v$, where k is the wavevector of the radiation. As the atoms decelerate this Doppler shift will change linearly with velocity. The aim of the Zeeman slower is to compensate for this Doppler shift by applying an opposing Zeeman shift to the atoms such that they remain on resonance with the radiation. This is known as the resonance condition and must be held for the entire slowing length of the decelerator. This condition can be written as $\hbar k.v_z = \Delta\mu\, B_z$, where $\Delta\mu$ denotes the difference between the magnetic moments of the excited state and ground state, $\mu_e - \mu_g$, of the chosen transition. To calculate the gradient of the magnetic field required to maintain the resonance condition the Zeeman shifts of the appropriate energy levels need to be calculated. The Hamiltonian for an alkali atom in a magnetic field is given by

$$\hat{H}_{tot} = A\ \underline{\hat{I}}.\underline{\hat{J}} + \mu_B(g_L\underline{\hat{L}} + g_S\underline{\hat{S}}).\underline{B} + \mu_N g_I\ \underline{\hat{I}}.\underline{B}\,, \tag{2.22}$$

where A is the hyperfine coupling constant, $\underline{\hat{L}}$, $\underline{\hat{S}}$, $\underline{\hat{J}}$, and $\underline{\hat{I}}$ are the operators representing the electronic orbital angular momentum, electron spin, total electronic angular momentum and nuclear spin, respectively and g is the associated Landé g-factor. $\underline{B}$ is the magnetic field vector and μ_B is the Bohr magneton. The nuclear magnetic moment, μ_N is $\sim$1880 times smaller than μ_B and the final term on the right hand side can be neglected. The first term in the Hamiltonian describes the hyperfine interaction while the second accounts for the Zeeman shift due to the

interaction between the applied field and the electron's angular momentum.

In order to maintain a closed two-level cycling transition the Zeeman slower may operate with either σ_+ or σ_- light acting on the $|g = 2s^2S_{1/2}\,, F = 2, M_F = 2\rangle \rightarrow |e = 2p^2P_{3/2}\,, F = 3, M_F = 3\rangle$ or the $|g, 2, -2\rangle \rightarrow |e, 3, -3\rangle$ transition, respectively. Note that unless stated otherwise all states written in Dirac notation are expressed in the $|g/e, F, M_F\rangle$ basis from here on. Where obvious, the clarification of the ground, (g), or excited, (e), state will be omitted. For a typical Zeeman slower the largest fields experienced by an atom will be on the order of 0.1 Tesla. Although typically at these fields mixing between M_F states would be expected, because the transitions of interest are between the stretched states there are no other states with which they can mix; figure 2.19. Consequently the Zeeman shifts for this transition are linear with magnetic field strength for all fields.

Rewriting the second and third terms in equation 2.22 as $\mu_B g_J B_z \hat{J}_z$ and applying the operator to the $|g, 2, \pm 2\rangle$ and $|e, 3, \pm 3\rangle$ states, we can show that

$$\Delta\mu B_z = ((g_J M_J)_e - (g_J M_J)_g)\, \mu_B B_z = \pm \mu_B B_z$$

for $\Delta M_J = \pm 1$, and the subscripts, g and e, denote the ground and excited state of the transition. The sign of the field depends on whether we wish to cycle from the $|2, 2\rangle$ (positive) or $|2, -2\rangle$ state (negative). The value for g_J is given as

$$g_J \cong 1 + \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)}, \tag{2.23}$$

under the approximation $g_S = 2$, and g_L is taken to be 1.

The magnetic field shape required to decelerate the atoms must therefore conform to the equation $k.v_Z = \pm\mu_B B_z/\hbar$, equating to a decreasing or increasing field Zeeman slower. Substituting in equation 2.21 gives

$$B(z) = \pm\frac{\hbar k}{\mu_B} v_0 \sqrt{1 - \frac{z}{L}}\,. \tag{2.24}$$

As the atoms slow down the laser frequency in their rest frame decreases and the transition frequency must decrease to compensate for this. On the σ_+ transition the frequency decreases with decreasing magnetic field, so B_z must be large at the entrance of the slower and taper down. The atoms then exit the slower at low speed and in low field such that the required laser frequency for the slower will be close to that of the MOT light. On the σ_- transition, the frequency decreases with increasing magnetic field and so B_z must be small at the slower entrance and taper up to large values. In this case the atoms are in low field when the Doppler shift is largest and laser frequency will need to be substantially red detuned from that of the MOT light. For example, for a capture velocity of 671 ms^{-1}, a 1 GHz detuning

is required. In either case (σ_+ or σ_-), in order to obtain a large capture velocity the maximum value of B_z must be large, typically about 0.1 T, and this can be quite challenging. A reduction in the field maximum can be obtained by offsetting the entire field profile such that it passes through zero and changes sign part way through the slower. An equivalent red detuning needs to be applied to the frequency of the cooling light to compensate the magnetic field offset. For the σ_- case this detunes the cooling light even further from the MOT light. This is disadvantageous since the frequency offset becomes difficult to obtain with AOMs. For the σ_+ case, the required offset can easily be generated and it becomes advantageous to do this since the slowing light is further detuned from the MOT light such that it does not disturb the MOT cloud. Including this additional frequency offset, the equation for the required magnetic field profile is

$$B(z) = \frac{\hbar k}{\mu_B} v_0 \sqrt{1 - \frac{z}{L}} + \frac{h\delta_0}{\mu_B} \quad \text{for } \Delta M_J = +1 \,, \tag{2.25}$$

where δ_0 is the laser detuning from the zero-velocity, zero-field frequency. By choosing to cycle with σ_+ light the size of δ_0 can be chosen to shift the Zeeman repump frequency to the locking frequency of our laser, thereby eliminating the requirement for an extra AOM; this can be seen in figure 2.9. The final parameters for our slower where chosen to be $\epsilon = 0.5$, $L = 0.51$ m and $\delta_0 = 401.75$ MHz, thus decelerating the atoms from a speed of 900 ms^{-1} to a target of 50 ms^{-1}; the value for L will be justified in the subsequent sections. The ideal magnetic field profile for these parameters is given in figure 2.18.

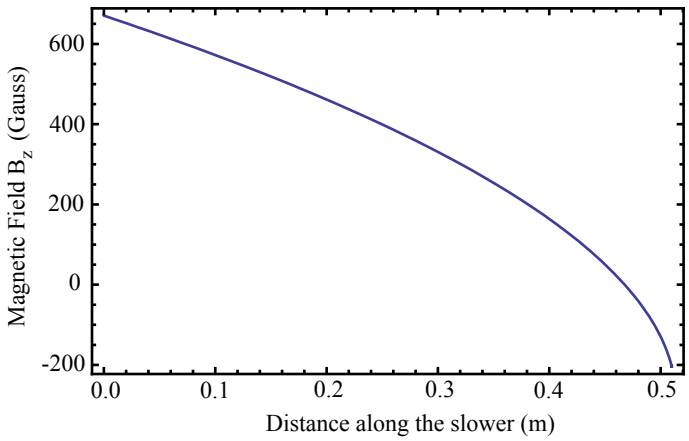

Figure 2.18: The ideal magnetic field profile required to decelerate ^{7}Li atoms from 900 ms^{-1} to 50 ms^{-1}, cooling with light detuned by 402 MHz from the $^2S_{1/2}\,|F = 2,\, M_F = 2\rangle \rightarrow {}^2P_{3/2}\,|F = 3,\, M_F = 3\rangle$ transition.

Losses from the Cooling Cycle

While the decision to laser cool on the $|g,2,2\rangle \rightarrow |e,3,3\rangle$ transition has certain practical benefits, as mentioned above, it is important to look at the behaviour of both possible cycling transitions as a function of magnetic field in detail. This is particularly true where decay rates into other, dark states are prominent.

Although both cooling transitions are closed for purely circularly polarised light, imperfections in the optics will distort this polarisation and lead to losses from the cooling cycle. In any real experiment the cooling light will never be purely circularly polarised, but instead contain a π-polarisation component polluting the state selection of a given excitation. To estimate the relative strengths of these loss mechanisms the coupling strengths between various energy levels must be calculated in both the small and large magnetic field regime. The full matrix elements for the Hamiltonian given by equation 2.22 give us the energy level diagram as a function of magnetic field for the $^2S_{1/2}$ and $^2P_{3/2}$ states, as displayed in figure 2.19.

For each case we need to look at the behaviour for small and large magnetic fields. For small fields, the Zeeman splitting is much smaller than the hyperfine splitting and the states are well represented by the quantum numbers F and M_F. For higher fields, the Zeeman splittings become much larger than the hyperfine structure and the states are better represented by the quantum numbers M_J, M_I. Each case is looked at individually for both the σ_+ and σ_- transitions.

Beginning with the $|g,2,-2\rangle \rightarrow |e,3,-3\rangle$ (σ_-) transition the presence of any π-polarised light allows excitations into either the $|e,3,-2\rangle$ or the $|e,2,-2\rangle$ state. For large fields, figure 2.19 tells us the $|e,2,-2\rangle$ state will be detuned far enough from the laser frequency that it is unlikely to be driven. Typical experimental values for the power broadened linewidth of the cooling light are approximately 15 MHz and consequently the $|e,3,-2\rangle$ state is unresolved, even for the highest fields. However, as the field approaches the strong field regime for both the ground and excited state ($\sim$ 600 gauss) the increased precession frequency of $\underline{\hat{J}}$ about $\underline{B}$ means F is no longer a good quantum number and the $\underline{\hat{I}}$ and $\underline{\hat{J}}$ operators begin to decouple. In this regime M_I and M_J are the good quantum numbers in the excited state. In the ground state there is still some degree of mixing between different M_I states, however, because we are pumping on the stretched state transition this state contains only one term when decomposed into the $|M_I, M_J\rangle$ basis meaning M_I and M_J are both still good quantum numbers and the transition into the $|e,3,-2\rangle$ state is forbidden by the nuclear selection rule, $\Delta M_I = 0$. A similar story holds for the σ_+ transition for which the $|2,2\rangle$ state is forbidden and the $|3,2\rangle$ state is significantly detuned.

For fields below $\sim$20 gauss M_I and M_J are no longer good quantum numbers, even in the excited state, and the power broadened linewidth of the laser allows

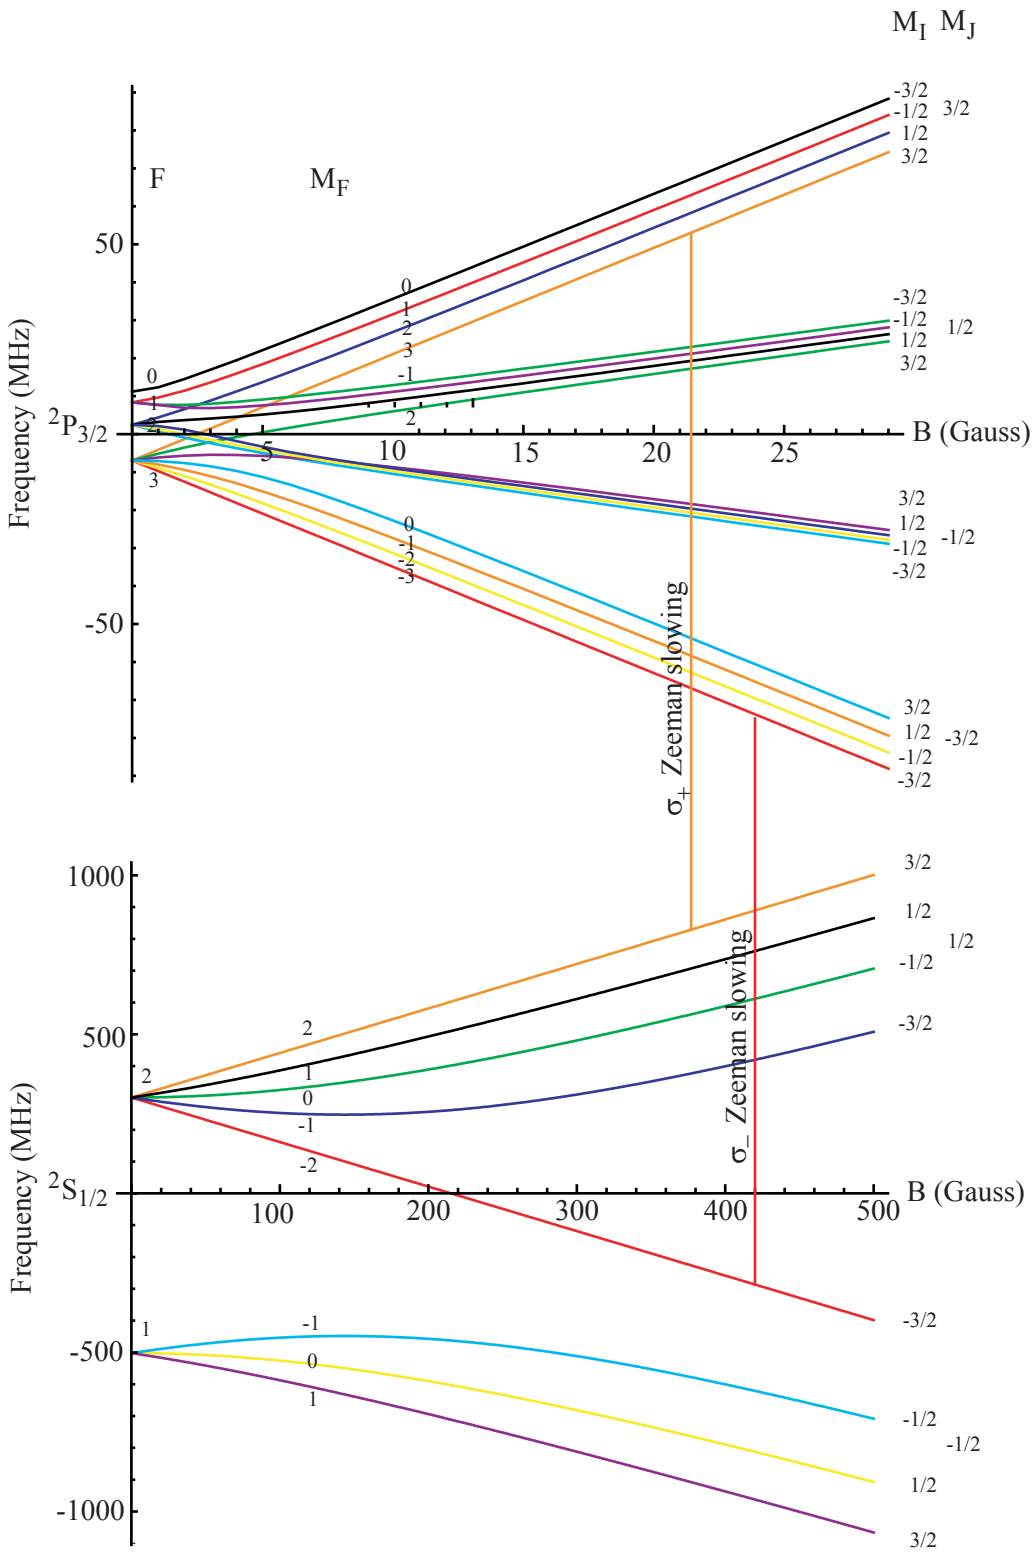

Figure 2.19: The full energy level diagram for the $^2S_{1/2}$ and $^2P_{3/2}$ lithium states as a function of magnetic field. The two possible cooling transitions for Zeeman slowing have been highlighted.

transitions to other states. For the σ_- slower this now allows excitations into the $|e,3,-2\rangle$ and $|e,2,-2\rangle$ states. From here decay routes to the $|g,2,-1\rangle$, $|g,2,-2\rangle$ and $|g,1,-1\rangle$ states exist. The first two states are still in resonance with the cooling light for such low fields, however, atoms in the $|g,1,-1\rangle$ will be permanently lost. To regain these atoms we must introduce a circularly polarised repump beam resonant with the zero-field $|g,1,-1\rangle \rightarrow |e,2,-2\rangle$ transition. However, these two states have almost opposite gradients for their respective Zeeman shifts and the linewidth of the repump laser is only large enough to maintain resonance for fields below 2 gauss. Figure 2.20(a) shows the detuning of the $|g,1,-1\rangle \rightarrow |e,2,-2\rangle$ transition as a function of magnetic field.

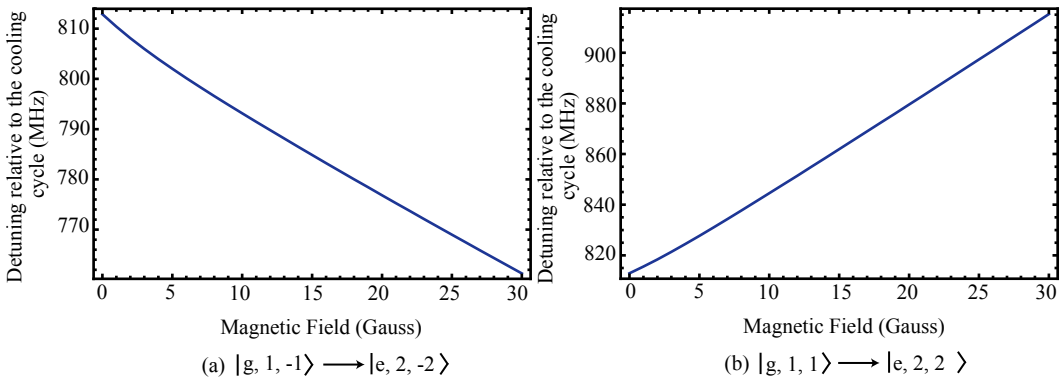

Figure 2.20: The transition frequencies for (a) the $|1,-1\rangle \rightarrow |2,-2\rangle$ transition and (b) the $|1,1\rangle \rightarrow |2,2\rangle$ transition as a function of magnetic field.

A similar story holds true for the σ_+ slower: excitations to the $|e,3,2\rangle$ and $|e,2,2\rangle$ states now fall within the natural linewidth of the transition for small fields and offer decay routes to the $|g,2,2\rangle$, $|g,2,1\rangle$ and $|g,1,1\rangle$ ground states, the latter of which is lost from the cooling cycle. The detuning of the $|g,1,1\rangle \rightarrow |e,2,2\rangle$ transition increases even more rapidly as a function of field such that for fields beyond even one gauss the repump light is no longer resonant; figure 2.20(b).

To determine how detrimental this rapid detuning is we need to look in more detail at the transition rates for the $|g,2,2\rangle \rightarrow |e,2,2\rangle \rightarrow |g,1,1\rangle$ and $|g,2,2\rangle \rightarrow |e,3,2\rangle \rightarrow |g,1,1\rangle$ transitions as a function of magnetic field. Figure 2.21 depicts these rates for σ_+ light with a 5% (red), 10% (yellow), 15% (green) and 20% (blue) π-polarisation component for the two transition routes. We see that the rate for the bad transition falls of rapidly as the field increases. This confirms the assertion given above that these transitions can be neglected for fields much greater than 20 G.

Note that the transition rate for the event $|g,2,2\rangle \rightarrow |e,3,2\rangle \rightarrow |g,1,1\rangle$ is strongly suppressed below 5 gauss. This is because the decay requires a change in the F quantum number of -2, which is forbidden by the selection rules. As the field increases mixing between the $|e,3,2\rangle$ and $|e,2,2\rangle$ occurs and the decay channel

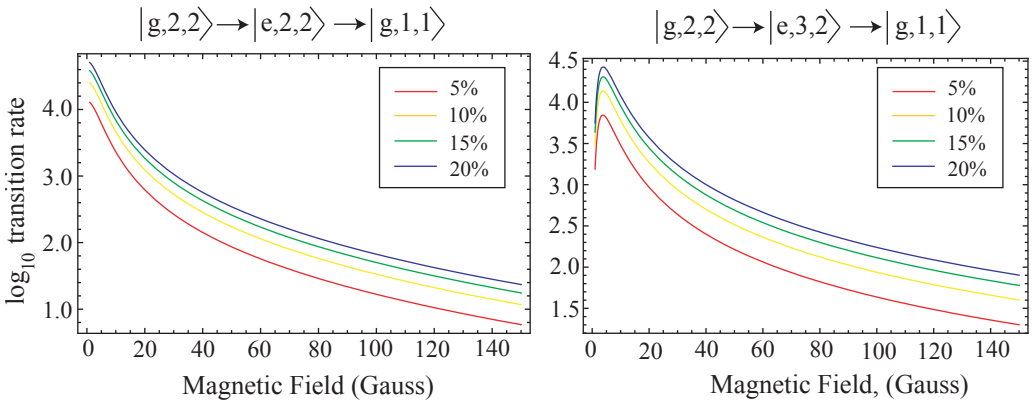

Figure 2.21: A plot of the transition rates as a function of magnetic field for the two decay channels pumping atoms into the $^2S_{1/2}\,|F=1,\ M_F=1\rangle$ state. The results are plotted for a total beam power of 25 mW with each line corresponding to various percentages of π-polarisation impurity.

to the $|g,1,1\rangle$ state is opened up. For larger fields the ΔM_I selection rule turns on and the transition rate is increasingly suppressed.

Given the transition rates one can calculate the probability of an atom reaching the state $|g,1,1\rangle$ as it traverses the Zeeman slower. If $R(B)$ is the total rate for the process, then the probability of an atom making a transition to the state, $|g,1,1\rangle$, is given by

$$P = 1 - \mathrm{Exp}\left[-\int_0^\tau R(B(t))\mathrm{d}t\right]. \tag{2.26}$$

As an example, for the σ_+ case we will model a slower able to decelerate atoms from $1\mathrm{kms}^{-1}$ to $50\mathrm{ms}^{-1}$ with a zero crossing at $269\ \mathrm{ms}^{-1}$, equating to a 401.75 MHz offset. Using equation 2.26 we find that for a beam with $1.5\mathrm{mWcm}^{-2}$ of π-polarisation (approximately 10% of the total intensity), $\sim$55% of all the atoms in the slower are pumped into the $|1,1\rangle$ state. In the experiment, in the absence of a Zeeman repump beam a reduction in the atom number of $\sim 75-85\%$ was observed. These discrepancies may be attributed to imperfections in the polarisation of the repump light or to oversimplifications in the models above.

Since the repump is only resonant at low field, we are interested in the fraction of this 55% lost in the lowest field region relative to the higher field regions. By choosing a window between two field values, labelled B and -B (i.e. symmetric about the zero-crossing) and normalising to the total number of atoms lost we can see where the atoms are lost in the slower. This is depicted in figure 2.22.

It is clear from figure 2.22 that, as expected, the majority of atoms are lost in the low field region, near the zero crossing. For 1.5 $\mathrm{mW\,cm}^{-2}$ of π-polarisation light,

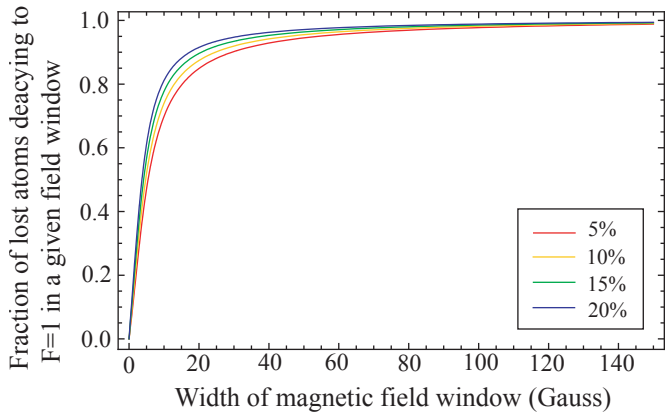

Figure 2.22: The probability of pumping an atom into the ${}^2S_{1/2}|F = 1,\ M_F = 1\rangle$ state within a given section of the Zeeman slower, bounded by a magnetic field window centred at the zero-crossing of the slower.

90% of the atoms lost to the $|g, 1, 1\rangle$ state are lost in the small window between -20 and 20 gauss. As mentioned previously, however, the $|1, 1\rangle \rightarrow |2, 2\rangle$ (σ_+) repump transition detunes rapidly from resonance with increasing field. The amount of repump power available is well below the saturation intensity meaning the linewidth for the transition is only 6 MHz, so at fields beyond 1.5 gauss the detuning is already larger than the linewidth. If, however, we chose the repump to have purely π-polarisation we are able to pump into the $|e, 2, 1\rangle$ state for fields as high as 20 gauss before detuning more than a linewidth off resonance, thereby recovering the majority of the atoms originially lost to the $|g, 1, 1\rangle$ state. From here the atoms may only decay to the $|g, 2, 2\rangle$, $|g, 2, 1\rangle$ or $|g, 1, 1\rangle$ state, all of which can be recovered into the cooling cycle.

The above arguments can also be followed through for the σ_- transition. For those atoms driven to the $|g, 1, -1\rangle$ state π-polarised light may be used to excite them back into the cooling cycle via the $|g, 1, -1\rangle \rightarrow |e, 2, -1\rangle$ transition. With careful tuning the transitions will also remain close to resonance with the repump beam for fields as high as 100 gauss, allowing a more efficient re-population of atoms into the cooling cycle.

For the $|g, 2, 2\rangle \rightarrow |e, 3, 3\rangle$ (σ_+) transition there is a small addendum to be noted for transitions into the $|e, 3, 2\rangle$ state in the large field regime. Although the state, $|3, 2\rangle$ appears to be detuned far from the laser for large enough fields, some slower atoms will find themselves falling into resonance with this transition due to their smaller Doppler shifts relative to those at the capture velocity of the slower. Atoms excited into this state may subsequently decay into either the $|g, 2, 2\rangle$ state(as required) or the $|g, 2, 1\rangle$ or $|g, 1, 1\rangle$ state. The $|g, 2, 1\rangle$ state is, as before, still in resonance with the cooling light so atoms in this state will not be lost. Some atoms will, however, be driven into the $|g, 1, 1\rangle$ state, which is not resonant with either the

cooling or repump light for fields larger than 20 gauss. For high fields the slowest atoms are lost. As we move along the slower towards smaller fields the velocity of atoms resonant with the $|g,2,2\rangle \rightarrow |e,3,2\rangle$ transition will increase. Taking into account the effects of the repump light it can be shown that those atoms with velocities below 250 ms^{-1} will be permanently lost. The fraction of atoms entering the slower in this velocity region, however, corresponds to less than 0.8% of the total flux and can be considered negligible.

In conclusion, while there are some disadvantages to choosing to cool the atoms with σ_+ light these are small and in total only decrease the atomic flux by less than 1% of the total number decelerated by the Zeeman slower. In contrast, the practical advantages for choosing a decreasing field slower greatly favour the decision to cool on the $|2,2\rangle \rightarrow |3,3\rangle$ transition. The following section discusses how one can maximise this flux through careful choice of the length of the slower.

2.5.2 Length Calculations

When designing the slower we wanted to choose a length that optimises the flux of decelerated atoms. The relationship between the efficiency of the slower and its length, however, is complicated and requires some detailed discussion. As the length of the slower increases we are able to remove increasing amounts of kinetic energy from the beam; a longer slower means we can capture a larger fraction of the forward velocity distribution. However, the atomic beam is divergent and so longer slowers will capture a smaller fraction of the transverse velocity distribution since there is no cooling in this direction. Consequently, there will be a length which optimises the number of atoms traversing the decelerator.

The transit time, τ, for an atom leaving the oven with an initial velocity, v_i, to the exit of the Zeeman slower is given by

$$\tau(v_i) = t_{OZ} + t_Z + t_d \,, \tag{2.27}$$

where t_{OZ} is the time taken to reach the slower entrance, t_Z is the time taken to travel in the slower before it is resonant with the radiation and t_d is the time taken to decelerate the atom to the capture velocity of the MOT. Each of the three times in equation 2.27 has some dependence on v_i that can be easily derived from the equations of motion. The time, t_{OZ}, is simply given by the distance from the oven to the slower divided by velocity of the undecelerated atom. The time, t_Z, is given by $(v_0^2 - v_i^2)/(2v_i\epsilon a_{max})$, where v_0 is the capture velocity of the slower. Finally, $t_d = (v_i - u)/\epsilon\, a_{max}$ is the time taken for an atom to reach the target exit velocity, u, undergoing constant acceleration, $\epsilon\, a_{max}$.

An atom will successfully exit the slower if its transverse displacement at the end of the slower is smaller than the radius of the slower, R. As shown by equation 2.13, the number of atoms emerging from the oven in the angular range from θ to $\theta + d\theta$ is proportional to $\cos(\theta)\sin(\theta)d\theta$. Integrating over θ between the limits 0 and Θ, where Θ is the maximum angle an atom can have and still exit the slower, and assuming Θ is small, such that $\sin(\Theta) \approx \Theta$, the total number of atoms emitted is proportional to Θ^2. The total flux then is given by the product of the Θ^2 and the normalised Boltzmann velocity distribution of the beam, integrated over the velocities the slower is able to capture,

$$N(L) \propto \int_0^{v_0=\sqrt{2\epsilon a_{max} L}} \left(\frac{M}{\sqrt{2}k_B T}\right)^2 \left(\frac{R}{v_i\ \tau(v_i, L)}\right)^2 v_i^3 \exp\left[-\left(\frac{M v_i^2}{2k_B T}\right)\right] \mathrm{d}v_i\ , \tag{2.28}$$

where a_{max} is derived from equation 2.3. In the presence of a collimated laser beam the transverse velocity of the atoms, v_r, remains unchanged along the length of the slower. For an atom to traverse the slower the maximum value v_r can take is $R/\tau(v_i, L)$. Thus, Θ is re-expressed in equation 2.28 as $v_r/v_i = R/v_i\ \tau(v_i, L)$.

So far we have not considered the diffusion of the atoms in the transverse direction due to the random nature of the spontaneous emission. This effect tends to increase the transverse displacement of atoms from the axis and is worse for atoms with high initial speeds due to a higher number of scattering events. This will tend to favour shorter slowers. The change in rms transverse velocity, v_T, after a time, t, resulting from spontaneous emission events is given by $v_T(t) = h/(3M\lambda)(\Gamma_s t)^{1/2}$, where Γ_s is the scattering rate. The transverse deviation resulting from this diffusion is then given by

$$\Delta x = \int_0^{t_d} v_T(t)\ \mathrm{d}t = \frac{2h}{9M\lambda}(\Gamma_s)^{1/2}\left(\frac{\Delta v_i}{\epsilon\ a_{max}}\right)^{3/2}. \tag{2.29}$$

To quantify the effects of the diffusion described by this equation we can calculate Δx for a set of typical parameters. The value for a_{max} is derived from equation 2.20. Taking a saturation parameter of 4 we find the the scattering rate for on-resonant light and will use a value of 850 ms^{-1} for the change in v_i. Given these parameters the increase in the radius of the atomic beam cross-section due to scattering-induced diffusion is, $\Delta x = 2.57$ mm. In order to modify θ and therefore $N(L)$ we can simply subtract the value of Δx from the slower radius in equation 2.28. The green curve in figure 2.23 shows the number of atoms that can successfully traverse the Zeeman slower as a function of the slower length including the effects of the transverse diffusion due to spontaneous emission. For very short lengths only the slowest atoms

can be decelerated to the MOT capture velocity and it is advantageous to increase the length. For lengths greater than 0.5 m this is no longer the case and as the length is increased, the number of atoms lost due to the beam divergence becomes greater than the number gained by increasing the slower's capture velocity.

In the above calculations, we have not taken into account the transverse cooling that comes about if the laser is converging towards the entrance of the slower. This is a useful method for increasing the flux of atoms traversing a slower. Other methods such as the employment of a 2D molasses before the slower have also been used to collimate the atomic beam [64] but in general are difficult to implement. We can account for such a converging beam properly with the use of a sophisticated numerical model, however, to keep things simple, we can think about the case where the convergence of the laser beam is matched to the divergence of the atomic beam. In this case the laser beam will fill the slower aperture at its exit and will converge to a waist near the oven. In this model the photon scattering events reduce both the longitudinal and the transverse velocities in such a way that the ratio of the two is kept constant. The atoms then just follow straight line paths and the equation for the flux is easily modified. The value of Θ becomes $R/(d_{OZ}+L)$ where d_{OZ} is the distance from the oven to the slower entrance. Figure 2.23 depicts a comparison of equation 2.28 with and without the application of a converging laser beam. In both cases the effects on the transverse displacement due to spontaneous emission have been included. The radius of the slowing exit aperture for this figure was chosen to be 19 mm.

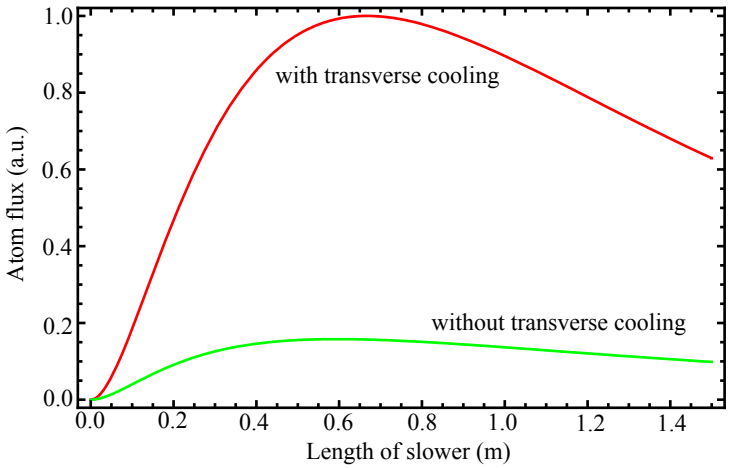

Figure 2.23: The normalised atom flux leaving the slower as a function of Zeeman slower length with and without the effects of transverse cooling from a converging laser cooling beam.

The effects of a transverse cooling component are clearly very significant, increasing the maximum flux by a factor of six and increasing the optimum decelerator length from 0.50 m to 0.65 m. Note, however, that this curve is not a universal curve

but will depend on the distance between the oven and the slower entrance and the diameter of the slower's aperture. The choice of diameter for the slower aperture is dependent on the trapping volume of the MOT. In practice, a typical trapping volume will only be approximately 20 mm in diameter due to practical limitations such as the available laser beam intensity and the size of the viewports. In addition, to maximise the number of trapped atoms, the trapping volume should ideally be placed as close to the exit of the slower as possible. Given that atoms exiting the slower will be diverging it is futile to choose a slower diameter any larger than the diameter of the trapping volume and detrimental to choose one considerably smaller. Given the availability of standard tube and flange sizes then, an inner diameter of 19 mm was selected. The distance from the source to the slower is mainly determined by the length of the two intervening 6-way crosses, however, an additional $15 - 20$ mm will be required before the start of the slower to allow for assembly of the chambers.

The peak of the red curve is in fact quite broad and we can still capture over 80% of the atoms for slower lengths between 0.35 m and 1.10 m. The choice of length for our slower was 0.51 m, corresponding to a capture velocity of 900 ms^{-1} and a target velocity of 50 ms^{-1}. This equates to a flux of 95% of the peak of the curve in figure 2.23. The following sections discuss some of the practical limitations that must be considered when implementing such a slower.

2.5.3 Vacuum Conductance

The conductance, C, in litres per second, of a tube of length, l, and radius, r, measured in mm, for a molecular effusive beam at temperature, T (in Kelvin), and molecular mass, M, is given by the equation [65]:

$$C(r, l, T, M) = \frac{r^3}{l} \frac{1}{1 + \frac{8r}{3l}} \sqrt{\frac{29T}{300M}}\,; \tag{2.30}$$

The operating temperature range for the lithium oven is typically between 700 and 800 Kelvin. Taking the upper bound of this range we can estimate from equation 2.12 a pressure inside the oven of $\sim 1 \times 10^{-2}$ mbar. The pressure in chamber 1, directly outside the oven, can thus be inferred from mass flow conservation if we assume that the quantity of the gas moving from chamber 1 to chamber 2 is negligible relative to that exiting via the pump. Equally, we will assume that there is no significant gas flow from the second chamber back to the first. Conserving mass flow into and out of the chamber we find a pressure, P_1, in chamber 1 given by $P_{oven}C = P_1 s$ where C is the conductance of the oven aperture and s is the pumping speed of the turbo pump attached to chamber 1. The dimensions for the

length and radius of the oven aperture are 2 mm and 0.35 mm respectively, giving an expected partial pressure for Li in chamber 1 of $\sim 5 \times 10^{-7}$ mbar. To avoid loading copious quantities of lithium into the turbo pump, a cold copper collection cup with an 8 mm diammeter aperture has been mounted outside the oven exit. Its dimensions are such that any atom colliding with the walls will have had too large an angle to reach the Zeeman slower. Since almost all the collected lithium will permanently adhere to the inner walls we can estimate the fraction of Li exiting the cup. Integrating over a $\cos(\theta)$ distribution (a good approximation for small angles [60]) over the angle subtended by the cup aperture we can show that only $\sim 8\%$ of atoms from the oven leave the cup, giving a partial pressure in chamber 1 of 4×10^{-8} mbar. Experimentally the total pressure in chamber 1 during initial operation measures at $2-3\times10^{-7}$ mbar at operating temperatures. Over the course of several months the chamber is baked and residual oil and water is removed from the oven. Over this time the pressure falls to $3 - 4 \times 10^{-8}$ mbar in accordance with our estimations.

In order to achieve a pressure of $\ll 5 \times 10^{-10}$ mbar in the MOT chamber we require a slower with low vacuum conductance. Typically to achieve pressures below 10^{-9} mbar it is necessary to use an ion pump; the lack of any moving machinery also greatly improves the mechanical stability of the chamber. For this experiment we have used a NOBLE DIODE ion pump from LEYBOLD with a pumping speed of 55 litres s^{-1} on the MOT chamber. However, with only this pump, and without any differential pumping between chambers 1 and 2 we would require a Zeeman slower with a length of 200 metres to achieve the desired pressure in the MOT chamber! To maintain a suitable vacuum in the MOT chamber for a slower whose length is of the order of 0.5 metres some differential pumping must be included before the slower. To this end a differential pumping tube is placed between the first and second chamber; see figure 1.2. The length and radius of the tubing chosen were 100 mm and 2.5 mm respectively in order to obtain a suitable pressure in chamber 2 and thus the MOT chamber. The dimensions were chosen such that the tube will not significantly impede any atoms that would have otherwise exited the Zeeman slower for slower lengths greater than 0.5 m. The effect of the pumping tube on the atom flux as a function of slower length is given in figure 2.24. The flux at 0.51 m is reduced from 95% to 80% when the differential pumping tube is included.

Given this additional differential pumping stage the pressure in chamber 2 is expected to be $\sim 3 \times 10^{-9}$ mbar which is reduced to below $\sim 3 \times 10^{-10}$ mbar in the MOT chamber for a Zeeman slower length of 0.5 m or more and a 55 $\mathrm{l\,s}^{-1}$ ion pump on the MOT chamber.

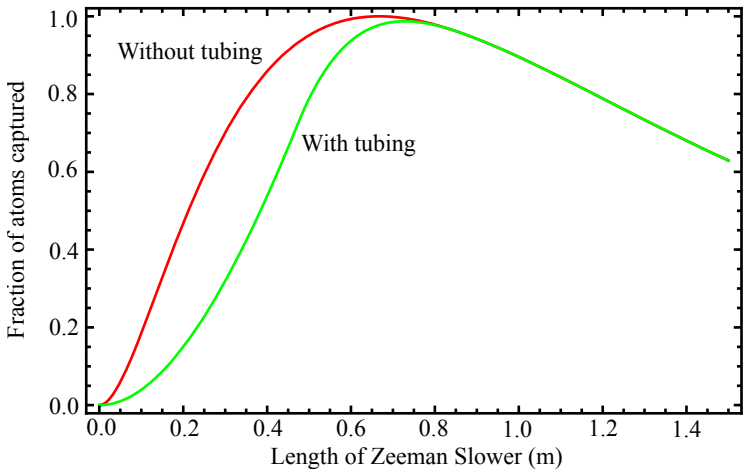

Figure 2.24: The normalised atom flux exiting the slower as a function of slower length with and without a differential pumping tube placed between chambers 1 and 2.

2.5.4 Magnetic Field Optimisation

When constructing a Zeeman slower there are two choices for producing the desired field shape, both with their own advantages and disadvantages. The first slower to be built [43] was wound using a single tapered coil, whose number of layers was gradually decreased along the length of the slower to provide the required field shape. Because the coil is wound from a single wire only one power supply is required, significantly simplifying the set-up. A single tapered coil, however, can only be wound once and does not have the same control over the field shape as a series of coils run in parallel. If the slower is divided into a number of individual solenoids each with a suitable current and number of windings there must be a dedicated power supply for each coil, however, each coil can be precisely tuned and optimised to maximise the atom flux. To achieve a satisfactory match with the ideal field shape ten to twenty coils may be required to cover the length of the slower meaning a large bank of power supplies is required. In our particular set-up we chose a hybrid of the above two methods which has proven to be a very successful design. The slower is divided into a set of coils approximately 40 mm in length. The coils are grouped into several segments run in series, each with a single dedicated power supply. The field in each segment is obtained by choosing an appropriate number of windings for each coil in the segment and a single current value delivered by the supply. Fine control of the field value at each coil is then set by adjusting a 10 Ω variable resistor in parallel with each coil.

In order to simulate the field produced by a given coil a simple program was written in *Mathematica* to calculate the contribution from each loop of wire for a coil with multiple turns and multiple layers. To estimate the parameters required to produce fields up to a maximum of 650 Gauss it is necessary to decide on a coil

wire diameter and an inner diameter for the coil former. The latter parameter is preferably as small as possible to reduce the amount of wire needed and hence the power dissipation, while also reducing the distance over the which the field drops off at the end of the slower. The diameter is restricted by several geometric constraints; outside of the Zeeman slower tube itself it is necessary to include room to wrap a heater strap to allow the option of baking the slower once under vacuum; this increases the diameter to 28 mm. It will also be beneficial to include a water cooling jacket for the coils, the design of which is described in detail in section 2.5.5. This increases the total diameter to 38.2 mm over which a coil former with a thickness, 1 mm, can slide, setting the coil inner diameter to 40.2 mm. For the wire itself, after testing several designs, it was decided a wire with rectangular cross-section would be preferential as it provides much tighter packing, allowing for better power dissipation, and produces a neater coil that is much easier to wind. The availability of wire with a square cross-section is limited, however, the Scientific Wire Company were able to provide us with polyester-enameled 2.85×4.08 mm copper wire rated to a temperature of 200 °C. Using these parameters it can be shown that to produce a field of 650 Gauss almost 4200 Amp-turns are required. Given that the length of a coil segment must be relatively short in order have enough control over the shape of the Zeeman field it was decided that a coil with 10 turns, 21 layers, and carrying ∼20 Amps would constitute sensible values for the size of the largest coil. It was shown in testing that such a coil at 20 Amps reaches a temperature of 120 °C with air cooling.

The coils are grouped into sections, with each coil in a given section carrying approximately the same current but wound with a different number of layers as appropriate for the field size. It was decided that the first and last coil of the slower should be assigned individual power supplies to allow for greater control and to accommodate for the larger currents necessary to meet the field specifications at these points; the remaining coils would be operated with as few power supplies as possible.

To simulate the magnetic field generated by a coil we begin with the equations for the radial and axial vector components of the field per amp for a single loop of wire of radius, r_0 and position z_0 relative to some origin [66]:

$$B_z = \frac{\mu_0}{2\pi}\,\frac{1}{\sqrt{(r+r_0)^2+(z-z_0)^2}} \times \Bigg(\text{Elliptic}K\Big[\frac{4\,r_0\,r}{(r+r_0)^2+(z-z_0)^2}\Big] + \Big(\frac{r_0^2-r^2-(z-z_0)^2}{(r_0-r)^2+(z-z_0)^2} \times \text{Elliptic}E\Big[\frac{4\,r_0\,r}{(r+r_0)^2+(z-z_0)^2}\Big]\Big)\Bigg), \qquad (2.31)$$

$$B_r = \frac{\mu_0}{2\pi\, r}\,\frac{z-z_0}{\sqrt{(r+r_0)^2+(z-z_0)^2}} \times \Bigg(-\text{Elliptic}K\Big[\frac{4\, r_0\, r}{(r+r_0)^2+(z-z_0)^2}\Big] + \Big(\frac{r_0^2+r^2-(z-z_0)^2}{(r_0-r)^2+(z-z_0)^2} \times \text{Elliptic}E\Big[\frac{4\, r_0\, r}{(r+r_0)^2+(z-z_0)^2}\Big]\Big)\Bigg). \quad (2.32)$$

Here the functions, Elliptic$K$ and Elliptic$E$ are the complete elliptic integrals of the first and second kind, respectively, and r and z are the radial and axial components.

Given these equations it is then simple to construct a coil layer built from sequential wire loops with the same radius, axially separated by the wire width. To calculate the field for a complete coil, several layers are summed, with the radius of successive layers increased by the height of the wire. A segment built from multiple coils is then constructed by calculating the total field produced by an array of neighbouring coils with varying layer numbers but a single current value, length and wire cross-section. The total field is the sum of all the segments in the slower. The process is somewhat a matter of trial and error, however, by fitting the simulated result to the ideal magnetic field shape, given in section 2.5.1, applying some initial guesses for the number of layers for each coil and leaving the currents for each coil as a floating parameter, a very good match can be obtained. With some experimentation the number of layers could be arranged such that many neighbouring coils could carry the same current to within at least 93% of the largest current for that section; this 7% deficit is then compensated for using the variable resistors placed in parallel with each coil.

The number of coils used to divide up the field profile is determined by the degree to which the simulated field and the ideal field are required to match. Because the field is not constructed from a single tapered solenoid but rather the summation of a finite number of individual coils there are ripples in the field that deviate from the desired theoretical value. Figure 2.25 compares the deviations from the ideal field case for a slower constructed from 8 (red curve) and 13 (blue) solenoids.

There is clearly a better match to the ideal field for a slower constructed from 13 solenoids over one constructed from 8. The maximum residual field in the central section of the slower for the 8 solenoid case is 14 gauss, compared to only 4 gauss in the 13 solenoid case. If we choose to build the field profile depicted in figure 2.18 these maximum residual fields correspond to 2.2% and 0.6% of the total field (650 Gauss) for the 8 and 13 solenoid case respectively.

In order for a Zeeman slower with a value of $\epsilon = 1$ to successfully decelerate the atoms the magnitude of these residual fields must be smaller than the power broadened linewidth of the transition. Although we have chosen a value of $\epsilon = 0.5$ we will aim to meet the magnetic field tolerances imposed by a value of 1 to allow

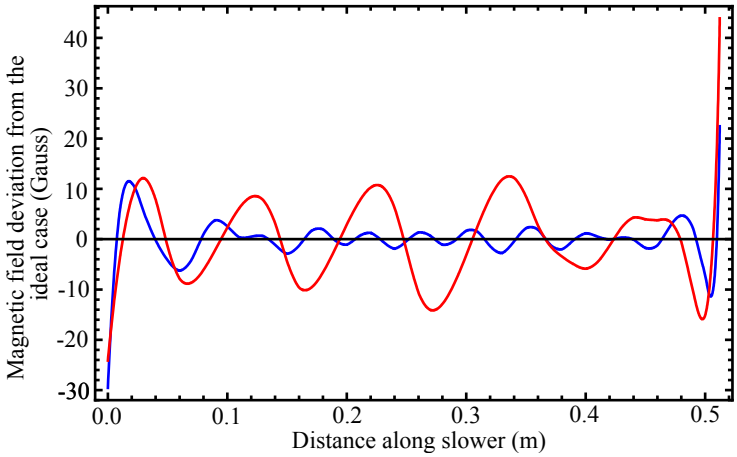

Figure 2.25: The magnetic field residues obtained by subtracting the ideal field from the field of a slower built from 8 (red curve) and 13 (blue curve) coils.

for other imperfections in the experiment, such as a reduced intensity in the wings of the cooling beam. The diameter of the laser cooling beam should match the diameter of the slower at the exit aperture and gradually focus as it approaches the oven. Because of the power lost to the AOMs we can provide no more than 27 mW of cooling power to the atoms, equating to a minimum intensity of $8-9\,\mathrm{mW\,cm^{-2}}$ at the exit of the slower. Calculating the power broadened linewidth one can show that these ripples should deviate by no more than $\sim 8-9$ gauss throughout the slower. To achieve this degree of agreement it was decided to divide the slower into 13 coils, decreasing the length of the coils towards the end of the slower to help compensate for the steeper field gradient, $\frac{\partial B}{\partial z}$.

In the case of a 13 solenoid slower the ripples are much smaller than the required tolerances for all but the fields at the ends of the slower. At the front end ($z = 0$) the intensity of the beam is much higher and the tolerance to the ripples is more substantial. For those atoms lost in the last centimetre of the slower due to field deviations their velocity is already very close to the target velocity and any atom loss will be negligible. Because of the detuning of the laser the field profile is shifted such that there is a zero-field point approximately 47 cm along the slower. Consequently, for the 13 coil slower there is a gap, 33.7 mm in length, close to the zero-field point between the 11^{th} and 12^{th} coil for which no coil is required.

The total field is depicted in figure 2.26 (red) with the desired field plotted in blue. Additional constraints were placed upon the size of the field beyond the end of the slower to try to minimise the field influence from the slower on the MOT fields. In later calculations the mutual effects of the MOT field and the slower field on each other are considered. To maintain a close match to the desired profiles of the two field profiles a compensation coil was introduced. The coil is positioned at the

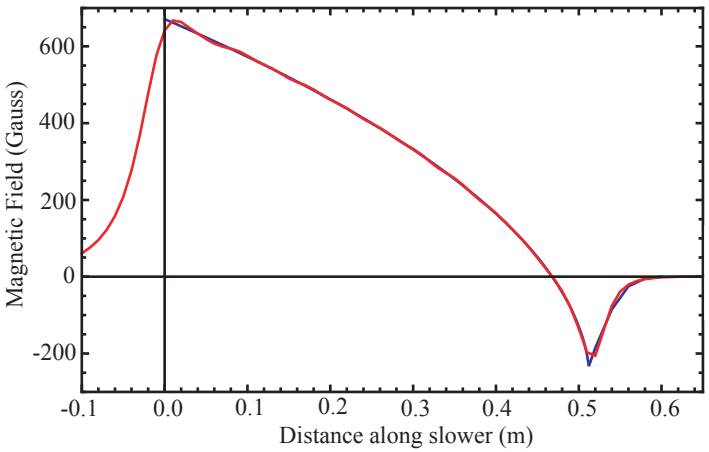

Figure 2.26: The total magnetic field profile for the Zeeman slower as calculated from equation 2.32 for a set of 13 coils run from 4 power supplies. The blue curve depicts the ideal field profile for comparison.

zero-crossing between the 11th and 12th coil with a larger inner radius of 28.5 mm. The coil carries a current of ∼1 Amp, depending on the size of the MOT field.

The relevant data for the solenoids used to produce the 13 coil slowing field is given in table 2.3.

Table 2.3: Data for the coils designed to produce the required magnetic field.

Segment No.	Coil No.	Number of layers	Current (A)	Former Length (mm)
1	1	21	18.57	40.80
2	2	16	11.72	40.80
	3	21	11.65	
	4	16	11.28	
	5	15	11.50	
	6	14	11.07	
	7	12	11.04	
	8	10	11.16	
	9	7	11.10	
3	10	13	5.83	40.80
	11	4	6.12	32.64
	Gap	0	0	33.72
	12	6	-5.65	24.48
4	13[2]	7	-10.75	10.00

Due to the large field required it was decided to dedicate a separate power supply

[2]This coil was wound with 1 mm diameter circular wire to increase current density.

to the first coil. Equally, the field gradient at the end of the slower is very steep so to ensure fine control over the field value at this point the final coil was also assigned a unique supply. This also allows for the dimensions of the final coil to be kept small to prevent the slower blocking any optical access into the MOT chamber. Thus, in total we were able to divide the Zeeman slower into 4 segments (table 2.3) and thus run all 13 coils from only 4 power supplies.

To understand in more detail how sensitive the Zeeman slower efficiency is to the exact field profile, a simulation was written to model the behaviour of an atom undergoing deceleration in a magnetic field. To numerically simulate the deceleration process we need to begin with the mean number of photons scattered by an atom in a given time interval, τ. This is simply the product of the time interval and the scattering rate, where the scattering rate is now a function of the laser power, the magnetic field and the velocity of the atom at a given position in the slower. This is valid provided these parameters do not change significantly during this time interval. Selecting a tolerance of a 1% change in any of the parameters in this time interval we can impose a limit on τ to be no larger than 5μs. The absorption process is a statistical one following a Poissonian distribution. Thus, if in this time interval, the mean number of photons scattered is n, we need to select a value at random from a Poissonian distribution, with a mean n, as a true representation of the number of photons scattered from an atom in order to account for the statistical nature of the process. For these simulations the distribution of the emission profile is assumed to be isotropic and the vector of the emitted photon is picked from a random distribution over a unit sphere. In reality, for an atom travelling along z in a field, B_z, towards counter-propagating radiation with a circular polarisation the emission profile is favoured along $\pm z$ and zero perpendicular to this axis. While this will decrease the transverse momentum diffusion of the beam the total divergence due to this random walk is small and therefore any changes to the emission profile have been neglected in these calculations.

As discussed in section 2.5.2 the atomic beam is diverging as it travels along the length of the slower. In an attempt to reduce the loss in flux due to this divergence the counter-propagating laser is gently converging such that the cones formed by the atomic beam and the light field overlap. The diameter of the Zeeman slowing light was chosen such that the $1/e^2$ radius of the beam matched the radius of the slower. While this means the intensity in the wings will be much lower than the central intensity it was found that by expanding the beam size further to deliver more power to the wings, the decrease in the peak intensity was too great to sufficiently decelerate the atoms. Given these parameters and the values quoted in table 2.3, simulations were then run for various beam powers. The results are given in figure

2.27.

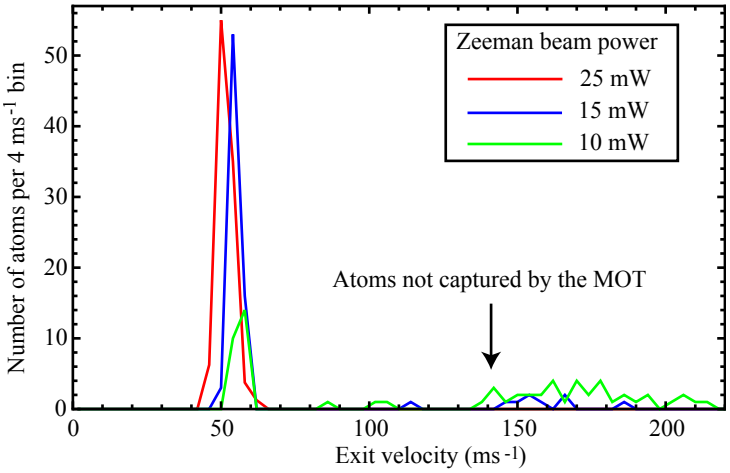

Figure 2.27: Numerical simulations showing the number of atoms as a function of velocity exiting the Zeeman slower, performed for 80 atoms. The simulations were run for beam powers of 25, 15 and 10 mW, for a $1/e^2$ diameter of 19 mm at the largest point. Atoms with a velocity greater than $\sim$60 ms^{-1} are not expected to be captured in the MOT.

The results of the above simulations match well with our experimental findings. One of the key results predicted by the simulations is the presence of a threshold power for the cooling beam, below which the number of atoms being decelerated to the target velocity begins to decrease. Simulations show that nearly all atoms are decelerated to 50 ms^{-1} for laser powers above 15 mW but are steadily lost from the cooling cycle if the power is decreased. The number of atoms exiting at the capture velocity for a power of 10 mW is less than 20% compared to the number at 15 mW. This is attributed to the decrease in the power broadened linewidth falling below the size of the field residues in the slower profile. The experimental data for these findings can be found in the section 2.7.

2.5.5 Construction

The photograph on page 79 shows the various components of the Zeeman slower prior to their final assembly. The compensation coil is not shown. The slower consists of an inner vacuum tube surrounded by a cooling jacket over which the formers for each coil rest. The stainless steel inner tube has an inner diameter of 19 mm and a wall thickness of 1 mm. The total length of this tube is 600 mm to accommodate the coils and allow sufficient room for assembly. Each end is capped by a DN25CF rotatable flange with one flange split to allow the cooling jacket to slide unobstructed over the inner tube. For baking purposes a heater tape was wrapped along the length of the tube adding approximately 8.8 mm to the total diameter.

The cooling jacket is a cylindrical steel tube with an inner diameter of 28 mm

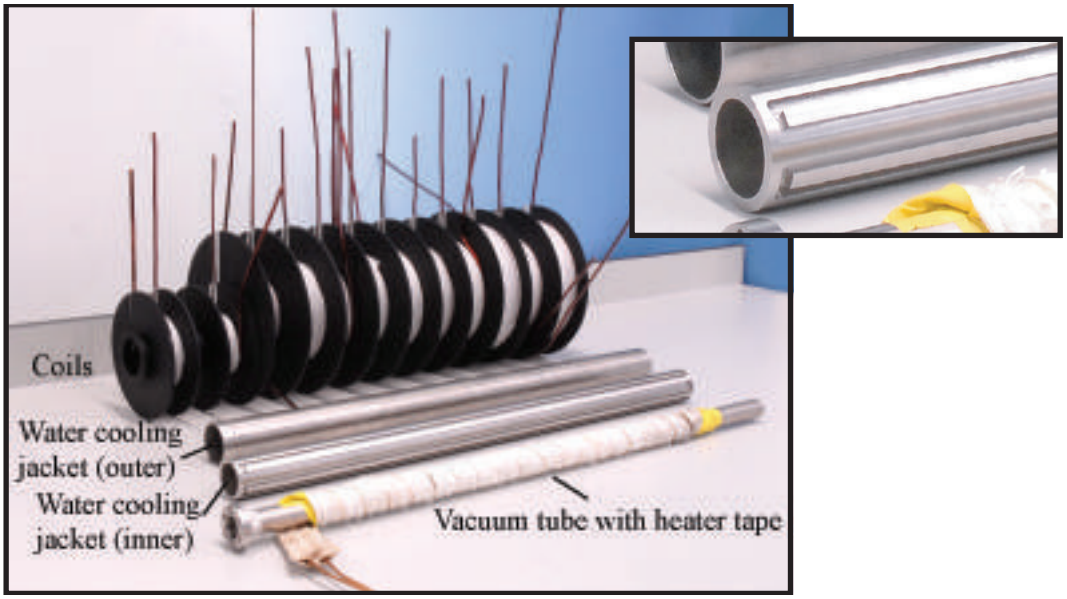

Figure 2.28: A photograph of the Zeeman slower components prior to construction. The inner tubing, the heater tape, the water cooling jacket and completed coil formers can all be seen. The compensation coil is not shown. Inset: Magnification of the cooling jacket, demonstrating the the water channel carved into the surface.

and a wall thickness of 3.55 mm. A 2.61 mm deep channel with a width of 4 mm is bored into the wall of the tube in a zig-zag pattern along the length of the jacket. The pattern can be seen in the inset of figure 2.28 and is machined such that the beginning and the end of the channel both finish at the oven end of the slower. To ensure the channel is water-tight a snug-fitting steel sleeve is slid over the jacket. The thickness of the sleeve is only 1 mm to maximise the thermal contact between the water and the coils. To ensure a tight fit the inner jacket is submersed in liquid nitrogen for approximately 30 seconds before the sleeve is slid over. This allows the channel walls to expand upon re-thermalisation to fill any gaps across which the water might leak. The sleeve is subsequently welded at both ends and two holes are drilled above the entrance and exit to the channel. Two 6 mm diameter lengths of piping are welded on to the holes to provide access for the cooling water. The cooling jacket only provides cooling on the inner side of the coils. For some of the larger coils additional cooling clamps were attached around the coil. The clamps were built from two thin-walled semi-circular hollow aluminium formers designed to fit around the chosen coil. The two halves were screwed together to ensure good thermal contact with the wire.

Each coil is wound on an individual former and tested before being slid onto the cooling jacket. A schematic for a typical former is given in figure 2.29. The formers are machined from a single aluminium piece and anodised to minimise the risk of a given coil shorting to ground. All the cylindrical walls are 1 mm thick to

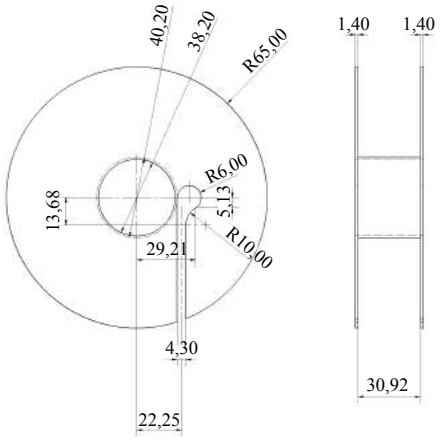

Figure 2.29: A technical drawing for a typical coil former used in the slower. Dimensions are quoted in mm.

improve thermal contact with the cooling jacket while the wings of the former are machined to be 1.4 mm thick for most of the formers. The thickness of the walls is determined by the dimensions of the wire. When winding a coil the wire must enter the former at the inner radius, requiring a hole in one of the side walls of the former. However, because one former should sit flush to its neighbour then the wire must run alongside the coil. The smallest dimension of the wire is 2.85 mm meaning this is the minimum gap allowed between two adjacent coils. Consequently the smallest possible wall thickness is limited to half this value.

The majority of coils were designed to accommodate the square cross-section copper wire with dimensions 2.85×4.08 mm, however, for two of the smaller coils a 1 mm diameter round wire was used to improve current density. The number of layers, currents and dimensions of each coil are given in table 2.3 in section 2.5.4. To wind the coils the former was placed on a lathe and the wire was fed onto the former as it was manually rotated. To push the wire into position a pair of soft teflon levers were used. Since the wire has a natural tendency to spring outwards several layers of high-temperature glass-fibre tape were wrapped around the coil once the final winding was completed. The tape is able to withstand temperatures up to 180 °C.

For a given segment the coils are linked in series with a coarse-graded ammeter placed between each one, as shown in figure 2.30. In addition, a 5 mΩ high-power resistor is placed in series with every coil, with an output socket at each end, across which a voltmeter may be placed. The exact resistance is calibrated by passing a known current through each resistor and measuring the voltage drop across it. A 10 Ω variable resistor is placed in parallel with these components creating a potentiometer that allows fine adjustment of the currents. By measuring the voltage drop across the 5 mΩ resistor the current can be determined with an accuracy of 10 mA. The size of the variable resistor allows a drop of up to 7% in the coil's current,

allowing small adjustments to be made to the local magnetic field value.

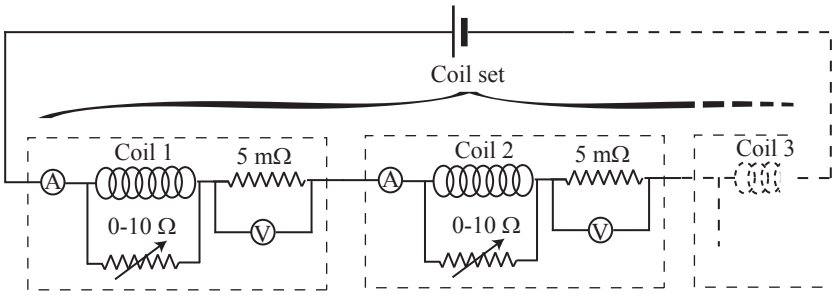

Figure 2.30: A schematic of the circuit diagram used to measure and adjust the current in each coil of a given coil set. Each set contains a different number of coils; see table 2.3

2.5.6 Performance and Optimisation

Before the Zeeman slower was assembled, the axial magnetic field of each coil was individually tested using a Hall probe and compared to the simulations performed previously. The results of these tests can been seen in figure 2.31. The data points indicate the experimental values and the red curves are the theoretical predictions calculated from equation 2.32. As can be seen, the results of these measurements agree very well with theory. Once the coils were tested the Zeeman slower was constructed and the appropriate currents were applied to each coil. However, when the field for the complete slower was measured, the field was smaller than expected with an overall discrepancy of 6% between the results and simulations. The field shape also implied that the slower was almost 20 mm longer than originally designed. It became apparent that these inconsistencies were due to small additional gaps between the formers. Each coil has a tendency to push out the thin wings of each former reducing the field overlap between neighbouring coils and increasing the total length of the slower. To correct for this four pushing screws were inserted between the front end of the slower and the neighbouring chamber in order to compress the coils against the wall of the MOT chamber. With these in place the new field now lay ≈2-3% short of the theoretical results. Figure 2.32 shows the measured field profiles for the compressed and uncompressed slower in comparison to the expected profile.

The results were measured using a Hall probe inserted into the vacuum tube of the slower. Uncertainty in the position and angle of the probe leads to a systematic uncertainty in the field measurement of up to $\sim 1\%$. This accounts for up to 50% of the discrepancy between measurement and calculation. Uncertainties in the currents for each coil and deformities in the coil former walls are thought to account for the remainder of any discrepancies between the measured and expected field. To check

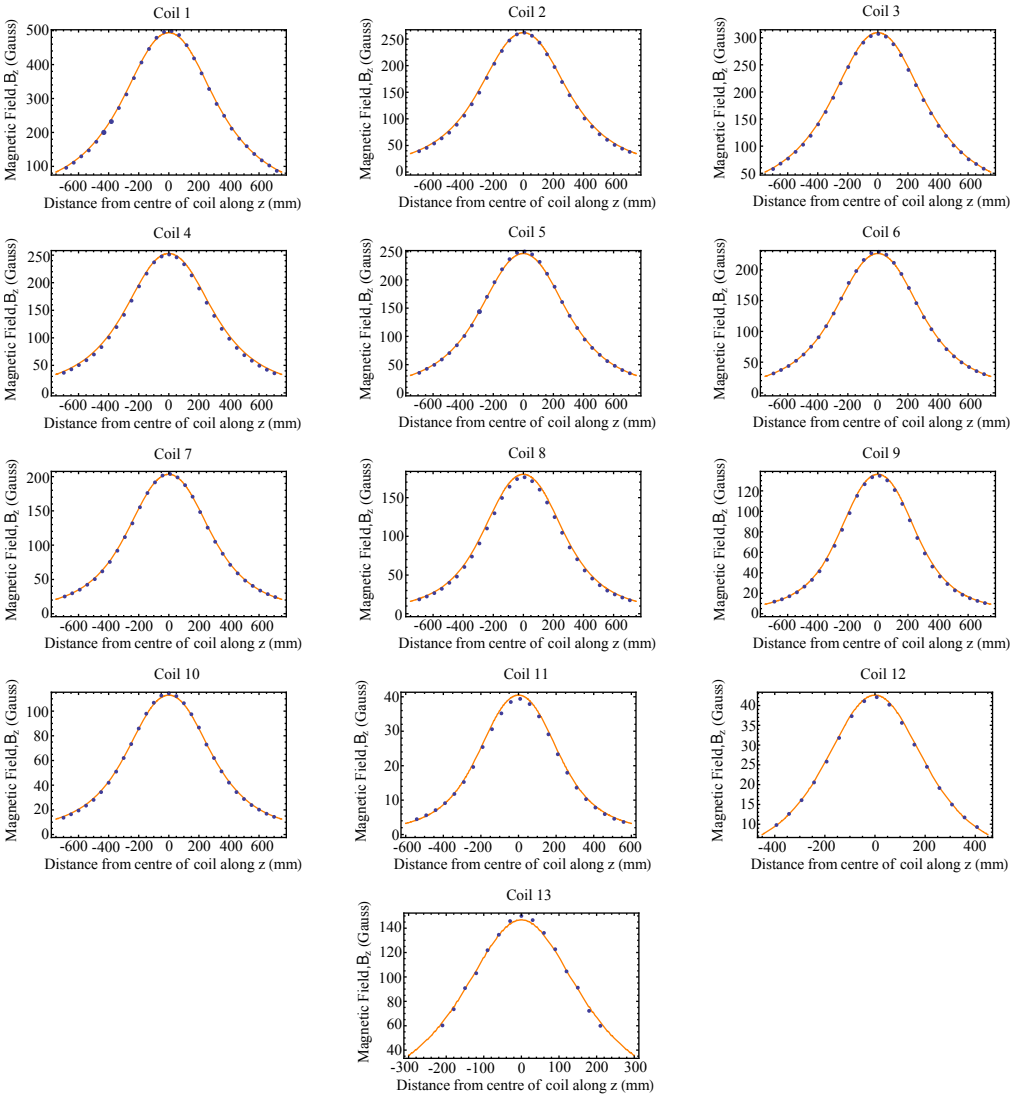

Figure 2.31: Comparison of the theoretical and experimental magnetic field data for each of the 13 coils used to construct the Zeeman slower

whether the above incongruities would be detrimental to the slower's performance the measured field was fed into the numerical deceleration simulations for various laser powers below 25 mW. For the measured field the number of declerated atoms was found to be approximately the same as for the ideal field case, however, the final velocity was a few metres per second faster. This is not thought to be detrimental to the success of the magneto-optical trap.

The performance of the Zeeman slower was not tested separately from the rest of the experiment. To infer whether the slower was working as expected the set-up for the MOT was completed and the trap was tested. Upon the successful trapping of lithium the Zeeman slower could be switched on and off, either via the laser cooling beams or the magnetic fields and the consequences to the trapped atoms

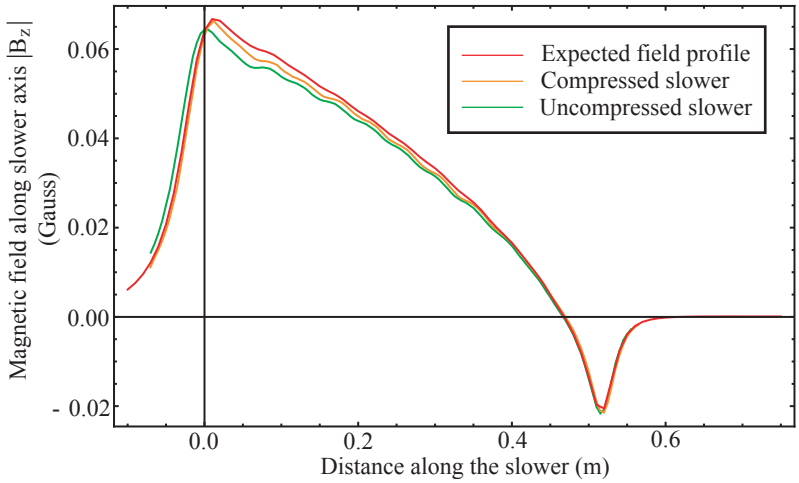

Figure 2.32: The magnetic field on the axis of each of the 13 coils used to construct the Zeeman slower. The red curve depicts the calculated expected field, given in figure 2.26, the green and orange curves are the joined data points for the measured field before and after compression of the slower, respectively.

observed. Unsurprisingly, without the Zeeman slower operating we were not able to trap any lithium. It was, however, possible to switch off some of the end coils while still maintaining a trapped cloud, albeit with one to two orders of magnitude fewer atoms, based on fluorescence measurements. This highlights the importance of decelerating all the atoms to below the capture velocity of the MOT, however, also implies that there is some probability for the trap to capture atoms travelling at velocities significantly larger than the anticipated capture velocity. The MOT number was also compared with and without the Zeeman repump light present. A drop in atom number of $\sim$45% is expected based on the theoretical discussions[3] covered in section 2.5.1.

To optimise the Zeeman slower the currents from each power supply were separately tweaked and the response to the number of trapped atoms was measured. The parameter space is too large to find the absolute optimum for each coil, particularly because the field from each coil influences the field of many of the neighbouring coils, providing no simple systematic method for optimisation. It was found that the trapped atom number was most sensitive to the field from the final coils. Thus, the current from the power supplies for the last coil, the compensation coil and the supply controlling the last group of three coils were scanned and the corresponding number of atoms optimised accordingly. These adjustments led to a factor of 2 increase in the atom number. The optimal currents for these supplies were significantly different to those expected from theory, the reasons for which are not fully understood. The currents supplied to the first two coil sets seemed much closer to

[3] In practice a drop of $75-85\%$ is measured; this discrepancy is thought to be due to the errors in the polarisation of the repump beam.

optimal and the atom number was much less sensitive to changes in the field. To calculate the field produced from these currents the current in each coil was measured by recording the voltage drop across the 5 mΩ resistor. The corresponding field profile is plotted in figure 2.33 (blue) and compared to the original (red). The insert plots the difference between the fields as a function of distance along the slower.

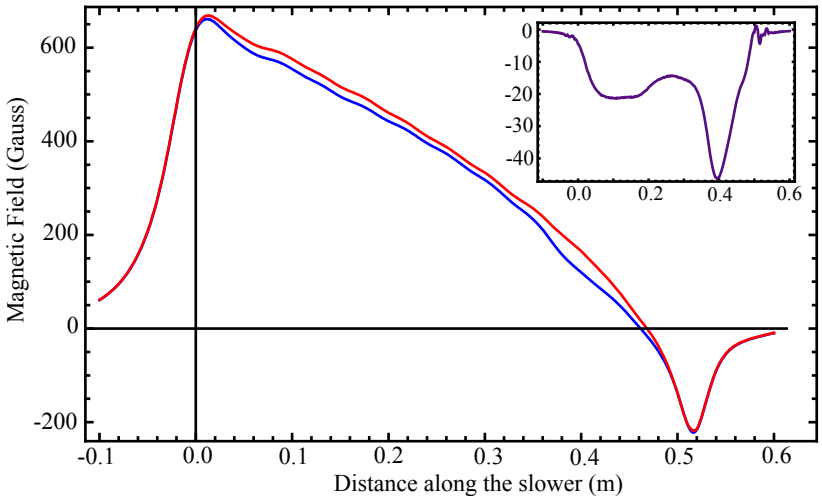

Figure 2.33: The experimentally optimised full magnetic field profile of the slower(in blue) as derived from the currents measured in each coil, compared to the field produced by the coils for the theoretical ideal currents. The insert depicts the field difference between the two curves.

Differences between the field profiles exceed 40 gauss in places making it difficult to reconcile the experimental success with theoretical expectations. The exact field shape, however, has only been estimated from measured current values rather than directly determining the magnetic field using the Hall probe. Errors in the calibration of current readings may account for these discrepancies. For example, at the start of the slower where the coil currents have not been adjusted there is still a $10 - 20$ gauss offset, suggesting the measured currents are slight underestimates. If this were true along the whole slower the largest offset is reduced to only 20 gauss. This may well be within the tolerance of the slower, given the small value chosen for ϵ and the power broadened linewidth of the transition line.

This idea is somewhat supported by simulation data. If this new field is fed into the deceleration simulator for laser powers used in experiments the result remains that all the atoms are still decelerated to below the capture velocity of the MOT. There is a small increase in the prediction for the exit velocity of the atoms from 50 to 56 ms^{-1} for the experimentally optimised profile compared to the ideal field, but this is still within the capture velocity range of the MOT. The final current values for the experimentally optimised field are given in table 2.4. Note that the currents for the coils belonging to segment 3 are 20 to 25% smaller than originally designed while the compensation coil has had its field direction reversed.

Table 2.4: The currents supplied to the coils of the Zeeman slower after optimisation.

Coil Number	1	2	3	4	5	6	7
Current, A	18.64	11.20	11.20	10.83	10.97	10.71	10.69
Coil Number	8	9	10	11	12	13	CC [4]
Current, A	10.71	10.70	3.80	4.33	-3.67	-10.87	-1.67

2.6 The Magneto-Optical Trap

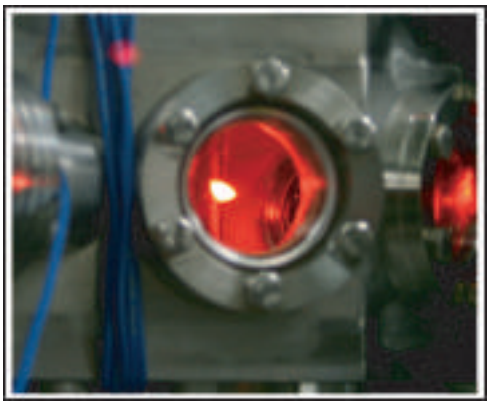

Figure 2.34: A true-colour photograph of a cloud of $\sim 1 \times 10^{10}$ neutral lithium atoms trapped in our MOT.

The following section discusses the design and construction of the magneto-optical trap. Section 2.7 onwards will report on the characterisation of the MOT and present the results for the trapped atom number and temperature. The largest number of atoms trapped in the F=2 ground state was measured to be 2×10^{10} atoms. The cloud is approximately spherical with a diameter of ∼7 mm equating to a maximum density of 1.1×10^{11} atoms cm^{-3}. The minimum temperature measured under these conditions was 0.85 mK at a detuning of four linewidths.

2.6.1 The Chamber Design

The chamber built to house the magneto-optical trap was designed to optimise the number of trapped atoms. To this end, a number of factors have been taken into account, including beam sizes, the magnetic field gradient and the distance of the trap from the end of the slower.

The flux of atoms entering the trapping volume depends on the divergence of the atomic beam exiting the slower. The divergence of the decelerated atomic beam is ∼86 mrad as it exits the slower. For a 19 mm diameter beam exiting the slower with a divergence of almost 100 mrad it is important that the MOT trapping volume is kept

[4] Compensation coil

as close to the end of the slower as possible. One parameter limiting this distance is the size of the trapping beams for the MOT since larger beams move the trapping centre further from the slower. The trapping region is closest to the end of the slower when two of the three perpendicular MOT beam axes lie at 45 degrees relative to the axis of the slower. The third beam axis then lies perpendicular to both of the beams and the axis of the slower; see figure 2.35(b). However, increasing the size of the beams increases the size of the trapping volume and the trap capture velocity, though, for a given power will reduce the intensity. Since all of these parameters influence the total number of atoms that can be trapped a compromise must be found. It has been shown experimentally [67], [68] that atom number increases linearly with intensity until it reaches $1-2\times I_{sat}$ per beam and subsequently begins to plateau as the transition saturates. However, equally, a strong dependence on beam diameter has been demonstrated, with atom number continuing to increase with the $1/e^2$ beam diameter even beyond the point where the intensity per beam drops below saturation. This trend is not expected to continue, however, for very large diameter beams. Although the results in [67] and [68] are reported for a trap loaded from a vapour cell the dependency laws are expected to be similar for a trap loaded from a decelerator. For this experiment we chose a $1/e$ beam diameter of 25 mm thereby maximising the beam size while avoiding having to purchase expensive polarisation optics larger than 1". With the power available and using three retro-reflected beams rather than six independent beams we were able to provide a peak intensity of 2.65 $\mathrm{mW\,cm^{-2}}$ per beam, slightly larger than the saturation intensity. To date we have not performed any measurements on the effects of beam diameter on the number of trapped atoms, however, section 2.7 discusses measurements of the atom number as a function of beam intensity for a constant diameter.

Having set a beam size, the ideal solution would be to design a glass cell that can be attached to the end of the slower and provide complete optical access. There are several problems with this, however. Glass cells are typically fragile and there is a significant force compressing the Zeeman slower against the walls of the cell. Secondly, without any pumping on the cell the pressure will be too high and adding a pump is difficult in terms of providing a seal and managing its weight. Finally, and perhaps most crucially, lithium attacks glass quickly, causing it to become opaque and eventually weakening its structural integrity. As a result of these impracticalities we chose to follow a more conventional route and design a steel chamber. To minimise any magnetisation of the chamber we chose to use 304L stainless steel; this material has the smallest magnetisability for any stainless steel. Other geometrical restrictions such as available flange sizes, wall thicknesses and so forth must also be taken in to account for the final chamber design, however, such idiosyncrasies are

too detailed for discussion here. In the final design the distance from the end of the slowing field to the centre of the trapping region was designed to be 110 mm. A simplified schematic of the chamber is given below, including the beam paths and the magnetic coils.

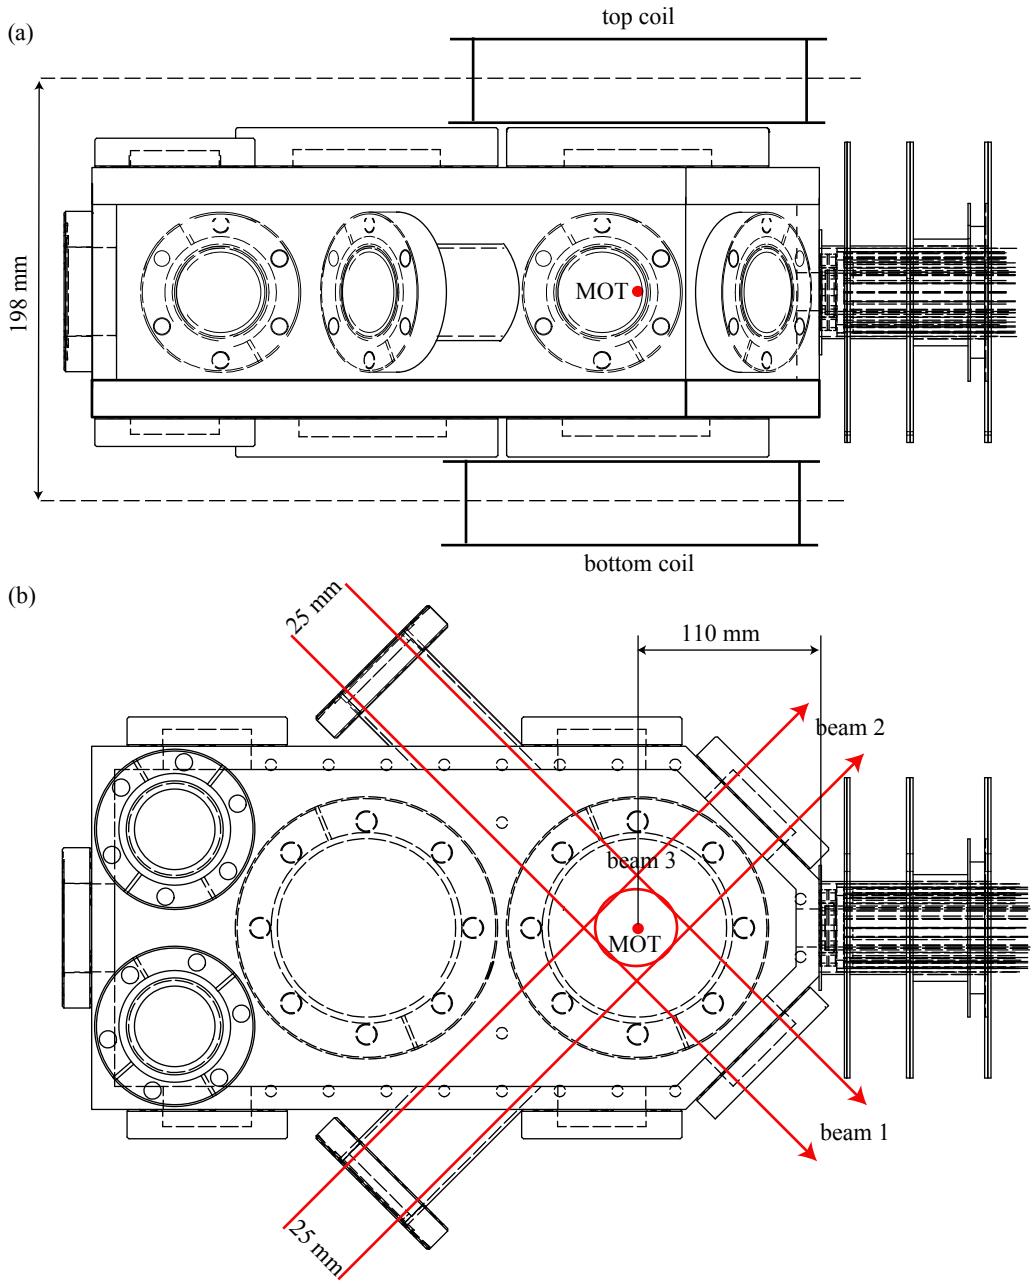

Figure 2.35: A schematic of the MOT chamber, illustrating the trapping beams, the position of the MOT cloud and the trapping coils, (a) from the side, (b) from above.

The effectiveness of the magneto-optical trap depends on both the radiation pressure and the size of the magnetic field gradient. In turn, the gradient depends on both the size of the field produced by each coil in the anti-Helmholtz pair configuration and the coil separation. It is therefore beneficial to keep the coil separation as

small as possible to reduce the amount of current required by the trap. Since, from a practical stand point it is much easier to keep the MOT coils outside the vacuum, the chamber height was kept as small as possible, again restricted by the diameter of the flanges and the thickness of the walls. The distance between the centres of the coils was set to be 198 mm, in order to produce field gradients up to 15 $\mathrm{G\,cm^{-1}}$ for a current of 20 Amps.

2.6.2 MOT coils

The optimal magnetic field gradient for a MOT is dependent on the trapping beam intensity and detuning and is typically experimentally determined due to the difficulty of calculating the forces on an atom in an environment with such complex polarisation gradients and a spatially varying magnetic field. The optimum field value is a consequence of the play-off between maintaining a large enough restoring force on the atoms to keep them trapped, without reducing the capture volume to a point that becomes detrimental to the atom number. Typical field gradients for lithium MOTs are between 5 and 15 $\mathrm{G\,cm^{-1}}$, so a coil capable of providing such a gradient was designed.

The coils for the MOT field are wound with the same rectangular cross-sectional wire used for the Zeeman slower in order to simplify the winding procedure and aid power dissipation. The inner diameter of the coil was set to 60 mm providing an almost constant gradient over 20 mm about the trapping centre, without restricting optical access to the chamber. Given the radius and separation of the coils, a value of 2400 Amp-turns is required to produce a magnetic field gradient of 15 $\mathrm{G\,cm^{-1}}$. Each coil in the pair is wound around an anodised aluminium former with 12 layers and 10 turns per layer. The coils are connected together in series and are run in an anti-Helmholtz configuration at typical operating currents of $12-15$ A, though they can comfortably maintain currents of up to 20 A without exceeding the temperature limit in order to achieve a gradient of 15 $\mathrm{G\,cm^{-1}}$.

Because of the proximity of the trapping fields to the end of the slower there is some interference between the two field profiles. The decreasing field at the exit of the slower, however, can be smoothly transformed into the profile from the MOT coils along the axis of the slower without distorting either of the desired field profiles. Off-axis this transition is not perfect, however, the distortion is small and should not significantly affect the atom flux reaching the trap. To cancel the influence of the MOT fields on the Zeeman field a compensation coil was placed at the zero-crossing point along the slower. The current in this coil may be changed to compensate for different MOT field gradients. The final field along the axis of the slower leading up to the trapping region is given in figure 2.36.

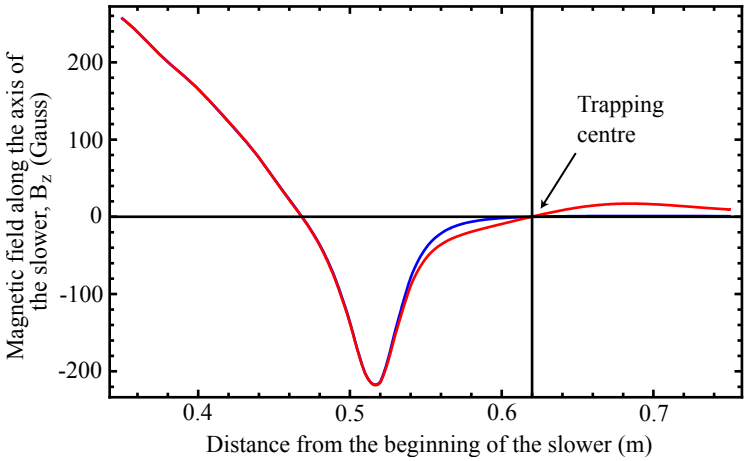

Figure 2.36: The magnetic field profile in the transitional period between the end of the slower and the centre of the trapping volume, in red. The blue curve gives the original field profile without the MOT fields or compensation coil.

Here both the profile of the Zeeman slower and the gradient for the MOT field remain unaltered for a field gradient of 15 G cm^{-1}. Similar profiles can be obtained for smaller field gradients by adjusting the current in the compensation coil.

2.7 Atom Number

In the following section the measurements of the number of trapped atoms will be discussed. I will report the results of both absorption and fluorescence imaging and the atom number dependence on oven temperature, laser intensity, laser detuning and the magnetic field gradient in the MOT. The largest number of atoms were found for a MOT beam detuning of $4.1 - 4.4$ linewidths and a magnetic field gradient of $6 - 7$ Gcm^{-1}, while the atom number continued to increase with oven temperature within the limits of our current oven design.

2.7.1 Fluorescence measurements

The first measurements of the atom number were recorded using the fluorescence of the atom cloud. The fluorescence from the cloud was imaged onto an amplified photodiode using a 100 mm focal length lens placed 120 mm from the MOT and 150 mm from the photodiode. The absolute atom number can then be determined. These fluorescence measurements are used to make estimates of the optimal trapping parameters to maximise the number of atoms. To convert the voltage reading from the photodiode into an atom number we need to make an estimate of the number of photons scattered by an atom per second into the solid angle subtended by the collection lens. The total number of atoms is

$$N = \frac{V}{h\ f\ R_s\ \Omega\ g\ R_{in}}, \tag{2.33}$$

where V is the output voltage from the photodiode, f is the frequency of the D2 transition, R_s is the scattering rate, Ω is the solid angle captured by the collection lens, g is the gain of the photodiode in $\mathrm{A\,W^{-1}}$ and R_{in} is the input impedance of the meter connected to the photodiode.

The scattering rate for any given atom in the MOT is not simple to calculate. The strength of the interaction of an atom with the light field depends on the velocity and direction of movement of the atom, the local magnetic field, the polarisation of the light and the internal state the atom. We can, however, make some good estimates by averaging over the ensemble of the cloud. The scattering rate for an atom is given by

$$R_s = \frac{2\pi\gamma}{2}\ \frac{s}{1 + s + (\frac{2\delta}{\gamma})^2}, \tag{2.34}$$

where δ is the total detuning from the transition frequency including Zeeman and Doppler shifts. To calculate the scattering rate averaged over the total ensemble we need to look at the dependence on each of the parameters in the equation, namely, s, γ and δ.

The saturation intensity in a MOT is not well defined as atoms are scattering on many transitions with different amplitudes. Some authors use a saturation intensity averaged over all states while others use the value for the stretched state only. The saturation intensity, however, can be measured experimentally, as demonstrated for caesium (see [69]). These results agree closely with the saturation intensity of the stretched state which will be used from here onwards for measuring the absolute atom number. For our MOT each beam has an intensity close to the saturation intensity, such that $s \approx 1$. The light intensity is independent of the direction of travel and is approximately uniform over the dimensions of the atom cloud.

The linewidth, γ, for the cooling transition is 5.92 MHz, however, within a MOT several broadening mechanisms are in place. The two dominating processes are magnetic broadening and Doppler broadening. At 1 mK the full-width half-maximum attributed to Doppler broadening is approximately 4 MHz. Typical values for the magnetic field gradient of the MOT are up to 10 $\mathrm{Gcm^{-1}}$. The diameter of our MOT is approximately $4 - 5$ mm implying a distribution of atoms across $4 - 5$ gauss; this equates to a width in frequency space of $\sim$7 MHz. Summing in quadrature, these broadening effects increase the linewidth of the atoms in the MOT to approximately 10 MHz. Thie width was experimentally determined to be between 10 and 11 MHz

using later absorption measurements.

The detuning of the light from the cycling transition depends on the local magnetic field, the polarisation of the light, the velocity of the atom and the detuning of the MOT beams from the field-free transition. While this dependency initially appears complicated the situation can be significantly simplified by considering how the atom is detuned for various circumstances. At the edge of the cloud the atom has exchanged all its kinetic energy for potential energy so the detuning of the light from the atomic transition is determined solely by the Zeeman shift, δ_z. From figure 2.2 one can show that the scattering from the counter-propagating beam will dominate and the Zeeman shift will bring the atom closer to resonance with the light. For an atom near the centre of the MOT there is no Zeeman shift, but a reasonable Doppler shift, δ_D. For an atom moving along one of the axes of the MOT beams the atom will interact with the counter-propagating beam more than the co-propagating beam, meaning that the majority of the scattered light will be from the counter-propagating beam which will again appear closer to resonance. Since for a temperature of $\sim$1 mK $\delta_D \approx \delta_Z \approx 5$ MHz then averaged over the MOT cloud the average detuning can be written as $\delta_L - 5$ MHz, where δ_L is the detuning of the MOT beams from the field free resonance, and assumed to be larger than δ_D and δ_Z.

For a MOT beam detuning of 4 natural linewidths and a transition with an intensity per beam matching saturation (2.56 mW cm^{-2}) one can show using equation 2.34 that the scattering rate, $R_s \approx 1.2 \times 10^6$ photons per second per atom.

The fraction of scattered photons reaching the photodiode is determined by the solid angle subtended by the collection lens. A 1" lens was placed 120 mm from the centre of the cloud providing a solid angle of $\Omega \approx 2.8 \times 10^{-3}$ srad, tantamount to 3400 photons per second per atom reaching the photodiode. Expressed as a power this equates to 1 fW per atom. The gain on the photodiode is 0.5 A/W equating to 500 kV/W for an oscilloscope with an impedance of R_{in} =1 MΩ. Thus, the expected voltage reading displayed by the oscilloscope is 0.5 nV per atom. Values measured by the oscilloscope were found to be as high as 1.5 volts, equating to as many as 3×10^9 atoms. The diameter of the cloud is approximately 7 mm suggesting a density of 1.7×10^{10}. Absorption measurements, discussed in the following section suggest a maximum measured atom number of 2×10^{10} suggesting a larger density of 1.1×10^{11} atoms/cm^3.

Fluorescence measurements are a very useful technique and provide a good estimate of the number of trapped atoms in a non-destructive manner. For high optical densities, however, fluorescence measurements are unreliable as multiple scattering processes become significant [69]. In this regime a scattered photon may be ab-

sorbed one or more times by other atoms in the cloud, reducing the fluorescence and hence providing an underestimation of the true atom number. On resonance, a photon escaping from the centre of the cloud sees an optical density of the cloud, given by $n\,\sigma\,d/2$, where $n \approx 1.0 \times 10^{10-11}$ is the cloud density, σ is the on-resonant absorption cross-section (table 2.2) and d is the cloud diameter. For the above parameters, with $n = 10^{11}$, the optical density is ∼75, however, the value quoted for σ is for the stretched state transition, while the light scattered by atoms in the MOT may be absorbed on other transitions, so the cross-section is reduced by a factor of 2 or 3. Furthermore, the detuning of the light from the atomic transition is not on-resonance, as discussed above. Taking into consideration Doppler and Zeeman shifts, the detuning of the MOT light, δ, is approximately $3\,\Gamma$, reducing the cross-section by a factor of $4\delta^2/\Gamma^2$. These modifications suggest a true optical density of less than 1. These results suggest a small underestimate of the atom number from fluorescence measurements. Absorption imaging of the cloud suggest a maximum number of $2{\times}10^{10}$ atoms recorded. This is in reasonable agreement with the fluorescence measurements, considering the uncertainties in the parameters contributing to equation 2.33, particularly for the value of R_s.

2.7.2 Absorption Imaging

To create an absorption image of a cloud of atoms a weak probe beam close to the frequency of an atomic transition is shone onto the cloud and imaged onto a CCD camera. The degree of absorption by the cloud is determined by the number of atoms in the cloud and can be calculated from Beer's law. In the case of a MOT this can be a non-destructive measurement providing the beam power and detuning are chosen appropriately. For non-optical traps where no light is scattered this method is essential for performing any measurements on a trapped species. Unlike with fluorescence images, the whole of the signal is collected by the camera and very little of the rescattered light reaches the detector providing a high level of contrast for the image. Experimentally, much greater control of the atom-light interaction can be undertaken during absorption imaging, thereby typically providing more reliable results than fluorescence measurements.

The absorption beam is picked off from the Zeeman cooling light. To tune the Zeeman light to the appropriate frequency for laser cooling the output of the laser is passed through two AOMs each applying approximately a 400 MHz shift. The absorption beam is picked from the zeroth order of the second AOM whose frequency coincides closely with the $F = 2$ to $F = 3$ transition as shown in figure 2.9. The detuning of the probe beam can be adjusted by altering the frequency shift of the first AOM, however, it is important to re-tune the second AOM in the opposite

manner in order to maintain the correct frequency for the Zeeman cooling light. In later experiments, a second, non-amplified diode laser at 671 nm was built to provide the absorption beam to achieve full control of the frequency and power of the beam without affecting the Zeeman cooling light. After the AOM the beam is spatially filtered by a 10 μm pinhole which is recollimated with a spot size of $\sim$7-8 mm FWHM; this spatial filtering significantly reduces any noise introduced in the data when subtracting a background image from a cloud image and provides a more uniform intensity distribution. Ideally a large spot size should be employed in order to ensure a more even illumination across the whole cloud. However, due to restricted optical access, the spot size could not be expanded further and was typically comparable to the $1/e^2$ diameter of the atom cloud.

The image is formed by a C-mounted lens unit fixed to the camera. The nominal focal length is set to approximately 22 mm providing a magnification of 0.2 - 0.3. The magnification was chosen such that the image of the cloud will fill as much of the image sensor as possible to improve resolution. MOT sizes, however, could be as large as 8 mm, significantly larger than size of the CCD, requiring magnification values smaller than 1.

The picture is recorded by an Allied Vision Technology MARLIN F-033B CCD camera. The pixel size is 9.9 μm per side and the chip size is 656×492 pixels. The camera provides an 8-bit image such that each pixel gives a number of counts between 0 and 255. The camera has a shutter with a minimum internal shutter speed of 32 μs. The CCD chip must be read out between frames, which takes 13.5 ms for a full-chip image, setting a limit on the rate at which images can be taken. If we neglect the finite resolution of the imaging system and assume the image is perfectly replicated by the camera then the absorption signal, $S(x, y)$, can be written as

$$S(x, y) = I(x, y)\, T(x, y) + C(x, y) \tag{2.35}$$

where $I(x, y)$ is the intensity profile of the probe beam, $T(x, y)$ is the transmission of the atom cloud and $C(x, y)$ is the background light. Sources of background light include ambient room light, scattered laser light and fluorescence from the cloud. The camera was positioned a significant distance away from the chamber apparatus such that any contributions from laser scatter and cloud fluorescence can be neglected. Images were taken in low light conditions to minimise noise from the ambient background light. Under these conditions, and in the limit of low optical density, such that $I(x, y)\, T(x, y) \gg C(x, y)$, the transmission signal of the atoms can be recovered with the help of a background image taken in the absence of any atoms. The ratio of these signals, $S'(x, y)$, is

$$S'(x,y) = \frac{I(x,y)\,T(x,y) + C(x,y)}{I(x,y) + C(x,y)} \approx T(x,y). \tag{2.36}$$

The absorption of monochromatic, collimated light whose intensity, I, fulfills the condition $I < I_{sat}$ is described by Beer's Law. The transmission profile, $T(x,y)$, can consequently be written as

$$T(x,y) = e^{-D(x,y)}, \tag{2.37}$$

where $D(x,y)$ is the optical density of the cloud. The optical density is related to the column density of the atoms, $\rho_{2D}(x,y)$ and the optical scattering cross-section, σ, via

$$D(x,y) = \sigma\rho_{2D}(x,y) = \sigma \int \rho(x,y,z)\,\mathrm{d}z. \tag{2.38}$$

For resonant light the optical density of the cloud is high and equation 2.36 does not hold. However, by detuning the absorption light several linewidths from resonance, the cross-section is reduced and the densest region of the cloud transmits between 10 and 50% of the beam depending on the atom number. Under these circumstances it is possible to ensure the background light is negligble relative to the transmitted laser light so that equation 2.36 remains valid.

Within the MOT the scattering cross-section is not well defined. The magnetic field vectors are dependent upon position so there is no well defined quantisation axis. Consequently for an absorption beam of a given polarisation there is no clear cycling transition with which to calculate a value for the cross-section. We can, however, provide an estimate relative to an upper bound. These measurements will only record atoms in the F=2 state and if we are cycling on the cooling transition as expected then one can assume these atoms are being pumped from the $|F = 2, M_F = 2\rangle$ level. The cross-section on resonance is $\sigma = \frac{3\lambda^2}{2\pi} f$, where f is a numerical parameter between 0 and 1, dependent upon the polarisation of the light and the driven transition. In an ideal case the absorption beam would carry circularly polarised light and only cycle atoms on the stretched state transition in a well defined magnetic field, parallel to the propagation of the beam. In these circumstances the absorption cross section is largest, such that $f = 1$. In reality, however, the field vectors are more complex. A reasonable estimate can be obtained, however, by averaging over the strengths of the available transitions if excited by an unpolarised beam. Although the frequency of the absorption beam is red detuned from the $|F = 2, M_F = 2\rangle \rightarrow |3,3\rangle$ transition there is significant mixing in the presence of small magnetic fields between the upper state M_F levels such that excitations to the $|3,2\rangle$, $|2,2\rangle$, $|3,1\rangle$, $|2,1\rangle$ and $|1,1\rangle$ states are possible. Averaging

over the transition strengths of these six states one can estimate a value of $f = 1/3$. For a two-level atom the scattering cross-section as a function of detuning, in the low intensity limit ($s \ll 1$) is given by

$$\begin{aligned}\sigma &= \frac{\hbar\omega\Gamma}{2}\frac{1/I_{sat}}{1+(\frac{2\delta}{\Gamma})^2} \\ &= \frac{3\lambda^2}{2\pi}\frac{1}{1+(\frac{2\delta}{\Gamma})^2}\ ,\end{aligned} \tag{2.39}$$

where Γ is the linewidth of the atoms in the MOT and δ is the detuning of the absorption beam. The absorption beam has a power of approximately 1 μW with a spot size of 8 mm ($I/I_{sat} \ll 1$) and a detuning of 2.7 natural linewidths. For the stretched state transition the cross-section is $\sigma = 2.0 \times 10^{-14}$ m^2. The transmission profile, given by the ratio of the absorption image signal and the background image is thus related to the total number of atoms in the $^2S_{\frac{1}{2}}(F = 2)$ state in the MOT:

$$N_{tot} = \frac{A}{f\,\sigma M^2}\sum_{i,j} -\ell n(T(x_i, y_j))\ . \tag{2.40}$$

Here A is the area of a pixel on the CCD, M is the linear magnification factor of the lens unit and the summation is performed over each pixel on the CCD with coordinates, $\{x_i, y_j\}$. Outside the region of the CCD illuminated by the probe beam there will be a significant number of pixels with zero counts. Consequently, the value, $\ell n(T(x, y))$, will give non-physical results for these pixels. To circumvent this issue a region of interest is post-selected from the background-free image. This region is a circle with a radius matching the probe beam spot size and chosen to contain only pixels with non-zero counts. The atom number is then calculated only using the pixels lying within this region.

Given the uncertainty in the value of the cross-section the absolute atom number cannot be quoted to better than a factor of 2, however, the relative atom number is experimentally reproducible to within an error of approximately 10%. The values calculated for absorption imaging are typically larger than the associated fluorescence imaging calculations.

The value for M is determined by calibrating the camera. The size of the cloud is calculated by calibrating the camera with an image containing an accurately known dimension. To achieve this a pair of calipers set to 3$\pm$0.005 mm is positioned in the absorption beam at various distances between the chamber and the CCD and the shadow is imaged onto the camera. The number of pixels corresponding to the width between the caliper blades is recorded for each distance and the degree of collimation of the absorption beam is determined. The collimation of the beam is

then adjusted until the number of pixels between the caliper blades recorded for each distance agree to within $2-3\%$, i.e. the beam is now very collimated. The number of pixels per millimetre is then calculated and the $1/e$ size of the cloud can be approximated to within a few percent.

In order to maximise the atom number the number of atoms as a function of laser power, detuning, magnetic field gradient and oven temperature have all been explored experimentally. These results are summarised and discussed below.

2.7.3 Characterisation

The dependence of the atom number on the oven temperature can be calculated using the basic gas kinetics reviewed in section 2.4.4. By integrating the total number of atoms given by equation 2.13 over the solid angle accepted by the Zeeman slower and the velocity distribution up to the capture velocity of the slower then the number of atoms leaving the slower per second within the capture velocity of the trap can be estimated. If we take the steady state solution for equation 2.18 the total number of atoms captured is given by $N_{tot} = R_c\tau$, where R_c is the capture rate and τ is the lifetime determined by the background pressure of the trapping chamber. Obviously τ will have some dependence on the temperature of the oven due to a larger background pressure at higher oven temperatures. However, to date no data has been taken to provide an empirical relationship between the lifetime and oven temperature. To make an estimate of τ measurements of the loading-time were made from observations of the MOT fluorescence during loading. For a temperature of 800 K a time-constant of 1.1 seconds was measured; this time is equivalent to the lifetime of the trap. If we assume that all the atoms leaving the slower per second within the capture velocity of the trap are indeed captured, then we can plot a curve for the number of expected atoms captured as a function of oven temperature. So far it has been assumed that all atoms below the capture velocity can be decelerated, however, we have already shown in section 2.5 that this is not true. The loss in atom number is estimated to be 10% as a consequence of cycling into dark states during deceleration. Including this in the calculations, figure 2.37 displays the theoretical predictions alongside the data recorded from absorption images. The data points in figure 2.37 have been multiplied by a factor of 8/5 to account for those atoms trapped in the F=1 state that do not participate in the absorption imaging. The value of 8/5 is given by the ratio of the total number of M_F levels in the $F = 1$ and $F = 2$ state to the number in the $F = 2$ state and rests upon the assumption that the atoms are evenly distributed amongst all the M_F states. This assumption is re-enforced by data taken in chapter 4. Note that to produce this theoretical curve no free parameters have been used; it is only a derivation based on the expected vapour

pressure inside the lithium oven as a function of temperature and the expected performance of the Zeeman slower and the MOT.

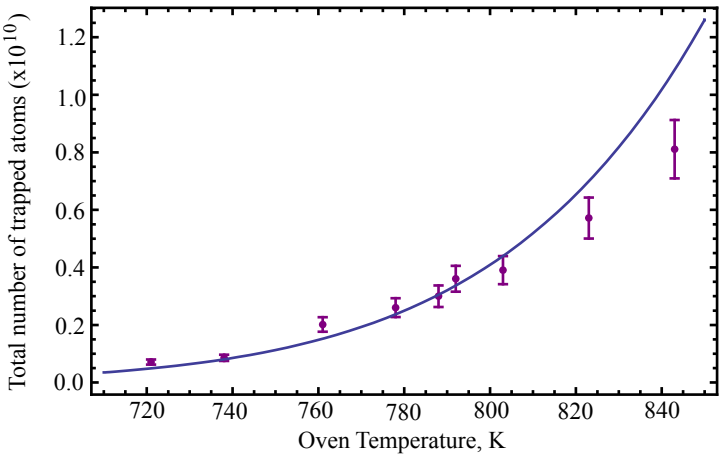

Figure 2.37: Experimental data depicting the dependence of the number of trapped atoms on the temperature of the oven source and a corresponding theoretically-derived curve. The error bars indicate the 10% uncertainty in the relative atom numbers after repeat measurements and do not indicate the uncertainty in the determination of the absolute atom number.

The error bars indicate the 10% uncertainty in the relative atom numbers after repeat measurements and do not indicate the uncertainty in the determination of the absolute atom number.

As the oven temperature and consequently the background pressure increases, the number of atoms is expected to eventually fall for higher temperatures as the increasing number of background collisions shortens the lifetime and thus reduces N_{tot}.

The atom number was also measured as a function of laser power for both the Zeeman slower and the MOT beams. Figure 2.38 summarises these results for the four frequencies used to cool and trap the atoms.

The largest number of atoms trapped was obtained for a MOT beam power of 13 mW per beam, a repump power of 7 mW, a Zeeman beam power of 27 mW and a Zeeman repump power of 3 mW. In these experiments one beam power is varied, while the other three are kept as close to these optimum values as possible.

The number of trapped atoms as a function of the slowing beam power (Fig. 2.38(a)) closely matches the expectations discussed in section 2.5.4. At approximately $17 - 18$ mW of power a plateau in the atom number is reached; this is close to the 15 mW threshold predicted by simulations. Below this threshold the atom number gradually decreases, though not as sharply as simulations suggest it should. At 10 mW the atom number has decreased by approximately a factor of 2.5, while simulations suggest a factor of 4 decrease. This suggests our Zeeman slower is more tolerant to very low laser powers than expected.

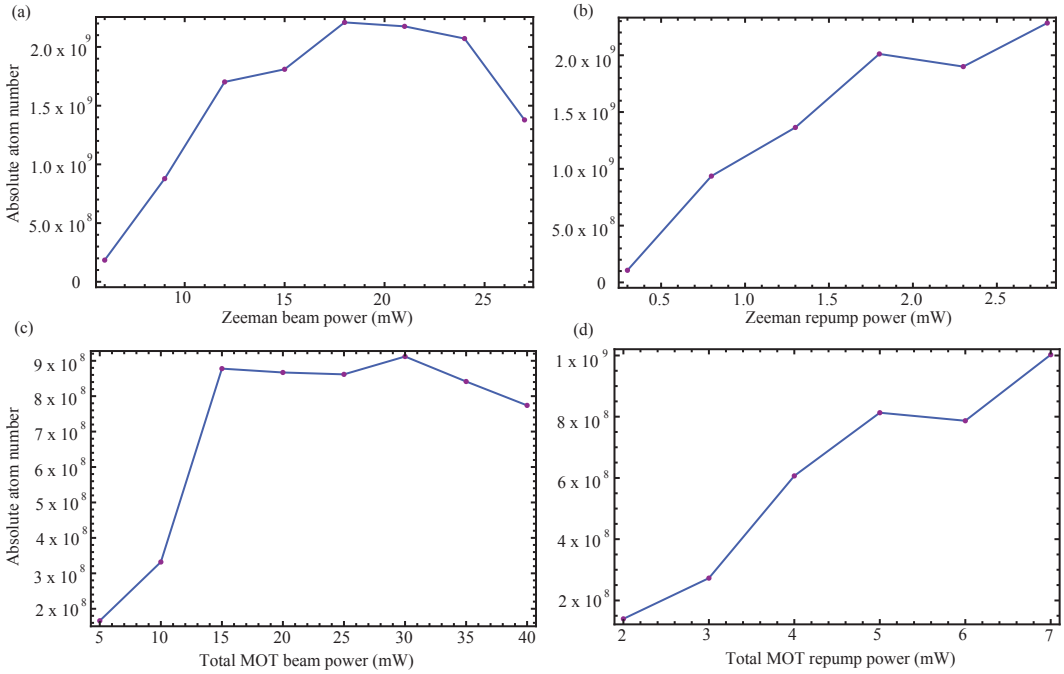

Figure 2.38: Experimental data for the dependence of atom number on the laser power for (a) the Zeeman beam, (b) the Zeeman repump, (c) the MOT trapping light and (d) the MOT repump light. The lines joining data points are used to guide the eye.

As can be seen in figure 2.38(b) the number of trapped atoms drops linearly with Zeeman repump power and is reduced by a factor of 20 in the absence of a repump beam. This is not consistent with the previous measurements referred to on page 83, where a factor of 4 decrease was observed. This discrepancy is thought to be due to a particular alignment of the MOT beams during these experiments; it is plausible that certain beam alignments favour a higher capture velocity than others. Thus, in the absence of a Zeeman repump beam, where exit velocities may be slightly higher, fewer atoms are captured in the MOT. There was not enough available power to measure a plateau in the relationship, however, as with the slowing beam there is expected to be a threshold power above which the number of atoms will not increase for the repump light also. This measurement suggests that higher atom numbers might be obtainable by diverting more power to the Zeeman repump beam.

The number of atoms as a function of MOT beam and repump power displays a similar relationship to the Zeeman slowing light and the Zeeman repump. For the MOT power dependence the number of atoms reaches a plateau for a total beam power of 18 mW, equivalent to 6 mW per beam, remembering that the 3 MOT beams are retro-reflected; this equates to an intensity, $I_{sat}/2$ per beam. This value is somewhat smaller than the results of [68] which predicts saturation at I_{sat} per beam. However, these results were performed for Cs, and without detailed simulations for Li atoms it is not clear what to expect.

Some researchers have found that the optimum detuning and magnetic field gradient are different for different laser intensities [68]. For the results collected in figure 2.38 both of these parameters remained constant at 3.8 Γ and $6 - 7\ \mathrm{G\,cm^{-1}}$, respectively; it would be interesting to retake the results, optimising the detuning and field gradient for each power. The dependency on the repump intensity is expected to have a similar saturation point, however, given the available power we were only able to explore the linearly dependent regime. As with the Zeeman repump light, these results suggest that more power to the MOT repump will lead to larger atom numbers. Note, however, that the light is close to resonance and increasing the intensity of the beam may mean sacrificing the cold temperatures achieved for the MOT.

Figure 2.39 portrays the number of atoms as a function of laser detuning for the Zeeman slowing beam, the MOT beams and the MOT repump light. The frequency of the Zeeman repump light is delivered directly from the laser and cannot be varied.

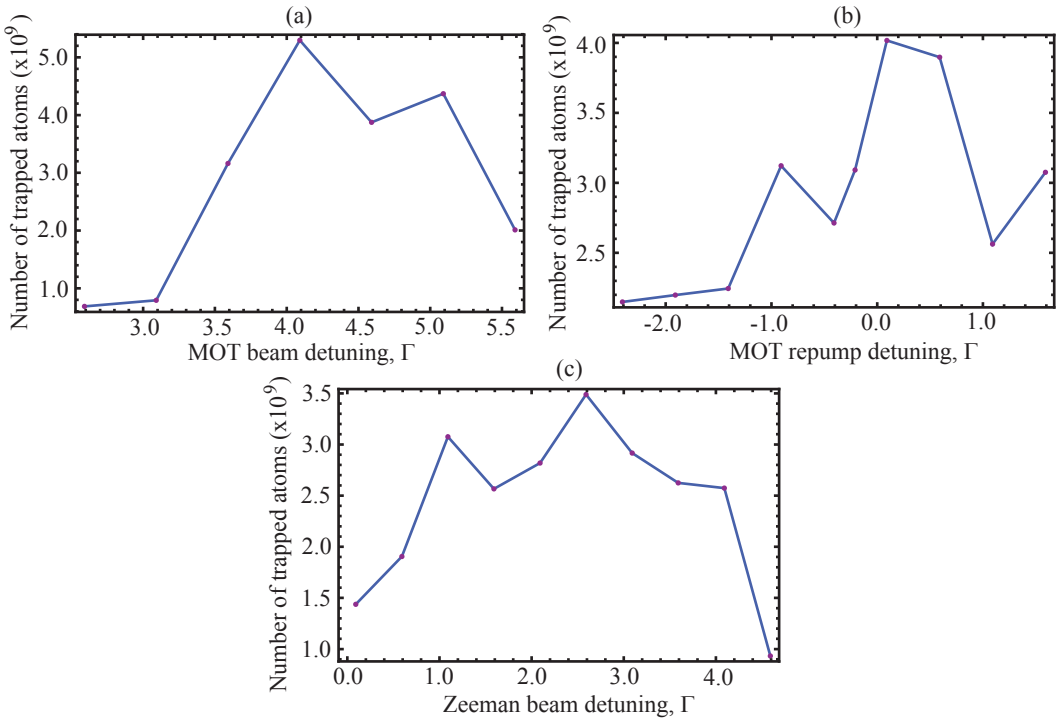

Figure 2.39: Experimental data for the dependence of atom number on the laser detunings for (a) the MOT trapping beams, (b) the MOT repump light and (c) the Zeeman cooling light; the detuning in (c) is measured relative to the transition frequency of an atom travelling counter to the beam at 269 $\mathrm{ms^{-1}}$ in zero field. The lines joining the data points are used to guide the eye. Greater positive values imply the frequency is more red detuned.

An optimum detuning for the MOT light frequency was found at approximately 4 Γ. For small detunings, as the detuning is increased the capture velocity of the MOT will also rise. The atoms leaving the slower are expected to have a velocity of approximately 55 $\mathrm{ms^{-1}}$ with a spread of 10 $\mathrm{ms^{-1}}$. As the capture velocity of the

MOT approaches this range the number of atoms is expected to rise sharply. Such an increase is observed for a detuning beyond 3 linewidths. However, for larger detunings the damping force is greatly reduced and the density of the MOT will fall; beyond a critical point the trapping force will be too weak to capture all the atoms. The optimum detuning will depend on the magnetic field gradient, however, a detuning of 4 Γ is in agreement with many other lithium experiments [34, 37]. The optimum detuning for the repump light was found to be zero, meaning that the light is on resonance with the field-free $^2S_{1/2}$ F=1 to $^2P_{3/2}$ F=3 transition. This is perhaps expected since the MOT light dominates the trapping processes and the forces on the atom are dependent on the detuning, whereas, the much weaker repump beam does not contribute to the trapping force but rather simply maintains a closed transition. The two peak features seen in figures 2.39 (a) and (b) are thought to be attributed to the resolved spectral lines corresponding to the $^2P_{3/2}$ F=3 hyperfine state and an unresolved peak for the $F = 0, 1, 2$ states, split by approximately 9 MHz. Because the repump is much weaker there is no power broadening and the spectral structure is much more apparent, however, without higher resolution data this interpretation is only speculative.

The structure in figure 2.39 (c) is more difficult to interpret. If the Zeeman slower is working to its specifications the highest atom flux should be found when the Zeeman light is resonant with the $F = 2$ component of the D2 transition for atoms moving at 269 ms^{-1}, (the zero-field point), corresponding to 0 Γ in figure 2.39(c). However, there is clearly a range between 1 and 2.5 Γ for which the flux is maximised. The difference may be a consequence of the change in magnetic field from the expected field, as discussed in section 2.5.6. In particular it suggests that the the speed of the atoms at the zero-crossing point of the Zeeman slower is slower than the design speed of 269 ms^{-1} by ~ 6 ms^{-1}. This discrepancy is most likely due to the mismatch between the ideal magnetic field and the experimentally measured field.

Figure 2.40 depicts how the atom number depends on the magnetic field gradient of the trapping fields for two different MOT beam detunings. For both detunings of 2.8 Γ and 4 Γ the optimum field gradient is found to be approximately 7 $\mathrm{G\,cm}^{-1}$. This is not in accordance with the simple Doppler theory [44, 70] which predicts a larger optimum field gradient for larger detunings. Without more data, however, no firm relationship between the parameters can be derived. We find that for the larger detuning case the number of atoms is much less sensitive to changes in the field gradient, with the atom number remaining approximately constant between 6 and 10 $\mathrm{Gauss\,cm}^{-1}$. For very steep gradients the atom number drops significantly. As the field gradient is increased the atoms must occupy a smaller volume for their given

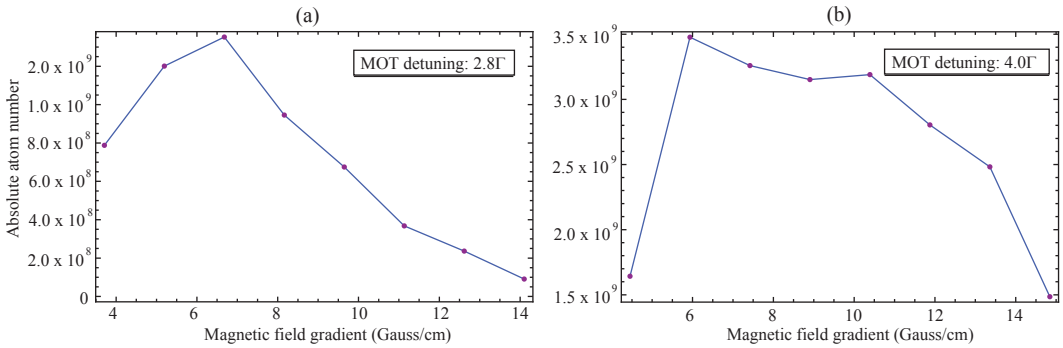

Figure 2.40: Experimental data for the dependence of atom number on the magnetic field gradient of the MOT for MOT beam detunings of (a) 2.8 Γ and (b) 4.0 Γ. The lines joining the data points are used to guide the eye.

temperature. However, since a cloud is density limited by factors such as radiation pressure the cloud cannot be compressed beyond a certain volume. Consequently for large field gradients the atom number will decrease to maintain the same atom density. For larger detunings the radiation pressure is weaker and the cloud is more tolerant to compression at larger field gradients.

2.8 Cloud Temperature

It is not immediately obvious how one should go about assigning a temperature to a cloud of atoms trapped in a MOT. The underlying assumption in the experimental determinations below is that the velocity distribution of the atoms in the cloud is a Maxwell-Boltzmann one, as described by equation 2.41. Each velocity component, v_i, follows a Gaussian distribution:

$$P(v_i) = \frac{1}{v_{i,rms}}\sqrt{\frac{2}{\pi}}\exp\Big(-\frac{v_i^2}{2v_{i,rms}^2}\Big)\,, \tag{2.41}$$

where $v_{i,rms} \equiv \sqrt{\langle v_i^2\rangle}$. The temperature, T, is defined by its relation to the mean squared velocity,

$$T \equiv \frac{m\langle v_i^2\rangle}{k_B}. \tag{2.42}$$

For most physical systems the Maxwell-Boltzmann distribution is typically derived by considering a classical, ideal gas of non-interacting particles in thermal equilibrium with a heat bath at a temperature, T; each particle is assumed to carry no intrinsic properties other than its mass. This is quite different to the case of laser cooled or magnetically trapped atoms which are isolated from any external environment by the vacuum. In the magneto-optical trap the atoms occupy several different

energy states and interact strongly with the radiation field, which under no circumstances can be described as a heat bath with a well defined temperature. To justify the assignment of a Maxwell-Boltzmann distribution to such a cloud one must look at the forces acting on the atoms, namely the competition between cooling and heating effects described by the damping force and the diffusion in momentum space, respectively [39]. The authors here pay careful consideration to the effects of these forces on the trapped atoms and the associated assignment of a temperature under different circumstances. The forces acting on the atoms in the MOT are described by a continuous damping force, proportional to the atom's velocity, and a statistical one, arising from the random nature of spontaneous emission. The mechanics of such an environment are well described by the Fokker-Planck equation. The steady state solution is a Maxwell-Boltzmann distribution. It should be noted, however, that this steady state solution is not, in the classical sense, a thermodynamical equilibrium state and in this strictest sense there is no thermodynamically definable temperature for a cloud of atoms in a MOT. The assignment of a Maxwell-Boltzmann distribution to such a trapped cloud may still be applied, however, and the temperatures quoted in the results below represent the usual characteristic temperature defined for such a distribution.

2.8.1 Release-recapture

Initial measurements of the MOT temperature were performed using the release-recapture method. The principle of the experiment rests upon the definition of an imaginary spherical surface surrounding the MOT cloud beyond which a stationary atom cannot be captured. By pulsing the trapping forces of the MOT off for a brief time interval, τ, the cloud is allowed to freely expand. The fraction of atoms recaptured when the trapping forces return will be related to the velocity of the atoms in the cloud and consequently the cloud temperature.

To calculate the temperature from these results, for the moment, let us assume that the cloud has no spatial distribution and that we are operating in only one dimension. The probability density of an atom occupying a position, x, at a time, t, for a cloud with a temperature, T, starting at the origin and expanding outwards is then given by

$$\rho(x,t) = \sqrt{\frac{m}{2\pi kTt^2}}\,\exp\left[-\left(\frac{mx^2/t^2}{2kT}\right)\right]. \tag{2.43}$$

For a given spatial distribution then, we can treat each point in the distribution as an individual expanding cloud described by equation 2.43. This treatment is simply given by the convolution of the velocity and spatial distributions of the cloud.

The force acting on the atoms in the MOT is linear with position over the central portion of the trap such that the potential can be approximated by a simple three-dimensional harmonic oscillator. For the moment we will assume an equipartition of the energy of the particles over the available degrees of freedom, such that the velocity spread and position spread, σ, are related by

$$k_B T = m\, v_{i,\,rms}^2 = \kappa\, \sigma_{i,\,rms}^2 \;, \tag{2.44}$$

where κ is the spring constant of the trap. This assumption implies a gaussian spatial distribution in the trap and has been confirmed from recorded images of the cloud. The probability of an atom occupying a given position, x, in such a distribution is thus simply given by

$$f(x) = \frac{1}{\sqrt{\pi}\sigma_x} \exp\left[-\left(\frac{x^2}{\sigma_x^2}\right)\right] \mathrm{d}x \;. \tag{2.45}$$

The convolution of equations 2.43 and 2.45 will therefore describe the 1D expansion of an atom cloud in a MOT with a Gaussian spatial distribution. If we now define some boundary beyond which an atom will not be recaptured, x_0, then the fraction of atoms still residing within this boundary at a given time, t, since release, is described by the integral of the convolution

$$N(t)_{1D} = \frac{1}{\sqrt{\pi}} \int_{-x_0}^{x_0} \frac{1}{\sqrt{\sigma_x^2 + q_x t^2}} \exp\left[-\left(\frac{x^2}{\sigma_x^2 + q_x t^2}\right)\right] , \tag{2.46}$$

where $q_i = 2k_B T_i/m$.

To modify the solution to three dimensions one must integrate over all three coordinates, x, y and z. To simplify the equation we will assume that both σ_i and q_i are equal in the three Cartesian coordinates. It should also be noted that in reality the recapture volume will be described by some irregular shape, perhaps the intersection of three orthogonal cylinders rather than one of a simple geometry. However, for the purpose of this calculation we will choose a cube whose volume is approximately equal to the real volume. It is clear from equation 2.46 in fact that the results are much more sensitive to the choice of x_0 than to any realistic choice for the shape of the volume. The total fraction of atoms then, occupying a cubic volume of side, $2\,x_0$, centred at the origin, at a given time, t, since release is described by the cube of equation 2.46:

$$N(t)_{3D} = \left(\mathrm{Erf}\left[\frac{x_0}{\sqrt{\sigma^2 + qt^2}} \right] \right)^3 . \tag{2.47}$$

Here Erf is the error function described by the integral of a Gaussian and σ is the spatial variance at $t = 0$. The calculations performed above assume the initial spatial

distribution of the cloud is isotropic. If this is not the case equation 2.47 cannot be applied. In the MOT the cloud shape is significantly determined by the laser beam alignment and can take on a number of shapes and modes [71]. In these experiments, however, the cloud was observed to be close to spherical, with a gaussian profile.

Below is a typical data set for our release-recapture experiments fitted to equation 2.47, floating the parameters T and x_0. The value for σ was determined experimentally to be 3 mm.

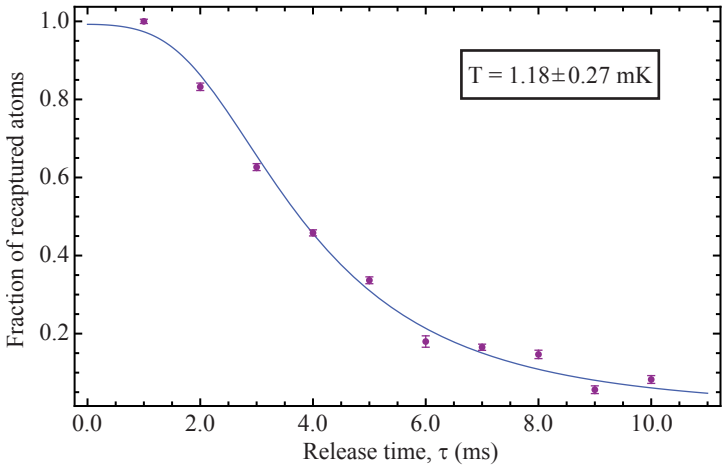

Figure 2.41: Release-recapture data for atoms in the magneto-optical trap. The data points show the number of atoms recaptured by the MOT after a release time, τ. The blue line is a fitted theoretical curve for a temperature of 1.18 mK and a recapture radius of 6 mm.

The above fit gives a temperature of 1.18±0.27 mK for a recapture radius of 6 mm. The temperature is very sensitive to the choice of the recapture radius, and very similar fits can be plotted for smaller temperatures with smaller recapture radii, and vice-versa.

2.8.2 Ballistic expansion

Figure 2.42 depicts a selection of absorption images taken for an expanding cloud at increasing times between 0 and 1.3 ms since release.

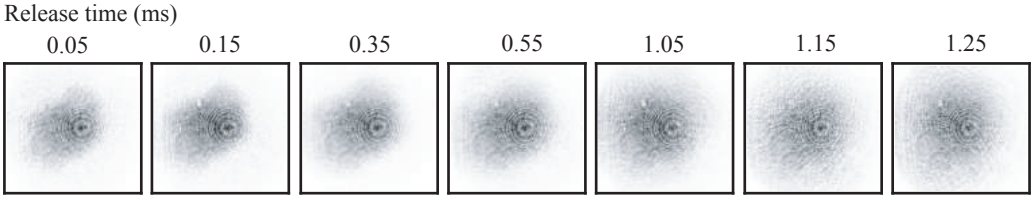

Figure 2.42: A series of absorption images of the atom cloud taken at successively longer times after release, illustrating ballistic expansion.

With the ability to produce high resolution absorption images the temperature of the atom cloud can be measured accurately and easily without requiring any

knowledge about the recapture volume of the MOT. If an atom cloud is released from the trap it will grow at a rate related to its temperature, given by the quadrature sum of the initial width, $\sigma(t=0)$, and the contribution from the velocity spread,

$$\sigma_{tot}(t) = \sqrt{\sigma(0)^2 + qt^2}\ . \tag{2.48}$$

Each image in figure 2.42 is an absorption image captured by the CCD camera as described in section 2.7. Because of the small times involved between successive images the collection of photographs above were not taken from a single MOT. Rather, the atoms were released from the trap and an image was taken at some later time, δt, followed by a background image without atoms. The experiment is subsequently repeated for increasing values of δt and a film strip can be built up, subtracting each background image from its corresponding cloud image. Figure 2.43 depicts the triggering schedule for the various components involved in the experiment to produce a single background-free image.

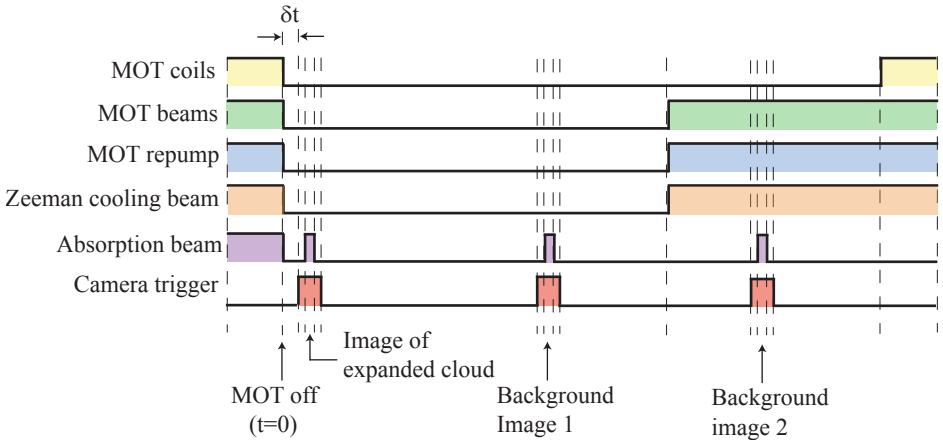

Figure 2.43: The digital triggering schedule delivered by the computer to the experimental apparatus to control the AOMs, the MOT magnetic coils and the camera shutter. The dashed lines depicts points of change to one of the parameters.

The first image taken is simply the image of the cloud shortly after it has been released. The second and third images are background images taken at least 50 ms after the atoms in the cloud have been released; after this time there should be no atoms remaining to image, only background light. The reason for taking two background images, one with the cooling and trapping beams off and one with them on, is because while running these measurements it can be useful to photograph the cloud before release, during which time the background light will be different due to scatter from the trapping beams. During image processing the appropriate background can then be chosen for subtraction.

To determine the temperature of the cloud the rms width of the cloud is measured for each image. The pixel matrix is collapsed both along the x and y axes to produce

two one-dimensional arrays for each image. Note that the choice of axes does not correspond to the axes of the MOT beams but to the arbitrary orientation of the camera and absorption beam. Each pixel array is fitted to a Gaussian distribution, allowing the amplitude, width and mean position to be determined from the fitting parameters. It is important to collect information for both axes as the spatial distribution of the cloud can lead to systematic errors in the measurements. For example, in some measurements one can see the cloud splitting along one axis but appears as a single expanding cloud if viewed along the orthogonal axis. In this case a Gaussian fit will be a poor estimation of the spatial distribution for the axis along which the MOT is splitting and will manifest as a warmer temperature in the results. The widths of the clouds are fitted according to equation 2.48 as a function of time and a temperature is estimated. Results for our lithium cloud are presented in figure 2.44, suggesting a temperature of 0.86 ± 0.13 mK in the x-axis and 0.85 ± 0.07 mK in the y-axis. The trap parameters prior to release were a MOT light detuning of 4.4Γ and a magnetic field gradient of 7 gauss/cm.

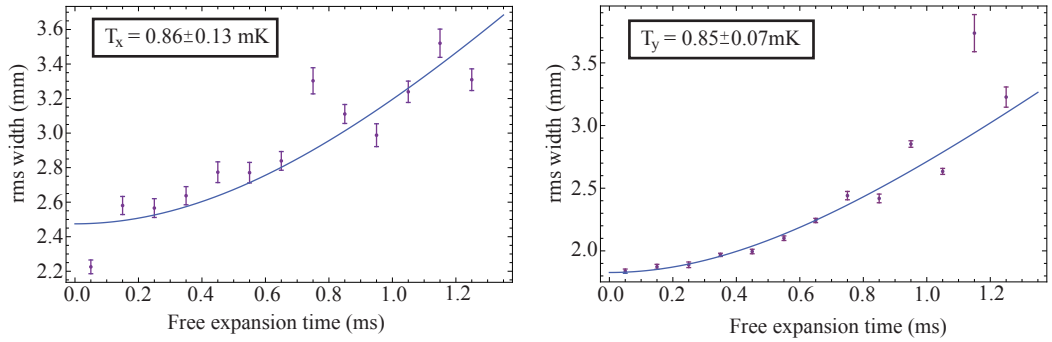

Figure 2.44: Experimental data for the rms width of the lithium cloud along two axes (x and y) as a function of time. A theoretical curve is fitted to the data to determine the temperature of the cloud. The results are given in the inserts.

The estimated temperatures for each axis agree well with each other within the allotted standard errors. The error bars for each data point are derived from the fitting errors of the Gaussian profile to the data arrays. They do not include any errors manifesting during the image collection, such as camera and probe beam vibrational instabilities and do not quantify any errors in the assumption that a Gaussian profile is the correct choice for the fitting function. The fitted lines in figure 2.44 are calculated according to the weighting of each error bar, with the temperature and initial cloud size left as floating parameters. Errors in the triggering times have been neglected as they have been measured to be less than 10 ns.

The temperature of the MOT is not trivially derived from first principles for a given set of laser detunings, intensities, magnetic field gradiet and so on. However, a

reasonable estimation can be derived from Doppler theory which assumes the atoms are sitting in a 1 dimensional optical molasses. The temperature as a function of detuning and beam intensity is given by equation 2.49, the derivation for which is covered in section 4.4.

$$T_{dopp} = \frac{1}{k_B}\frac{\hbar\Gamma^2}{8\delta}\left(1 + s + 4\frac{\delta^2}{\Gamma^2}\right) , \tag{2.49}$$

where $\Gamma = 2\pi\gamma$ and δ is in $\mathrm{rad\,s^{-1}}$. For $s \ll 1$ the minimum of this curve falls at 142μK, which is the expected Doppler temperature. For our set-up the intensity per beam is approximately the saturation intensity, so for a 1D molasses we have $s = 2$. The MOT beams are detuned by 4 linewidths to the red. Substituting these values into equation 2.49 gives a predicted temperature of approximately 600μK. This is in very good agreement with our experimental observations, considering the simplicity of the calculations. For a calculation in 3 dimensions, taking into account the heating and cooling rates of the repump beam as well as the effects of the magnetic field and light polarisations a better match may be found, however, this requires very complex and detailed simulations. Other effects such as multiple scattering events may also lead to further heating.

2.9 Summary

We have trapped up to 2×10^{10} neutral ^{7}Li atoms in the $^2S_{1/2}\,F = 2$ state in a magneto-optical trap with a density of 1.1×10^{11} atoms $\mathrm{cm^{-3}}$ and a minimum temperature of 0.85 ± 0.10 mK. The number of atoms has been determined using both fluorescence and absorption measurements, and their results are in reasonable agreement considering measurement uncertainties. We have made measurements of the dependency of the atom number as a function of oven temperature, laser detuning, beam intensity and magnetic field and find qualitative agreement with theoretical expectations. In addition we have measured the temperature of the cloud using both release-recapture and ballistic expansion techniques and find good agreement between the two. At the time of writing, a literature view of previous lithium MOTs was made, including traps for both isotopes. In light of this review it is believed that a trap containing 2×10^{10} atoms is the largest lithium MOT to date and is an encouraging sign for success in the experiments outlined in this thesis and beyond. Experiments such as BEC condensation and sympathetic cooling with molecules are known to be severely detrimental to the remaining number of trapped atoms and it is thus clearly beneficial to undertake such projects with an initially large number of atoms.

The following chapters will discuss the transfer of the atoms to a magnetic trap

and the transportation to an isolated region of the chamber for measurements of the cloud under large electric fields. Firstly, chapter 3 will discuss the set-up and results of a separate project investigating the feasibility of producing electric fields approaching the 1 MV cm^{-1} regime required to polarise a cloud of lithium atoms.

Chapter 3

Production of Very Large Electric Fields

In the regime of classical physics the only limit to the maximum electric field sustainable between two parallel, infinite capacitor plates with perfectly flat surfaces is defined by the ionisation energy of the metal. The energy required to dissociate an electron from the nucleus of an atom is governed by the Coulomb interaction. For stainless steel this is estimated to require a field strength of approximately 400 MV/cm. Upon the inclusion of quantum effects, however, for fields exceeding a value of 10 MV/cm significant currents between electrodes may be observed. This is a consequence of electrons tunneling through the potential barrier present at the surface of the metal. This process is known as field emission. However, even these fields appear to be unattainable and few experiments achieve fields larger than a few hundred kV/cm. Such discrepancies are thought to be due to imperfections in the electrode's surface, such as microscopic protrusions and impurities. The following chapter will cover some of the work we have done, through various polishing treatments and high voltage conditioning techniques, to push the size of these fields towards the 1 MV/cm milestone. We also explore the use of transparent indium-tin-oxide (ITO) coated glass plates for electrodes, as a technique to achieve a high level of parallelism using interferometric techniques, while providing optically flat surfaces. Section 3.1 provides a basic background for understanding the processes underlying electrical breakdown. Sections 3.2 to 3.4 will describe the apparatus constructed as a general purpose test rig for exploring breakdown between two broad area stainless steel electrodes and the results achieved for numerous surface treatments. Finally I will describe the design and construction of ITO-coated glass plates and an adjustable mounting rig for aligning the plates parallel with each other.

3.1 Theoretical Overview

3.1.1 Field Emission

In 1928 Fowler and Nordheim published a paper [72] deriving an expression for the field emission of electrons from a metal surface based on quantum tunnelling. The paper describes a potential barrier near the surface of the plate which is modified by the presence of an electric field, as described schematically in figure 3.1. The size of the barrier is determined by the Coulomb attraction of an electron to its nucleus at the surface of the metal and may be approximated as a step potential.

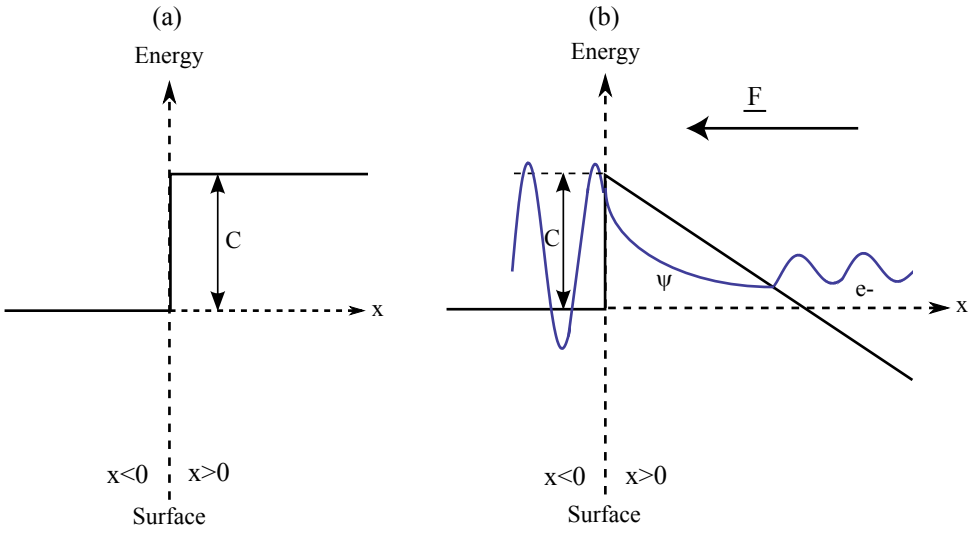

Figure 3.1: A simplified representation of the potential barrier at the surface of a metal at $x = 0$ as seen by an electron (a) without and (b) with an electric field, **F**, present. In the presence of an electric field the barrier potential is lowered outside the metal ($x > 0$) and the electron has a finite probability of tunneling out of the metal.

In reality the corners at the peak of the potential barriers in figure 3.1(b) will be rounded off as a consequence of an image-charge effect, however, these modifications are of little consequence to the calculations, as noted by Fowler and Nordheim. To solve the transmission probability for an electron with energy, W, tunnelling through a barrier, of height C, whose form has been modified by an external field, F, (figure 3.1(b)) we only have to solve the wave equations each side of the barrier

$$\frac{\mathrm{d}^2\psi}{\mathrm{d}x^2} + \kappa^2(W - C + Fx)\psi = 0 \qquad (x > 0), \tag{3.1}$$

$$\frac{\mathrm{d}^2\psi}{\mathrm{d}x^2} + \kappa^2 W\psi = 0 \qquad (x < 0), \tag{3.2}$$

subject to the conditions that ψ and $\mathrm{d}\psi/\mathrm{d}x$ are continuous at $x = 0$ and that for $x > 0$, ψ represents a stream of electrons *only* co-propagating with x. The constant κ is given by $\kappa^2 = 8\pi^2 m_e/h^2$, where m_e is the electron mass. The calculations for this problem are a little in depth for the purpose of this outline, but can be found

in detail in [72]. The final solution for the fraction of electrons, D, (energy, W), penetrating the boundary and emerging under the influence of the external field is given in equation 18 of [72] and written as

$$D(W) = \frac{4\sqrt{W(C-W)}}{C}\, e^{-4\kappa(C-W)^{\frac{3}{2}}/3F}, \tag{3.3}$$

where C is the barrier height and F is the electric field strength.

To express this as a current we need to know the number of electrons, *N(W)*, on a surface of unit area per unit time with momentum normal to the surface. Following through the arguments outlined in [72] we reach a final solution for current density in Am^{-2}

$$I = \frac{e}{2\pi h}\frac{\mu^{\frac{1}{2}}}{(\chi+\mu)\chi^{\frac{1}{2}}}\, F^2\, e^{-4\kappa\chi^{3/2}/3F}\,, \tag{3.4}$$

where μ is the chemical potential and $\chi = C - \mu$ is the work function of the material. This formula is valid for all temperatures for which the condition $\mu \gg k_B T$, corresponding to temperatures up to 5000 K for metals. In the experiments outlined below the electrodes were stainless steel cylindrical buttons, 13mm diameter, sitting at room temperature. The chemical potential for most metals lies between 5 and 10 eV; the exact value chosen from this range has little effect on the value of I. The work function for stainless steel is $\sim 4.7 - 5.6$ eV. Thus, putting in some values to equation 3.4 for the electrodes used in the experiments a graph for the current between two electrodes as a function of the electric field between them can be plotted, figure 3.2.

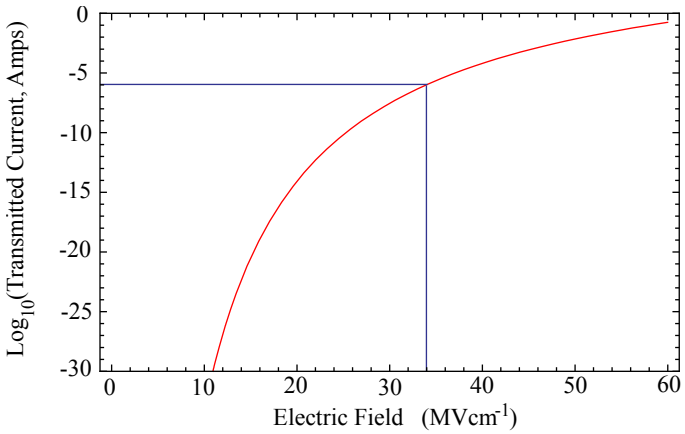

Figure 3.2: A theoretical curve of the current drawn between two electrodes as a function of electric field due to field emission, as described by the Fowler-Nordheim theory. The formula for the curve is given in equation 3.4 and is plotted for values of $\mu = 10\,eV$ and $\chi = 5.3\,eV$. The blue lines indicate the field strength required to produce a field-emission current of 1 μA.

Experimentally breakdown tends to occur when the current between the electrodes exceeds $\sim$ 1 μA. According to figure 3.2 this should not occur until field values of around 34 MVcm^{-1} while the current at 1 MVcm^{-1} is vanishingly small. In practice, however, achieving over 100kVcm^{-1} is experimentally difficult. There is a lot of literature on the subject of electric field breakdowns in a vacuum, beginning with work by Wood in 1897 and continuing throughout the next century. Here I will bring attention to the main mechanisms that lead to premature breakdowns.

Early studies by Millikan and Sawyer [73] established that a vacuum gap has a small but finite electrical conductivity, even prior to breakdown, inferred from the existence of 'pre-breakdown' currents across the gap. These currents increased in magnitude with increasing voltage until breakdown occurred. Furthermore, it could be shown that if the voltage was increased slowly enough such that all pre-breakdown currents were allowed to decay, then the breakdown threshold of the gap could be increased. This technique is known a current conditioning; other conditioning techniques, such as gas conditioning, can be found in [74]. Further investigation into the sources of these pre-breakdown currents showed that they were due to field electron emission at isolated sites on the electrode. Milikan and others [75, 76] reasoned that microscopic imperfections in the material lead to regions of enhanced electric field strengths around protrusions from the surface of the electrode. If this enhancement exceeded the field threshold breakdown would incur. Chemical impurities were also recognised as a contributing factor by locally lowering the electrode's work function.

3.1.2 Breakdown Mechanisms

Premature electrical breakdown across a gap may be attributed to one of two main causes. (A particularly good introduction and theoretical analysis of these processes can be found in [77]). The first of these, known as cathode-initiated breakdown was first recognised by Dyke *et al.* [78] who identified a critical emission current at which the cathode becomes thermally unstable due to resistive heating processes as a result of field emission currents. They modelled protrusions from an electrode surface by a cone whose apex is capped by a hemisphere of radius, r. Above electron densities of 10^8 A cm^{-2} [79], where the area is defined by the radius, r, of the tip, the induced heating leads to vaporisation of the surface and subsequently breakdown is initiated. Figure 3.3 depicts the modification of an electric field between two plates in the presence of a protrusion. The protrusion locally enhances the field strength thereby increasing field emission. Higher electron densities in this region may also result in cathode- and anode-initiated breakdown (see below).

The second decay mode is anode-initiated breakdown. If a surface protrusion is short and dull electron emission can occur without resistive heating effects increasing

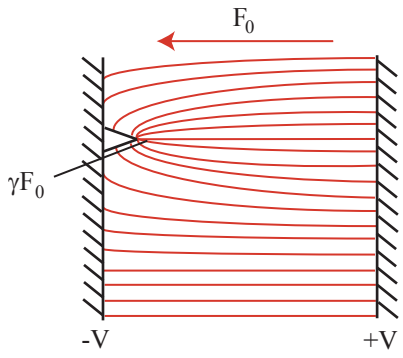

Figure 3.3: The modification of an electric field between two plates due to the presence of a microscopic protrusion.

to the threshold required for cathode-initiated breakdown. The electrons are emitted from the protrusion and accelerated across the vacuum gap by the electric field. As the field strength increases the electron-beam power density at the anode can reach a critical threshold value before the cathode vaporises, causing vaporisation of anode material and thus leading to anode-initiated breakdown. The boundary between cathode and anode-initiated breakdowns can be characterised by a critical value, γ_0, [77] of the local field enhancement factor at the cathode, $\gamma = F/F_0$, which is the ratio of the enhanced local field F at the tip of a given cathode protrusion to the macroscopic field F_0 in the gap. For DC fields equation 25 of [77] gives

$$\gamma_0 = r\left(\frac{J_c\, F_c\, d^{1/2}}{2K\Delta T_M}\right)^{2/3}, \tag{3.5}$$

where r is the radius of the protrusion, d is the distance between the electrodes, K is the thermal conductivity of the anode, ΔT_M is the temperature change that occurs when the voltages are switched on and the temperature has reached a steady state. J_c and F_c are the critical field emitted current density and the corresponding field at the cathode projection. J_c may be calculated from the steady state solution of the basic heat conduction equation applied to the dimensions of the protrusion [77] and is related to F_c through equation 3.4. For values of $\gamma > \gamma_0$ cathode-initiated breakdown processes will dominate, while for $\gamma < \gamma_0$, anode-initated breakdown is the most important mechanism.

In an accompanying paper [80] which presents supporting experimental evidence for the theories outlined above, it is shown that protrusions with local field enhancement factors greater than γ_0 are normally found on even carefully processed electrode surfaces. Typical values for γ_0 were calculated to be between 30 and 50 for a variety of metals while values of γ were measured to be between 15 and 400. These authors suggest that in the case of DC fields, where the anode has an oppor-

tunity to reach thermal equilibrium, then in a large number, if not a majority of cases, breakdown is ignited by anode-initiated processes. However, for short-pulse voltages the value of γ_0 decreases and cathode initiated arcs become more probable.

Other papers suggest a different enhancement factor, $\beta = F/F_0$, which describes the largest local electric field present at the electrode, due to sharp protrusions from the surface. Here, the local enhancement does not immediately lead to cathode or anode-initiated breakdown, but allows for large currents to flow between the plates through the process of field emission. Consequently, currents measured between the plates may be much larger than those expected from equation 3.4 for the measured field, however, fall in accordance with this equation when the true field is given by βF_0. Eventually, the large currents will trigger a full cathode or anode-initiated breakdown. In this respect the parameters γ and β can be considered synonymous. Experimental [81, 82, 83] and theoretical results [84] suggest values of β in the range of 1-100. From figure 3.2, for a field enhancement factor of $\sim$10, a current of 1 μA due to field emission is expected to occur at values of F of a few MVcm^{-1}, in congruence with our own and other [81] experimental observations.

3.2 Experimental Set-up

In an effort to achieve electric field strengths approaching 1 MVcm^{-1} an experiment was designed to apply high voltages to pairs of broad area electrodes separated by a gap of 0.5 mm. Note that 'broad area' is a term reserved for electrodes whose surface area is much greater than the electrode spacing. We measure the resulting currents drawn across the plates and investigate a number of surface preparation techniques to increase the electric field at which breakdown is incurred. We present results for pairs of stainless steel, cylindrical electrodes and discuss the benefits and pitfalls of different polishing techniques. Detailed reports of the experimental set-up can be found in references [83] and [85].

3.2.1 Outline

Two electrodes are placed under vacuum with a separation of 0.5$\pm$0.05 mm and isolated from the ground. Two contact probes provide the required voltage and receive any current passing between the electrodes. Each probe is connected to an ammeter which converts any measured currents into a voltage which is subsequently translated into a digital frequency. This is fed into a data acquisition board and read by an on-board clock. Computer software provides a user interface to display the data in real-time and allow control of the voltage supplies. Breakdown is recognised as any event for which a sudden, large and persistent current above a preset baseline

is recorded by the ammeters. In general such large currents often destroyed the ammeters so the values were rarely recorded. Breakdown was simply marked at the field at which this occurred. Breakdowns are divided into two types; those for which an electrode shorts to ground, indicated by current measured by only one ammeter and those for which current travels across the vacuum gap, indicated by equal and opposite currents measured by both ammeters. The details of the apparatus are described below.

3.2.2 Vacuum System

The main chamber is a $4\frac{1}{2}$" 4-way cross with two DN40CF flanges welded at the junction to allow for the inclusion of two observation windows, used to help in the assembly of the electrode contacts in the chamber, figure 3.4.

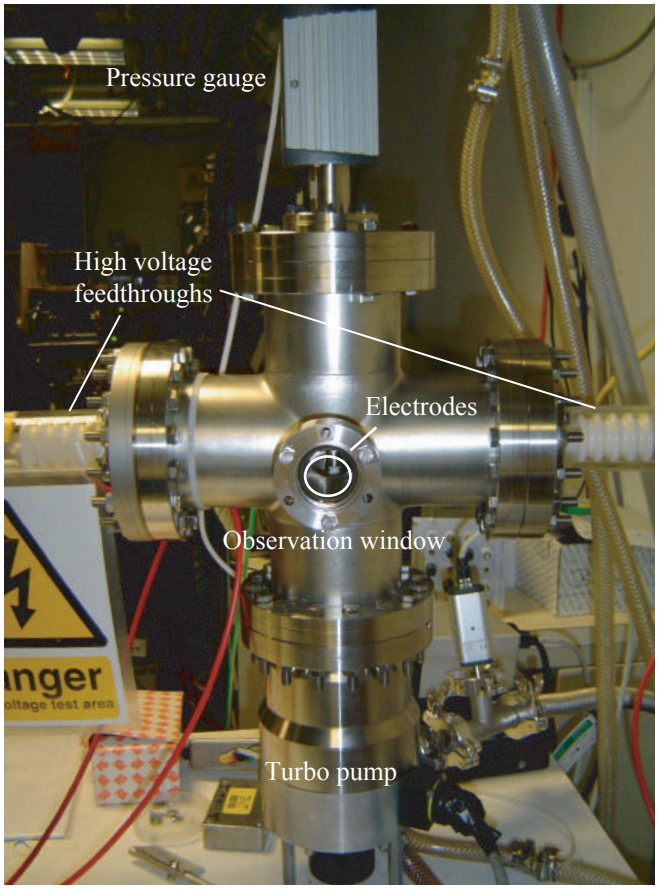

Figure 3.4: A photograph of the experimental apparatus used to collect the data discussed in section 3.4. The main attributes of the experiment are highlighted. A detailed schematic of the electrode support rig is given in figure 3.5.

The chamber is pumped with a 150 l/s turbo pump backed by a roughing pump. Without baking, the chamber typically reached a pressure of 5×10^{-8} mbar after 1-2 weeks of pumping. To ensure that breakdown is a consequence of field emis-

sion alone rather than ionsisation processes between the electrodes it is important to achieve a suitable level of vacuum. Particles occupying the space between the electrodes are ionized by electrons released into the gap through field emission. The higher the applied potential difference the more energy these electrons gain as they accelerate across the gap and thus the greater the probability of an ionisation event occurring. An analysis of the required vacuum was made by F. Paschen [86]. It was shown empirically that the breakdown voltage is dependent on the product of the gas pressure and the gap size. To provide sufficient vacuum insulation it is considered sufficient for the mean free path of the electrons to be long in comparison to the inter-electrode gap. For air the mean free path is on the order of several metres for pressures below 10^{-4} mbar, well above the pressures achieved here. The top arm of the four-way cross is used to support the electrode assembly described below and each side arm houses a contact probe used to supply the voltages to each electrode. The pressure in the chamber is recorded with a Leybold Penning gauge. A photograph of the apparatus is given in figure 3.4.

In general the electrode support assembly will be designed specifically for a given pair of electrodes, depending on their shape and size. In its simplest form the electrodes are mounted to a stainless steel rigid rig suspended from the top flange of the chamber and are contacted via the electrical feed-throughs at each side. For our tests the electrodes used were cylindrical stainless steel plates (see section 3.2.4) mounted to the rig shown in figure 3.5. The electrodes are electrically isolated by ceramic rods.

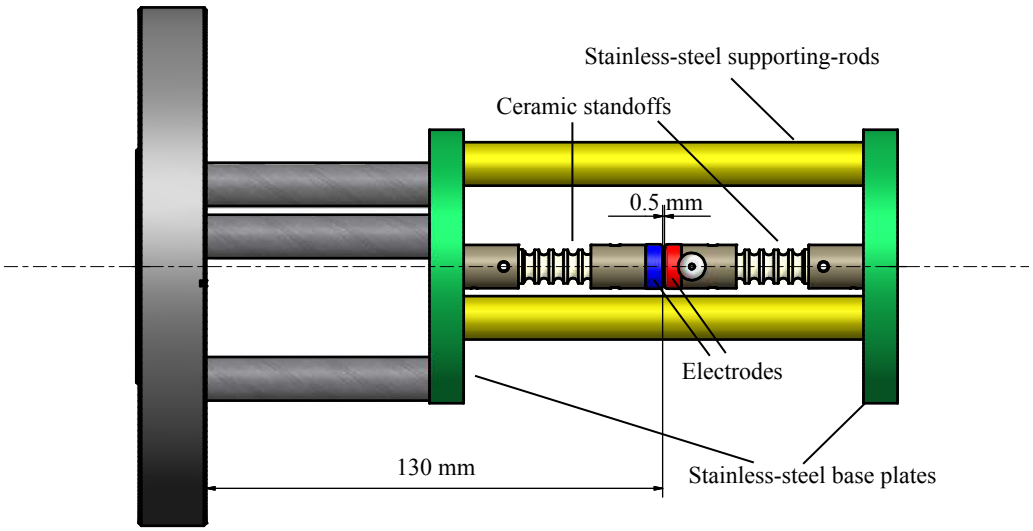

Figure 3.5: A scaled diagram of the electrode test rig used to support the stainless steel electrodes at a separation of 0.5 mm. The rig is suspended from the upper-most flange of the 4-way cross test chamber.

3.2.3 Measurement Circuitry

Two ammeters were built to measure any small currents drawn across the electrodes or between the electrodes and ground. Floating the inputs of an ammeter to many kilovolts is not possible for conventional designs thus necessitating the construction of a custom designed, carefully isolated ammeter, capable of recording currents in the micro-amp and nano-amp regime. Over time the circuitry used for our ammeters has been steadily improved. A current design for the circuitry and overall construction of the ammeter can be found in reference [87], however, below I will describe the key components of the design used for this experiment. The general process for collecting the data has not changed. The basic principle for our ammeters is no different to the standard you may find in a typical laboratory; the current drawn is passed through a sense resistor (1M Ω in this case) and the voltage dropped across it is measured. To read out this voltage, without electrically coupling the ammeter to any external readout circuitry, the recorded voltage is amplified to a value between 0-10 V and sent to a voltage-to-frequency converter. The converter outputs a TTL square-wave whose frequency falls between 0 and 10 kHz, proportional to the input voltage in a linear fashion, which in turn drives an LED at the same frequency. The light from the LED is coupled by an optical fibre to a photodiode which in turn reproduces the digital voltage signal, providing a measure of the voltage, and thus the current drawn, to the computer. A data acquisition board (DAQ) is used to collect the readings at a frequency of 10 Hz. The DAQ board has only a single counter operating at 20 Hz; in order to read information from both ammeters a multiplexer is used to read the voltages from one ammeter on the rising edge of the clock signal and the other on the falling edge. This 100 ms resolution is ample for monitoring any interesting events during the experiment.

A number of protective features were installed in the ammeter circuitry in an attempt to minimise any damage caused by breakdown. In series with the 1 MΩ sense resistor is a chain of resistors totalling to 130 MΩ. These resistors reduce the maximum current that can be drawn at $\pm$20 kV to 150 μA. For currents of 150 μA, however, the voltage drop across the sense resistor still exceeds 150 V which is too high an input for the op amp. To prevent voltages exceeding 1 V a pair of oppositely biased diodes are placed in parallel with the sense resistor. For potential differences greater than 1 V the diodes short and the voltage input of the op-amp is reduced to $\sim$0 V.

The data collected by the DAQ unit were interpreted by LabVIEW using a simple program outlined by the block diagram on the following page. This allowed real-time readouts of the current drawn across the two ammeters and the voltages on each electrode relative to ground. The data were also recorded for later analysis.

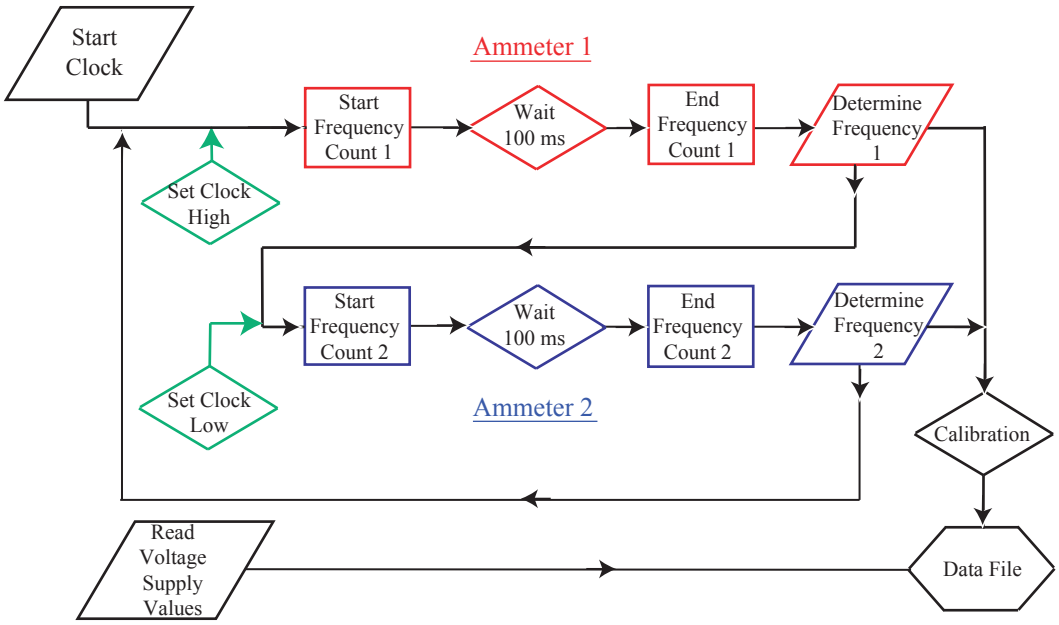

Figure 3.6: A flow chart highlighting the basic operations of the program implemented in LabView to record the current and voltage data from the experiment.

The power supplies are connected to the electrodes in series with the ammeters. Each electrode is charged by a BERTAN 602C 200P,N supply, applying a bias between 0.5 and 20 kV$\pm$1% at 400 μA $\pm$ 2%. Given a separation of 0.5 mm between the electrodes this allows fields up to 800 kVcm^{-1} to be reached.

3.2.4 The Electrodes

The electrodes are stainless steel cylinders with a height of 5 mm and and a diameter of 13 mm. The circumference of the face of each electrode is rounded with a 1mm radius in an attempt to reduce regions of enhanced field. Each electrode is mounted via a tapped hole on the reverse to allow easy assembly onto steel rods, isolated from ground by the ceramic insulators separating the electrodes from ground. The electrodes plus ceramic insulators are mounted onto a rigid support rig maintaining a separation of 0.5 mm between the electrodes; figure 3.5. Each electrode also has a small indentation in the side wall to accommodate the tips of the contact probes. The capacitance between the electrodes is

$$C = \frac{\epsilon_0 A}{d} \approx 2.35 \text{ pF}. \tag{3.6}$$

Each electrode is connected to the power supply using high voltage feedthroughs. The feedthroughs were mounted in a ceramic casing to prevent any shorts to the grounded chamber and entered the chamber through the side arms. The pins of the

feedthrough were spring loaded to allow some degree of compression improving the stability of the contacts to the electrodes.

3.3 Polishing Techniques

As discussed in section 3.1 the surface quality of the electrodes is integral to achieving large electric fields. We have investigated a number of polishing techniques in an effort to achieve a qualitative understanding of the effects of different surface conditions under high fields. For each of the polishing methods employed the surfaces of the electrodes were examined under an optical microscope. To provide a control test untreated electrodes were examined under an optical microscope at 50× and 100× magnification; figure 3.7.

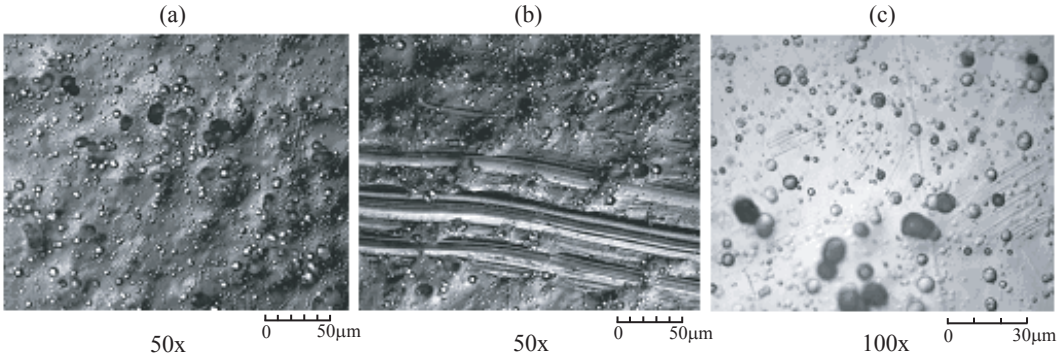

Figure 3.7: Photographs of unpolished electrodes viewed through an optical microscope. Figures (a)&(b) are taken at 50× magnification, and figure (c) at 100× magnification.

Clearly the surface is highly scratched and contaminated with residue, most likely consisting of water and grease. These residues will trap pockets of air whose lower work function may attribute to premature breakdowns.

Tests were performed on hand-polished, machine-polished, electro-polished and optically-polished electrodes. Each of these were tested under high electric fields and their surface quality was assessed. These results together with details of the polishing procedures are presented here.

3.3.1 Hand-polishing

All sets of electrodes were initially hand polished before being subjected to more advanced techniques. This involved manually polishing the surfaces with 600, 800 and 1200 grit silicone carbide paper, followed by further polishing with alumina paste. The electrodes were rinsed in cold water between each polishing stage. Before being placed in the chamber each electrode was also cleaned in an ultrasonic bath using a 5% solution of Decon-90 detergent, followed by an isopropanol wipe with

powder-free tissues. The quality of the electrode surfaces were examined under an optical microscope and are depicted in figure 3.8

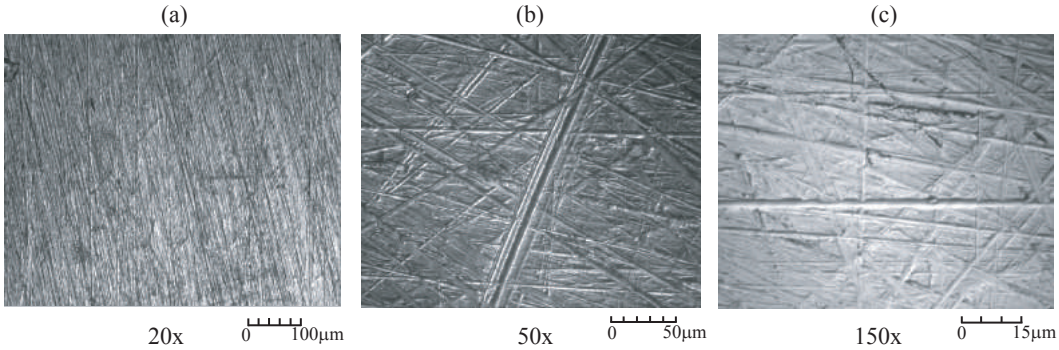

Figure 3.8: Photographs of hand-polished electrodes viewed through an optical microscope. Figures (a), (b) and (c) are taken at 20×, 50× and 150× magnification, respectively.

The most significant improvement between these electrodes and the untreated pair is the lack of residue evident on the surface; this is most likely a consequence of our cleaning procedures. Polishing has removed all of the largest scratches and has provided an overall smoother appearance.

3.3.2 Machine-polishing

Machine polishing, as the name suggests, uses a machine to apply equal force from the polishing cloth to the whole surface of the electrode. This has an obvious advantage over hand-polishing which is subject to human error and can be applied for many hours to achieve the desired finish. The electrode is rotated off-axis during polishing in order to prevent the formation of a grain on the surface in any particular direction. The electrodes are first polished with 1000 grit silicone carbide paper. Afterwards a 6 μm, 2 μm and finally 1 μm finish was subsequently achieved using progressively finer diamond polishing cloths. Between each stage the surface of the electrodes were rinsed and the final set were subjected to the same ultrasonic and isopropanol cleaning techniques as above. Machine polishing typically leaves a very fine grit on the surface of the plate which is not easily removed by standard cleaning techniques. Despite machine polishing providing an otherwise very smooth surface this grit is thought to significantly reduce the breakdown capacity of these plates. In addition the method employed to secure the plates during polishing applies a mechanical stress to the plates that may cause the surface to warp. Photographs of the surfaces achieved are given in figure 3.9 at magnifications of 10, 20 and 150 times.

Drawing comparisons between figures 3.8 and 3.9 the benefits of machine polishing over hand polishing are immediately obvious. The vast majority of the scratches

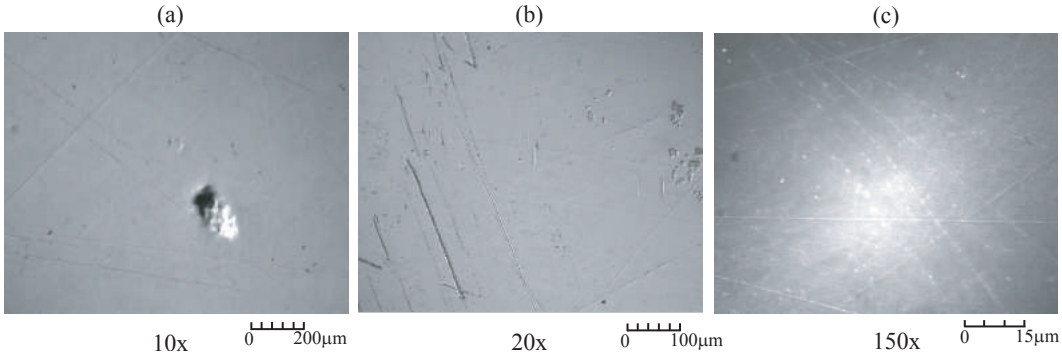

Figure 3.9: Photographs of machined-polished electrodes viewed through an optical microscope. Figures (a), (b) and (c) are taken at 10×, 20× and 150× magnification, respectively.

and whiskers populating the surface previously are now removed. Larger dents such as those in figure 3.9(a) still remain, however, these are less likely to create local areas of high current density, detrimental to high field capacities.

3.3.3 Electro-polishing

Electro-polishing is a common technique, used particularly in industry, for achieving a high quality finish to a large array of metals, including stainless steel. It is an electrochemical process whereby protrusions are removed from a metal surface by chemical etching. The metal test object is used as an anode and submerged in an electrolytic bath. The cathode is made from a different material, typically copper or aluminium, and is submerged in the same electrolytic bath. The bath is filled with a mixture of acids and a voltage is applied between the anode and cathode. As a current between the electrodes is drawn, metal ions are pulled from the anode surface and quickly saturate the solution near the surface of the test object. This provides a protective film over the anode with resistive properties. As electrons from the cathode bombard the anode the areas of the metal that protrude from the insulating film the most receive the most bombardment and are consequently eroded the quickest. In addition, due to electric field enhancement at the sharpest of protrusions a higher current density forms locally, further increasing the speed of erosion. As a result the test piece is gradually polished, with the most prominent protrusions being etched away first.

The first set of electrodes were polished in our laboratory using the recipes documented in [88] and [89]. Here, the electrolyte solution is composed of one part concentrated sulphuric acid, 3 parts concentrated phosphoric acid and one part of distilled water. The copper cathode is a copper sheet with a surface area several times larger than the surface of the electrodes which form the anode. A current of 5 A is applied across the electrolytic solution for up to 7.5 minutes and the solution

is stirred throughout the procedure. The procedure was repeated for polishing times of 1, 3.5, 5.5 and 7.5 minutes, the results for which are presented in figures 3.10(a-d).

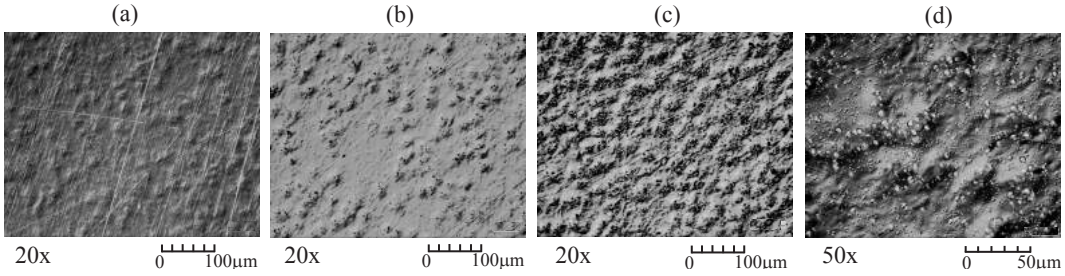

Figure 3.10: Photographs of electro-polished electrodes viewed through an optical microscope. These electrodes were polished in our labs using recipes documented in [88, 89]. Figures (a) to (c) are taken at 20× magnification, and image (d) is viewed at 50× magnification.

In all the images, significant pitting of the electrode surface is observed; the surface could certainly not be considered smooth. The pitting is thought to be a consequence of over exposing the surface to the polishing process. In light of later experiments performed by A. Kovaleva and J. Millen [82, 83], this is almost certainly the case, as described below. Evidence of any improvements to the surface of the electrodes are obscured by the degree of pitting. Large quantities of oxygen and hydrogen gas expelled during the procedure also appear to permeate the surface of the metal, particularly in image (d), and could not be removed by cleaning with alcohols. Alkali chemicals were also used in case the residue was acid-based, but were also unsuccessful. This residue may lead to a lower work function at local points on the surface of the electrode and may be detrimental to the breakdown capacity of this surface. Despite these results these electrodes were tested under high fields.

Later experiments performed by A. Kovaleva and J. Millen on a similar pair of electrodes employed professional electro-polishers to try to improve the surface quality. The electrodes were mechanically polished before electro-polishing. The electrolyte solution was comprised of two parts concentrated phosphoric acid, one part concentrated sulphuric acid and one part distilled water. A current of 0.2 A was passed between the anode and copper cathode for a period of 20 seconds. This recipe is significantly gentler than the one described above. The currents and exposure times are an order of magnitude smaller and the concentration of acids is weaker. Photographs of the resulting surfaces are given in figure 3.11.

These photographs depict a much smoother surface, almost completely free from scratches and whiskers, even when viewed at 150× magnification. There is a general roughness to the surface, most apparent at 50× magnification, however, the size of these protrusions are sub-μm and hopefully do not significantly increase the likelihood of breakdown.

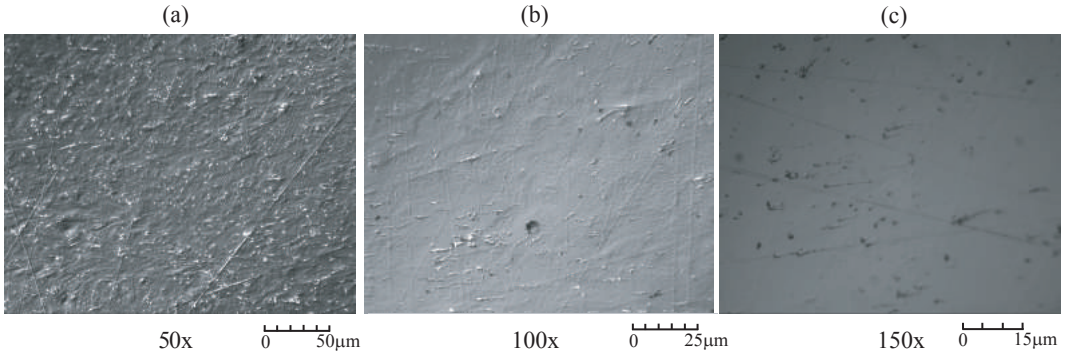

Figure 3.11: Photographs of electro-polished electrodes viewed through an optical microscope. These electrodes are polished in a much gentler manner than those in figure 3.10. Figures (a),(b) and (c) are taken at 50×, 100× and 150× magnification, respectively.

3.3.4 Optical-polishing

Optical polishing is normally a procedure reserved for glass surfaces and can provide a very uniform and level surface. The technique was applied to the electrodes to see if a similar quality could be replicated. The procedure is similar to the machine polishing described above but a finer grain diamond paste is used. After polishing, the electrodes were washed with warm ultra-pure water (18.0 MΩ cm), as suggested by C. Suzuki *et al.*, [90] to remove many of the impurities residing on the surface. The images below were taken after polishing but before treatment with ultra-pure water.

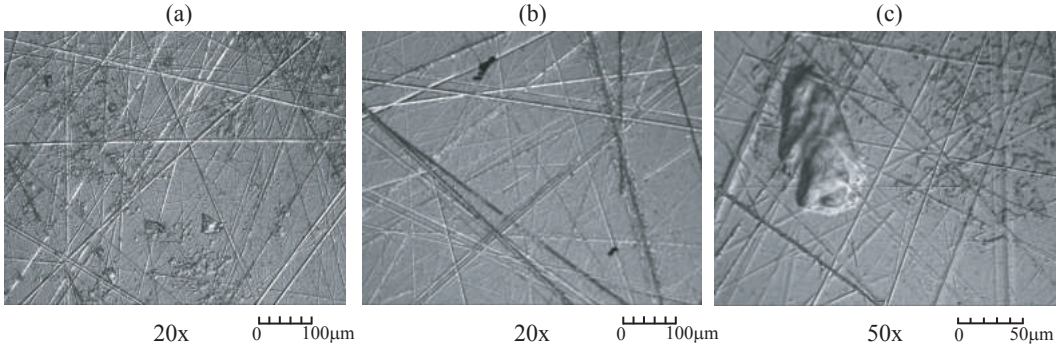

Figure 3.12: Photographs of optically polished electrodes viewed through an optical microscope. Figures (a)&(b) are taken at 20× magnification, and (c) at 50× magnification.

The results are a little disappointing, with many scratches and whiskers still pervading the surface. The surface quality lies approximately mid-way between the results returned for hand-polished and machine-polished electrodes. The poor quality surface may be a consequence of polishing for too short a time, and longer machine times may lead to higher quality results.

3.4 Results

For each test performed on a given pair of electrodes the current across the electrodes and the voltages at each electrode were recorded via the DAQ instrumentation. Figure 3.13 shows the typical results for the current recorded by one ammeter and the voltage at the corresponding electrode.

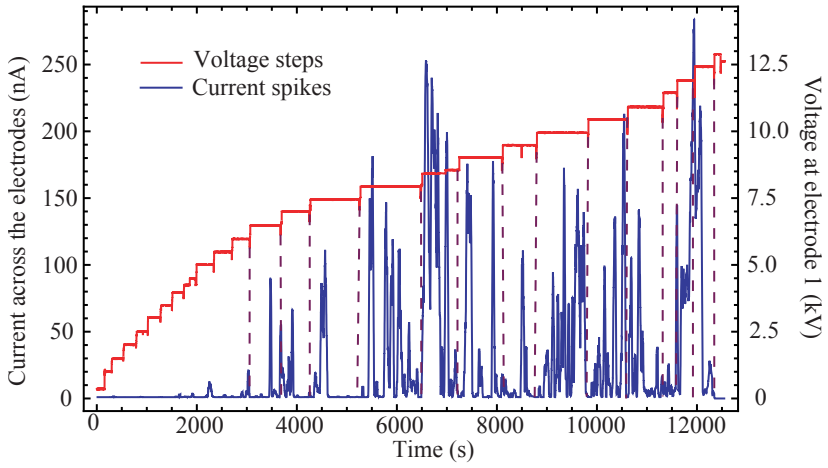

Figure 3.13: A graph presenting the current and voltage data collected by the DAQ board for a typical experimental run. The red line depicts the voltage difference applied to the cathode relative to ground as a function of time. The blue data presents the current measured by the ammeter, connected between the cathode and ground. A similar set of data is collected at the anode. Dashed lines are presented to guide the eye and illustrate the relationship between current spikes and changes in voltage.

The voltage was gradually increased over many hours in steps of 500 V. For the majority of results the voltage at each electrode was increased in an alternating fashion. For example, electrode 1 would first be floated to 0.5 kV. After some interval, electrode 2 would be also be increased to -0.5 kV, followed by another pause followed by an increase at electrode 1. This pattern would be repeated until breakdown occurs and is referred to as symmetric testing. Asymmetric testing whereby only one electrode is floated was investigated to measure if this affected the breakdown capacity of the plates, however, no significant changes were observed. As the voltages become larger it is common to observe large current spikes immediately after a voltage increase; some of these correlations are notable in figure 3.13 (dashed lines). These currents gradually decay and the voltage is increased again. These observations are consistent with the current conditioning techniques discussed in section 3.1.

One proposed explanation for the current spikes is the clump hypothesis, put forward by Cranberg [91]. After each voltage increase, loose adhering particles on the electrode surface are removed and accelerated towards the opposite electrode. By increasing the voltage slowly the kinetic energy of the removed particle is minimised at the time of release, reducing the potential damage to the other electrode where

the particle collides. Once the particles that can be removed for a given voltage are gone the voltage can safely be increased and a new set of particles with slightly more adhesion to the electrode can be safely removed.

Not all the currents observed are due to the removal of such particles. Both anode- and cathode-initiated breakdown processes can occur in a controlled manner if the voltage is increased incrementally. The current density at the anode or cathode gradually approaches its critical density and vaporisation of the material forming the protrusion can occur. By gradually increasing the current density the energy released at the point of vaporisation can be minimised, thus reducing the probability of an avalanche effect causing breakdown.

Observations from other data sets also suggested an ohmic relationship present, indicated by a constant and persistent current which increases linearly as a function of voltage. This is not consistent with the data expected for field emission but may be attributed to leakage currents from the ammeters at high voltages. This explanation does not account, however, for the symmetry in the current data between the two ammeters that suggests current really is passing directly between the plates.

A summary of all the results collected is given in table 3.1. Tests 1-3 were performed on the same pair of electrodes and the benefits of repeated current conditioning is reflected by the continually increasing maximum electric field values obtained for each test. Tests 4 and 5 were performed with a second pair of electrodes that were electro-polished using the procedure described in section 3.3.3. Although electro-polishing was expected to increase the field value at breakdown, no such observations were made for this electrode pair. Because the electro-polishing procedure was new to us at this time it is thought that the lower results are likely to be a consequence of poor polishing procedures rather than a fundamental problem with the technique. This is supported by later experimental data as well as by measurements of the surface quality of the plates used in tests 4 and 5 under an optical microscope, figure 3.10.

The electric field values for tests 6 to 9, quoted in table 3.1, are the maximum values obtained after a series of repeat tests on the same electrode pair for a given polishing procedure. For each electrode pair the largest field value obtained typically increased for each test performed on them, indicating an improvement in the surface quality of the electrodes as a result of current conditioning.

The images of the electrodes taken by the optical microscope for the various polishing techniques indicate that the electro-polished plates had the highest quality surface, followed by the mechanical, optical and hand-polished electrodes in that order. This is not, however, reflected in the results for the electric field shown above. To explain these discrepancies it is necessary to look at other contributing

Table 3.1: A summary of the results for the largest obtained electric field strengths for various polishing techniques.

Test No.	Polishing Technique	Largest Positive Voltage (kV)	Largest Negative Voltage (kV)	Plate Separation (mm)	Largest Electric Field ($MV\,cm^{-1}$)
1	Hand-	6.5	7	0.5	0.27
2	Hand-	18	0	0.5	0.36
3	Hand-	13	13	0.5	0.52
4	Electro-	6.5	6.5	0.5	0.26
5	Electro-	6.5	6.5	0.5	0.26
6	Hand-	13.2	13.2	0.6	0.44
7	Machine-	8.0	8.2	0.45	0.36
8	Electro-	13.8	13.8	0.35-0.40 [5]	0.79
9	Optically-	16.2	16.2	0.4	0.81

factors. The additional treatment of the optically polished plates in warm, ultra-pure de-ionised water is thought to remove a significant quantity of any residue remaining on the plates after polishing. This leads to a reduction in currents caused by Cranberg's clump hypothesis, leading to a larger, more stable electric field at high voltages. Contrariwise, the surface of the mechanically polished electrodes, although measured to be smoother than hand or optically polished plates, has a thin layer of metallic grit left on the surface after polishing. These loosely adhering particles will encourage significant currents as a result of clump transport during the application of large electric fields, increasing the probability of premature field breakdown. Lastly, it was noted after testing that the electro-polished plates were misaligned due to a mistake in their machining. This created a non-uniform gap spacing between the plates with the centres of the plates lying off-axis. It is thought this poor alignment is responsible for achieving comparatively smaller fields despite the improved surface quality. These results perhaps demonstrate the importance of maintaining a clean as well as smooth surface that must be carefully maintained prior to testing. The following sections will discuss the extra difficulties entailed when one wishes to place atoms between such plates in order to measure the effects of electric fields on cold atomic gases.

[5] The plates were not parallel during this test. A value of 0.35 mm was used to determine the size of the electric field.

3.5 Atoms Under Electric Fields

To be able to study trapped atoms under high electric fields the fields must be extremely uniform in order not to apply a force opposing the trapping potential confining the atoms. The force vector acting on an atom sitting in a non-uniform electric field is proportional to the gradient of the Stark shift experienced by the atom in the field. In this case this force must be smaller than the trapping forces supplied by our quadrupole magnetic trap, described in chapter 4. This gives an approximate criterium to be maintained, described by equation 3.7.

$$\mu \frac{\mathrm{d}B}{\mathrm{d}r} > \frac{\alpha}{2} \frac{\mathrm{d}(E^2)}{\mathrm{d}r} \; ; \tag{3.7}$$

where $\mathrm{d}B/\mathrm{d}r$ is the magnetic field gradient of the trap along r, α is the polarisibility of the atom, given in table 2.1, and E is the size of the electric field. In fact the conditions required for the plates are slightly stricter than this. The above equation only applies to stationary atoms, however, our atoms have a finite energy corresponding to $1/2k_BT$ per degree of freedom which must also be considered. The atoms will experience an electric field gradient inside a parallel plate capacitor in two cases. Firstly, if the atoms are close to the plate edge they will experience a gradient due to the fringe fields present here. In the second case a field gradient will be formed if the two plates are not perfectly parallel. Each of these cases must be considered carefully.

Until now we have discussed the use of broad-area electrodes, defining broad-area plates to mean those for which the plate diameter is much greater than their separation. However, in our case we are trapping atoms between the plates, requiring a slightly stricter definition for broad-area plates. These atoms are confined by relatively weak forces so only relatively small electric field forces present due to fringe-field effects are required to counter the trapping force. By calculating the gradient of the field as we approach the edge of the plates we can deduce how far in the atom cloud must sit between the plates to maintain a sufficiently uniform field across the atoms. This limitation will then serve as a definition for what it means to use broad-area plates in these circumstances.

The electrostatic potential, ϕ, between two parallel, infinitesimally thin, circular plates, of radius a, separation d and voltage $\pm V$ is [92]

$$\phi = \left(\frac{2V}{\pi}\right) \sum_{0}^{n} \left[\sin^{-1}\left(\frac{a}{x_-}\right) - \sin^{-1}\left(\frac{a}{x_+}\right) \right], \tag{3.8}$$

where

$$x_\pm = \frac{1}{2}\left(\sqrt{b_\pm^2 + (\rho + a)^2} + \sqrt{b_\pm^2 + (\rho - a)^2} \right)$$

and

$$b_{\pm} = n\,d + |z \pm \frac{d}{2}| \,.$$

The coordinates ρ and z are the radial and axial coordinates, respectively, whose origin is positioned along the central axis of the plates, midway between them. Circular plates have been used here as the circular symmetry simplifies the calculations necessary to determine the function, ϕ. Nevertheless, the solutions for the potential are not thought to differ significantly if rectangular plates with a small aspect ratio are used instead. The summation in equation 3.8 can be taken for n *ad infinitum*, however, since there is no simple integral solution for this formula the sum is normally truncated. We find that values of $n \geq 5000$ provide sufficiently accurate results (an error of 1%) for what is required here.

The force on an atom in a non-uniform electric field is given by the gradient of the Stark shift, $E_{stark} = -\frac{1}{2}\alpha F^2$, where α is the atom's polarisability and F is the electric field strength given by the gradient of ϕ. For the electric field sizes we will be applying the atoms confined by the magnetic trap will be strong (electric-) field seeking. Because the field between the plates is obviously stronger than the outside, we can ignore any forces in the radial direction, ρ, as they will always be restoring the atoms towards the centre. The force along z has been calculated as a function of distance from the edge of the plate. Figure 3.14 shows the maximum acceleration along the z-axis felt by a lithium atom as a function of distance from the plate edge, for a pair of plates with a separation of 0.5 mm, a radius of 12 mm and held at a voltage of ±25 kV (equating to a field between the plates of 1 MV/cm).

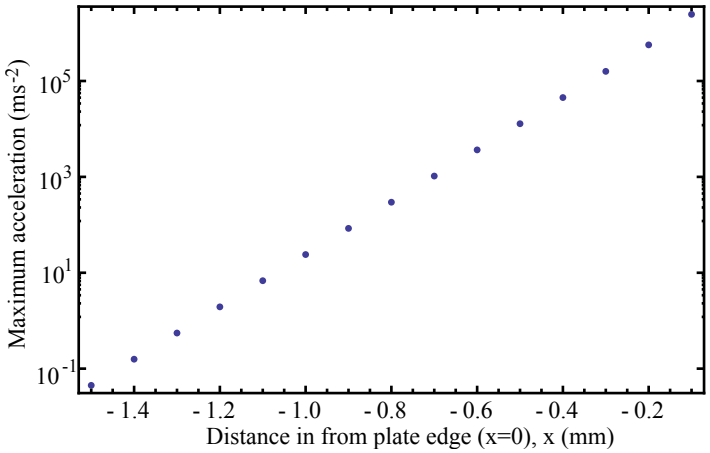

Figure 3.14: The maximum acceleration along the axis normal to the plate surface experienced by a lithium atom as a function of distance from the plate edge. The plate separation is 0.5 mm with a diameter of 24 mm. A field of 1 MV/cm is held between the plates.

The electric field very quickly becomes uniform as we move from the plate edge and the maximum acceleration drops off exponentially. Consequently, once we have

moved in from the edge by as little as 1 mm the forces on the atom are already much weaker than the magnetic trapping forces, which apply accelerations of the order of 300 ms^{-2}. Such a small distance will almost certainly be inconsequential in any future electrode designs.

Let us now consider the case where the plates are tilted along one axis; we will define the local coordinate system for each plate such that the plate surface lies in the x-y plane and they are separated along z. If the plates are misaligned such that their respective x-axes are no longer parallel then we create a field gradient, $\mathrm{d}E_x/\mathrm{d}x$; there is no gradient along y. Gauss' law then tells us that $\mathrm{d}E_x/\mathrm{d}x = -\mathrm{d}E_z/\mathrm{d}z$, implying a force along z. For now let us consider only those forces in the x-y plane.

If one simulates the force due to the presence of a non-uniform electric field as a tilt in the magnetic trap potential, this will lower the trap depth at one edge. If the force is too great the trap will tilt so the warmest atoms will be lost. For a temperature of 100 μK, a magnetic field gradient in the radial direction of 30 G/cm and an electric field of 1 MV/cm we can look at how the cloud changes as a function of angle between the plates[6].

Using the dimensions of the magnetic trap coils described in chapter 4 we begin with a trap depth of 1.5 mK. Restricting the problem to one dimension, for a misalignment in x of 0.1 mrad the cloud will spread out from a diameter of 1 mm to 1.3 mm in x and the trap depth is reduced to 0.8 mK. For a 0.15 mrad tilt from parallel the potential is tilted further; atoms have spread to a diameter of 4.7 mm and the trap depth is only 400 μK. By 0.2 mrad the trap depth is less than 40 μK and all but the coldest of atoms are lost. Such a degree of parallelism is extremely difficult to achieve, and in fact, if one takes into account the forces along z these restrictions become even tighter. Along z the field gradient is twice as strong, approximately 60 $\mathrm{G\,cm}^{-1}$. Despite this, already at 100 μK, the atoms are already spread over 0.5 mm filling the gap between the plates. Any degree of non-uniformity in the field will cause the cloud to spread out and atoms will be lost by collisions with the plate surfaces. It should be noted that the force due to these electric fields scales as the square of the field strength. Reducing the field to 0.8 $\mathrm{MV\,cm}^{-1}$ already reduces the forces on the atoms by more than 30%, such that for angles of 0.15 mrad, a large proportion of the atoms may be maintained.

The strength of these electrical forces at such high fields highlights the importance of achieving very parallel plates. From an engineering point of view, to achieve

[6]The magnetic field gradient chosen here is a realistically achievable value given the constraints of the experimental set-up and is discussed further in chapter 4. The temperature of the cloud is required to compress the cloud to a size that can sit comfortably between the plates given the size of the magnetic field gradient. Proposed methods to reach this temperature range are discussed in section 4.2.

even a 1 mrad tolerance over plates several centimetres in length is extremely difficult. For this reason we wanted to design a rig that can be adjusted post-assembly while offering a way to measure the degree of parallelism without contacting the delicate surfaces of the plates.

3.5.1 ITO coated electrodes

By using glass plates interferometric techniques can be employed to measure the degree of parallelism of two neighbouring plates. By using the plates as a Fabry-Pérot interferometer the parallelism can be determined to better than 0.1 mrad. In addition we can use lasers to accurately determine the separation of the plates without requiring the use of feeler gauges that can damage the surface. To test these properties we used borosilicate crown (BK-7) glass plates, with a refractive index of 1.52±0.01. While glass is not particularly conductive it can be coated with a transparent conductive medium as used for many industrial purposes. Indium-tin-oxide (ITO)and aluminium doped zinc-oxide (AZO) are the most common coatings for these applications, while ITO is more suitable for redder wavelengths [93]. Although ITO is not especially transparent at optical wavelengths, for thin coatings it is more than suitable. Comar were able to supply glass plates with a thin ITO coating on one side with a sheet resistance of $\sim$20 Ωm^{-2} and a transmission of $\sim$88% at 635 nm. The glass plates were cut to 42×24×1.1 mm with an optically buffed radius of 0.5 mm on all edges and a buffed radius of 1 mm on all corners. This additional polishing is to reduce any field enhancement at sharp edges. It was not possible for Comar to coat over these edges and on to the reverse of the plates; this is necessary in order to make contacts to the ITO coating without placing any contacts inside the gap between the electrodes. To circumvent this problem we applied our own coating around the edges of each plate and onto the back. We have the capabilities to coat a number of metals onto a surface, including aluminium, gold, platinum and silver. For gold and silver coatings a 5 nm under-layer of chromium must be applied first to help the metals to cohere. We chose to use platinum as it has a very high conductivity and does not require a time-consuming under-coating. A coating thickness of 20 nm was deemed sufficient to maintain a good level of conductance without impeding the electric field capacity of the electrodes. To coat the plates a soft teflon mask was created to cover the areas that needed to be protected; the mask was built from two teflon blocks, clamped together around the plate, leaving two edges exposed during each coating. Once all four edges were coated the plate was reversed and the process was repeated. The coating creates a border, $3-4$ mm from the edge that connects the front and back of the plate. The coating is formed using an RF field to produce a plasma on the surface of a platinum disk. The plasma slowly

sputters metal onto the piece which slowly rotates above the platinum plasma. Typical coating speeds are approximately 0.1 Å/s. A photograph of the finished plates in their assembly is given in figure 3.15.

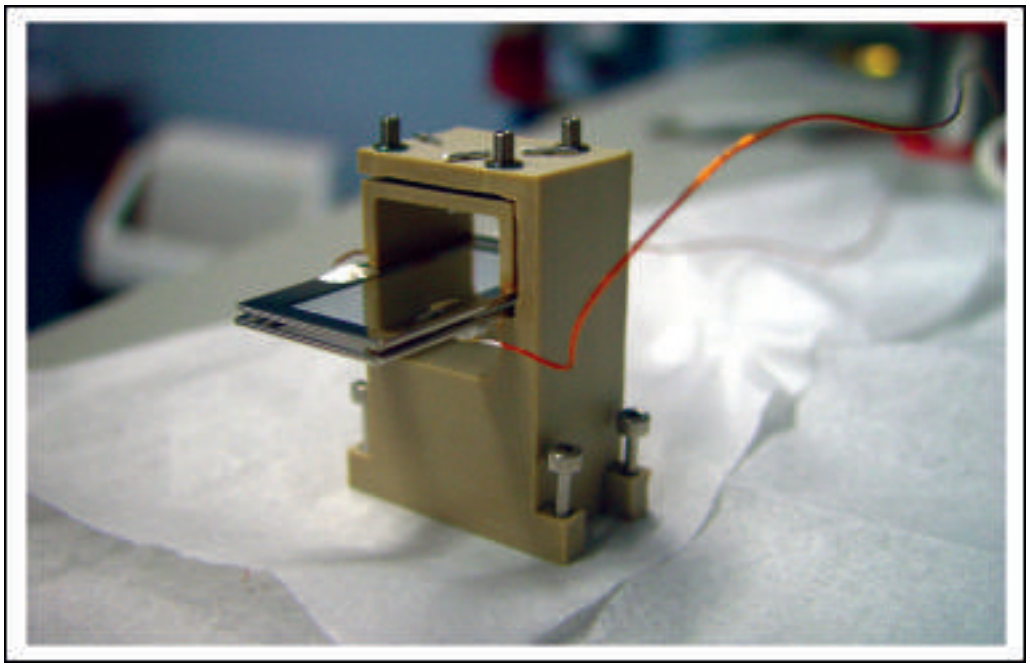

Figure 3.15: A photograph of the full assembly of the ITO coated electrode plates and their support rig. The platinum frame borders each plate on both sides, electrically connecting the front and back of the glass plate. This allows contacts to be made to the backs of the plates in order to apply a voltage to the ITO coating.

To connect to the plates a 0.8 mm diameter polymer-coated copper wire is stripped at each end and attached to the electrode plates at the platinum border, as depicted in figure 3.15. The wires are attached using an electrically-conductive silver epoxy, H27D, supplied by EpoTek. The other end of the wire is left loose, to be later attached to an electrical feedthrough.

The plates are glued onto two supports machined from a polyether-ether-ketone (PEEK) compound.

PEEK has a very low electrical permittivity as well as a very large breakdown voltage, making it ideal for isolating the electrodes from ground. The glue is supplied by EpoTek; we use H77 epoxy, a thermally and electrically-insulating glue, vacuum compatible down to the 10^{-10} mbar range. The mounting bases of each support have a 1 mm deep step machined into them along which the glue can be applied to connect the plates to the supports; a schematic is given in figure 3.16. The plate and support are carefully strapped together using a clean, soft nylon wire so as to avoid damaging the ITO surface. The glue is then applied into the resulting 1×1 mm channel, using a fine needle point. The glue is cured at 150°C for 90 minutes.

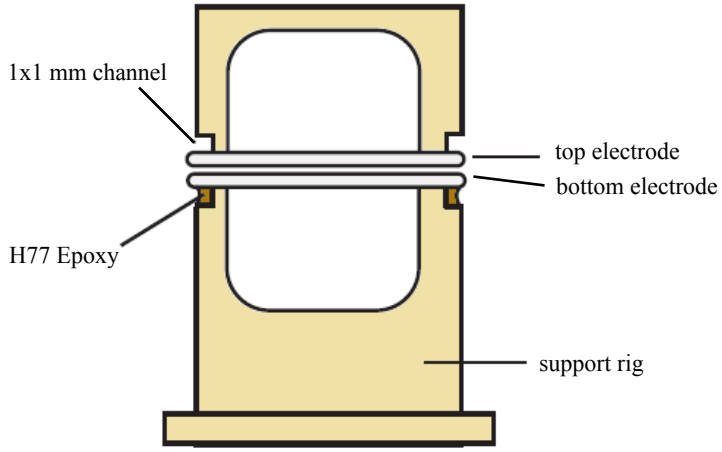

Figure 3.16: A basic schematic illustrating how the plates are secured to the support rig. The plates are tied to their separate supports using a clean, soft, nylon wire and H77 epoxy is applied in the 1 mm channels running along side the plates. The glue is cured at 150°C for $60 - 90$ minutes.

As will be discussed in chapter 4 the design of the support rig must allow optical access for the Zeeman slowing beams as well as meet several other geometrical constraints imposed by the size of the chamber. To accommodate these restrictions the plates are held along only two edges and protrude 21 mm from the rig. Figure 3.17 shows detailed drawings of the plates and the supporting rig.

The rig is composed of two supports. The lower support is fixed to the chamber floor and the lower plate is firmly attached to the PEEK arms. The upper support, carrying the second electrode, is suspended from the lower support (see figure 3.17) by three springs whose tension is maintained by three extra-fine threaded pushing screws (not shown) residing in the roof of the lower support. Each screw is capped with a 1 mm diameter metal sphere that sits in a depression made on the top side of the upper support. In a similar fashion to a mirror mount, these screws can be adjusted to tilt the upper plate relative to the lower plate as well as adjust the plate separation.

To align the plates, the support was fixed to an optical breadboard. The beam from a portable diode laser at 635 nm was expanded and collimated to a 1/e diameter of 10 mm and directed onto the plates, approximately perpendicular to the surface of the plate. We can treat the plates as a Fabry-Pérot interferometer and use the resulting interference pattern to measure their degree of parallelism. If we assume the front and back surfaces of each plate are parallel then any interference pattern imaged onto our screen can be assumed to be the result of a misalignment of the two plates from parallel. As the plates approach parallel the fringe spacing increases until the fringes can no longer be distinguished. For a collimated beam incident on a pair of wedged plates we expect to see a series of straight line fringes. Each fringe

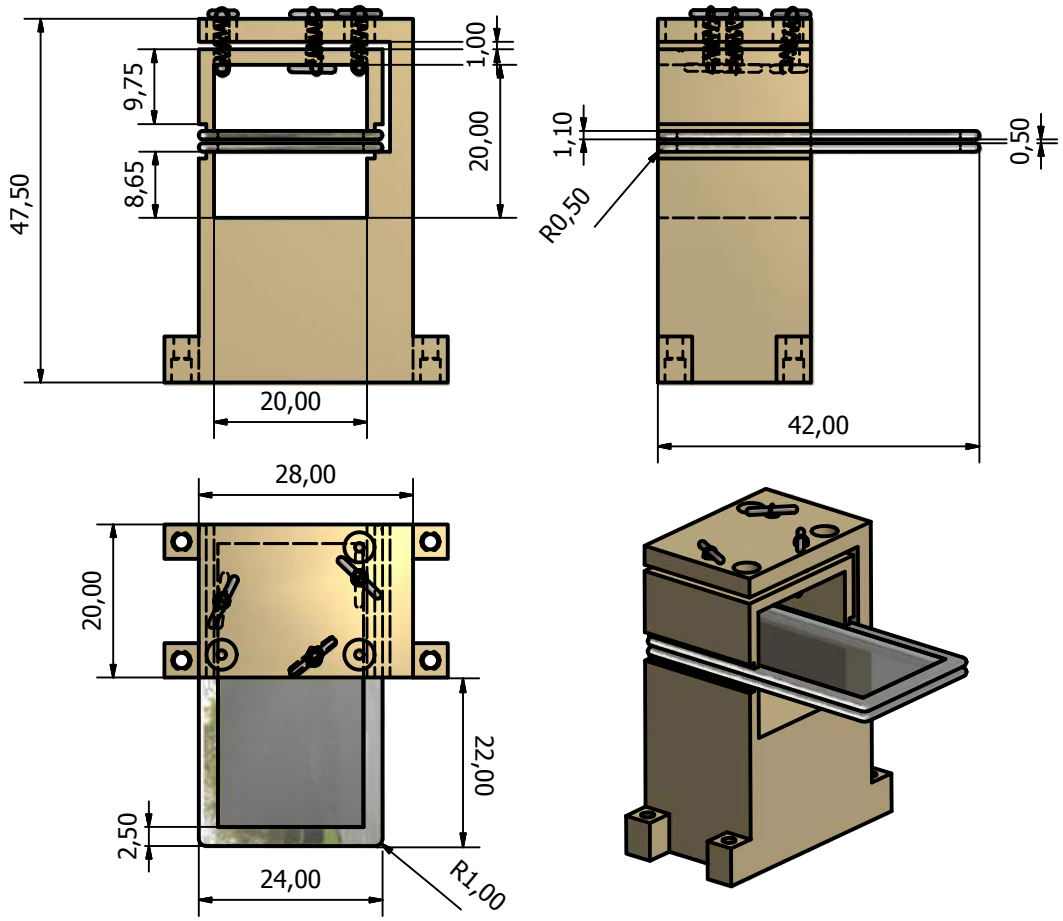

Figure 3.17: Detailed drawings of the adjustable capacitor assembly design for the ITO coated electrodes.

corresponds to an increase in the size of the plate spacing by $\lambda/2$. If there are N fringes over a distance L then we can write $b-a = N\lambda/2$, where $b-a$ is the change in gap spacing over the length, L. The angle, θ, between the plates can thus be written as

$$\theta = \frac{b-a}{L} = \frac{\lambda}{2}\frac{N}{L}, \tag{3.9}$$

where N/L are the number of fringes per unit length. To determine the limit to which we can measure the angle between the plates we have to take into account practical limitations. Assuming we can expand the laser beam to be larger than the plate area, then the largest value we can assign to L is the plate width, 24 mm. To observe any fringes, the number of fringes, N, must be at least 2. Substituting these numbers into equation 3.9 it can be shown that the angle between the plates can be measured to better than 25 μrad. This is well below the tolerances imposed by the forces deforming the trapping potential. When the above techniques were applied in practice the precision of the results were limited by the formation of concentric rings

as the plates approached parallelism. It is thought that these rings are a consequence of a relative curvature between the plates. This has been modelled in figure 3.18 as a single curved plate next to a flat plate for simplification.

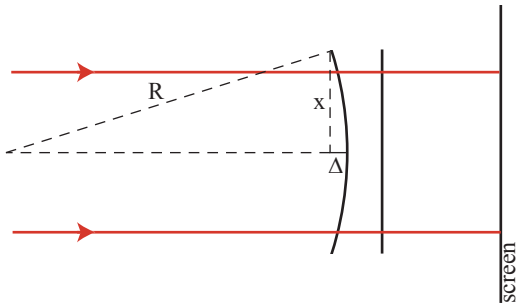

Figure 3.18: A schematic illustrating how a relative curvature between the electrode plates will lead to interference fringes, separated by a distance x.

If x is the distance between two fringes and we assume the bowing is circularly symmetric then the value Δ is equal to $\lambda/2$ and can be re-expressed in terms of the radius of curvature as $x^2/2R$. Substituting x for L/N we can then write $R = (N/L)^2(1/\lambda)$. The measured fringe spacing of the rings was approximately 4 mm, giving a value for R of 25 m. This implies our plates are still very flat. The curvature does limit our ability, however, to determine the parallelism of the plates. If we insert a fringe spacing of 4 mm into equation 3.9 we find we can measure the angle between the plates to no better than 80 μrad. Although this value is approaching the tolerances calculated at the beginning of this section this degree of parallelism is much better than can be engineered between two metal plates using standard machining techniques.

To measure the plate separation, a second much finer laser beam of diameter 300μm is passed through the plates at a 45 degree angle relative to the surface of the plate and parallel to the long plate axis as shown in figure 3.19.

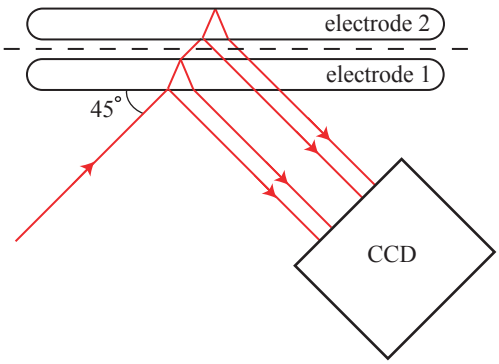

Figure 3.19: An illustration of the optical set-up used to measure the separation of the electrodes. The red lines indicate the beam paths that reach the CCD after reflections from the plate surfaces, as presented in figure 3.20.

A CCD camera with a pixel size of 9.9 μm is then positioned to pick up the reflections from the various surfaces of the plates. The face of the CCD is positioned to lie perpendicular to the beams. Given the approximate refractive index of the plates it is then possible to determine the spacing between the plates. There is a complicated array of reflections picked up by the camera from both the front and back surfaces of each plate and from multiple reflections inside each plate. However, by carefully varying the distance between the plates and measuring the distance between the numerous reflections each reflection can be matched with a particular surface. A thin piece of paper can be slid between the plates to help determine which reflections correspond to electrode 1 and which to electrode 2. Figure 3.20 gives the results for a particular plate separation. The upper image is the 2D intensity profile collected by the CCD. The fact that the beams lie in a horizontal plane indicates the beam is very parallel to the long axis of the electrodes and that the value calculated for the spot separation is an accurate measure of the plate spacing. The graph below is the 1D intensity profile obtained by integrating over the image above. The blue curve depicts the profile obtained for electrode 1 only and the purple curve is the profile collected from both electrodes.

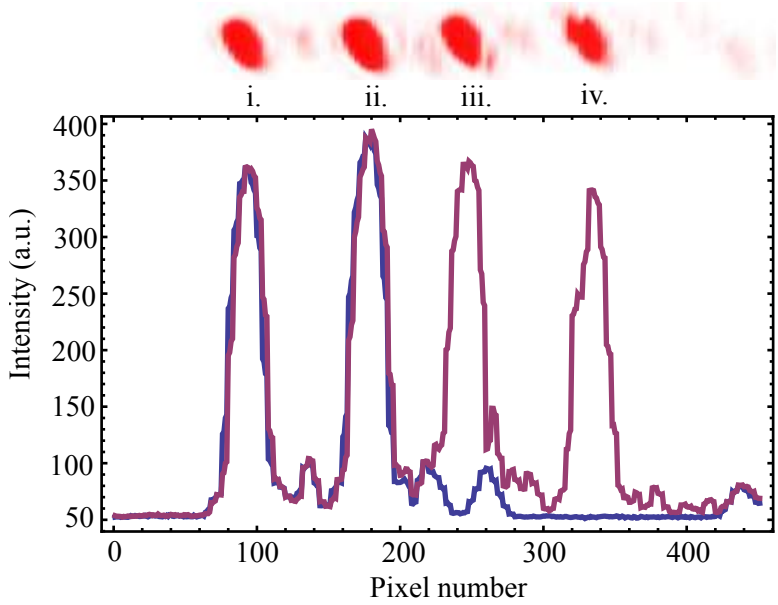

Figure 3.20: The upper image is the 2D intensity profile collected by the CCD. Below is the 1D intensity profile obtained by integrating over the image above. The blue curve depicts the profile obtained for electrode 1 only, and the purple curve is the profile collected from both electrodes.

By comparing figure 3.19 with the graphs above it is easy to see that, providing the plates have the same thickness, the distance between peaks (i) and (ii) and peaks (iii) and (iv) should match. Their separation suggests a plate thickness of 1.17 mm, in good agreement with the quoted nominal value of 1.1 mm. Given the pixel size

the gap between the plates can also be calculated from the distance between peaks (ii) and (iii). In the above figure the gap between the plates was measured to be 483±10 μm. This is significantly better than the typical 50 μm error associated with feeler gauges.

Summary

Our aims for these experiments were to show that it is possible to achieve electric field strengths of up to 1 MV/cm. To date we have created fields of 0.81 MV/cm for stainless steel electrodes; a significant step towards this goal. Our experiments also suggest that achieving higher fields should be possible with greater care given to cleaning the electrodes after polishing, using warm, ultra-pure water. In future experiments other materials for the electrodes will also be tested. An electric field of 1.08 MV/cm was achieved by Alpert *et al.* [81] for tungsten plates at a separation of 50.1 μm. These plates were baked before any field was applied, but, were not subjected to the same polishing techniques we have available here. These improvements may be associated with the much larger electron affinity of tungsten (78.6 kJ/mol) compared to iron (15.7 kJ/mol). To bake our plates under vacuum considerable efforts must be undertaken to create a functioning design and thus have not been employed to date. Future designs may consider this possibility and the possibility of baking the electrodes before they are placed under vacuum, which is a trivial task in comparison, albeit less effective.

We have also demonstrated the ability to produce an electrode support unit that may be adjusted to provide ITO-coated glass electrodes with a very high degree of parallelism, using interferometric techniques. Such plates also allow us to accurately measure the spacing between the plates to within 10 μm with very little effort and without making physical contact with the plates. With improved equipment accuracies of $\pm 1\mu$m are very achievable. The practicalities of using such a set up for producing high electric fields will be discussed in the following section.

3.6 Polarising the Atoms

The experiments discussed in the preceding chapters have been aimed at producing an ultra-cold source of lithium atoms that can be transported and held in a magnetic trap. From this point we may chose to either transport our atoms to a trap over-lapping a molecular trap, or instead place the atoms directly under a uniform electric field. The experiments discussed in this chapter were performed with an aim in mind to study our cloud of Li atoms under electric fields approaching 1 MV/cm. At low fields we can study the effects of the combination of magnetic

and electric fields acting on the atoms in preparation for our proposed sympathetic cooling experiments. At intermediate fields we also have the opportunity to detect a previously unobserved shape-resonance predicted to lie at an electric field value of 540 $\mathrm{kV\,cm^{-1}}$, [15]. Lastly, at the highest fields we can search for phase transitions in an increasingly polarised ultra-cold gas. To achieve these goals the electrodes and their support rig must be designed with a list of challenging criteria in mind. The most important of these are presented below:

1. The electrode plate surfaces must be extremely smooth.
2. The electrodes must be parallel to within at least 0.1 mrad.
3. The electrodes must be able to maintain stable DC electric fields up to 1 MV/cm.
4. The electrode rig must be stable and rigid enough to sustain both a constant separation and the required degree of parallelism against electric field strengths of up to 1 MV/cm.
5. The electrodes must sit over the magnetic trap and the support rig must conform to the volume constraints of the chamber so as to avoid blocking any optical access to the MOT and Zeeman slower.

While individually challenging, each of the above criteria have been addressed and where the goals have not been met entirely, significant progress has nonetheless been made. The all-metal electrode pairs tested in section 3.4 were used to determine the effects of a number of polishing techniques designed to optimise the surface quality of the plates and fields as large as 0.81 MV/cm were achieved. Several suggestions have been put forward to extend these fields beyond the 1 MV/cm milestone. The support structure of the high-field test rig was very rigid, but was not designed to meet criteria two or five. In light of this, the prospect of using ITO coated electrodes was put forward. The results from these tests meet both the first and second criteria, while the design of the support rig was specifically chosen to meet the conditions imposed by the fifth criterium. Unfortunately, due to time constraints the ITO electrodes were not tested under high electric fields and their performance under these conditions is not known. However, a number of papers documenting the tests on ITO coated electrodes under high electric fields show promising results [94]. A test rig for the ITO electrodes, compatible with the apparatus described in section 3.2, has been constructed but to date has not been put into operation.

Unfortunately the current design for the ITO electrode support rig does not meet the conditions outlined in item four. When a potential difference is applied

between the plates the electrostatic forces will pull the plates together. For very large electric field strengths these forces are very significant. If the separation between the electrodes is required to be maintained to better than a few microns the strength of the supporting structure must be equally sufficient to compensate. The force acting on one plate due to the field generated by the other is given by the charge of one plate multiplied by the electric field. The total electric field between the plates is $E = \sigma/\epsilon_0$, where σ is the charge per unit area of surface. The force is then

$$F = \frac{Q\,E}{2} = \frac{Q\,\sigma}{2\epsilon_0}. \tag{3.10}$$

where the factor of 1/2 is to include the field strength from only one plate. Using $Q = CV$, $C = \epsilon_0 A/d$ and $E = V/d$, where V is the potential difference between the plates, d is the size of the gap, A is the plate area and C is the plates' capacitance, we can re-express 3.10 as

$$F = \frac{1}{2}\epsilon_0\,A\,E^2. \tag{3.11}$$

The force between the plates scales quadratically with the field and linearly with the area. For plates with an area, $A = 10.08\,\mathrm{cm}^2$, with an electric field of $E = 1$ MV/cm the force acting on each plate is 45 N.

The spring constant for the springs used to retain the upper electrode support rig in the design described in section 3.5.1 is given to be $0.66\,\mathrm{Nmm}^{-1}$. Once aligned in the correction position their extension is approximately 1.6 mm, equating to a restoring force of $\sim$1 N per spring. Even with three springs this restoring force is much smaller than the force between the plates. Thus for large enough fields the plates will be pulled together. To produce a field of only 3 N, given the same plate dimensions we can only maintain a total electric field of $260\,\mathrm{kV\,cm}^{-1}$ between the plates. This support rig then is only suitable for relatively small fields. The force between the plates grows as the square of the total electric field so without a significant redesign this method of support is unsuitable for our applications. However, our plate design does allow us the possibility to accurately measure the parallelism and separation of rigidly supported glass plates to test their suitability before installation. While the possibility of constructing adjustable plates, capable of withstanding the forces quoted above, has not been ruled out, it is also worthwhile exploring other design options.

One point of significant importance omitted from the above calculations is the assumption that the plates are supported uniformly across their surface. In fact, in the designs discussed in section 3.5.1 the plates are suspended over the magnetic trap centre, unsupported for approximately half their surface area. Even if this design was adjusted such that the plates where rigidly held by their support, the

over-hanging portion will bend and even break for sufficiently high electric fields. We are thus presented with two possible solutions to overcome these problems. Firstly, the plates and support rig can be mounted entirely between the trap coils such that no over-hang is necessary. The plates can then be constructed from ITO-coated glass and supported rigidly or otherwise, depending on the merits of each design. However, given the extremely stringent geometrical constraints available inside the chamber such a design may be very difficult to implement. The alternative is to construct an electrode whose shape provides a suitably large moment of inertia, from a material with a sufficiently high Young's modulus such that the elastic forces resisting deformation exceed the electrostatic forces pulling the plates together.

To calculate the degree of displacement, u, of our electrodes under their electrostatic forces one must solve the Euler-Bernoulli beam equation. Our model is describe by a rectangular plate, of length, L; width, W; and thickness, t, rigidly fixed at one end. A schematic for our plates is given in figure 3.21.

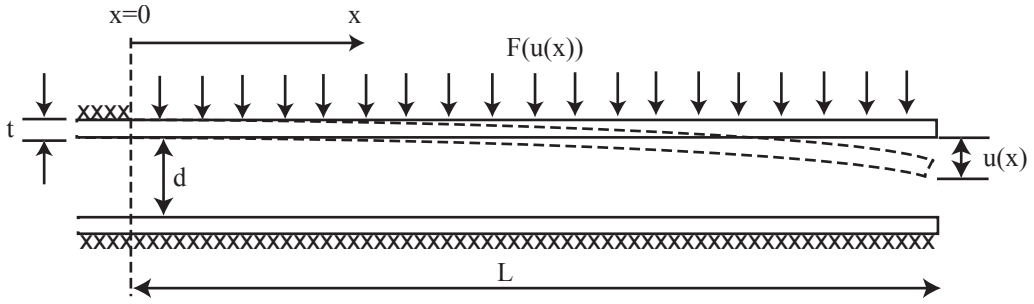

Figure 3.21: A schematic depicting the model used to calculate the effects of an applied electric field between two plates. One plate is rigid, the other is fixed along only one edge. The plates have dimensions $L \times W \times t$ and experience a displacement $u(x)$, where x is the distance along the plate from the fixed edge. The applied force is a function of $u(x)$.

The distance from the fixed end to a point along the length of the plate is labelled x. The force on the plate is initially uniform over the surface and is given by equation 3.11. The second plate has identical dimensions but is fixed along the complete length; rather than model two plates bending towards each other the plates can be treated as a single free plate bending with twice the expected displacement. As the displacement, u, increases the electrostatic force increases and the free plate will bend further. The force per unit length, w, is therefore a function of the displacement of the plate from its initial position, u, which itself is a function of x. The degree of bending is dependent upon the Young's modulus, Y_M, of the electrode material and the second moment of inertia given by $I = t^3\,W/12$. For electrodes with an initial separation, d, and a potential difference, V, the load per unit length is described by the differential equation

$$w\big(u(x)\big) = Y_M I \frac{\mathrm{d}^4 u}{\mathrm{d}x^4} = \frac{\epsilon_0 W L V^2}{2(d - 2u(x))^2}. \tag{3.12}$$

The boundary conditions for a plate fixed at one end and free at the other are $u(0) = 0$, $\frac{\mathrm{d}u(0)}{\mathrm{d}x} = 0$, $\frac{\mathrm{d}^2 u(L)}{\mathrm{d}x^2} = 0$ and $\frac{\mathrm{d}^3 u(L)}{\mathrm{d}x^3} = 0$. Given a suitable parameter set, equation 3.12 converges to a steady-state solution for $u(x)$. In cases where the plate is too weak or the applied electric field is too strong the differential equation diverges from a stable solution; in these cases the plates continue to bend until they touch (or snap) and the distance between the plates becomes negative, leading to a non-physical solution. The Young's modulus for BK-7 glass is 82 GPa. If we try to solve equation 3.12 for the plate design discussed in section 3.5.1 at a field strength of 1 MV/cm, we find there is no stable solution. We can find a stable solution if we decrease the length of unsupported plate, L, from 22 mm to approximately 16 mm, however, the degree of bending is too severe to be able to trap atoms between the plates due to the non-uniformity of the electric field. If we reduce L to 5 mm the situation is much more favourable, with the maximum angle of deflection being 0.1 mrad. A protrusion of 5 mm, however, is not sufficient to extend beyond the centre of the magnetic trap so the support rig arms must be extended to compensate for the lack of reach. While these designs are possible to implement the rig begins to significantly impede upon the optical access required for the MOT beams.

It is also possible to increase the thickness of the plates, with a thickness of 4.8 mm being sufficient to prevent the plates from deflecting by more than 0.1 mrad for an unsupported length, L, of 22 mm. The use of thicker plates, however, significantly decreases the transmission of the Zeeman and Zeeman repump light into the slower and may consequently reduce the number of atoms reaching the MOT. Other materials have much higher Young's moduli but do not have the optical transmission properties of glass which allow us to determine the angle between the plates.

These calculations show that with the current designs it is not possible to study our atoms in the presence of very large electric fields. Some design tweaks may allow one to achieve these goals but each design is pushing both the machining capabilities of most companies and testing the tolerances imposed by the five main criteria required. The difficulties arise from two problems. The first is that the constraints imposed by the size of the chamber make it extremely difficult to place the plates in the chamber such that they are supported along their entire length and sit above the magnetic trap. The second problem is that any design for an adjustable rig will typically not be rigid enough to withstand the large forces pulling the plates together. To meet these requirements a significant redesign of the experimental chamber and electrode support rig is necessary.

Having demonstrated the ability to transport atoms over several centimetres (chapter 4) there is no foreseeable limit, other than the background-pressure induced lifetime, to transporting them much further using either a series of over-lapping anti-Helmholtz coils, or a single coil pair mounted to a translation stage. To solve the first of the above problems, the most suitable scenario is one in which the MOT and the electric field plates occupy different chambers. This has a number of advantages over the current design. Firstly, optical access to the MOT and the Zeeman slower is unrestricted, as is access to the electrode plates in the neighbouring chamber. Secondly, the chambers can be arranged such that our magnetic transport coils can be positioned outside the chamber, allowing us to operate at significantly higher currents, and thus free to choose the gradient of our trap to suit the nature of the experiment. With the possibility of making larger field gradients the tolerance on the angle between the plates is also relaxed. This design allows us to place the plates directly between the trapping coils and remove the necessity to have the plates protrude beyond their support rig. In addition, there is no restriction on the electrode plate thickness making them more stable while making it easier to bridge electrical contacts to. Lastly, by separating the chambers by some distance, the vacuum tubing bridging the two chambers will act as a differential pumping stage allowing one to achieve excellent background pressures in the second chamber and thus extended trap lifetimes.

One particularly appealing design to solve the second problem is to use small piezoelectric actuators to adjust the alignment of the plate. Using ITO-coated glass plates, small actuators can be attached to the corners of one plate while the other plate remains rigidly fixed. Before placement in the chamber the interferometric techniques discussed previously can be used to determine the plate separation and adjust their degree of parallelism to within the required tolerances. Once in the chamber imaging techniques can be used to monitor the position of the atoms, and the atoms themselves can be used to gauge the degree of parallelism between the plates, allowing for unprecedented precision. With the new chamber designs these changes are relatively simple to implement and a polarised gas may be realised in the near future.

Chapter 4

Transporting the Atoms

The aims of this thesis are to make progress towards the realisation of an ultra-cold dipolar gas. To achieve this a source of Li atoms has been loaded into a magneto-optical trap and cooled to 0.85 mK. With further cooling, these atoms may either be polarised directly by the application of an extremely large electric field, or used as a refrigerant for the sympathetic cooling of a sample of polar molecules, which may be then polarised by relatively weaker fields. To meet either aim it is necessary to load the atoms from the MOT into a second trap and to be able to efficiently transport them, either to a position between a pair of electrodes or to a molecular trap. This chapter details the experiments performed to demonstrate the principle of loading a new trap and transporting the atoms in a controlled fashion. The designs for the trap and transport mechanism are chosen with the expectation of building an experiment to polarise the atoms directly, however, the techniques may be easily applied to overlapping the atoms with a molecular trap.

In order to study the atoms under strong electric fields it is preferable to isolate them completely from their environment. In a magneto-optical trap the atoms are very strongly coupled to the radiation and the conditions are not suitable for creating an ensemble of atoms in a single hyperfine state, necessary for these types of experiment. In addition, forming a MOT between two electric field plates is difficult and unnecessary. The atoms must therefore be transported from the MOT to a second trap, centered around the electric field plates. There exist several ways to move atoms under vacuum, using either optical or magnetic forces. Each method carries its own set of advantages and disadvantages and should be chosen based on the details of the experiment in hand. Optical forces are normally employed for moving atoms short distances on the millimetre scale using optical tweezers, but have also been successfully operated over large distances [95]. The technique makes use of the dipole force to trap the atoms, by tightly focusing a high-powered, far-detuned laser on the trapping region. The focal point can then be moved to the

desired location and providing the movement is slow enough, the atoms will follow. The use of magnetic forces requires a moving trapping potential. This is achieved by either constructing a sequential array of overlapping traps, whose potentials can be modified to translate the trapping center and thus transport the atoms, or a single trap that can be physically transported.

Because of the restricted optical access into the trapping chamber the possibility of using optical forces was ruled out. The choice then became whether to use magnetic trap coils, integrated into a translation stage to move the cloud, [96, 97], or to build a series of traps, arranged in tandem and choreograph the currents such that the centre of the trapping potential moves smoothly from one trap to the next, [98]. The geometry of our set-up meant that neither method could be employed if the trapping coils were to remain outside the chamber, while building a clean, low-vibration, moving coil pair inside a vacuum chamber would be an expensive and unnecessarily difficult task. It was therefore decided the simplest and most inexpensive solution would be to construct an array of overlapping in-vacuum quadrupole traps to create a 'magnetic conveyor belt' for the atoms.

Current cloud temperatures are of the order of 1 mK. At these temperatures a significant field gradient is required to sufficiently compress the cloud such that it can be positioned between two electrodes with a 0.5 mm separation. However, the size of the fields achievable with in-vacuum coils are relatively weak. Although it is possible to position the coils reasonably close to the atom cloud, the stringent vacuum requirements mean that the amount of current that can be supplied to the coils is not sufficient to produce the desired gradients. This will be discussed further in section 4.2.2, however, it is sufficient here to simply state that without cooling and thus compressing the cloud further a significant fraction of atoms would be lost to the electrode surfaces. In addition, it is clearly beneficial to have a cold atomic cloud if one wishes to create a BEC or sympathetically cool molecules to the sub-mK regime.

The Doppler temperature for lithium is 142 μK, which if simple Doppler-theory applies (see section 4.4) is achieved in a MOT at low trapping beam intensities and a laser detuning approaching $\Gamma/2$. In most cases, while adjusting the trapping beams to meet these requirements is feasible and will significantly lower the temperature of the MOT cloud, other experimental limitations, such as heating from repump beams prevent one reaching the Doppler limit. Realistic ambitions may be to achieve temperatures between 200 and 300 μK. In order to drop below these temperatures and move beyond the Doppler temperature limit other methods must be employed. In the case of Li, as discussed in 2.1.3, polarisation gradient cooling is not a feasible option, however, it is possible to use evaporative cooling techniques

[50]. By carefully tuning the trap dynamics it is possible to expel those atoms with the greatest velocity, thereby lowering the velocity spread and thus the temperature. Although this method incurs some atom loss, the current atom numbers achievable in the MOT afford us some tolerance to losses, without hindering future, proposed experiments. This cooling technique is typically required for achieving Bose-Einstein condensation and will likely form one the steps in the experiment's future.

4.1 Magnetic Trapping of Neutral Atoms

Magnetic confinement of neutral atoms is very useful in many ways. Magnetic traps isolate the atoms from the environment, providing long lifetimes in the dark, leaving the atoms available for interaction with electromagnetic field probes, ideal for spectroscopic measurements. The magnetic confinement of atoms relies on the interaction between an inhomogeneous field with the magnetic dipole moment, $\underline{\mu}$, of the atom. This interaction provides a force for a given magnetic field gradient

$$\underline{F} = \nabla(\underline{\mu}.\underline{B}) = g_F \, M_F \, \mu_B \, \nabla|B| \; . \tag{4.1}$$

Many trapping geometries exploiting this force have been studied in the literature. The simplest of these designs is a quadrupole trap, formed from a single pair of anti-Helmholtz coils, identical to those employed to create the field gradient for the MOT. The quadrupole trap has a single minimum, at the centre between the coils. Since it is impossible to create a magnetic field maximum, only minima, the atoms must be driven in to a weak-field seeking spin-state in order to be confined. If possible the atoms are pumped into the most positive M_F state to maximise the trapping force; see figure 2.19. It is still possible to trap strong-field seeking atoms by using AC fields to create a rotating saddle point potential, providing the magnetic moment can adiabatically follow the changing field gradient. One of the simplest designs for such a trap is an AC magnetic trap based on the same dynamical principle as the Paul trap for ions and was first demonstrated by E. A. Cornell *et al.* [99]. Such traps are typically 1 to 2 orders of magnitude shallower, however, than DC traps.

Typically, for ground state atoms, the Zeeman energy-level shifts associated with the practical field sizes obtainable in the lab are small relative to $k_B T$, where $T = 1 \; K$, making the traps relatively shallow. Consequently, atoms are normally cooled to the mK regime before being confined in a magnetic trap. The small trap depth also dictates stringent vacuum requirements as even gentle collisions with the thermal energy background gas are sufficient to eject an atom from the trap. This will be discussed in more detail further on.

The absolute ground state of any atom or molecule is always strong-field seeking. For DC traps with a field minimum, such as the basic quadrupole trap, the atoms are confined in the weak-field seeking component of their electronic ground state and there is thus an opportunity for atoms to decay to the strong-field seeking component. This decay to a non-confined state is most probable at the zero-field point where the energy eigenvalues of the strong and weak-field seeking components are degenerate. For non-zero field points such transitions can still occur if the magnetic field in the atom's frame rotates with an angular frequency comparable to the Zeeman splitting between the two states. These spin-flip transitions are known as Majorana transitions and the probability of a transition occurring has a dependence on the temperature of the cloud; larger temperatures imply a greater spatial variance for an atom in the cloud and the atoms will have a smaller probability of passing through the zero-field point. For lower temperatures the atoms are more tightly confined to the centre of the trap and hence pass through the zero-field point more frequently. In most experiments involving alkali atoms the decay rate due to these spin-flip transitions is far smaller than the rate of collisions with background particles until the temperature of the cloud approaches the low μK or nK regime. Trap designs with a non-zero minimum have been developed to circumvent these problems, however, they have not been employed here because of the relatively warm temperatures of our clouds.

For our experiment the distribution of atoms amongst the eight $^2\mathrm{S}_{1/2}$ M_F states in the MOT is unknown, however, in the worse case scenario the atoms are expected to be uniformly distributed amongst the eight states. Of these states those atoms in the $|g, 2, 2\rangle$ state will be trapped with the maximum possible force for a given field gradient, while those atoms in the $|g, 1, -1\rangle$ and $|g, 2, 1\rangle$ states will be trapped with half this force and experience a trap with half the depth. Atoms in all other states will be lost. The prospect of creating a spin-polarised cloud using an on-resonant repump beam has been considered as a method to increase the atom number in the magnetic trap. This is discussed briefly in section 5.2.

4.2 Designing the Magnetic Traps

The principle behind a magnetic travellator is very simple and has been employed successfully by M. Greiner *et al.* [98]. The atoms are initially loaded from a MOT into a magnetic trap; the method for doing this is described in section 4.3. The magnetic trap is formed by two quadrupole coils in the usual anti-Helmholtz configuration. A second pair of coils is then positioned such that there is some overlap between the circumferences of the trapping coils and the centres of the traps lie

mid way between the coils in the same horizontal plane. To transfer the atoms the current is lowered in the first trap and simultaneously increased in the second. By carefully adjusting the currents during transfer the zero point of the field will move smoothly from the centre of one coil pair to the next. In principle, an indefinite number of coils can be linked together in this way and atoms can comfortably transported many tens of centimetres with almost no loss. The dimensions of our coils and the distance the atoms are to be transported are discussed below.

Because of the volume occupied by the Zeeman and MOT beams, finding a suitable location for the coils and the electrodes inside the chamber is challenging. Obviously any partial obstruction of the MOT beams will cause an imbalance in the radiation pressure of the trapping beams and the atoms cannot be trapped. It is also important to be able to image the cloud in its final position and preferably throughout transportation for characterisation purposes. It was decided to move the atoms along z (see figure 1.2 for a coordinate axes reference) thereby moving clear of the MOT beams while still allowing for optical access through the side viewports. Figure 4.1(a) gives a schematic drawing of the transport coils, illustrating the main dimensions of the support structure. The circle in the centre depicts the available optical access. Image (b) is a rendered drawing of the coil formers in their position in the chamber showing the relative positions of the MOT and Zeeman beams. Here, the electric field plates can be seen in their proposed position in the chamber. The designs for the electrode unit and their associated problems are discussed in detail in sections 3.5.1 and 3.6. The Zeeman beam must pass through the electrodes so the support rig was designed to allow most of the beam to pass unobstructed; see figure 3.17. The electrodes are kept thin to minimise the amount of Zeeman light obstructed. Many factors have been taken into account before completing the final design. In the following sections I will discuss the final apparatus and highlight each of the main potential problems and their solutions.

4.2.1 First Coil Pair

The first coil pair must sit centred around the MOT in order to capture as many of the atoms as possible during transfer from the MOT to the magnetic trap. In addition, to place the atoms between the electrode plates the cloud must be compressed to a diameter of a few hundred microns, requiring a potential with a steep field gradient. Because the coils are in vacuum we are significantly limited by the amount of current that can be applied, meaning the coils must be closely spaced to reach the desired field gradients. The inner diameter of the coil, however, must be larger than the MOT beam diameter and separated by at least the same amount, as the vertical MOT beam passes through the centres of the two coils and the horizontal

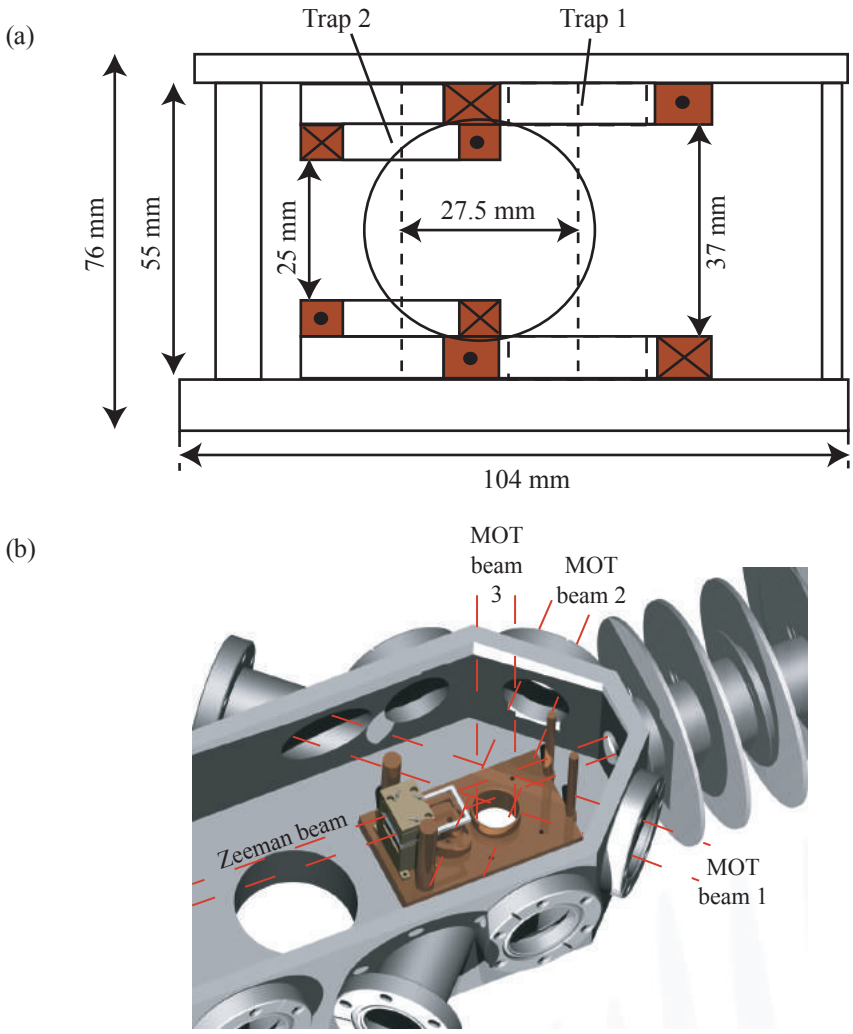

Figure 4.1: (a) A simple schematic of the magnetic traps and the supporting assembly, (b) Computer generated illustration of the MOT chamber and the magnetic trap support rig, showing the position of the traps and field plates in the chamber relative to the MOT and Zeeman beams.

MOT beams pass between them. This can be seen in figure 4.1. The choice for an outer diameter is more complicated, but in general a larger coil means a larger field and thus a steeper field gradient. The exact choice of the diameter will be discussed in section 4.2.2.

Each of the two coil formers for the first trap and their supporting structure are machined from a single piece of copper. The inner wall of the former is a 1 mm thick, hollow, copper ring whose inner diameter matches the diameter of the vertical MOT beam, so as to maintain optical access for the MOT. The former has only one supporting face, formed by the base (/top) plate of the supporting structure for the lower (/upper) coils. This maximimises the heat conduction from the coils to the chamber floor. The dimensions of the lower and upper plates are $104 \times 60 \times 10.5$ mm and $100 \times 60 \times 4.5$ mm, respectively. The wire for each coil has a diameter of 0.9 mm

in addition to a $30 - 40\ \mu$m thick polyester enamel used for insulation, suitable to temperatures up to 200°C. The dimensions of the square cross-section wire used for the Zeeman slower were too large to conform with the small size of the coil and could not be used. To hold the coil together under vacuum the outer layer was initially potted in an electrically-insulating, vacuum-compatible epoxy glue. The make of glue used in these experiments was H77, supplied by Epotek. It is cured at $100 - 150$°C for approximately 1 hour.

4.2.2 Second Coil Pair

The former for the second coil pair is a two-tier design, machined from a single copper piece; figure 4.2.

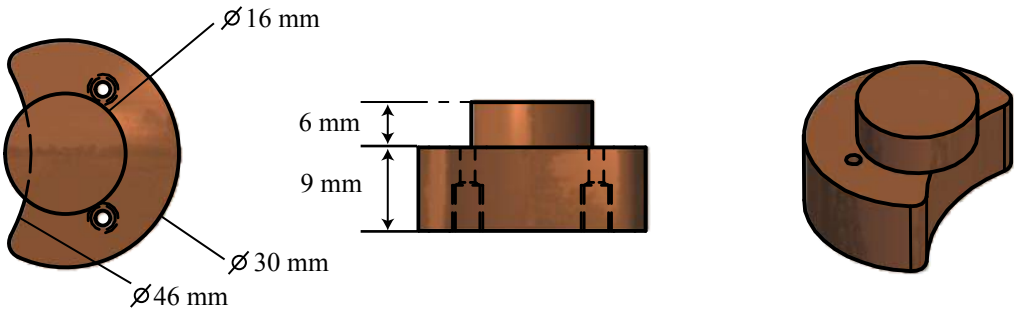

Figure 4.2: Diagrams illustrating the shape and dimensions of the coil former used for the second trap.

The lower tier has a crescent shape machined from it, cutting up to 3.5 mm underneath the top tier. The radius of the lower tier matches the outer diameter of the lower coil and forms one supporting face of the former. The top tier forms the inner wall of the former. The height of the bottom tier matches the thickness of the first coil such that the second coil pair sits between the first pair and the coil circumferences overlap. Figures 4.1(a) and 4.3 illustrate how the coils are positioned. The radius of the crescent machined from the first tier is 46 mm to accommodate the outer diameter of the coils forming the first trap; figure 4.1(b). The formers are sunk into a 2 mm recess cut into each of the copper plates forming the larger former. The recesses are cut to form a tight fit around the former to improve heat conduction away from the coils. Screws are inserted through counterbored holes in the reverse of each plate to secure the formers to the copper base. The coils are wound with the same wire as the first pair and cemented in the same fashion using the epoxy.

Optical access for imaging is through the two DN40CF viewports that flank the MOT cloud; figure 4.1. The MOT sits at the centre of the first magnetic trap. Unfortunately, for large MOT clouds this position means the entire cloud cannot

be imaged through this viewport and the usual vertical imaging beam must still be used. To maintain a view of the cloud in the second trap the atoms should not be moved beyond the line of sight provided by these windows. The diameter of each window is approximately 35 mm so the distance between the centres of the two traps should be smaller than this. The atoms must also move a minimum distance beyond the MOT beams in order to sit between the electric field plates without the plates obstructing the laser light. To meet these constraints the separation between the trap centres was chosen to be 27.5 mm, as depicted in figure 4.1(a).

The outer diameter of the second coil pair can be no larger than ∼30 mm otherwise it will overlap with the vertical MOT beam. The diameter of the crescent recessed from the lower tier of the second coil former, and thus the outer diameter of the first coil, is limited by the amount of material that can be removed from the lower tier of the second former. If too much material is removed from this tier, the second coil may not have sufficient integrity to stay together once wound. It was estimated that the supporting face of the coil former should be present for at least 50% of the surface area, so a diameter of 45 mm was chosen for the first coil. The crescent has a diameter of 46 mm to allow some leeway during construction.

Since compression of the cloud is most crucial in the second trap the second coil pair was placed to sit within the coils of the first trap to maximise the field gradient about the centre of the second trap. The separation between the upper and lower coil of the second trap should be as small as possible, but is restricted by the diameter of the MOT beams. The heights of the coils, as with everything, are limited by the constraints imposed by the chamber size and the required optical access. The distance from the roof to the floor of the chamber is 76 mm. The separation of the second coil pair is defined by the MOT beam diameter and thus must be greater than or equal to 25 mm. Equally, the total structure can be no taller than 70 mm, to allow 6 mm of clearance to the chamber roof for assembly. The coils and their formers must subsequently fit into the remaining space. The heights for the coils were chosen to be 9 mm and 6 mm for the first and second coil pair, respectively. This ratio was chosen to provide a steep enough gradient in the second trap to compress the atoms sufficiently enough to fit in the 0.5 mm gap between the field plates, for a cloud temperature of 100 μK: The volume occupied by the atoms in the trap is bounded by the walls of the field plates. The depth of the trap is thus limited to $U = \mu_B A\, d/2$ where A is the field gradient, normal to the plate surface, and d is the distance between the plates that can be safely occupied by the cloud without under going collisions with the plates. If one takes d to be $\sim 500\ \mu$m and given a field gradient of $\sim 65\ \mathrm{G\,cm^{-1}}$ the trap depth in Kelvin is $U/k_B \approx 110\ \mu K$. If the lithium cloud can be cooled to the Doppler temperature, $T_D = 142\ \mu K$, then

a significant fraction of the atoms in the MOT can be captured, transported and polarised.

The coils were initially tested separately outside the vacuum chamber to ascertain their current-temperature relationship. A temperature limit of 120 °C was imposed on the coils because it is known that at higher temperatures outgassing of the polyester coating increases the pressure in the vacuum chamber above 10^{-10} mbar. We then measured the effects of convection and conduction cooling processes at atmosphere in the lab, by running two tests, one with the coils heat sunk to the optical table and the second with the coils thermally isolated. By comparing the rate of heating with and without the effects of conduction cooling present we were able to estimate the effect of cooling due to convection and thus extrapolate the corresponding current limits in a vacuum environment. In practice this led to a substantial overestimation of the maximum currents sustainable in vacuum. For the first coil pair it was estimated that the coils could withstand currents up to 7 Amps, however, under vacuum conditions at 4.5 Amps outgassing at the 10^{-8} mbar level was already observed. During these in-vacuum tests it was initially unclear whether the epoxy glue, the insulating coating or both were responsible for the observed level of out-gassing recorded. To investigate, a pair of identical coil formers were built that could maintain the coil integrity without the use of any epoxy. The design is almost identical to the original, however, in addition a 1 mm aluminium plate was screwed to each of the coil formers to provide a second supporting face for the copper formers. The plates provide enough compression to the coil to hold it in place. The wires from each coil were tightly wrapped around the steel screws securing the aluminium plates in place to maintain the necessary tension in the coils. Some of the screws, however, bit into the polyester enamel and caused electrical grounding issues that had to be addressed in the electronics at a later date. Once the epoxy coated coils were removed from the experimental chamber and dismantled, it became clear that the polyester coating had been severely damaged. This suggested the coils had become significantly hotter than anticipated, destroying the integrity of the enamel and outgassing copious quantities of CO_2. This suggests that the epoxy was not responsible for the outgassing and could be used in later designs.

After the lessons learnt from the first coil set, the new coils were tested more carefully in a separate vacuum chamber to determine a more accurate estimation of the current limits for the coils. During winding, a thermocouple was embedded in the centre of each coil. The four coils were then secured in a test chamber and each pair was tested separately. The coils were mounted carefully to try and mimic the degree of heat sinking the coils would experience inside the MOT chamber. The test chamber was baked and pumped down to a pressure of 1×10^{-9} mbar by a 70 ls^{-1}

turbo pump, similar to the pumping speed expected from our ion pump. The coils and thermocouples were connected to the outside via an electrical feed-through and the temperature and pressure were monitored as a function of current. Choosing a safe limit of 120 °C the maximum current allowed for each coil pair was determined. The pressure readings were found to raise by no more than 3×10^{-10} mbar in testing. The results are summarised in table 4.1, including data on the coil geometries and field gradients.

Table 4.1: Coil data for the magnetic travellator

Trap	Coil i.d. (mm)	Coil o.d. (mm)	Coil height (mm)	Current at 120 °C (A)	Corresponding field gradient ($\mathrm{G\,cm^{-1}}$)
1	27	45	9	3.75	40.2
2	16	30	6	5.35	64.4

Once in the main chamber the current for each coil was supplied through an electrical feed-through. Each coil was operated by a dedicated supply to give full control over the position of the zero-field point of the trap along the coil axis. The thermocouples were removed before the assembly was placed in the MOT chamber. To date, operating at these current limits has not significantly impinged on the vacuum quality, supporting the above measurements.

A photograph of the assembly is given in figure 4.3. The bottom two coils have

Figure 4.3: A photograph of the full trap assembly showing both magnetic traps and their supporting structure. The bottom two coils have been wound and cemented in epoxy.

been wound and cemented in epoxy. These coils were later replaced by a second

set without the epoxy, but held down by aluminium plates. The electrode plates are not shown in this image. Instructions for the installation and assembly of the support rig can be found in appendix A. The appendix includes details about the various components forming the support frame and the order in which the rig must be assembled inside the chamber. Detailed drawings of the assembly and its position in the MOT chamber can also be found.

4.3 Realising the Magnetic Trap

With the trapping coils in place we proceeded to make a MOT to try and recreate the atom number and temperature documented in sections 2.7 and 2.8. Initial measurements of the atom number and temperature measurements were repeated for atoms in the MOT using the usual outer-vacuum coils. To simplify nomenclature any MOT formed using only the outer-vacuum coils will be referred to as an outer-MOT while any formed using the in-vacuum coils will be referred to as an inner-MOT. To create an outer-MOT the same parameters for the laser intensity and detuning used to make the MOTs described in section 2.6 were used. These detunings were then optimised, but there were no significant changes between the original and final atom number. An increase in atom number was recorded by using a steeper field gradient of ~ 12 G cm^{-1}, however. Under these conditions the total atom number was measured by absorption imaging to be $\sim 8 \times 10^8$ atoms, some 10 times smaller than previous MOTs. The cloud has an approximate volume of 0.014 cm^3 suggesting a density of 5.7×10^{10} atoms cm^{-3}, similar to earlier results. The lower atom number was not expected as the trapping parameters have not changed. It is thought that the local pressure near the MOT may now be higher due to outgassing from the magnetic trap installation. The larger required field gradient is not understood; it is possible that the Zeeman slower was not performing at its optimum during these experiments and that the exit velocity was slightly larger that the intended speed of 50 ms^{-1}. Thus, by increasing the field gradient one may serendipitously increase the capture velocity of the MOT, and hence improve the atom number.

Temperature measurements indicated that the MOT was significantly warmer than previously. For MOT beams with 4 linewidths detuning we find a temperature of 1.80 ± 0.12 mK using ballistic expansion measurements. This is more than twice the temperature measured in the MOTs created prior to the installation of the magnetic traps. The reasons for this are unclear but may be a consequence of the change in field gradient. In an attempt to cool the MOT the detuning for the MOT beams was decreased to 3 linewidths. Although poor quality images lead to large errors in the measurements we did not observe a detuning dependence

as predicted by the empirical laws discussed in section 2.8. A significant amount of work is still required in order to fully explore and understand the cloud's temperature dependence. Recent measurements performed at the time of writing suggest there is a significant heating effect from the MOT repump light and that lower intensities and greater detunings may result in considerably lower temperatures. Temperatures of 500 μK have been observed for repump intensities of 2.5 mW, in line with the results predicted by Doppler theory (equation 2.49).

In spite of the higher temperature we attempted to load a magnetic trap. The simplest method to load a magnetic trap is to load a MOT and switch off the trapping light, retaining those atoms with the correct spin states at the time of release in the magnetic trap. To capture a significant number of atoms from the MOT and maintain a reasonable atom density it is beneficial to form a trapping potential with a large depth and a steep field gradient. The size of the outer MOT coils means the trap is sufficiently deep, however, the small field gradient means the cloud is spread over several centimetres. These low atom densities make imaging difficult and such a cloud size not amenable for transportation. To create field gradients above 15 G cm^{-1} much larger currents are required, however, ultimately we require the atoms to reside in the in-vacuum trap in order to be able to transport them. The inner trap is capable of providing up to 40 G cm^{-1}, although trap depths are significantly shallower. It was clear from initial attempts to create a MOT using the in-vacuum coils that there was a poor overlap between the position of the cloud and the centre of the magnetic trap. This was not due to any design failures but rather due to the alignment of the MOT beams. To move the MOT a bias coil was wound around the chamber down stream for the MOT. The coil is wound with a 2 mm diameter wire, with 8 turns carrying a current of 10 Amps. The alignment of the MOT beams was adjusted to optimise the atom number in the new position. In the final position the optimised atom number was approximately 4×10^8.

Once the cloud was in position the outer coils were switched off and the in-vacuum coils were used to make an inner-MOT. In preparation for loading a magnetic trap the field gradient was increased to 18 G cm^{-1}, forming a 0.85 mK deep magnetic trap. 80-90% fewer atoms were observed in the inner-MOT even with a 12 G cm^{-1} field gradient to the outer-MOT, suggesting the atom number has a strong dependence on the field profile between the trap and slower. As the field gradient is increased to 18 G cm^{-1} and beyond the atom number is reduced further. The smaller geometry of the inner coils creates a significantly different field profile in the region between the slower and the MOT. The final four Zeeman coil currents were retuned to optimise the atom number. This successfully increased the atom number by a factor of four. To date a complete calculation of the magnetic field

shape required to merge the Zeeman field and new trapping fields optimally has not been performed and a better configuration may be found than the profile currently used. It is clear from our observations, however, that the new slower profile increases the atom number significantly.

The largest atom number recorded in the inner-MOT was 5.4×10^7. Due to the high field gradient the cloud is compressed and the atom number density remains high despite the reduced atom number. The cloud is therefore still optically thick and consequently fluorescence measurements will tend to provide an underestimate of the true value. Reliable absorption measurements have not been performed to date due to poor quality imaging apparatus resulting in small signal to noise ratios.

The magnetic trap is loaded from a MOT by pulsing off the light for the MOT beams, the MOT repump beams and the Zeeman slowing beams, for a period δt, without modifying the magnetic field profile. When the laser beams are turned back on, the atoms that remain are recaptured into the MOT and a fluorescence image of the atoms is recorded. The number of re-captured atoms as a function of δt determines the lifetime in the magnetic trap. Two additional experiments were also performed. The first of these is a repeat of the above experiment with the magnetic field gradient set to 40 $\mathrm{G\,cm^{-1}}$ for the MOT and magnetic trap. The increased field gradient in the MOT means the initial atom number is about a factor of two smaller and the cloud is compressed to a volume of 9×10^{-4} $\mathrm{cm^3}$, equating to a density of $\sim 3\times10^{10}$ $\mathrm{atoms\,cm^{-3}}$. The second experiment measures the free expansion of an 18 $\mathrm{G\,cm^{-1}}$ MOT by switching off the magnetic trap. Although inductance in the coil prevents the field switching off extremely rapidly, the magnetic field was measured to have fallen by 90% after 1.6 ms, sufficiently quick for the experiments described here. This data is used as a control experiment to determine the decay constant for a cloud expanding under no external forces. This provides a useful comparison to the predicted decay rates for strong-field seeking atoms actively expelled by the magnetic field gradient discussed further on. The results for the number of atoms as a function of release time are given in figure 4.4 for all three experiments. For the 40 $\mathrm{G\,cm^{-1}}$ data a longer camera exposure was used to offset the smaller atom number, and this allowed us to observe the trap lifetime for up to 30 ms.

To produce the lifetime data each image from the camera is integrated to compress a 2D array of single-value pixels into a 1 dimensional list. A Gaussian function is then fitted to this data, keeping the amplitude, mean, variance and background as free parameters. The error bars for the data depict the usual fitting errors associated with the ability to match a Gaussian profile with the MOT cross section. The amplitudes are then plotted as a function of release time and an exponential decay is used to model the decay of the cloud. The amplitude and decay constant

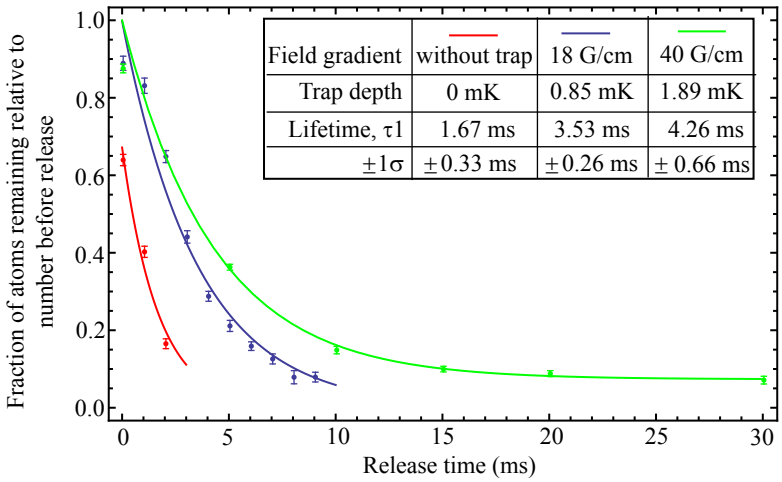

Figure 4.4: A graph depicting the number of atoms remaining in the magnetic trap once released from the MOT as a function of the time since release. The table summarises the decay time constants for traps with different depths and field gradients. The red and blue curves are fits to a single exponential decay, the green curve is a fit to a double exponential decay.

are free parameters and the fit is weighted using the individual error bars. The 1/e lifetimes for a free expanding cloud and for a cloud in a trap with a field gradient of 18 G cm^{-1} are $\tau = 1.67 \pm 0.33$ ms and $\tau = 3.53 \pm 0.26$ ms respectively. It was evident from the match between the results of the fitting function and the 40 G cm^{-1} data that a double exponential function of the form A Exp[-t/τ1]+ (1-A) Exp[-t/τ2] fits this data much more accurately. The second decay constant, τ2, could not be determined accurately, however, as it is significantly longer than the 30 ms period over which the data was taken. An explanation for the double exponential decay will be discussed further on in this section.

The sub-10 ms decay times measured above are the result of two process occurring at the time of release. The first of these is the expulsion of those atoms trapped in spin states not confined by the trap. An atom in the $|g, 2, -2\rangle$ state is expelled with an acceleration of 320 ms^{-2} for a trap gradient of 40 G cm^{-1}. Given the recapture radius of 6 mm, determined in section 2.8, the estimated time taken for such an atom to be expelled from the centre of the trap is 2.1 ms. This is comparable to the free expansion lifetime. Although faster expulsion times may be expected the above estimation is derived from a very simplified calculation and more detailed simulations may be required. Those atoms in more weakly coupled M_F states will experience a smaller acceleration and take longer to exit the recapture volume.

The second decay process is due to the loss of atoms whose kinetic energy is greater than the trap depth at the time of release. To judge the loss rate for these atoms in a trap of a given field gradient we can the find the average time an atom in a given confined state will take to leave as a function of velocity, if its kinetic

energy is greater than the trap depth. The average time is given by the integral of the product of the Maxwell-Boltzmann distribution and the function, derived from the equations of motion, describing the time taken for an atom to leave the trap for a given initial velocity. The integral is performed for velocities corresponding to a kinetic energy greater than the trap depth. For atoms in the $|g, 2, 2\rangle$ state beginning in the centre of the trap the average time is approximately 3 ms. There will be a distribution about this value for atoms starting elsewhere in the trap and for those in more weakly confined M_F states.

The decay times estimated above match well with the short time constants measured for the initial loss rate of atoms from the trap. We can thus explain the two stage decay observed in the green curve in figure 4.4. There is an initial rapid loss of atoms with a decay constant of several milliseconds due to those atoms whose kinetic energy is too great to be contained or whose spin state is not a weak-field seeking state. There is then a much longer decay time for those atoms which are slowly removed from the trap by collisions with background particles. This will be discussed in detail in section 4.5.

Currently our imaging techniques are limited by mechanical instabilities between the table and the experiment that move our imaging beam, leading to a poor signal to noise ratio. Consequently, lifetimes beyond 30 ms cannot be recorded. To measure these longer lifetimes we need to either improve the atom number that survive the first decay stage or improve our imaging.

4.4 Lowering the Cloud Temperature

To increase the number of atoms loaded into the magnetic trap the cloud may be cooled. The steady-state kinetic energy of atoms in a MOT is given by equating the heating and cooling rates. The cooling rate, $C = F.v$, is determined by the optical molasses, where F is given by equation 2.4. The heating rate is determined by the momentum transfer imparted from the light field to the atoms as they undergo spontaneous emission in a random direction. The atoms are heated by the recoil experienced after each absorption and emission of a photon of momentum, $\hbar k$. The associated recoil energy is given by $E = \hbar^2 k^2 / 2M = \hbar\omega_{rec}$. Thus, for each scattering event an energy of $2\hbar\omega_{rec}$ is lost from the light field. This occurs at a rate $2R_s$, (see equation 2.34), since the atoms interact with only 2 beams on average. The total cooling and heating rates, C and H, are thus given by

$$C = \frac{8\hbar k^2 \delta s v^2}{\gamma(1+s+4\frac{\delta^2}{\gamma^2})^2}$$
$$H = 4\hbar\,\omega_{rec}\,R_s\ . \tag{4.2}$$

To find the equilibrium energy one can equate the above equations, substituting ω_{rec} for $\hbar k^2/2M$ and rewriting v^2 as $2\,E_{eq}/M$. By rearranging for E_{eq} one arrives at

$$E_{eq} = \frac{h\gamma^2(1+s+4\frac{\delta^2}{\gamma^2})}{16\delta}\ . \tag{4.3}$$

Equating this to the thermal energy, $\frac{1}{2}k_BT$, and arranging for T we find the temperature as a function of the detuning for a 1D optical molasses (equation 2.49). The minimum temperature is found at the minimum in this curve when $\mathrm{d}T/\mathrm{d}\delta$=0. The minimum temperature and the detuning corresponding to this temperature are given by

$$\delta_{min} = -\frac{\gamma}{2}\sqrt{1+s}\ ,$$
$$T_{min} = \frac{h\gamma}{2k_B}\sqrt{1+s}\ . \tag{4.4}$$

For low intensities these equations reproduce the well known Doppler temperature at a detuning of half a linewidth. Thus, by ramping the detuning of the MOT beams towards half a linewidth in the red and reducing the intensity of the beams, a cooler MOT phase can be reached. This is a common method employed to cool a MOT and will typically lead to a temperature reduction of 50%, e.g. [100]. There are clearly a number of parameters to be explored in order to optimise this cooling stage, including the rate of change of intensity and detuning, their final values and the length of cooling period at these values. The experiments below look at the effects of decreasing the detuning of the MOT to $\Gamma/2$ and reducing the intensity to 20% for varying lengths of time prior to release. To date only the length of the cooling period has been explored. The effects on atom number, trap lifetime and cloud temperature as a function of cooling time are discussed below.

Figure 4.5 shows the number of atoms remaining after a release time of 10 ms for various lengths of cooling time performed immediately prior to release. The length of the cooling period was increased in steps of 2 ms up to 10 ms. The field gradient is kept at 40 $\mathrm{G\,cm^{-1}}$ with a trap depth of 1.9 mK. To change the intensity and detuning of the light a pulse was delivered from the computer to the AOM drivers with a very small ramping period of the order $10-100$ ns. The initial values for the intensity and detuning were restored before recapture took place.

Without cooling only $4-5$% of the atoms remain after 10 ms in the magnetic trap. With only a 2 ms period of cooling this number is almost doubled. As the

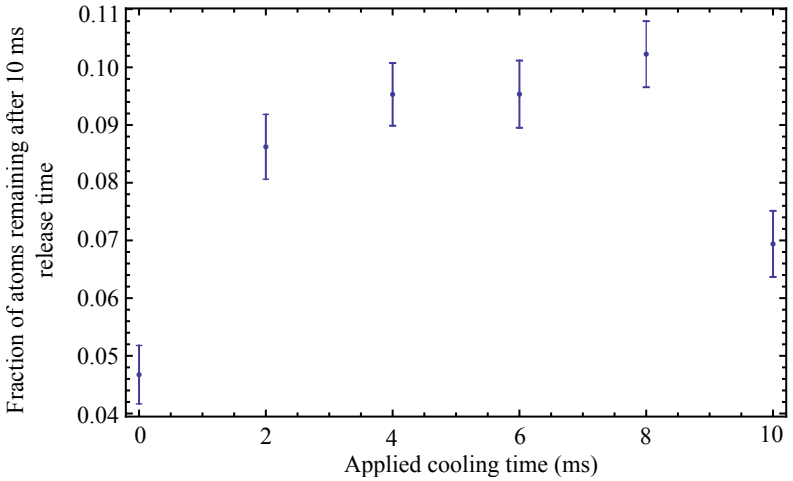

Figure 4.5: A graph presenting the number of atoms remaining in the magnetic trap, 10 ms since release from the MOT for various lengths of cooling pulse applied prior to release.

length of the cooling period increases towards 8 ms there is no further increase in the number of captured atoms. There is a drop in atom number for a period of 10 ms, however, evidence collected from the measurements below suggest this is an anomaly and the effects at 10 ms are comparable to 8 ms of cooling. This improvement in atom number suggests that the cooling mechanism works as expected, cooling the cloud and thus increasing the number of atoms with an energy less than trap depth. To fully understand the data it is useful to measure the lifetime of the cloud released into the magnetic trap as a function of cooling time. Figure 4.6 portrays the trap lifetimes for cooling periods of 0, 4, 8 and 12 ms.

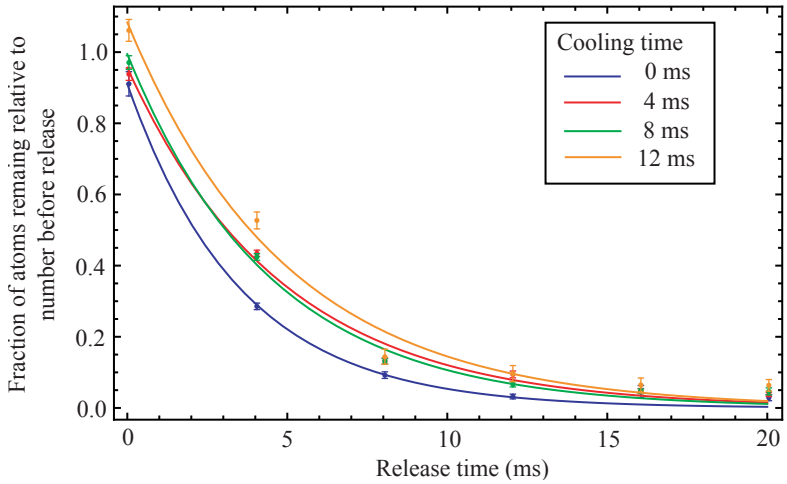

Figure 4.6: A graph depicting the decay curves for atoms in a magnetic trap as a function of time since release for different lengths of cooling pulse applied prior to release.

Applying the cooling for 4 ms increases the lifetime from 3.5±0.2 ms to 4.8±0.4 ms, with no further change for longer cooling times. The difference in time constants between cooling and no cooling reflect the factor of two increase in atom number

observed after 10 ms of release, as shown in figure 4.5. It also supports the hypothesis that the initial fast decay of atoms from the trap is, in part, due to the higher energy atoms escaping the trap.

To measure the exact temperature changes as a result of cooling, ballistic expansion measurements were performed on the MOT cloud with and without cooling prior to release. Unlike in the ballistic expansion results taken in section 2.8.2 the density distribution of the atoms is measured by turning the MOT beams back on, while leaving the magnetic field off, and imaging the resulting fluorescence onto a CCD camera. The application of the 3D molasses effectively freezes the motion of the cloud allowing accurate measurements of the temperature to be taken despite a long camera exposure time relative to the release time. The analysis of the data is performed in an identical manner to the experiments discussed for the earlier ballistic expansion results. To calibrate the pixel size the camera is repositioned to view an image of known size, positioned at the focal plane, such that the distance from the object to the lens is the same as the MOT to lens when the camera is imaging. The object we used was a pair of calipers adjusted to a separation of 1±0.005 mm. By measuring the number of pixels corresponding to 1 mm an accurate measure of the width of the cloud can be deduced. To determine the width of the cloud as a function of time a region of interest is selected from each of the 2D arrays collected by the CCD. This array is integrated over to give a 1D list and a Gaussian profile of the form described by equation 2.45 is fitted to the resulting profile. The rate of expansion can then be calculated and a temperature inferred. Figure 4.7 provides temperature data for a MOT without cooling (in blue) and draws comparisons to both 4 ms (red) and 8 ms (green) of cooling, immediately prior to release.

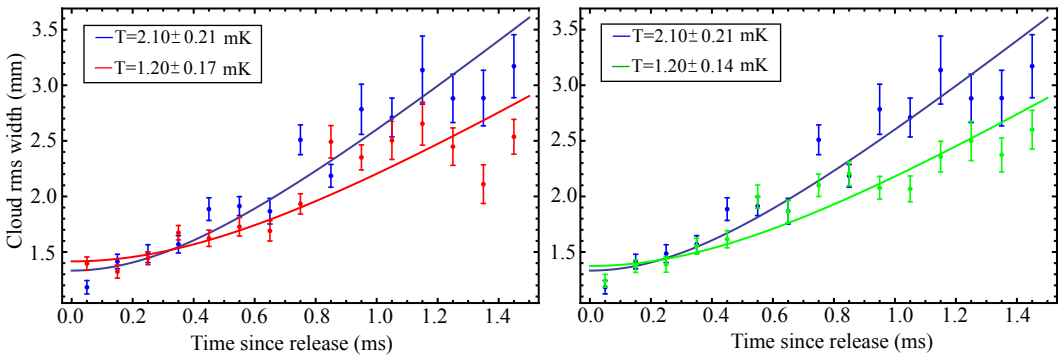

Figure 4.7: Ballistic expansion data used to determine the temperature of the MOT. The graph shows the Gaussian width of the cloud as a function of time since release. The left hand figure compares the expansion rates for a MOT cooled for 4 ms prior to release to an uncooled cloud. The right hand graph compares a cloud cooled for 8 ms with the uncooled cloud.

While there is no apparent benefit in cooling for 8 ms over 4 ms, in congruence with the previous data, there is over a 43% decrease in the cloud temperature when

the cloud is cooled for 4 ms compared to an unaltered cloud. This drop in temperature agrees with the results quoted in [100]. Comparing the Maxwell-Boltzmann distributions at the two temperatures above and integrating up to the maximum velocity an atom can have in a trap of depth 1.9 mK, a factor of two increase in atom number is expected when loading from a MOT at 2.1 compared to 1.2 mK. These results further reinforce the measurements shown in figures 4.5 and 4.6.

4.5 Trap Lifetime Measurements

Given the above improvements in atom number it is now possible to measure the atom number over a longer period to try to determine the longer time constants corresponding to losses due to background collisions. To measure these decay times the lifetime of the trap was observed for up to 300 ms. Beyond 300 ms the signal to noise ratio was too small to be able to accurately determine the fraction of atoms remaining in the trap. The data is given in figure 4.8 and was recorded at a trap gradient of 35 $\mathrm{G\,cm^{-1}}$, equating to a trap depth of 1.65 mK.

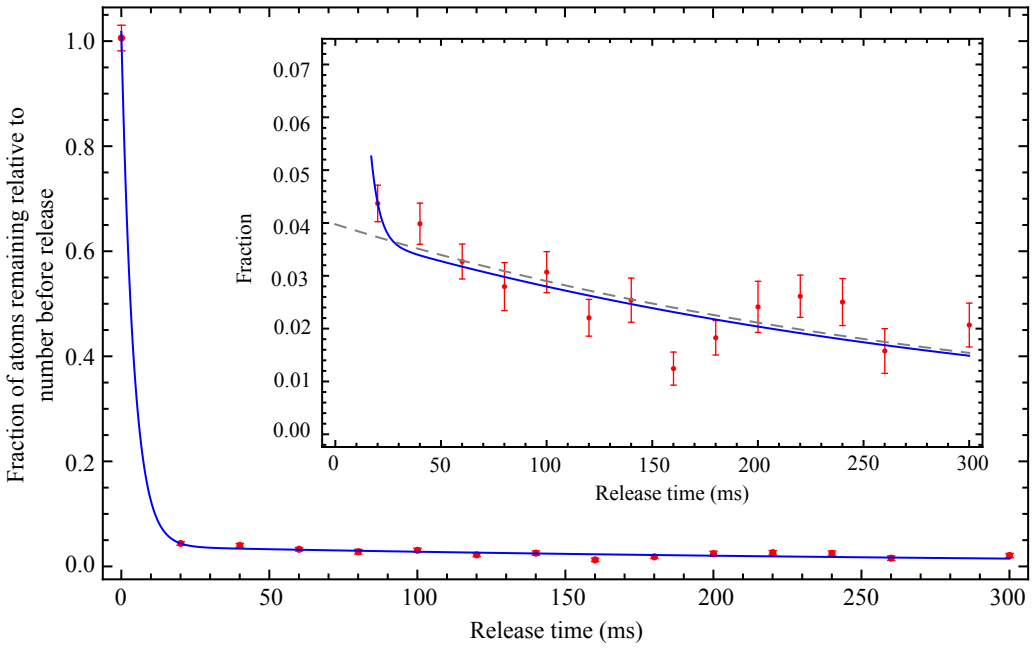

Figure 4.8: A graph presenting the number of atoms remaining in the magnetic trap as a function of time since being released from the MOT. A double-exponential decay is fitted to the data to reflect the physical model describing the results. The inset highlights the data showing the slow decay rate for atoms removed from the trap by background collisions. The grey curve depicts the function describing the slow decay only.

The data in figure 4.8 is fitted to a double exponential decay. In accordance with expectations there is a fast decay corresponding to the processes described earlier and a much longer lifetime in agreement with the loss rate due to background collisions. The 1/e lifetimes for the two decay processes are found to be 4.14±0.76 ms and

317±98 ms respectively. The large uncertainty in the second result is a consequence of the short period over which the data was recorded relative to the lifetime itself.

The atom density in the magnetic trap is too low for three body collision rates to contribute significantly to the decay time. Thus, the vast majority of atom-particle interactions will be due to collisions with the background gas. Since the trap is shallow relative to the thermal energy of the background gas even gentle collisions will propel an atom from the trap. To calculate the loss rate due to background gas it is assumed that every collision with a trapped atom will carry enough momentum to displace the atom from the trap. The 1/e decay time, τ, for atoms in the trap due to background collisions is dependent upon the background particle density, n, the mean velocity, v, and the Li-background scattering cross section, σ_{Li}

$$\frac{1}{\tau} = n\, v\, \sigma_{Li} = \sqrt{\frac{8}{m\pi k_B T}}\, P\, \sigma_{Li}\ , \tag{4.5}$$

where P and T are the pressure and temperature of the background gas. The exact composition of the background gas is unknown, however, the cross-sections between Li and the most likely candidates are all very similar [101]. The cross-section is estimated from an average of values found in [101] to be $\sigma_{Li} \approx 2 \times 10^{-18}$ m^2. While pressure readings suggest the chamber is about $1{\times}10^{-9}$ mbar, out-gassing from the coils in the chamber may increase the local pressure, since they are by far the dominant source of background gas in the chamber. A pressure of $1{\times}10^{-9}$ mbar corresponds to a lifetime of 20 s. To obtain a lifetime of 300 ms a background pressure of $\sim 6 \times 10^{-8}$ mbar is required; this is 2 orders of magnitude higher than the pressure recorded by the ion pump. However, the close proximity of the atom cloud to the outgassing coils and the coils' high temperature (∼100°C) means such local pressures are not unreasonable.

We can compare the data in figure 4.8 to some theoretical expectations. By extrapolating the curve corresponding to the long lifetime decay in figure 4.8 (grey line) back to the time of release before background collision losses take effect we find that 4% of atoms are loaded into the magnetic trap. This is congruent with the expected losses incurred due to atoms being too fast or in the wrong spin state to be trapped: We shall assume that the initial state distribution of atoms in the MOT at the time of release is equally spread over the eight ground state M_F sub-levels. Those atoms in the $|g, 2, 2\rangle$ state will see the full trap depth of 1.65 mK. Thus, beginning with a cloud with a Boltzmann velocity distribution at a temperature of 1.8 mK we would expect approximately 18% of these atoms to have a kinetic energy smaller than the trap depth. For those atoms in the $|g, 1, -1\rangle$ and $|g, 2, 1\rangle$ states the atoms only experience a trap with half the field gradient and thus half the depth. If we select from the same Boltzmann distribution those atoms with a kinetic energy less

than half the trap depth we expect only 7% of these atoms to remain. Those atoms in strong-field seeking states or with $M_F = 0$ will not be trapped. In total then, of all the M_F states one would expect only 4% of the atoms to have a kinetic energy small enough to be trapped and reside in the correct spin state. These calculations match well with our data and imply the state distribution of the atoms in the MOT is spread evenly amongst the ground state M_F sublevels.

4.6 A Deeper Trap

To improve the fraction of atoms remaining in the trap further several options can be explored. The first option is to try to skew the distribution of atoms amongst the M_F states towards higher M_F values using a circularly polarised repump beam immediately prior to release. We may also try to further decrease the temperature of the MOT cloud, or equivalently increase the depth of the magnetic trap. To demonstrate the improvements of a deeper trap the outer MOT coils were used in conjunction with the inner coils to increase both the field gradient and the trap depth. Figure 4.9 presents a plot of the total absolute magnetic field profile along the axis of the Zeeman slower, including the Zeeman coils, the compensation coils and the in-vacuum magnetic trap with a 35 $\mathrm{G\,cm^{-1}}$ gradient, with and without the outer MOT coils present; the atoms are trapped at 0.62 m from the start of the slower. The second trap potential formed at 0.66 m (blue curve) is due to the combined fields from the compensation coil wound around the MOT chamber and the magnetic trap and is not intended to be used as a trap in itself.

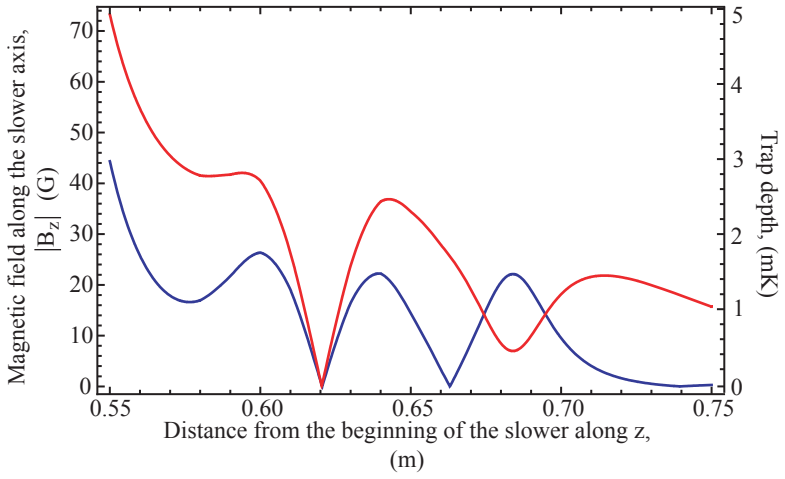

Figure 4.9: A graph showing the absolute magnetic field profile along the axis of the slower, including field effects from Zeeman slower, magnetic trap 1 and the compensation coils. The red curve shows how the magnetic field is modified if the field from the outer-MOT coils is included. The equivalent trap depth is also given. The trap centre is found at 0.62 m.

The outer MOT coils are operated at 20 Amps, corresponding to an axial gradi-

ent of 15 G cm^{-1} and a trap depth of 2 mK. By adding these coils to the in-vacuum trap the total axial gradient and trap depth is increased to 54 G cm^{-1} and 2.48 mK, respectively. Figure 4.10 shows the decay lifetime in this deeper trap between release times of 50 and 1000 ms in intervals of 50 ms. The atoms are cooled for 4 ms prior to release by reducing the intensity of the MOT beams to 20% and the detuning to -Γ/2.

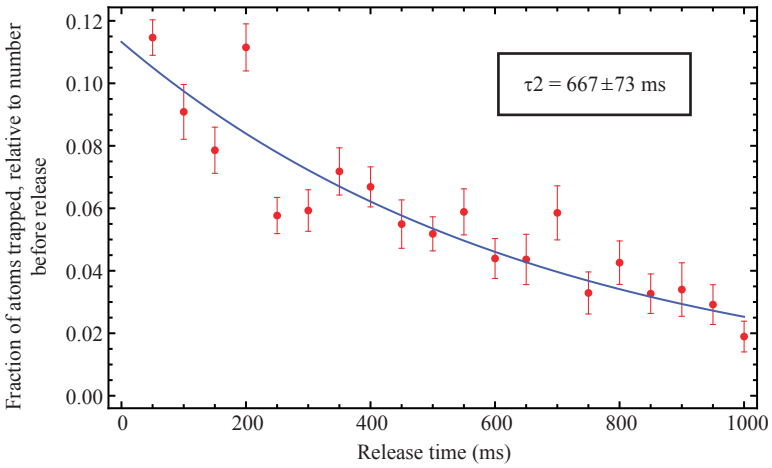

Figure 4.10: A graph depicting the increase in atom number collected by the magnetic trap for an increased trap depth. Only the data for the long term decay are shown. The longer trap lifetime reflects the improved pressure conditions in the chamber.

Figure 4.10 clearly shows a significant increase in the atom number in comparison to the data presented in figure 4.8. The initial fast decay of atoms from the trap is not shown in the above figure as no data points were taken in the first 50 ms. By extrapolating the decay curve in the above data back to the time of release we now find that approximately 11% of the atoms are loaded into the magnetic trap. The initial MOT cloud temperature is estimated to be 1.2 mK based on the results given in figure 4.7. If we again begin with a uniform distribution of atoms amongst the M_F states we find the expected fraction of atoms remaining to be 10.7%, consistent with the above data.

It should be noted that the lifetime of the trap measured in figure 4.10 is significantly longer than previous results. After several weeks of operation the level of outgassing from the in-vacuum coils decreased and accordingly an increase in the trap lifetimes was recorded. This data was collected some time after the data presented in figure 4.8 and thus was performed under lower pressure conditions. Recent data shows we now find trap lifetimes exceeding 1 second, corresponding to a local pressure near the cloud of 2×10^{-8} mbar.

Even prior to these improvements in trap lifetimes, lifetimes of 300 ms were already thought to be adequate to test the transportation of atoms into the second

trap. The following section will discuss the results recorded for transporting to the second trap, holding them there for a length of time and transporting them back where the fluorescence of the remaining atoms is measured.

4.7 Transferring Between the Traps

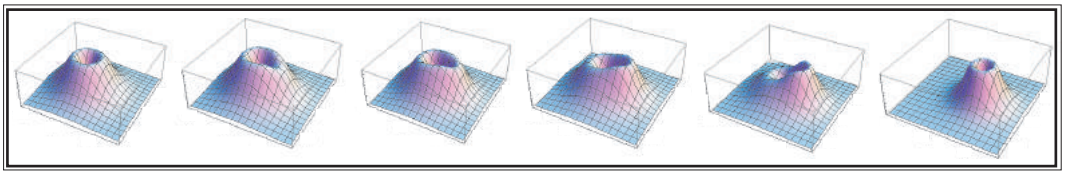

Figure 4.11: A magnetic trapping potential as it is transferred between the centres of two quadrupole traps.

Although the electrode plates are still to be implemented it is valuable to show that the position of the atomic cloud can be accurately controlled through the manipulation of magnetic fields. The set-up outlined in section 4.2 describes a pair of magnetic quadrupole traps in the anti-Helmholtz configuration. The second trap is positioned such that the coils overlap with those forming the first trap. The centre of the first trap overlaps with the centre of the MOT and the centre of the second trap is positioned in the same plane, displaced 27.5 mm along the z-axis. By modifying the currents in each trap the local minimum of the total field can be adiabatically displaced from the centre of one coil pair to the next. The technique can be applied to many sequential traps allowing very large distances to be covered and even transport the atoms around corners, [98]. In principle the distance is only limited by the lifetime of the atoms in the trap governed by collisions with the background gas.

To transport the atoms the currents supplied to the coils of the first trap are reduced while the currents in the second are increased, and conversely upon return. How exactly these currents are modified during transport can have significant consequences on the efficiency of the transport and the temperature of the cloud. To transport the atoms with minimum loss, aside from the effects of background pressure, it is necessary to maintain as deep a trap as possible during transport. A deep trap, however, is futile unless the cloud can be kept cool. Rapid changes in the acceleration of the trap potential will heat the atoms during transport. The increase in temperature is proportional to the rate of change of acceleration, $\mathrm{d}^3x/\mathrm{d}t^3$, also known as the jerk.

For initial experiments the current ratios, $I_1(t)/I_1(0)$ and $I_2(t)/I_2(T)$, supplied to the two traps were varied such that the quadrature sum remains constant. This is described by the $n = 2$ case of equation 4.6, where T is the total transport time

and will determine the maximum acceleration and jerk experienced by the atoms.

$$
\begin{aligned}
I_1(t) &= I_1(0)\left(1-\left(\frac{t}{T}\right)^n\right)^{\frac{1}{n}}, \\
I_2(t) &= I_2(T)\,\frac{t}{T}\,, \\
\Rightarrow\ \left(\frac{I_1(t)}{I_1(0)}\right)^n &+ \left(\frac{I_2(t)}{I_2(T)}\right)^n = 1, \quad \text{for } n = 2, 3\,.
\end{aligned}
\tag{4.6}
$$

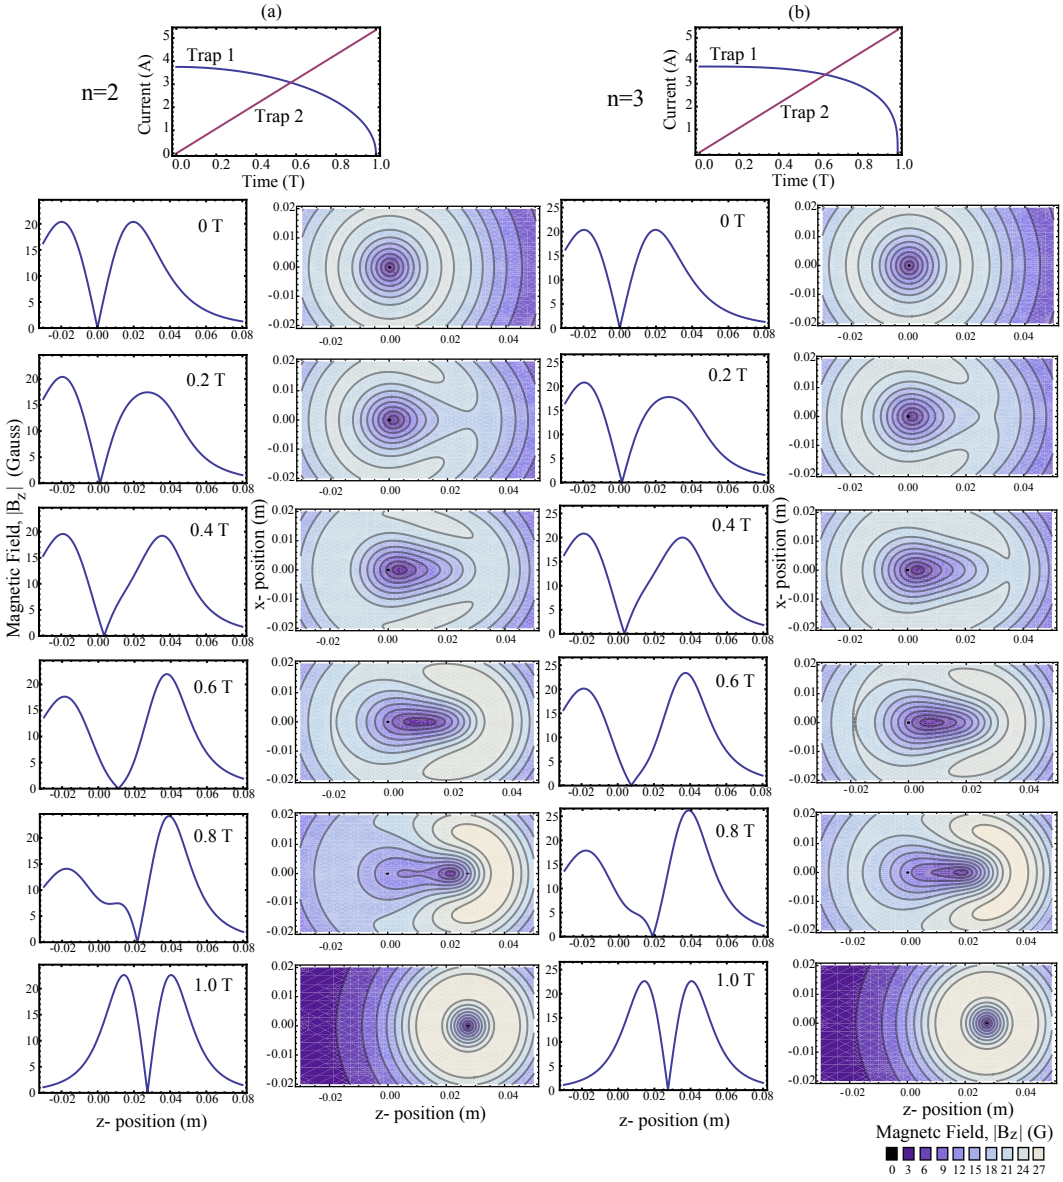

Figure 4.12: A series of 1D and 2D plots of the magnetic field profile created by traps 1 and 2 during transport as calculated in the frame of the laboratory. The 1D profiles are plotted along the direction of transportation and the 2D contour plots are in the xz-plane. (a) depicts the profiles for the $n = 2$ case of equation 4.6. (b) depicts the profiles for the $n = 3$ case. The two graphs at the top show how the currents in the coils of traps 1 and 2 change during transportation, for the $n = 2$ and $n = 3$ cases.

The graph at the top of figure 4.12(a) shows how the currents for the two traps

change in time for the $n = 2$ case. The currents in the top and bottom coils of trap two are linearly increased to a value of 5.35 A while the currents for trap 1 are reduced from 3.75 A to zero. The graphs in the left hand column depict the trap potential along the z-axis for increasing times from 0 to T in steps of $0.2\,T$. The right hand column illustrates the contour plots for the field profile in the xz-plane for the same time intervals. These plots do not include any effects due to the presence of the Zeeman or compensation coils.

From figure 4.12(a) it is clear that at a time of $0.8\,T$ the trap depth is 25% smaller in the $n = 2$ case on the left-hand side, compared to the initial depth. This is due to the trap 1 potential decreasing too much, too soon. To counter this the current ratios were modified such that the cubic sum remains constant; this is described by equation 4.6 for $n = 3$. The currents for traps 1 and 2 in the $n = 3$ case can be seen in the graph at the top of figure 4.12(b) and the same field profiles and contour plots are depicted in an identical manner to the $n = 2$ case. In the $n = 3$ case the trap depth at $0.8\,T$ is only 10% less than the initial trap depth.

The effective trap potential experienced by the atoms is distorted during acceleration. To provide an analogy, consider a full bowl of water transported laterally. If the acceleration applied to the bowl is too great the water will spill over the side; this is physically equivalent to tipping the bowl and can be modelled by the same equations. Applying these arguments to atoms sitting in a magnetic field trapping potential it is clear that greater accelerations of the field potential during transport will lead to a more tilted trapping potential and consequently a shallower trap. Given the same function for the acceleration, $a(t)$, the maximum acceleration will increase quadratically as the total transport time is decreased. To estimate the type of transport times that can be tolerated by the atoms we can compare the maximum acceleration experienced during transport to the trapping forces. The radial force in trap 1, expressed as an acceleration, is approximately 160 ms^{-2}; this acceleration is reached for transport times below 80 ms. To minimise the number of atoms lost due to background collisions we would like the transport procedure to be as quick as possible without spilling atoms from the trap. To evaluate the effects of transport over these kind of times, the effective potentials (viewed from the reference frame of an atom in the trap) for 50, 70 and 100 ms were investigated. The results are given in figure 4.13.

For the 100 and 70 ms case an atom at the centre of the trap with zero initial velocity will remain trapped throughout transport, albeit very weakly confined at times. For the 50 ms case the atom will not be confined around the 30 ms mark, however, may still remain in the trap after transport if the rate of change of potential is sufficient to push the atom back towards the centre of the trap shortly afterwards.

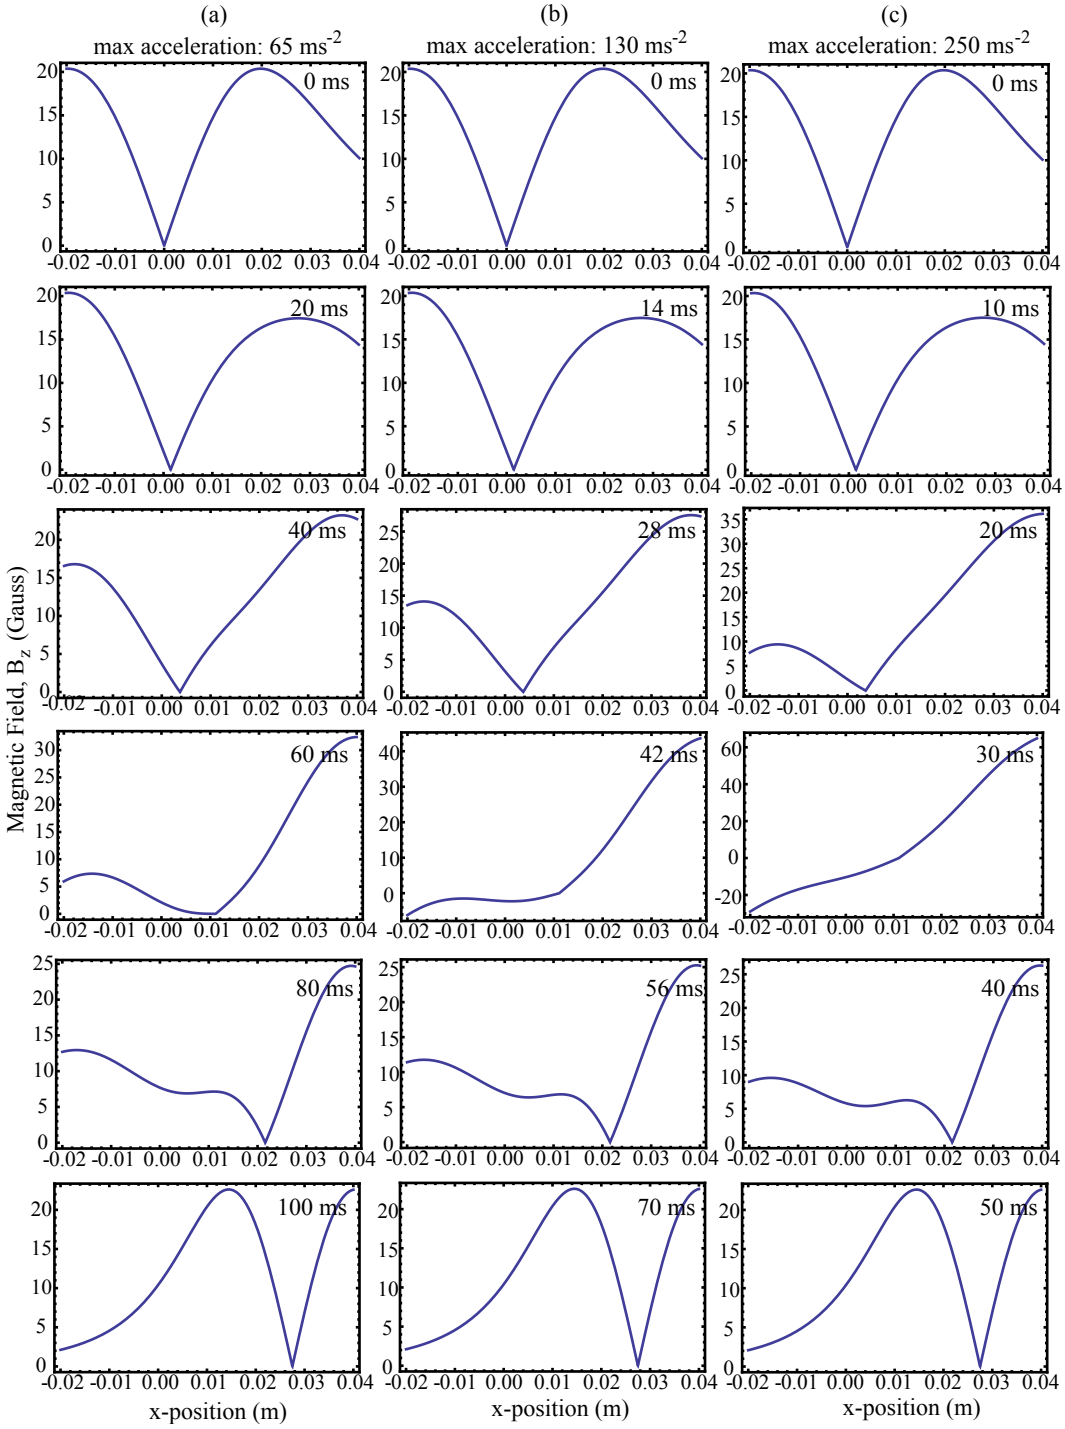

Figure 4.13: The effective potentials during transport in the frame of an atom beginning with zero velocity at the centre of the trap. The results are plotted for a transport time of (a) 100 ms, (b) 70 ms and (c) 50 ms, corresponding to maximum accelerations of 65 ms^{-2}, 130 ms^{-2} and 250 ms^{-2}, respectively.

From these plots we cannot easily determine the dynamics of an atom with different initial conditions; one might expect that an atom with a velocity in the positive z-direction may be more likely to be trapped than one beginning with a velocity in the negative z-direction, but the exact set of atoms that will remain trapped

cannot be extrapolated by looking at the data in figure 4.13 alone. To establish a better estimate of the number of atoms in the trap after transport, one-dimensional computational simulations were written to estimate the transport efficiency, η, for transport times of 50 and 100 ms. Figure 4.14 shows the phase-space diagrams for these two cases, before (red) and after (blue) transport.

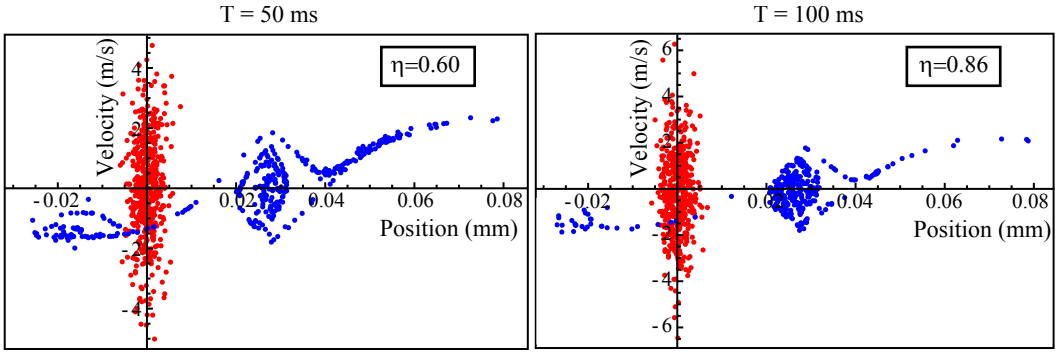

Figure 4.14: Phase-space plots showing the results of 1-D simulations used to calculate the number of atoms expected to remain trapped after a single transport from trap 1 to trap 2, for transport times of (a) 50 ms and (b) 100 ms. The red points mark the initial positions and velocities of the atoms in the MOT before release, the blue points mark the final coordinates of the atoms in trap 2 after transport. The simulations were performed for 500 atoms. The atoms are held in the trap for 30 ms before and after transport.

The initial conditions for the atoms were approximately matched to the conditions expected in the MOT at the time of release. The spatial and velocity distributions are initially gaussian with the widths chosen to match the observed size and temperature of the MOT cloud. The atoms were allowed to expand in the trap for 30 ms before transportation, in concordance with the experiments described below. During this 30 ms interval some of the atoms are lost from the trap as their energy is greater than the trap depth. In the experimental runs those atoms in non-trappable spin states will also be expelled in this 30 ms window. The simulations are only concerned with the atoms in the strongly confined $|g, 2, 2\rangle$ state.

These results demonstrate that a significant fraction of atoms remain trapped after transport, for even relatively fast times, despite the large deformations to the trapping potentials shown in figure 4.13. To experimentally determine the transport efficiency and the lifetime of trap 2 the atoms were transported to trap 2, held for a period, τ, and transported back to trap 1 where they are imaged. Currently the apparatus are not in place for measuring atoms in trap 2 directly. The sequence of events for the 100 ms (50 ms)case is as follows: the atoms are loaded into the first magnetic trap and held there for 30 ms. They are then transported to trap 2 in 100 ms (50 ms), held for an interval, τ, transported back in 100 ms (50 ms) and held in trap 1 again for 10 ms. The laser light is then turned back on and an image of the cloud is recorded after 1 ms. Figure 4.15 shows the number of atoms transported

back to trap 1 as a function of τ between 10 and 100 ms, for transport times of 50 ms (blue) and 100 ms (red).

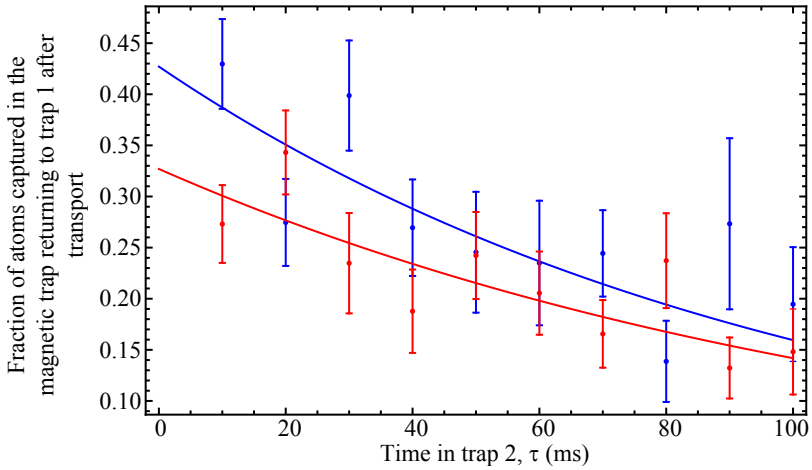

Figure 4.15: The fraction of atoms remaining relative to the number loaded into the magnetic trap, after transport to trap 2 and back as a function of time held in trap 2. The results were performed for a transport time of 50 ms (blue) and 100 ms (red) per trip.

To satisfy ourselves that the experiments are successful many measurement runs are repeated, recording images of the first trap either while the atoms are in the second trap or once they have returned. These results consistently show that the atoms leave the first trap as expected and that some fraction do return. Figure 4.16 shows 3 images, taken before transport, while the atoms are held in trap 2 and after transport for a transport time of 100 ms and a hold time of 30 ms.

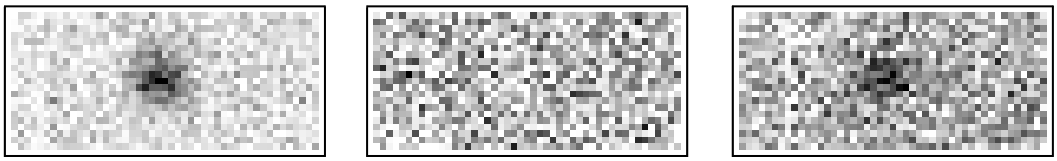

Figure 4.16: Photographs of trap 1 (a) before transport, (b) while the atoms are held in trap two, and (c) after transport back to trap 1.

The lifetimes for the two curves determine the lifetime of trap 2 and are found to be 102±26 ms and 120±29 ms for the 50 ms and 100 ms case, respectively. These values are consistent with each other within their uncertainties and can be averaged to find a value for the trap 2 lifetime of 111±19 ms. The short lifetime relative to that measured in trap 1 is thought to be due to local pressure effects; the heat from coil 1 warms up coil 2, which in turn outgasses material. Since trap 2 has rarely been used to date, the level of outgassing is expected to be much larger than from trap 1; likewise it is expected that the lifetime of trap 2 will increase with use.

To calculate the efficiency of the transport we need to know the fraction of atoms

that would have remained had they not undergone transportation. This is determined by measuring the decay curve for trap 1, which was performed immediately prior to the above transport measurements. The fraction of atoms loaded into the magnetic trap was found to be 8.1±0.5% of the inital MOT number and decayed from this value with a lifetime of 839±78 ms. The number of atoms remaining after transportation can now be compared to the trap 1 data. From figure 4.15 we can extrapolate back to a hold time of zero seconds and thus infer the fraction of atoms lost due to transfer alone. From figure 4.15 we find we recover 42.7±4.9% and 32.7±3.7% of the atoms initially loaded into trap 1, for transport times of 50 ms and 100 ms, respectively. These values represent the number of atoms remaining 140 and 240 ms since their release from the MOT. In the absence of transport the atoms are lost from trap 1 due to background collisions. For the 100 ms case the number of atoms remaining in trap 1 after 240 ms is expected to be $75.3 \pm 6.2\%$ without transportation, and 85.2±6.2% for the 50 ms case after 140 ms. These values imply transfer efficiencies of $\eta_{50} = 0.71 \pm 0.07$ and $\eta_{100} = 0.66 \pm 0.07$. The value for η_{50} is larger than the above simulation predicts, whereas the value of η_{100} is smaller.

During these calculations we have not considered the effects of additional background collisional losses during transport due to the increased proximity to the trap 2 coils. This loss rate will act for longer on the cloud during a transportation time of 100 ms than for 50 ms. It is not possible to assign a value to the lifetime, τ_{bg}, due to these additional collisions. Nor can the effects of background collisions be decoupled from heating effects brought about by accelerating the cloud. It is not unreasonable, however, to expect the loss rate to remain the same for the two transport times. Under this assumption, however, and making appropriate modifications to the calculations there is no value of τ_{bg} that can reconcile the measurements with the predictions of the simulations. To marry these findings one must conclude that either transport in 100 ms is less efficient than predicted by theory or that transport in 50 ms is more efficient. Currently the data does not exist to distinguish between these two possibilities, however, with more detailed simulations and a direct investigation of where the atoms are during the entire transportation time, it is hoped that these findings can be rectified.

Ultimately, however, the aims of these investigations have been met. A cloud of ^{7}Li atoms has been successfully transported across the chamber and back with a good efficiency, demonstrating good control over the cloud position. Future experiments will investigate the dynamics of the cloud, including the effects of transport on the cloud temperature.

Chapter 5

Conclusion and Outlook

5.1 Final Conclusions

The aims of this thesis were to build and characterise a source of ultra-cold lithium, suitable for the sympathetic cooling of molecules and to investigate the feasibility of directly polarising this atomic gas. The production of such strongly correlated quantum gases is of much scientific interest to both the cold atom and condensed matter communities. Because of the high degree of purity of atomic and molecular gases and the precise control over many of the parameters defining the nature of their interactions, such a system is ideal for understanding the physics of strongly correlated many-body systems. By testing these it will hopefully provide insight into the more complex dynamics underlying condensed matter systems that are difficult to study directly.

The experiments undertaken in this work were designed to keep open the possibility of creating such a gas using either an atomic or molecular species. For both lines of investigation an ultra-cold atomic source was required; for these experiments we chose to use bosonic lithium. Lithium is a good candidate for direct electrical polarisation and is also thought to be a good sympathetic coolant for light polar molecules. In concurrent experiments a molecular source is being investigated with the aim to trap molecules from a Stark decelerated beam, sympathetically cool them and finally, electrically polarise them. Both LiH and CH molecules are being studied, as their small mass to electric dipole ratio makes these molecules excellent candidates for Stark deceleration and for the study of dipolar gases. While the feasibility of sympathetically cooling these molecules with any atomic species is still unknown, it is known that choosing an atomic to molecular mass ratio close to one is beneficial in reducing the number of collisions required for the molecules to thermalise with the atoms. Thus, since Li is very amenable to laser cooling, it is an obvious choice as a refrigerant for such light molecules. It is also known that to avoid inelastic collisions in sympathetic cooling, it is best to choose a light collisional partner with a small

Van der Waals coefficient [27] for the coolant, even for relatively larger diatomic molecules; in this respect lithium is the best choice.

In this thesis I have described the formation of a large, ultra-cold source of ^{7}Li atoms, confined in a magneto-optical trap. The lithium atoms are heated in an oven to create a thermal effusive beam. The atoms subsequently enter a 60 cm long Zeeman slower where a large fraction are laser cooled to a velocity of 50 ms^{-1} and loaded into a magneto-optical trap. The Zeeman slower delivers approximately 4×10^9 slow atoms per second, leading to a maximum recorded atom number in the MOT of 2×10^{10} atoms with a minimum temperature of 0.85 mK.

To transport the atoms collected in the MOT a pair of in-vacuum magnetic traps were designed and built. The magnetic traps are each formed from a pair of anti-Helmholtz coils forming a quadrupole magnetic field with a zero-field point at the centre. The first trap is centred directly over the MOT and was used to form a new MOT with a magnetic field gradient of 40 $\mathrm{G\,cm}^{-1}$. In this compressed MOT a maximum atom number of 3×10^7 was recorded at a density of 3×10^{10} atoms/cm^3. The magnetic trap was loaded by switching off all light into the MOT chamber. The magnetic trap has a depth of 1.9 mK and atoms were loaded from the MOT with a 4% efficiency. A trap lifetime of 300 ms was measured, though longer lifetimes exceeding 1 second were later recorded under better background pressure conditions. A deeper trap, with a depth of 2.5 mK, can be formed by combining both the outer and inner vacuum coils, capturing 11% of the atoms from the MOT. An increase in the atom number transferred into the magnetic trap was also demonstrated by applying a cooling stage to the MOT immediately prior to release into the magnetic trap. This cooling stage reduced the temperature of the MOT cloud by approximately a factor of 2.

To electrically polarise the atoms directly, or to sympathetically cool a molecular sample it is necessary to transport the atoms away from the MOT region. To demonstrate this process the second magnetic trap was positioned such that the trapping coils partially overlapped with the coils forming the first magnetic trap. By reducing the current for the first trap while increasing the current to the second the zero field point can be smoothly transferred from the centre of one trap to the other. For gentle enough accelerations the atoms will adiabatically follow the zero-field point and move from one trap to the next. This procedure was used to transport the atoms from trap 1 to trap 2 and back, for various transfer times. Fluorescence images showed that the atoms are transferred in a controlled manner with an efficiency of approximately 70% per trip for transfer times of both 50 and 100 ms. Computational simulations performed for our transport parameters suggest the transfer efficiency should be higher for the 100 ms case than the 50 ms one. The

discrepency between theory and the experimental results is not fully understood and further investigation is required.

In parallel to these projects an investigation into the feasibility of creating extremely large electric fields has been performed. To electrically polarise ultra-cold ^{7}Li sufficiently, such that dipole-dipole interactions dominate over the usual s-wave scattering processes, a field strength of at least 0.8 MV/cm is required. We have shown that between two broad-area, optically-polished, stainless steel plates a field of 0.81 MV/cm can be realised without electrical breakdown occurring. The feasibility of placing atoms under such electric fields is then discussed. Calculations show there is a strict tolerance on the size of the electric field gradients allowed between the plates in order to keep the atoms confined in their magnetic trap. These tolerances equate to a required parallelism between the plates of better than 0.1 mrad. To this end we investigated the construction of optically-polished, indium-tin oxide coated glass plates mounted to an adjustable rig capable of achieving a high degree of parallelism between the plates. The results suggest an angle between the plates better than 0.08 mrad can be achieved. This mounting rig also allows us to accurately control and measure the plate separation to within 10 μm, without making contact with the delicate plate surfaces. Unfortunately, no work on the electric field strengths maintainable between two ITO coated plates has yet been carried out.

This work has paved the way for the realisation of an ultra-cold dipolar gas of ^{7}Li in the near future. Such a gas will allow one to study many of the interesting quantum phases found in such a strongly correlated quantum system.

5.2 Outlook

To date we have demonstrated the controlled manipulation of the position of a magnetically confined cloud of ultra-cold ^{7}Li atoms. In addition we have shown that the electric field strengths required to polarise a gas such that the dipole-dipole interactions dominate the collisional processes are achievable. The next stage for this work is to bring these two components together, in order to magnetically confine our atoms between the electrodes under field strengths approaching 1 MV/cm.

The first requirement to achieve these goals is the realisation of a magnetically confined cloud close to the Doppler temperature. This temperature is required to compress the cloud sufficiently in our magnetic trap such that a large fraction of our atoms can be positioned between the electric field plates. Further cooling has been achieved by lowering the intensity and the detuning of the MOT beams immediately prior to release into the magnetic trap. Currently the MOT detuning is reduced to $\Gamma/2$ to cool the cloud, however, a greater effect is expected if the MOT detunings

can be increased to much larger detunings of the order of 10 Γ. Our most recent work suggests that considerably cooler MOT temperatures can be achieved using a lower intensity repump beam. Temperatures of 500 μK have been recorded in the MOT for repump powers as low as 2.5 mW. In addition to cooling the cloud, lower temperatures will lead to greater atom numbers loaded into the magnetic trap. Further increases in atom number may also be obtained by creating a strongly spin-polarised cloud. At present it is thought that the atoms in the MOT are uniformly distributed amongst their eight ground M_F states. If the atoms are irradiated by a beam with σ_+ polarised light, with respect to the direction of a weak, uniform magnetic field, parallel to the propagation of the beam, and resonant with the $F = 1$ component of the D2 line a significant fraction of the atoms can be pumped into the $|g, 2, 2\rangle$ and $|g, 2, 1\rangle$ states, both of which are confined by the trap. Larger atom numbers are also achievable by loading the magnetic trap from a larger MOT. Currently the size of the MOT is significantly limited by the use of very large field gradients. These gradients have been applied to simplify the loading process into the magnetic trap, however, there is no reason why in the future the field gradient cannot be optimised for the MOT and strengthened at the point of release to create a deeper, steeper trap. Beginning with 10^{10} atoms in the MOT we could expect to capture more than 10^9 atoms in a magnetic trap through the application of these proposals.

The second stage is to design an electrode pair with a support rig capable of meeting the various criteria discussed in section 3.6. As was discussed, it may be possible to modify our current designs to meet these goals, however, the experiment is ultimately constricted by the stringent volume constraints imposed by the small chamber and the required optical access, making such an attempt unnecessarily difficult to realise. It is therefore proposed that the experiment be divided into two vacuum chambers, one for housing the MOT and the other housing the field plates. The atoms may be transported between the two chambers using magnetic field coils mounted to a translation stage outside the chamber. This transport mechanism has been demonstrated for rubidium already [96] and we can see no reason why the technique cannot be applied equally well to lithium. Because the coils would be outside the vacuum we are not limited by the currents that can be applied, allowing us to build deep traps with steeper field gradients. Larger gradients mean warmer cloud temperatures can be tolerated for atoms placed between the electrodes, while larger trap depths equate to more atoms collected in the magnetic trap from the MOT. By dedicating a chamber solely to the electric field plates an appropriate design can be constructed that meets all the required criteria while providing plenty of optical access for the imaging required. The current design for the electrode

support rig is not rigid enough to withstand the electrostatic forces acting between the plates, while more rigid designs are very difficult to build to meet the severe tolerances and the required parallelism between the plates. One proposed solution, to build a rigid, but adjustable support rig, is to use small piezoelectric actuators to adjust the alignment of the plates. The plates can be adjusted outside the chamber to be close to parallel using the interferometric methods described in section 3.5.1. Once under vacuum the actuators can be further adjusted through an electrical feedthrough, using the atoms themselves as a guide to determine the degree of parallelism with unprecedented precision. In addition, by separating the chambers by a long thin tube, differential pumping can be employed in order to significantly improve the vacuum quality in the second chamber, thereby extending the lifetime of the trap and thus the ease of the experiment.

With these improvements we are confident that a dipolar atomic gas is well within reach, with very few experimental obstacles to be overcome along the way. Once such a gas has been demonstrated the experiment has many avenues that can be explored, including the realisation of a polarised Bose-Einstein condensate and the coupling of such a cloud to an optical lattice. In addition, the proposed modifications allow us to easily couple our atomic trap with a molecular trap for the purpose of sympathetic cooling.

Appendix A

Installation and Assembly of the Magnetic Traps

A major obstacle to be circumvented during the design of these coils is that the apparatus must fit in the chamber. The largest access points into the chamber are through the two DN63CF flanges at the top of the chamber, providing a 66 mm diameter circular aperture for parts to enter. Both the height and length of the structure, however, are larger than this. Consequently, the final design must be constructed from separate pieces each of which can fit into the chamber after the previous pieces are inside, and be simple enough that it can be constructed via access through a hole with a diameter smaller than a hand width. Simple plastic models representing the coil structure and the parts of the chamber involved were machined from nylon to test various recipes for assembly and to judge which pieces could fit and which must be assembled once inside.

The final design was constructed from 8 main components. The coil formers for the first trap were machined from a single slab of copper whose thickness was chosen to position the coils at the correct height. The upper slab has a reduced thickness to allow room for assembly. The second coil pair was wound onto a separate copper former that is tightly bolted to the main slab, such that the coils from each trap overlap. The upper and lower slabs are connected by four copper pillars which slot tightly into four counterbored holes both in the upper and lower copper slab. This allows thermal conductance from the upper slab to the lower slab and so to the floor of the chamber. To maintain good thermal contact four holes were drilled through the upper slab, through each of the pillars and finally part way into the lower slab. The holes in the lower slab were tapped to hold four supporting threaded rods, inserted through the centre of each pillar and protruding a short distance above the upper slab. Nuts could then be threaded onto each rod and the upper and lower slab could be pinched together. This improved both stability and thermal conductivity.

To maintain the position of the structure in the chamber and to obtain good thermal contact with the chamber floor, four springs were inserted between the roof of the chamber and the upper side of the top slab. Shallow holes were machined into the roof of the coil structure for the springs to sit in. The springs were compressed as the structure was put in place to firmly secure the coil structure against the chamber floor.

The order of placement and assembly in the chamber is very important and not necessarily intuitive so has been presented here.

- The larger coil is wound first. The aluminium plate is pinched against the copper base using a small clamp inserted through the centre of the former. The securing screws cannot be inserted until after winding is complete. Once the coil is wound the aluminium plate is secured by four screws connecting the plate to the copper base. Tension in the wire is maintained by wrapping the wire around these screws.

- The second coil is then wound with the aluminium plate in place. Once wound the former is bolted onto the copper slab, flush with the first coil. The process is identical for both the upper and lower copper plate.

- The lower plate is inserted into the chamber first through the upper DN63CF flange closest to the slower. Two of the four threaded rods are screwed into the back two holes set for the thicker support pillars. These two pillars are then slid over the rods and pushed firmly into the base plate. The plate is then slid to the back of the MOT chamber to allow room for the upper slab to be inserted.

- The upper plate is then placed in the chamber. The plate is held against the roof of the chamber while the lower plate and pillars are slid underneath. The back pillars are positioned into the corresponding holes on the upper plate.

- The thin pillars are then placed at the front of the plate. To do this the top plate must be angled up enough to allow the pillars to be pushed between the two copper slabs before the two pillars are sunk into their corresponding holes. The final two threaded rods can then be inserted from above and the four nuts are screwed onto each threaded rod to compress the upper and lower plates together.

- Finally, the two front springs are inserted into the top plate, compressed, and the structure is slid to the front of the chamber allowing access to the two recesses cut to accommodate the two remaining springs. The last two springs

are then inserted and the structure is slid back into position such that all four springs now compress the structure against the chamber floor. Finally each coil is connected to the electrical feedthrough and the wires are positioned to avoid obstructing any of the MOT beams. The MOT beams are switched on to check no light is blocked and the electrical connections are tested before the chamber is sealed.

Additional room on the slab was also incorporated to house the electrodes and their supporting unit. A small area was recessed into the lower copper plate for the unit to sit as close to the second trap as possible. The details of the electrode design are given in sections 3.5.1 and 3.6.

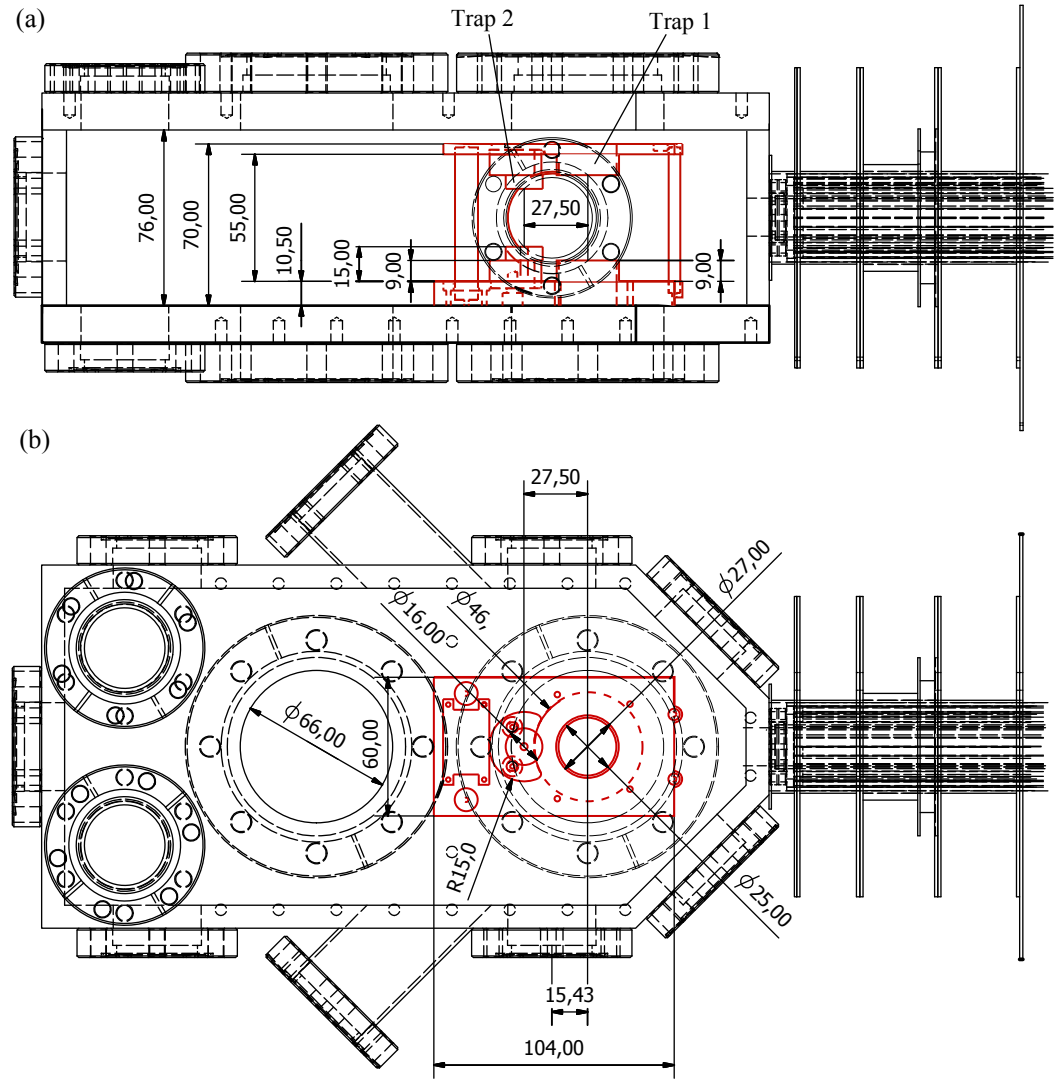

Figure A.1: Detailed drawings of the MOT chamber and the in-vacuum magnetic traps. The support structure for the coils and the main flanges on the chamber are shown. The electric field plates have been omitted. All units are in mm.

Bibliography

[1] A. Micheli, G. K. Brennen and P. Zoller, *Nat. Phys.* **2**, 341 (2006)

[2] J. L. Roberts *et al.*, *Phys. Rev. Lett.* **81**, 5109 (1998)

[3] C. A. Regal *et al.*, *Nature* **424**, 47-50 (2003)

[4] K. Góral, L. Santos and M. Lewenstein, *Phys. Rev. Lett.* **88**, 170406 (2002)

[5] D. DeMille, *Phys. Rev. Lett.* **88**, 067901 (2002)

[6] A. André *et al.*, *Nat. Phys.* **2**, 636 (2006)

[7] D. H. J. O'Dell, S. Giovanazzi and C. Eberlein, *Phys. Rev. Lett.* **92**, 250401 (2004)

[8] A. Griesmaier *et al.*, *Phys. Rev. Lett.* **94**, 160401 (2005)

[9] J. Stuhler *et al.*, *Phys. Rev. Lett.* **95**, 150406 (2005)

[10] J. Kim *et al.*, *Phys. Rev. Lett.* **78**, 19 (1997)

[11] E. Narevicius *et al.*, *Phys. Rev. A* **77**, 051401 (2008)

[12] P. A. Ruprecht *et al.*, *Phys. Rev. A* **51**, 4704 (1995)

[13] F. Dalfovo and S. Stringari, *Phys. Rev. A* **53**, 2477 (1996)

[14] Y. Kagan, G. V. Shlyapnikov, and J. T. M. Walraven, *Phys. Rev. Lett.* **76**, 2670 (1996)

[15] M. Marinescu and L. You, *Phys. Rev. Lett.* **81**, 4596 (1998)

[16] Faraday Discussions 142, *Cold and Ultracold Molecules*, RSC Publishing, (2009)

[17] J. Doyle *et al.*, *Eur. Phys. J. D* **31**, 149 (2004)

[18] D. Wang *et al.*, *Phys. Rev. Lett.* **93**, 243005 (2004)

[19] J. Deiglmayr *et al.*, *Phys. Rev. Lett.* **101**, 133004 (2008)

[20] A. J. Kerman *et al.*, *Phys. Rev. Lett.* **92**, 033004 (2004)

[21] K. -K. Ni *et al.*, *Science* **322**, 231 (2008)

[22] E. Kuznetsova *et al.*, *New J. Phys.* **11**, 055028 (2009)

[23] M. D. Di Rosa, *Eur. Phys. J., D* **31**, 395 (2004)

[24] B. K. Stuhl, B. C. Sawyer, D. Wang and J. Ye, *Phys. Rev. Lett.* **101**, 243002 (2008)

[25] E. S. Shuman, J. F. Barry, D. R. Glenn, and D. DeMille, *Phys. Rev. Lett.*, **103**, 223001 (2009)

[26] T. E. Wall *et al.*, *Phys. Rev. A* **78**, 062509 (2008)

[27] M. Lara, J. L. Bohn, D. Potter, P Soldán, and J. M. Hutson, *Phys. Rev. Lett.* **97**, 183201 (2006)

[28] E. Urban *et al.*, *Nat. Phys.* **5**, 110 (2009)

[29] T. Pohl and P. R. Berman, *Phys. Rev. Lett.* **102**, 013004 (2009)

[30] I. Mourachko *et al.*, *Phys. Rev. Lett.* **80**, 253-256 (1998)

[31] M. P. Robinson *et al.*, *Phys. Rev. Lett* **85**, 4466-4469 (2000)

[32] S. Tokenaga *et al.*, *New J. Phys.* **11** 055038 (2009)

[33] M. Mudrich *et al.*, *Phys. Rev. Lett.* **88**, 253001 (2002)

[34] S. Jochim, *Diploma Thesis* - Bose-Einstein Condensation of Molecules, Institut für Experimentalphysik (2004)

[35] J. Kawanaka *et al.*, *Appl. Phys. B* **57**, 113 (1993)

[36] Z. Lin *et al.*, *J. Appl. Phys.* **30**, 1324 (1991)

[37] U. Schünemann *et al.*, *Optics Communications* **158**, 263 1998

[38] Erling Riis and Charles S. Adams, *Laser Cooling and Manipulation of Neutral Particles,* To appear in The New Optics, Cambridge University Press, (2010)

[39] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* Springer-Verlag, (1999)

[40] J. Prodan and W. Phillips, Chirping the Light Fantastic — Recent NBS Atom Cooling Experiments, *Prog. Quant. Elect.* **8**, 231 (1984)

[41] R. Watts and C. Wieman, *Opt. Lett.* **11**, 291 (1986)

[42] R. Gaggl, L. Windholz, C. Umfer and C. Neureiter, *Phys. Rev. A* **49**, 1119-1121 (1994)

[43] W. Phillips and H. Metcalf, *Phys. Rev. Lett.* **48**, 596 (1982)

[44] E. L. Raab *et al.*, *Phys. Rev. Lett.* **59**, 23 (1987)

[45] P. Lett *et al.*, *Phys. Rev. Lett.* **61**, 169 (1988)

[46] J. Dalibard and C. Cohen-Tannoudji, *J. Opt. Soc. Am. B* **6**, 2023-2045 (1989)

[47] Y. Castin, J. Dalibard, and C. Cohen-Tannoudji, *"The limits of Sisyphus cooling"*, in *Proceedings of Light-Induced Kinetic Effects on Atoms, Ions and Molecules*, L. Moi, S. Gozzini, C. Gabbanini, E. Arimondo, and F. Strumia, eds. ETS Editrice, Pisa, Italy, (1991)

[48] D. Felinto *et al.*, *J. Opt. Soc. Am. B* **18**, 10 (2001)

[49] C. Fort *et al.*, *Eur. Phys. J. D* **3**, 113 (1998)

[50] W. Ketterle and N. J. van Druten, *Adv. Atom. Mol. Opt. Phys.* **37**, 181 (1996)

[51] M. Mudrich *et al.*, *Phys. Rev. Lett.* **88**, 25 (2002)

[52] C. Wieman and T. W. Hänsch, *Physical Review Letters* **36**, 20, (1976)

[53] M. L. Harris *et al.*, *Phys. Rev. A* **73**. 062509 (2006)

[54] W. Demtröder, Laser Spectroscopy *Basic Concepts and Instrumentation* Chapter 7. p463-476 (1981)

[55] J. Fortagh *et al.*, *J. Appl. Phys.* **83**, 6499 (1998)

[56] B. P. Anderson and M. A. Kasevich, *Phys. Rev. A* **63**, 023404 (2001)

[57] L. Dunoyer, *Comptes Rendus* **152**, 594 (1911).

[58] L. Dunoyer, *Le Radium* **8**, 142 (1911).

[59] L. Dunoyer, *Le Radium* **10**, 400 (1913).

[60] N. F. Ramsey, *Molecular Beams* Oxford Classic Texts (2005, Orig. 1956)

[61] E. H. Kennard, *Kinetic Theory of Gases* McGraw-Hill Book Company, Inc. (1938)

[62] P. Claussing, *Physica* **9**, 494 (1942) and *Zeits f. Physik* **6**, 471 (1931)

[63] M. R. Tarbutt *et al.*, *arXiv*:0803.0967v2 (2009)

[64] F. Lison *et al.*, *Phys. Rev. A* **61**, 013405 (1999)

[65] J. F. O'Hanlon, *A User's Guide To Vacuum Technology* A Wiley-Interscience Publication, (1989)

[66] J. V. Stewart, *Intermediate Electromagnetic Theory*, World Scientific Publishing Co. Pte. Ltd., pp.275-276 (2001)

[67] K. E. Gibble *et al.*, *Opt. Lett.* **17**, 526 (1992)

[68] K. Lindquist *et al.*, *Phys. Rev. A* **46**, 7 (1992)

[69] C. G. Townsend *et al.*, *Physics Review A* **52**, 1423, (1995)

[70] A. M. Steane *et al.*, *J. Opt Soc. Am. B* **9**, 12 (1992)

[71] D. W. Sesko *et al.*, *J. Opt. Soc. Am. B* **8**, 946-958 (1991)

[72] R. H. Fowler and L. Nordheim, *Proceedings of the Royal Society of London, Series A, Containing a Mathematical and Physical Character* **119**, 781, pp.173-181, (1928)

[73] R. A. Millikan and R. A. Sawyer, *Phys. Rev.* **12**, pp.167-170, (1918)

[74] R. V. Latham, *High Voltage Vacuum Insulation: The Physical Basis* Academic Press Inc. (1981)

[75] R. A. Millikan and B. E. Shackelford, *Phys. Rev.* **15**, 239, (1920)

[76] R. A. Millikan and C. F. Eyring, *Phys. Rev.* **27**, 51, (1926)

[77] F. M. Charbonnier, C. J. Bennette and L. W. Swanson, *(Field Emission Corporation, McMinnville, Oregon) Journal of Applied Physics* **38**, 2, 627 (1967)

[78] W. P. Dyke *et al.*, *J. Appl. Phys.* **24**, 570 (1953)

[79] W. W. Dolan, W. P. Dyke and J. K. Trolan, *Physical Review* **89**, 799-808, (1953)

[80] C. J. Bennette, L. W. Swanson and F. M. Charbonnier, *(Field Emission Corporation, McMinnville, Oregon) Journal of Applied Physics* **38**, 2, 634 (1967)

[81] D. Alpert *et al.*, *J. Vacuum Sci. Technology* **1**, 2, pp.35-50, (1964)

[82] A. Kovaleva, *Masters Thesis, unpublished*, The Production of Extremely Large Electric Fields, Imperial College London (2007)

[83] J. Millen, *Masters Thesis, unpublished*, The Production of Extremely Large Electric Fields, Imperial College London (2007)

[84] P. A. Chatterton, *Proc. Roy. Soc.* **88**, pp.231-245, (1966)

[85] J. Stack, *Masters Thesis, unpublished*, The Production of Extremely Large Electric Fields, Imperial College London (2006)

[86] F. Paschen, *Wied. Anal. Phys. Chem.* **37**, 69-96, (1889)

[87] B. E. Sauer, D. M. Kara, J. J. Hudson, M. R. Tarbutt and E. A. Hinds, *Review Of Scientific Instruments* **79**, 126102 (2008)

[88] A. Guthrie, *Vacuum Technology*, (1985)

[89] F. Roseburg, *Handbook of Electron Tube and Vacuum Techniques*, (1965)

[90] C. Suzuki *et al.*, *The Fourteenth International Spin Physics Symposium* **570**, pp.1009-1011 (2001)

[91] L. Cranberg, *J. Appl. Phys.* **23**, 5 (1952)

[92] W J Atkinson, J H Young and I A Brezovich, *J. Phys. A* **16**, 2837 (1983)

[93] S. H. Brewer and S. Franzen, *Journal of Alloys and Compounds* **338**, 73-79 (2002)

[94] J. He *et al.*, *Thin Solid Films* **363**, 240-243 (2000)

[95] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Görlitz, S. Gupta, D. E. Pritchard and W. Ketterle, *Phys. Rev. Lett.* **88**, 2 (2002)

[96] S. Schlunk, A. Marian, P. Geng, A. P. Mosk, G. Meijer and W. Schöllkopf, *Phys. Rev. Lett.* **98**, 223002 (2007)

[97] H. J. Lewandowski, *Coherences and correlations in an ultracold Bose gas*, Doctoral Thesis from the University of Colorado, (2002)

[98] M. Greiner, I. Bloch, T. W. Hänsch, and T. Esslinger, *Phys. Rev. A* **63**, 031401 (2001)

[99] E. Cornell, C. Monroe and C. E. Wieman, *Phys. Rev. Lett.*, **67**, 18, (1991)

[100] F. Schreck, *Doctoral Thesis* Mixtures of ultracold gases: Fermi sea and Bose-Einstein condensate of lithium isotopes, Laboratoire Kastler Brossel, (2002)

[101] N. F. Ramsey, *Molecular Beams* Oxford Classic Texts (2005, Orig. 1956), pp.34