A PERMANENT MAGNET TRAP FOR BUFFER GAS COOLED ATOMS

Didier Marcus Marcel Nohlmans

Department of Physics
Imperial College London

Thesis submitted in partial fulfilment of the

requirements for the degree of

Doctor of Philosophy in Physics

of

Imperial College London

July 2015

DECLARATION

I declare that this thesis is my own work. Where I have used the work of others, the sources are appropriately referenced and acknowledged.

ABSTRACT

Achieving precise control over an array of ultracold molecules would provide a unique tool-set for carrying out quantum simulations and quantum computations, as a result of the molecules' rich internal structure. To realise this aim, the molecules have to be cooled and trapped. This is much more difficult for molecules than for atoms due to their complex internal structure. This thesis presents preliminary work towards realising a versatile, permanent magnet trap for buffer gas cooled molecules. Atoms are used throughout to test the feasibility of the trap, as they are easier to produce and detect.

Two novel methods for trapping buffer gas cooled atoms in a permanent magnet trap are investigated. The first of these involves trapping the atoms directly from a cryogenic buffer gas cooled ablation plume. Dy atoms, with a magnetic moment of $10\mu_B$, are trapped with a lifetime of $800 \pm 30\,\mu\mathrm{s}$, thought to be limited by collisions with a high density of background buffer gas atoms remaining in the trap region.

Information gained from the direct trapping experiments motivated the design of a second trapping set-up. Here, a beam of Dy atoms is first extracted from a cryogenic buffer gas source, and when this beam reaches the trapping region, a fraction of the atoms are stopped through collisions with cold helium gas present in the trapping region. This second method reduces the density of buffer gas required in the trap region. The trap lifetime achieved in this arrangement of $810 \pm 40\,\mu\mathrm{s}$ is no longer than in the direct trapping experiments, but this arrangement is much more stable and repeatable. The lifetime here is also thought to be limited by collisions with background buffer gas atoms.

ACKNOWLEDGEMENTS

The work in this thesis would not have been possible without the support of a great number of people. The first and foremost of these is Mike Tarbutt, who has been a brilliant supervisor, offering expert guidance on the experiments and always making time to help, despite his busy schedule. I also want to thank Ed Hinds, Ben Sauer and Danny Segal for stimulating discussions and advice during weekly meetings, and Ben for always knowing where to find a particularly niche piece of equipment in the lab.

None of this work would have happened without the support of Sarah Skoff, particularly during long days of data taking in the lab. She has been great at helping to come up with solutions to numerous experimental problems, and her knowledge of buffer gas cooling has been instrumental in advancing the experiments. I also want to thank the rest of team Buffer Gas and CCM for making my time in the lab a thoroughly enjoyable experience, and Sanja Maricic, without whom CCM would not function. The expertise of Jon Dyne and Stephen Maine in the mechanical engineering workshop made possible the construction of all the parts in the apparatus, without which none of these experiments would have been built.

I have been lucky enough to have been part of the Centre for Doctoral Training (CDT) for my PhD, and I want to thank everyone in the CQD CDT for providing great support, and for the good times during the parties and conferences. I am also extremely grateful to my family and friends, in particular my parents, who have always supported me and provided me with the opportunity to pursue my studies; and of course Helen, for putting up with me throughout this PhD and for always being there at the end of a long day.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

Cold molecules are poised to provide a wealth of new science, even exceeding that of their atomic counterparts. This is predominantly due to their vastly richer structure as a result of rotational and vibrational energy levels. It is this complicated internal structure that also makes molecules much more difficult to cool than atoms. As a result, the field of cold molecules is still in its infancy compared to cold atoms, with the majority of the cooling and trapping techniques described in this chapter having only been developed in the last decade, and many of them much more recently than that.

With so much having already been achieved with ultracold atoms, such as precise control in optical lattices [1] or the creation of BECs [2], one might wonder what exactly can be accomplished with cold molecules that can not already be done with atoms. Some of these reasons are summarised in section 1.1. Essential to the implementation of these are the cooling and trapping of molecules, which are discussed in more detail in section 1.2. There are two temperature regimes that are referred to when it comes to cold molecules: cold (1 mK–1 K) and ultracold (< 1 mK). Different temperature regimes and molecule densities typically require different cooling methods and allow for many different applications.

1.1 Applications of cold molecules

There are numerous research directions where cold molecules can be a real asset, some of which are already being explored and others which will likely be possible in the near future with the advancement of molecular cooling and trapping techniques. Two of the major research directions for cold molecules are discussed in some detail in this section, which also highlights some other applications.

1.1.1 Precision spectroscopy

One of the main uses of cold molecules is for precision spectroscopy and the testing of fundamental theories in laboratory-based experiments [3]. In some modern theories of particle physics, the fundamental constants, such as the fine structure constant, are predicted to vary with time, position or the local density of matter [4]. Cold molecules can be used to search for such variations by carrying out precise spectroscopy to measure the frequency of molecular transitions in the lab, and comparing those to astronomical measurements. For example, ground state microwave transitions of OH molecules were measured precisely by using cold, Stark decelerated OH molecules [5]. Comparing these to astrophysical measurements of the same transitions, constraints could be put on the time evolution of the fine structure constant. Similarly, very accurate measurements of microwave transitions in CH [6] were compared to data from the interstellar medium in the Milky Way to constrain the variation of the fine structure constant between the high density environment of the earth and the low density environment of the interstellar medium. Additionally, there are suggestions that if the fine structure constant were different in the past then there should be an associated change in the electron-proton mass ratio, m_e/m_p, which may be a lot larger. Atomic structure has almost no dependence on m_e/m_p, but molecular structure does through rotational and vibrational energy levels. There have been a number of proposals to use molecular transitions to probe the time variation of m_e/m_p [7, 8, 9] by comparing laboratory measurements with high redshift astronomical data. Cold molecules would allow very accurate laboratory measurements of these transitions, putting ever tighter constraints on the variation of these fundamental constants. Several of these measurements have already been made [10, 11].

Cold molecules can also be used to measure the electric dipole moment (EDM) of the electron much more sensitively than measurements made using atoms [12]. Unless time reversal symmetry (T) is violated, the electron EDM should vanish entirely. Observing a non-zero electron EDM would therefore be a measure of T violation, and by extension also CP violation. With CP violation being necessary to explain the matter-antimatter imbalance in the universe, it is clearly of great interest to make accurate measurements of the electron EDM. In the last few years, first YbF molecules [13] and subsequently ThO molecules [14] have allowed the most accurate measurements of the electron electric dipole moment to date, putting a new upper bound on the electron EDM. The YbF experiment [13] used a room temperature supersonic beam of YbF molecules, whilst the ThO experiment [14] used a cryogenic buffer gas beam of ThO molecules, helping them to improve upon the electron EDM measurement from [13] as the slower beam increased their interaction time. One future possibility to improve the accuracy of the electron EDM measurement is to create a molecular fountain by laser cooling YbF molecules from a cryogenic buffer gas source [15], thereby increasing the interaction time even further. With the Standard Model of particle physics predicting an almost zero value for the electron EDM, searching for a non-zero value is a very powerful tool in determining the validity of extensions to the Standard Model, each of which predict different, much larger, values of the electron EDM [16].

1.1.2 Quantum information and quantum simulations

Cold *polar* molecules are of particular interest, as the large dipole-dipole interactions between them could potentially allow them to be used for quantum simulations [17, 18] and quantum computation [19]. An applied electric field polarises the molecules through the mixing of rotational states. As these are close in energy, the field needed to fully polarise the molecules is relatively small. This consequently allows these dipole-dipole interactions to be manipulated with external DC electric fields and AC microwave fields [20] leading to tunable interactions between the molecules. DeMille's proposal for a molecular quantum computer [19] consists of polar molecules arranged in a 1D optical lattice, with one molecule per lattice site. An external electric field is applied to the 1D array, consisting of a constant bias field on top of a linear field gradient, which allows the molecules to be addressed individually as they all experience a different Stark shift. The two qubit states are the electric dipole moments (EDMs) of the molecules, aligned either with or against the external electric field. The qubits are coupled together through the dipole-dipole interaction, and transitions between qubit states to carry out quantum computing operations can be driven by microwave pulses.

There are many quantum aspects of ultracold many-body systems that would be interesting to simulate, but that are intractable on classical computers. Instead of simulating these quantum processes classically, ultracold molecules can be used to directly simulate a complex quantum system by engineering the system of molecules to have a Hamiltonian that matches that of an interesting, but not fully understood many-body system. The long range dipole-dipole interactions between polar molecules would provide a novel and versatile tool set to carry out these simulations. Micheli *et al*. [17] have shown how such a toolbox of ultracold polar molecules in an optical lattice could be used in practice to carry out quantum simulations.

One of the major requirements to carry out quantum computation and quantum simulations using ultracold polar molecules, is that the molecules must be confined to a one or two-dimensional array such as an optical lattice [21]. This was recently achieved for the first time in both a 2D and 3D optical lattice [22], starting to pave the way for some of these quantum information and simulation applications to be realised.

1.1.3 Other applications

As well as the two major applications of cold molecules discussed above, there are several other areas where cold molecules can be of great importance. One of these is in ultracold chemistry, where cold molecules will allow accurate studies of chemical reactions at low temperatures, and the role that quantum effects have on the reaction rates. Another, related direction is to manipulate collisions of molecules at ultracold temperatures using external electromagnetic fields, thereby achieving control over chemical reactions [23].

Nanodeposition of molecules onto solid materials has the potential to greatly advance lithographic processes used to produce transistors on chips. Ordinarily, optical lithography is carried out using a laser beam to transfer a pattern onto a light sensitive chemical on the substrate, with subsequent chemical treatments then depositing the desired material onto the pattern. The size of the features that can be created in this way is limited by the diffraction of the laser light. Atomic/molecular lithography would deposit the atoms/molecules directly on the surface from a cold beam, thereby removing diffraction as a limiting factor and allowing the creation of much smaller features. Atomic lithography is already well established, for example with laser cooled Cr [24] or Cs [25] atoms. Extending this to cold molecules would allow for much greater versatility.

1.2 Overview of molecular cooling and trapping

The field of molecular cooling and trapping is fast expanding, with new techniques continually being developed and improved. The complicated energy level structure that makes cold molecules so appealing for new research, at the same time means that the cooling and trapping techniques that are already well established for atoms (predominantly laser cooling) are often not readily applicable to molecules. Laser cooling in atoms works because they contain a closed multilevel system that allows many absorption and emissions cycles, and hence momentum transfer from a large number of photons from the laser beam. Molecules on the other hand, have no such closed system as emission can occur to a number of different vibrational states. This makes laser cooling molecules extremely difficult. Nonetheless, several methods are currently in existence that are able to cool molecules and, as described in [26], they can be divided into two categories: direct and indirect. The vast majority of these molecular cooling and trapping techniques are currently carried out with diatomic molecules, however polyatomic molecules are also increasingly being investigated (see for example [27] and [5]), particularly as they may allow the study of interesting phenomena in cold chemistry [28].

1.2.1 Cooling methods

Indirect cooling

Indirect methods involve creating cold molecules from their constituent atoms that have already been cooled and trapped using laser cooling– the main advantage here being that the methods for the cooling and trapping of atoms are already very well established and are capable of producing ultracold atoms with relative ease. Simply confining the constituent atoms to a trap will not cause them to form molecules naturally as the density in the traps is too low, resulting in a lack of 3-body collisions that would be required to form molecules. Additionally, forming molecules from 3-body collisions would heat the molecules, which is undesired. As such, there are two main ways to form molecules from ultracold trapped atoms: photoassociation [29] and magnetoassociation [30].

Photoassociation (PA) is achieved by tuning a laser beam to a resonant transition from the ground state of two free atoms to an electronically excited bound state of the molecule. This produces a molecule in a weakly bound, highly excited vibrational state. From this state there are then two possible decay channels: one to the bound electronic ground state of the molecule and another to a pair of free atoms, which are lost from the trap. Decay to a deeply bound molecular state is unlikely, due to unfavourable wavefunction overlap, and this makes it very difficult to form vibrational ground state molecules in this way. Nevertheless, it is possible to find suitable PA transitions with favourable decay channels that allow ground state molecules to be formed. Some examples for these include LiCs created in the rovibrational ground state [31] from a dual Li-Cs MOT, or more recently the formation of RbCs molecules in the rovibrational ground state [32].

Magnetoassociation occurs via a Feshbach resonance, which is a resonance in the collision cross-section that occurs when the energy level of a weakly bound molecular state crosses that of the free atoms. Usually, the bound molecular state and the free-atom state have different magnetic moments and so a magnetic field can be applied to tune the states through the resonance. There is an avoided crossing at the resonance, and by sweeping the magnetic field through this avoided crossing, the free atoms can be turned into bound molecules. Contrary to PA, the weakly bound molecule that is formed after a Feshbach resonance is in the electronic ground state, albeit still in a highly excited vibrational state. Stimulated Raman adiabatic passage (STIRAP) can then be used to coherently transfer the weakly bound Feshbach molecules into the rovibrational ground state, as was done in 2008 with KRb molecules [33] and more recently with RbCs molecules [34, 35]. The difficulty with forming molecules through Feshbach resonances is that a suitable Feshbach resonance has to be found that can be swept with realistically achievable magnetic fields, and these do not necessarily exist for every potential molecule.

Typical temperatures that can be achieved using these indirect cooling methods are hundreds of μK for PA and hundreds of nK for magnetoassociation, mainly because the constituent atoms are already at these temperatures before the molecules are formed. Whilst both of these methods are currently capable of producing colder molecules than direct cooling methods, they are limited to very specific species of molecules– those whose constituent atoms can be laser cooled and trapped, and where suitable Feshbach or PA resonances can be found. For more versatile and widely applicable techniques, we need to look at methods that cool the molecules directly.

Direct cooling

Whilst direct methods are so far not capable of reaching ultracold temperatures, they are much more versatile and have managed to create a wide variety of cold molecular species. Perhaps one of the most versatile direct cooling techniques is through elastic collisions with a cryogenically cooled buffer gas– buffer gas cooling. As this technique is used in this thesis, it is explained in more detail in section 2.1. It typically involves introducing the molecules into a cold cell of (often helium) buffer gas, and allowing the molecules to thermalise to the temperature of the buffer gas through elastic collisions. This was first demonstrated for molecules in 1998 [36], achieving a temperature of 300 mK, and has since been applied to numerous molecules.

Typically, temperatures achieved using buffer gas cooling are in the region of 1 K, depending on the type of buffer gas used, and as such it is generally seen as a first stage cooling technique, with further cooling required to get down to lower temperature regimes. It is, however, extremely versatile and can, in theory, be applied to any molecule as it only relies on elastic collisions with a buffer gas. A cold beam of buffer gas cooled molecules can also be created by having an aperture in one side of the buffer gas cell [37]. Such a buffer gas beam has, for example, been used to measure the electron EDM (see section 1.1.1), or it can be used as the starting point for Stark deceleration (see below).

One of the major techniques for the direct cooling of molecules involves slowing down molecular beams using the Stark shift that molecular energy levels experience in an electric field [38, 39]. These Stark decelerators consist of an inhomogeneous electric field that causes molecules in a low field-seeking state to decelerate. The fields are then switched appropriately to ensure that the molecules in the beam are always travelling up the 'potential hill' and can thus be slowed down. This technique has been used to decelerate a range of polar molecules including CO [38], YbF [40] and OH [39], all of which were decelerated from a supersonic molecular beam. It is also possible to Stark decelerate molecules from a buffer gas cooled molecular beam [41, 42]. As these buffer gas beams have a lower starting velocity compared to room temperature supersonic beams, Stark deceleration to rest is made significantly easier and therefore allows much heavier molecules to be decelerated. Once these molecular beams have been slowed down they can be loaded into a variety of molecular traps, discussed in section 1.2.2.

There has also been considerable progress towards the direct laser cooling of molecules [43], despite the difficulty of doing this compared to atomic laser cooling. The problem with laser cooling molecules is that molecules have multiple vibrational energy levels, with no selection rules governing which vibrational state an electronically excited molecule decays down to. This makes it difficult to find a closed cycling transition that allows enough absorption and emission cycles for the molecules to be laser cooled, as the molecule will readily decay into a different vibrational state that is not addressed with the cooling laser. The branching ratios for the different vibrational levels an excited molecule can decay down to are determined by the Franck-Condon (FC) factors. Therefore, molecules have to be found with favourable FC factors that ensure that the vast majority of the excited molecules decay down to the same vibrational state. Even then, repump lasers are still required for some of the other vibrational levels that a significant fraction might decay down to. Despite these requirements, suitable molecules have been found, and laser cooling was first achieved in 2010 for SrF [43]. A beam of SrF molecules has since also been slowed using laser beams [44], and a beam of CaF molecules has been laser cooled and slowed [45]. Two-dimensional magneto-optical compression of a beam of YO molecules has also been demonstrated [46], and most recently a 3D magneto-optical trap of SrF molecules has been achieved [47]. A number of other molecules have also been earmarked as potential candidates for laser cooling [48, 49]. The ability to laser cool molecules effectively has huge potential for the applications discussed in section 1.1, however it will likely remain applicable to only a very specific set of molecules as they require a very favourable set of transitions.

Collisional cooling, which includes sympathetic and evaporative cooling, can be used to cool the molecules down to the ultracold temperature regime, once they have been confined to a trap (see section 1.2.2). This can only take place once a sufficiently high density of molecules has been trapped and elastic collisions dominate over inelastic collisions. Evaporative cooling involves successively removing molecules with the highest energies, leaving the lower energy molecules, and thereby reducing the overall temperature of the sample. Elastic collisions constantly re-thermalise the sample, ensuring that the high energy tail of the distribution gets repopulated and can continue to be removed. Until recently, this technique had only been applied to trapped atomic distributions, but it has now been demonstrated for OH radicals in a magnetic trap [50]. Sympathetic cooling would cool a sample of trapped molecules through elastic collisions with an ultracold sample of trapped atoms. The ultracold atoms would essentially act as a refrigerant for the molecules, bringing the molecular sample down to the ultracold temperature regime. A number of schemes have been proposed to achieve this [51, 52], although it has not been realised to date.

1.2.2 Trapping methods

In order to study molecules over longer timescales, they can be held in traps made of electric and/or magnetic fields. Electrostatic traps make use of the Stark effect that molecular energy levels experience in electric fields. The Stark shift may be positive or negative, corresponding to low field-seeking and high field-seeking states respectively. If the molecule is in a low field-seeking state, a static trapping arrangement with an electric field minimum at the centre confines the molecules. This was first demonstrated by Gerard Meijer's group for ND_3 [27], loaded from a Stark decelerated molecular beam. There are, however, a number of disadvantages to trapping low field-seeking molecules: the ground state of every molecule is high-field seeking and has the largest Stark shift; for heavy molecules, all the low-lying states are high-field seeking; and non-adiabatic transitions can put molecules into high-field seeking states (discussed further in section 2.2.1). As a result, it is often advantageous to be able to trap high field-seeking molecules. This can be done in an AC electric trap, also first achieved by Gerard Meijer's group in 2005 [53]. This trap works by creating an electric field that has a maximum in one direction and a minimum in the other and then rapidly switching the two directions so that a high field-seeking molecule is trapped. Another proposed method for trapping ground state, high field-seeking molecules is to confine them to an anti-node maximum of a standing wave microwave field [54].

Magnetic trapping works on the same principle as electrostatic trapping, but this time makes use of the Zeeman shift of the energy levels rather than the Stark shift. Creating a magnetic field arrangement with a minimum at the centre would again allow low field-seeking molecules to be trapped. As magnetic trapping is used throughout this thesis, the principles behind it are explained in more detail in section 2.2. The first demonstration of the magnetic trapping of molecules was in 1998 by J. Doyle's group [36], where buffer gas cooled molecules were trapped in a magnetic field generated by superconducting coils. This same method has since been used to trap a variety of molecules, as shown in table 2.1. There are also other ways of loading molecules into magnetic traps, for example from a Stark decelerated beam [55].

In addition to the two main magnetic and electric trapping methods, magneto-optical trapping of molecules has also been realised in the last two years, first in 2D for YO [46] and then in 3D [47, 56], loaded from a laser cooled buffer gas beam of SrF. The most recent of these achieved a trap lifetime of 136 ms [56]. Since its first demonstration in 1987 [57], magneto-optical trapping has become one of the most widely used trapping methods in atomic physics, but had previously been limited to atoms for the same reasons as described above for laser cooling. Demonstrating it for molecules therefore marks a significant step towards realising some of the applications in section 1.1.

The ultimate aim in order to realise many of the applications described in section 1.1 is to confine molecules to an optical lattice so that they can be used for quantum simulations and computations. The interaction of atoms with light causes a 'light shift' of the atom's ground state by an amount equal to $\hbar\Omega^2/(4\delta)$, where Ω^2 is proportional to the light intensity and δ gives the detuning of the light from resonance. If the laser light is detuned below resonance then the light shift is negative, with the highest intensity light causing the largest negative light shift. This results in a force on the atoms attracting them towards regions of high intensity, called the optical dipole force. This can be used to confine atoms to regions of high intensity laser light, such as in the tight focus of a Gaussian laser beam in optical dipole traps [58]. First demonstrated in 1987 [59], an optical lattice generates a standing wave pattern of light using counter-propagating laser beams. This creates an 'egg box' potential of maxima (and minima) of the light field, to which the atoms can be confined through the optical dipole force. By having this array of 'traps', atoms can be confined individually, allowing unprecedented levels of control over each atom [1]. It is therefore also the ultimate aim to be able to readily confine molecules in the same way, such as first achieved for polar molecules in 2D in 2011 [60] and in 3D in 2012 [22]. The molecules trapped in these cases were created from ultracold atoms through magnetoassociation and STIRAP, as in [33], with lifetimes in the 3D optical lattice up to 25 s [22]. Reaching the ultracold temperature regime using direct cooling methods would provide an even more versatile way of loading molecules into optical lattices.

1.2.3 More novel cooling and trapping methods

As well as the more universally used cooling and trapping techniques for molecules described above, there have been a few more novel approaches that are highlighted here. The first of these is by Lu *et al.* [61], where a cold, buffer gas cooled beam of CaF molecules is loaded directly into a magnetic trap without any further laser cooling or Stark deceleration of the beam. The beam is produced from a two-stage buffer gas cell [62], with a velocity as low as 30 m/s. A superconducting magnetic trap is placed downstream that creates a quadrupole trapping potential for low field-seeking (LFS) molecules. As the LFS molecules approach the trap, they will experience a potential hill that they have to get over to make it into the trap region, which slows the molecules. As the molecules reach the maximum of this potential hill, they are optically pumped into a high field-seeking (HFS) state. This results in them experiencing a new potential hill on the approach to the trap centre where the magnetic field is zero, slowing the molecules further. As they reach the trap centre, the molecules are pumped back into the LFS state, leaving them trapped with a lifetime of 500 ms [61]. This optical pumping cycle allows the molecules to be slowed using only the magnetic field, and therefore makes it a fairly versatile method that could be applied to other species.

Centrifugal forces have also been used recently to slow molecules [63]. A beam of molecules is guided by a spiral shaped electric quadrupole guide mounted on a disk. The molecules enter the guide at the outer edge of the disk and the guide finishes at the centre of the disk, where the molecules emerge. By spinning the quadrupole guide, the molecules experience a centrifugal force barrier as they approach the centre of the disk, slowing them down. This allows the authors to slow beams of polar molecules from 200 m/s to 15 m/s [63]. Another method that has been used is to photo-dissociate a beam of NO_2 molecules so that the NO fragment recoils with a velocity that is exactly equal and opposite to the velocity of the beam, leaving the NO molecules stationary [64]. These more novel trapping methods have all been developed in the last few years. Their full potential has therefore not yet been fully investigated, and they may well become more widely used in the future.

1.3 A permanent magnet trap for buffer gas cooled atoms and molecules

The work presented in this thesis is aimed at realising a simple, versatile permanent magnet trap for buffer gas cooled molecules, from which further cooling can then take place to reach the ultracold temperature regime. To investigate the feasibility of the trapping arrangements in this thesis, atoms are used throughout as they are easier to produce and detect. Once it has been established that the trapping arrangement works for atoms, it is then a fairly simple step to use the same arrangement for molecules, since the methods we use are applicable to both atoms and molecules.

The experiments in this thesis share some similarities with traditional buffer gas magnetic trapping experiments [36], and some major differences. We use permanent magnets rather than superconducting coils as this leads to a much simpler experimental set-up, and the geometry and arrangement of the magnets can be easily adapted in order to suit the requirements of the trap. However, permanent magnets do not provide as high a trapping field as the superconducting coils. The traps used in this thesis also have a much more open geometry, rather than the traditional, fully enclosed buffer gas cell, to allow the buffer gas to leave the trap region more quickly once the atoms have thermalised.

Magnetic trapping of neutral atoms goes back a long way, with the first such experiment carried out in 1985 [65], where laser cooled neutral atoms were confined in a quadrupole trap made up of two anti-Helmholtz coils (see figure 2.3 for an example of such a set-up). Magnetic trapping of *buffer gas cooled* neutral atoms has since been achieved in the Doyle group [66], where superconducting anti-Helmholtz coils were used to create the quadrupole trapping field. Trapping atoms using *permanent* magnets is not a new venture either [67], however the atoms here had to be loaded into the trap from an optical molasses set-up. More recently, photodissociation of a molecular beam of Br_2 molecules has been used to trap Br atoms using permanent magnets [68]. Molecules have also been loaded into a permanent magnet trap from a Stark decelerated beam [55]. Both of the atomic species used in this thesis, dysprosium and lithium, have previously been buffer gas cooled and magnetically trapped using superconducting coils [69, 70]. However, *buffer gas* loading of a *permanent* magnet trap has not been achieved before.

This thesis begins with an overview of the theory behind buffer gas cooling and magnetic trapping, presented in chapter 2. Chapter 3 then discusses different permanent magnet arrangements that were considered to provide the trapping field. Two different methods for trapping buffer gas cooled atoms are subsequently presented. The first of these involves directly trapping the atoms from a buffer gas cooled cloud by placing two permanent magnets around the region where the cloud is formed. This is presented in chapters 4 and 5 for Li and Dy respectively. Knowledge gained from these experiments motivated the second method of trapping Dy atoms from a buffer gas cooled beam, presented in chapter 6. The results are summarised in chapter 7, where some future improvements to the set-up are also discussed.

CHAPTER 2

BACKGROUND & THEORY

2.1 Buffer gas cooling

2.1.1 The principle

Buffer gas cooling works by allowing molecules to thermalise with a cryogenically cooled helium buffer gas at 4 K through elastic collisions. This was first demonstrated, for CaH molecules, by J. Doyle's group [71, 36], and has now been used to cool a number of molecular species (see table 2.1). Figure 2.1 shows a schematic for a typical buffer gas experiment. Helium is used to cool the cold plate (and hence cold cell) of a cryocooler to 4 K. The helium buffer gas is then allowed to flow, either continuously or pulsed, into the cold cell through the inlet, and thermalises to 4 K through collisions with the cold plate and the walls of the cell. The desired molecules are then loaded into the cell and, through elastic collisions with the helium buffer gas, thermalise down to 4 K. It is not possible to use the cold plate to cool the molecules directly as the vapour pressure of molecules at low temperatures is essentially zero, meaning that the molecules would stick to the surface of the cold plate [72].

A common method for introducing the desired molecules inside the cell is by laser ablation of a suitable target (see section 2.1.2), which releases the molecules directly into the buffer gas. Whilst this is probably the most widely used method, there are several other methods that have been used to load molecules into buffer gas cells. The two main alternative methods are beam injection and capillary filling. Beam injection involves creating an aperture in one side of the buffer gas cell which allows a room temperature beam of the molecules to enter the cell and thermalise with the buffer gas, as used for NH molecules in [73]. This requires finely balancing the helium density inside the cell and the size of the aperture to make sure

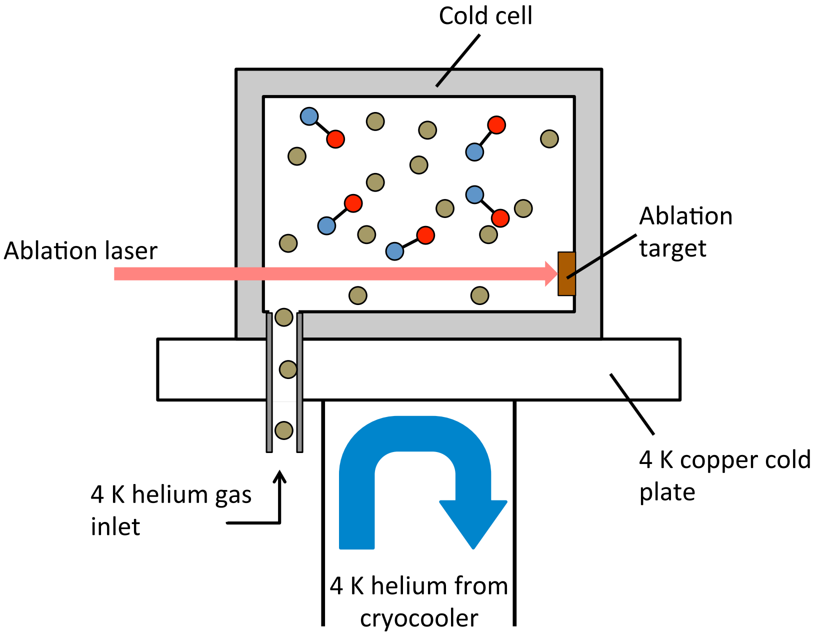

Figure 2.1: Diagram of a typical buffer gas cooling set-up. Cold helium is used to cool the cold plate, and hence the cold cell, to 4 K. Helium gas is then pumped into the cold cell and allowed to thermalise to the temperature of the cell walls. Laser ablation of the solid target creates the desired molecules inside the cell. These then thermalise with the cold helium gas through elastic collisions.

that enough molecules make it into the cell with a high enough helium density for them to thermalise, without the outflow of helium from the same aperture pushing the molecules back out of the cell. These problems are avoided in capillary filling, where a capillary tube connects a room temperature molecular gas reservoir with the cold buffer gas cell, with the walls of the capillary tube guiding the molecules into the cell. This was first demonstrated for CO molecules in [74] and has since been demonstrated for a number of molecules, including O_2 [75] and ND_3 [76]. Capillary filling has its own drawbacks, predominantly due to the fact that the capillary has to be kept at a temperature above the freezing point of the molecules. This means that a significant heat load might be deposited onto the buffer gas cell, raising its temperature, as well as the possibility of the end of the capillary tube becoming cold enough to freeze the molecules before they reach the buffer gas. As suitable solid targets were readily available for both of the atomic species buffer gas cooled in this thesis, laser ablation was used for both of these.

The main advantage of buffer gas cooling over other molecular cooling techniques is that it is incredibly versatile, as it only relies on elastic collisions between the molecules of interest and the buffer gas. Therefore, in theory, it can be applied to many different species of molecules. As well as this, it is capable of producing higher density molecular clouds than Stark deceleration

($10^9\,\mathrm{cm}^{-3}$ as opposed to $10^6\,\mathrm{cm}^{-3}$) [26]. Once the molecules have been loaded into the buffer gas cell, they can either be studied in detail by allowing them to diffuse to the cell walls; they can be extracted from the cell to form molecular beams; or they can be magnetically trapped. These are discussed in more detail in sections 2.1.2, 2.1.3 and 2.2 respectively.

2.1.2 Gas dynamics inside a buffer gas cell

Laser ablation

Laser ablation is a process that involves removing material (in this case the particles of interest) from a solid target through absorption of a laser pulse. A high enough intensity laser beam will cause the formation of a plume of material from the target [77], which is the desired effect to introduce molecules into the buffer gas cell. The higher the intensity of the laser beam, the more molecules will be released from the target, however the precise number of molecules released during ablation is very hard to determine theoretically due to the complicated interaction between the solid and the laser light. Even though the process has a highly complicated nature and is poorly understood, some estimates can be made for the number of molecules released from the target during the ablation process by considering the absorption of laser light by a target material [78]. In order for molecules to be released from the target, the target material has to be vaporised. The energy required to convert a mass m of material into a vapour can be given by [78],

$$E = m(C_s(T_m - T) + C_L(T_v - T_m) + L_f + L_v), \tag{2.1}$$

where C_s and C_L are the solid and liquid specific heat capacities, T_m and T_v are the melting and boiling points of the material, L_f and L_v are the latent heats of fusion and vaporisation and T is the initial temperature of the material. Assuming that $C_s \approx C_L = C$, the mass of material released from a target is given by

$$m = \frac{E}{C(T_v - T) + L_f + L_v}. \tag{2.2}$$

The energy that is absorbed by the material, E, can be written as $E = E_p(1 - R)$, where E_p is the energy of the laser pulse incident on the target and R is the reflectivity of the target material. Hence, the number of atoms/molecules released from the ablation target during a single laser pulse is given by

$$n = \frac{E_p(1-R)}{m_{mol}C(T_v - T) + L_f + L_v}, \tag{2.3}$$

where m_{mol} is the mass of an individual atom/molecule. This gives an upper limit to the number of atoms/molecules that can be produced through ablation as it assumes that all of the absorbed energy goes towards bringing the particles to boiling point, whereas in reality part of the energy will heat the surrounding target material. It also assumes constant heat capacities and ignores the fact that the mass released from the target will consist of other ablation by-products as well as the desired atoms/molecules. The ablation pulses used in the experiments in this thesis have a pulse width of 5-7 ns and typically have an energy of 10 mJ (see section 4.2.4). Taking Li as an example, inserting the appropriate constants and assuming an ablation pulse energy of 10 mJ, equation 2.3 puts an upper limit of 10^{16} on the number of Li atoms that can be produced from a single ablation pulse.

Collisions and thermalisation

After the ablation, the molecules must thermalise with the helium buffer gas before they diffuse to the cell walls. This occurs through elastic collisions [1] between the molecules and the helium which cools both the translational and rotational energy of the molecules [80]. Thermalisation of the translational temperature can be modelled using a hard sphere model, where m denotes the mass of the buffer gas atom and M the mass of the molecule being cooled. Using energy and momentum conservation (the conditions for an elastic collision), the difference in temperature of a molecule before and after a collision with a buffer gas atom is given by [80],

$$\Delta T = (T_I - T_B)/\kappa, \tag{2.4}$$

where T_B and T_I denote the temperature of the buffer gas and the initial temperature of the molecule respectively and $\kappa = (M+m)^2/(2Mm)$. Rewriting this as a differential equation,

$$\frac{\mathrm{d}T_i}{\mathrm{d}i} = -(T_i - T_B)/\kappa, \tag{2.5}$$

---

[1] An elastic collision is a collision in which the total kinetic energy of the particles involved does not change, however their individual kinetic energies can change. This is in contrast to inelastic collisions where the internal energy of the particles, and hence total kinetic energy, can change. [79]

where T_i is the temperature of the molecule after i collisions with the buffer gas. The solution of the above equation is as follows,

$$T_i/T_B = (T_I/T_B - 1)e^{-i/\kappa} + 1. \tag{2.6}$$

Doing a rough order of magnitude calculation assuming $T_I \approx 10\,000\,\mathrm{K}$, $T_B \approx 1\,\mathrm{K}$ and $M/m \approx 50$, around 100 collisions are necessary in order to cool the molecules to within 30% of the buffer gas temperature.

In order to study the molecules at low temperatures, or extract a cold beam from the cell, it is essential that the molecules thermalise with the helium buffer gas before they reach the cell walls. To ensure this is the case, the buffer gas density must be high enough. The rate of elastic collisions between the molecules and the buffer gas depends on the elastic collision cross section. Assuming a typical elastic collision cross section, σ, between molecules and the buffer gas of around $10^{-14}\,\mathrm{cm}^2$, and a mean free path of the molecules of $\lambda = 1/(\sqrt{2}n\sigma)$, the minimum buffer gas density required for molecules to thermalise before they reach the cell walls of a typical 1 cm diameter cell is of the order of $10^{16}\,\mathrm{cm}^{-3}$. This also sets a lower limit on the temperature of the buffer gas as the helium will have a certain saturated vapour density at low temperatures. In order to achieve a density of $10^{16}\,\mathrm{cm}^{-3}$, the helium can be no colder than 700 mK [80].

There have been several studies of the formation, diffusion and thermalisation of molecules produced by laser ablation of a target inside a buffer gas cell. One is the work of Skoff *et al.* [81], who studied the dynamics of YbF molecules using absorption imaging and spectroscopy. The authors found that the YbF molecular plume expanded ballistically after being formed at the target and that this expansion was subsequently arrested by the helium buffer gas. The higher the helium density inside the cell, the closer to the ablation target the expansion was arrested, with the entire cell being filled with molecules if the helium density is low enough. The subsequent diffusion of the molecules through the helium was modeled and experimentally observed to find the diffusion cross sections for YbF in He at several temperatures. They also found that both the rotational and translation temperature of the molecules thermalise with the buffer gas on a timescale shorter than 50 μs when the helium density is $10^{22}\,\mathrm{m}^{-3}$. However, the initial ballistic expansion leads to some heating of the buffer gas, which was subsequently observed to cool back to the temperature of the cell walls on a longer timescale. Another study is that of Lu and Weinstein [82], who investigated the dynamics and collisions of TiO molecules in a helium buffer gas at 5 K. They found that the decay lifetime of molecules inside the cell initially increases linearly with increasing helium density, as expected if the lifetime is dominated by diffusion to the cell walls, where the molecules are lost as they condense onto the cold surface. However, above a helium density of about $10^{22}\,\mathrm{m}^{-3}$, the lifetime actually starts to decrease with increasing helium density as a result of inelastic collisions which cause the molecules to decay into the lowest vibrational state, which the authors do not detect. A similar pattern was found in [83], where gold, silver, lithium and rubidium atoms were studied in several different buffer gases: helium, nitrogen, neon and argon. However, in this case the decrease in lifetime above buffer gas densities of approximately $10^{24}\,\mathrm{m}^{-3}$ was attributed to dimer formation, atom loss on clusters or atom loss on impurities in the buffer gas, rather than inelastic collisions. Additionally, the authors found that the decay time did not depend on which type of buffer gas was used.

The calculations in this section have provided an indication of the typical buffer gas densities required inside a buffer gas cell and have given an upper limit to the number of atoms/molecules that can be produced from an ablation pulse. As the experiments in this thesis have an open geometry, rather than the typical closed buffer gas cells, the helium will dissipate over time rather than remaining at a constant density. This has to be taken into account in our experiments, and this helium dissipation is studied in more detail throughout the experimental chapters 4, 5 and 6.

2.1.3 Beams from a buffer gas cell

Another application of buffer gas cooling is the creation of atomic or molecular beams. This is achieved by slightly modifying a standard buffer gas cell to include a small aperture in one face, as shown in figure 2.2. Molecules are then introduced to the buffer gas cell as normal, and after they have thermalised with the buffer gas, both the molecules and the buffer gas are allowed to escape the cell through the aperture. Usually, the aim is to produce a high-flux beam of cold, slow-moving molecules. Hence, care has to be taken to ensure that the helium density inside the cell is high enough for the molecules to efficiently thermalise, but low enough to prevent the molecules exiting the cell at too high a speed [84]. The density of the helium inside the cell is kept constant by allowing a continuous flow of helium into the cell to replace that lost through the aperture [80]. In order to ensure that only the desired molecules remain

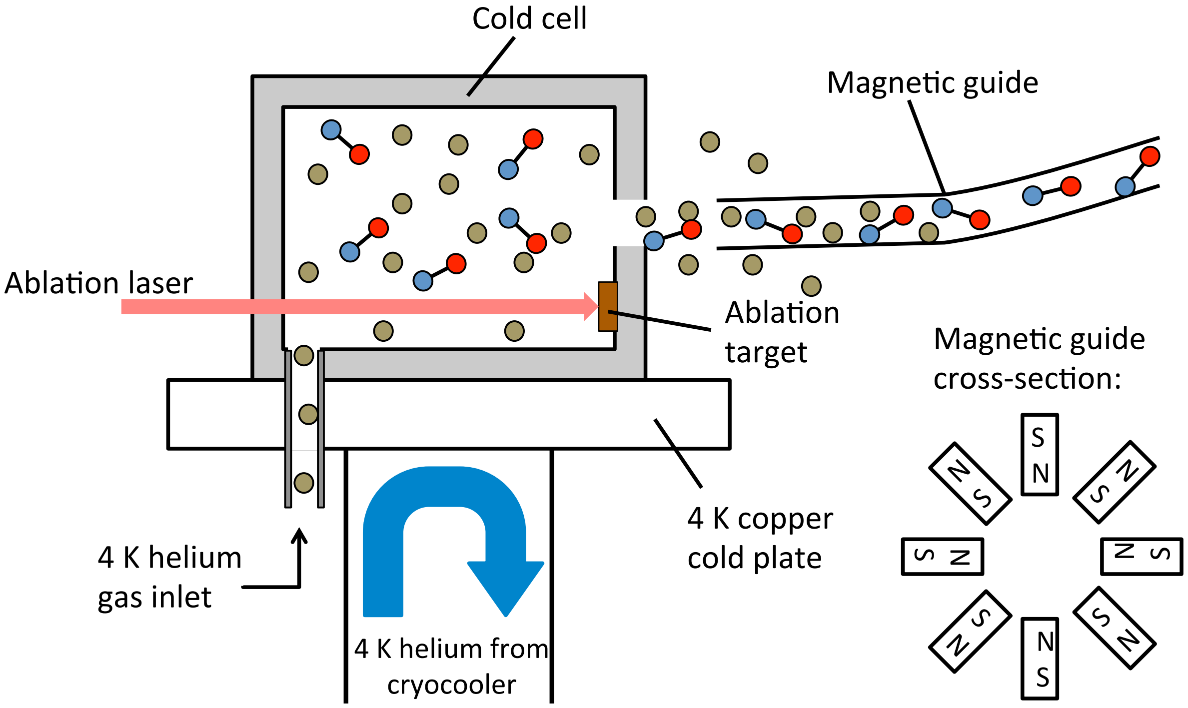

Figure 2.2: Diagram of a typical buffer gas beam set-up, showing how the aperture in the cell allows the molecules and buffer gas to escape. The flow of helium through the inlet is balanced with that escaping from the aperture to keep a constant helium density inside the cell. The magnetic guide placed near the aperture selects only the desired molecules from the beam.

in the beam, a curved magnetic guide can be placed at the exit of the cell [75]. A fraction of the molecules follows the guide, whereas the non-magnetic helium atoms do not. The magnetic guide only works for paramagnetic molecules, however a source of cold polar molecules can also be guided using an electrostatic guide. This has been demonstrated in [85] with a room temperature effusive source of ammonia molecules, and later from a cryogenic buffer gas source of ND_3 molecules in [86] and [76].

2.2 Magnetic trapping

The aim of this work is to demonstrate magnetic trapping of buffer-gas cooled atoms and molecules. The phenomenon that makes magnetic trapping of both atoms and molecules possible is the Zeeman shift of the energy levels in a magnetic field. This shift arises due to the interaction between the magnetic field and the magnetic dipole moment of the atoms or molecules, and leads to an energy shift that, in low fields (where the Zeeman splitting is much less than the splitting between neighbouring energy levels), is proportional to the magnitude of the magnetic field,

$$\Delta E(x, y, z) = g_F\, m_F\, \mu_B\, B(x, y, z), \tag{2.7}$$

where g_F is the Lande g-factor, m_F denotes which magnetic sub-level the atom or molecule is in and μ_B is the Bohr magneton. The $g_F\, m_F\, \mu_B$ factor is called the magnetic moment of the atom or molecule in this particular quantum state. It is also made explicitly clear that the magnetic field in a typical magnetic trap is a function of x, y and z. Since each energy level, labelled by F, is split into m_F sub-levels from $-F$ to F, it is clear from equation 2.7 that the energy of some m_F levels will increase with magnetic field and the energy of others will decrease. Particles in a state with positive magnetic moment are *low* field-seeking as their energy increases with magnetic field and hence their energy will be lower in a weaker field. For the same reason, particles in a state with negative magnetic moment are *high* field-seeking.

For higher fields, where the Zeeman splitting becomes comparable to the difference in energy between neighbouring levels, this linear dependence no longer holds. An example of this can be seen in figure 4.2, which shows the energy level dependence with magnetic field for the ground state of lithium. In this regime, some states turn over from low to high-field seeking states, or vice versa, due to the mixing of states of different F but same m$_F$. Only the 'stretched' states (those with biggest F and biggest |m$_F$|) remain unaffected and continue to shift linearly as they have no other states to mix with. When the Zeeman splitting becomes much larger than the energy difference between neighbouring levels – the so called strong field regime – the Zeeman shift again becomes linear, however some states will have turned from high to low-field seeking or vice versa (see figure 4.2). In this strong field regime the energy shift is now given by [87],

$$\Delta E = (m_J g_J - m_I g_I)\mu_B B + A\, m_I m_J, \tag{2.8}$$

where g_J and g_I are the g-factors corresponding to the total electronic angular momentum and to the nuclear spin angular momentum, and A is the hyperfine constant. In the strong field regime the energy shift no longer depends on which F level the atom is in, but instead on the I, J, m_I and m_J sub-levels. These are now the better quantum numbers to describe each energy level, as illustrated in figure 4.2.

From equation 2.7, the force on the atom or molecule that arises due to the Zeeman shift is given by,

$$\text{Force} = -g_F\, m_F\, \mu_B\, \nabla B(x, y, z). \tag{2.9}$$

In order to confine atoms or molecules to a trap, their kinetic energy must be converted to potential energy. This is achieved by designing a trap that has a magnetic field minimum at the centre, with the field increasing in all directions away from the centre. One such design that is commonly used is the anti-Helmholtz coil arrangement, shown in figure 2.3. This has a magnetic field zero at the centre, with the magnitude of the field near the centre of the trap increasing linearly in all directions away from the centre as [88],

$$B = A\sqrt{\rho^2 + 4z^2}, \tag{2.10}$$

where A is the field gradient and $\rho = x^2 + y^2$. Particles in a weak field-seeking state can then be trapped, as their internal energy levels (and hence potential energy) will increase away from the trap centre, thus removing kinetic energy. From equation 2.9, it can be seen that this positive magnetic field gradient results in a force on the particles back towards the centre of the trap. For the particles to remain in the trap, their initial kinetic energy must be low enough that they do not go beyond the confines of the trap before they are brought to a halt. Particles in a strong field-seeking state are automatically lost from the trap. [89]

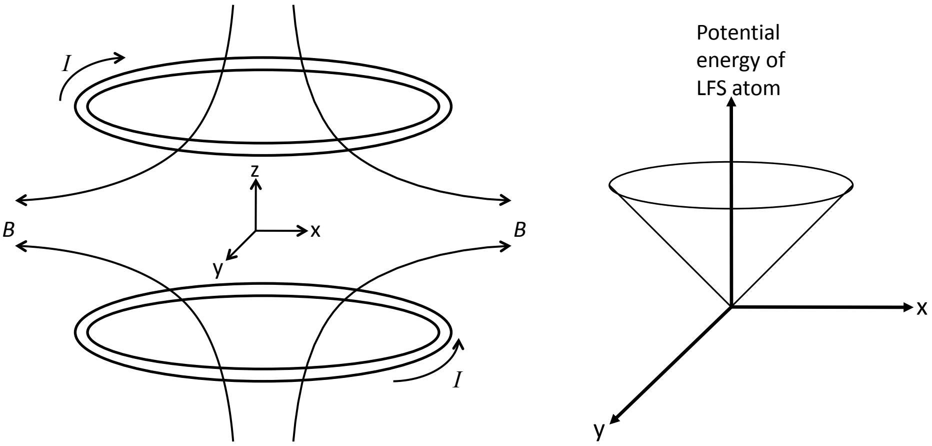

(a) (b)

Figure 2.3: (a) Anti-Helmholtz coil arrangement for a magnetic trap. With the currents in the two coils in opposite directions a magnetic field zero is created at the centre of the trap. Near the centre of the trap the magnitude of the magnetic field increases linearly with distance from the centre. (b) The potential energy surface of a low-field seeking atom or molecule in the trap.

In order for the kinetic energy of the molecules to be low enough that they can be trapped in this way, they must first be cooled. This must be done within the confines of the trap, which is where buffer gas cooling comes in. The first demonstration of this by J. Doyle's group [36] used superconducting coils in an anti-Helmholtz configuration similar to that in figure 2.3(a) around the buffer gas cell. This led to a magnetic trap 3 T deep, which, for low-field seeking states of 1 μ_B magnetic moment, corresponds to a 1.3 K trap depth. This is more than deep enough to trap a significant fraction of the buffer gas cooled molecules. Unfortunately however, the presence of the buffer gas brings with it the added problem of trap loss through inelastic collisions between the molecules and the buffer gas, which is explained in detail in section 2.2.1.

2.2.1 Trap loss

Spin-flip collisions

There are two main channels through which molecules can be lost from the trap. These are in addition to the fact that some molecules are simply travelling too fast to be confined. The first of these loss mechanisms is inelastic collisions between the molecules and the buffer gas. Whilst elastic collisions are essential to ensure the molecules are fully cooled, inelastic collisions in which the internal energy changes can result in the molecules ‘spin flipping’ into high field-seeking states, thereby losing them from the trap. This is possible because a low field-seeking molecule is in a higher energy state than a high field-seeking molecule, and therefore a collision between the molecule and the buffer gas can result in the molecule ‘relaxing’ into a high field-seeking state.

In atom-atom as well as atom-molecule collisions these spin flip transitions can occur through three major channels: spin-exchange, dipolar relaxation and interaction anisotropy [80]. Spin exchange involves the direct trading of spins between two collision partners, often as a result of the interaction between the electron spin of one atom with the nuclear spin of the other [90, 91]. However this is not an issue for collisions with ^{4}He as it has zero electronic and nuclear spin. Dipolar relaxation only occurs between two particles with magnetic moments– the spin-spin interaction between the magnetic moments can then lead to angular momentum exchange between the electrons and the relative motion of the particles [80]. This is again not applicable to helium atoms, where the long-range interaction is instead mediated by the van der Waals potential. Interaction anisotropy occurs as a result of the helium atom distorting the spatial charge cloud around the trapped atom. For atoms with a non-zero orbital angular momentum (i.e. non S-state atoms), a collision with a helium atom can induce an admixture of different orbital angular momentum, M_L, states. The spin-orbit interaction in the atom then leads to an admixture of different M_S states, and can therefore result in a spin-flip transition. On the other hand, atoms that are in the $M_L = 0$ orbital angular momentum state have a symmetric angular charge distribution, and hence require the helium atom to first mix in a state with higher orbital angular momentum to induce a spin-flip transition. This is very rare due to the large energy difference between the $M_L = 0$ state and higher orbital angular momentum states [80] and therefore S-state atoms are much less susceptible to inelastic collisions than non S-state atoms.

The three mechanisms described above apply as much to atom-atom collisions as atom-molecule collisions, however there is an additional mechanism that allows a helium atom to cause a diatomic molecule to spin-flip into a high field-seeking state. This is the interaction between the helium atom and the molecule's rotational wavefunction. Much in the same way that an atom's orbital angular momentum determines the angular charge distribution around the atom, the rotational state of a molecule determines the lab-frame charge distribution around the molecule. In the molecule's rest frame, the charge distribution is likely to be highly asymmetrical given that the two constituent atoms are nearly always different. However in the lab frame, the rotation of the molecules has to be taken into account, which results in this asymmetry being averaged out so that the helium atom essentially does not see it. In the rotational ground state ($R = 0$), this rotation is spherically symmetric and therefore leads to a symmetric charge distribution analogous to that of an S-state atom. This makes it difficult for a helium atom to exert a torque on the molecule. Higher rotational states, on the other hand, are not spherically symmetric, leading to asymmetric charge distributions around the molecule in the lab frame. A collision with a helium atom can then easily couple to the rotational angular momentum of the molecule by exerting a torque on it. There is then an interaction internal to the molecule that couples the spin to the rotational wavefunction, the spin-rotation interaction [92]. These two mechanisms can then combine to cause a spin-flip of the molecule. Due to the symmetry of the molecule's charge distribution in the lab frame, spin-flip collisions of rotational ground state molecules are much less likely than molecules in a higher R state, and therefore makes it desirable to trap molecules in the rotational ground state.

These loss processes mean that it is vital to ensure that the rate of elastic collisions between the molecules and the buffer gas is much higher than the rate of inelastic collisions. As well as this, to increase trapping lifetimes, it is important to remove the buffer gas from the cell after thermalisation before the molecules have a chance to undergo spin-flip collisions. This can be done by rapidly pumping out the helium through a hole in the side of the cell [93]. If this is achieved successfully then trap lifetimes are likely to be limited by effects such as Majorana transitions (see below) or molecule-molecule collisions [80].

Majorana transitions

The second major loss mechanism in magnetic traps is through Majorana transitions near the field minimum, which can flip the spin of the atom into a high field-seeking state. Near the centre of the trap, where the field is very close to zero, the energy difference between the low and high field-seeking states is very small. If, in addition to this, the field is changing rapidly with position, this can lead to a non-adiabatic transition from the low to the high field-seeking state [94, 95]. These transitions are often called Majorana transitions. The condition for the atom to pass through the centre of the trap adiabatically is [89],

$$\frac{\mu B}{\hbar} \gg \frac{1}{B}\frac{dB}{dt}, \tag{2.11}$$

where μ is the magnetic moment of the atom and B is the magnetic field. If the fractional rate of change of the magnetic field becomes comparable to the Zeeman splitting between the low and high field-seeking states then Majorana transitions become likely. As a result, it might be advantageous to have a non-zero trap minimum so that there is a much larger energy difference between the low and high field-seeking states, and hence these Majorana transitions are suppressed. This has been realised in several different magnetic trapping configurations such as the Ioffe-Pritchard [96] and TOP [97] traps.

2.2.2 Molecular species trapped to date

Several different species of molecules have now been cooled and trapped using the technique of buffer gas loading, details of which are summarised in table 2.1. This table shows only those

Species	Temperature (mK)	Trap Lifetime (s)	Reference
CaH	400	0.5	[36]
CrH	650	0.12	[98]
MnH	650	0.18	[98]
ND*	550	0.5	[72]
NH*	500	20	[99]

Table 2.1: The different species of molecules that have been buffer gas cooled and magnetically trapped, showing the temperature and trap lifetimes (1/e lifetimes) achieved. *ND and NH were loaded into the buffer gas cell using a molecular beam-loading technique rather than laser ablation (see the references for details).

molecules that have been cooled *and* trapped, however there are many additional species of molecules that have been cooled using a buffer gas. Most of the trap lifetimes so far achieved are of the order of 500 ms and are limited by the loss channels described in section 2.2.1. More recently however, trapping of NH has been achieved with lifetimes of more than 20 s [99]. This was accomplished by creating an aperture at the back of the cold cell and sending in a pulse of helium rather than a steady flow as was the case in previous experiments. This results in a rapid decrease in the helium density after the molecules have been loaded into the trap, and hence fewer spin-flip collisions leading to the much longer lifetimes that they observe [99]. In addition to molecules, buffer gas loading has also been used to magnetically trap atoms, for example in [66, 100]. All these species were trapped in a superconducting magnet arrangement, with trap depths in the region of 3 T, as opposed to the permanent magnets used in this thesis.

CHAPTER 3

MODELLING PARTICLES IN THE TRAP

Two different magnet arrangements were investigated in order to assess their suitability for the trapping of Li atoms. As Li atoms are very easy to make (from ablation of a solid Li target) and detect, they were deemed a suitable candidate to attempt the first trapping experiments. As well as this, the inelastic scattering cross section between ground-state Li and He is much smaller than the elastic scattering cross section so trap loss through spin-flip collisions should be negligible [69]. The trajectories, phase space acceptance [1] and trap loss of Li atoms in the two different trapping arrangements were modelled, the results of which are presented in this chapter.

3.1 Six-magnet trap

The first magnet arrangement that was investigated is shown in figure 3.1(a), with six NdFeB magnets arranged with their north poles facing inwards to create a magnetic field zero at the centre. The holes inside each magnet allow easy access to the trap for the ablation and detection lasers. The Radia plug-in [101] for Mathematica was used to determine the shape and magnitude of the magnetic field produced by this magnet arrangement, with a cross section of it shown in figure 3.1(b). It is clear that this arrangement fulfils the requirement of having a magnetic field minimum at the centre with the field increasing in all directions away from the centre. Figure 3.1(b) also shows the dip in the magnetic field around the edges at $x = 0$ and $y = 0$ due to the presence of the holes in the magnets. This results in a slightly lower trap depth than would have been achieved without the holes, however the holes are a necessary requirement to allow laser access to the trap. The choice of this trap was motivated by the small magnetic field over a large volume, which minimises the Zeeman broadening of the Li transition, making it easier

---

[1] The initial positions and velocities of particles that can remain trapped

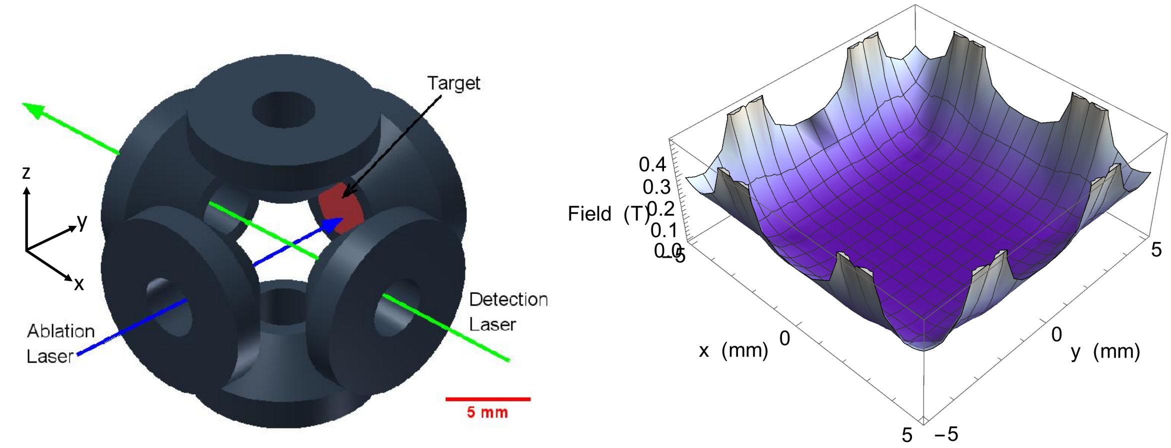

(a) (b)

Figure 3.1: (a) First proposed arrangement of the six permanent NdFeB magnets making up the trap, each with their north poles facing inwards to create a magnetic field zero at the centre. The holes inside the magnets allow the ablation and detection lasers to enter the trap, with the ablation target positioned inside one of the holes of the magnets. (b) Magnitude of the magnetic field inside the trap with this magnet arrangement. A cross section of the field is shown in the x-y plane at z = 0 (the centre of the trap). This magnet arrangement leads to a trap depth of approximately 0.25 T.

to detect. Additionally, it means that a large volume of the cloud of atoms buffer gas cooled in the trap region will be at small magnetic field, and hence at the bottom of the trap potential.

3.1.1 Trajectories of particles in the trap

The trajectories of the particles inside the trap were modelled in Mathematica. By inserting the magnetic field from figure 3.1(b) as a function of x, y and z into equation 2.9, the force on a particle in the x, y and z directions can be found. The magnetic moment of the particle was assumed to be 1 μ_B, which is correct for the most favourable state in Li. A cloud of particles was then created with random initial positions and a Maxwell-Boltzmann distribution of velocities at 4 K. The force calculated from equation 2.9, along with the initial positions and velocities, were then used to solve the equations of motion for each particle in the distribution. This gives a trajectory as a function of time for all the particles in the trap. Particles whose trajectories went beyond the boundaries of the trap because they were going too fast were deemed to be lost from the trap and were ignored in any further calculations. Figure 3.2 shows a typical trajectory of a trapped particle. Due to the ‘bucket’ like nature of the trapping field, the edges essentially act as walls which the particles bounce off. This results in a very irregular trajectory in which the particles spend most of their time in regions where the magnetic field is low– this will prove to be an issue when it comes to Majorana transitions (see section 3.1.2).

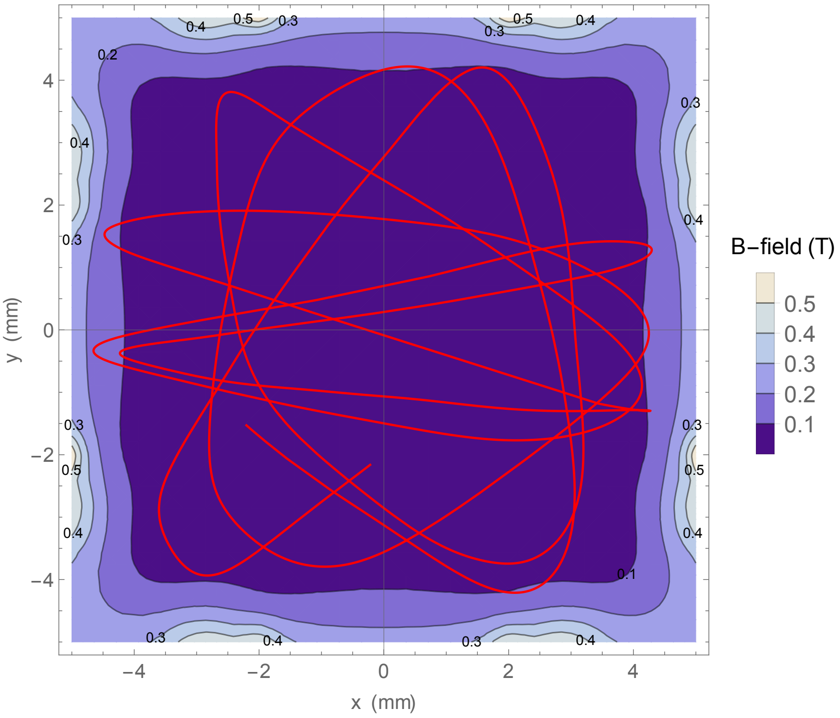

Figure 3.2: 2D projection into the x-y plane of the typical trajectory of a particle in the six-magnet trap (red line). A contour plot of the magnitude of the magnetic field is also shown, with darker areas indicating low magnetic field and lighter areas high magnetic field. Contours are labelled by their magnetic field values. The initial conditions for this trajectory are: $(x, y, z) = (-2.2, -1.5, -2.1)$ mm and $(v_x, v_y, v_z) = (12.8, -11.2, -10.5)$ m/s.

An initial distribution of Li atoms assumed to be at a temperature of 4 K due to thermalisation with the buffer gas was modelled. The fraction of particles lost over the edge of the trap because they were travelling too fast was calculated to be $\gtrsim$ 99.5%. This is a large fraction, but is not unexpected since the trap depth of 0.25 T corresponds to a temperature of 0.11 K, far smaller than the initial temperature of the atomic distribution. The Maxwell-Boltzmann velocity distribution gives the distribution of particle velocities at a particular temperature and is given by

$$f(v) = \sqrt{\left(\frac{m}{2\pi k_B T}\right)^3} 4\pi v^2 e^{-\frac{mv^2}{2k_B T}}, \tag{3.1}$$

where k_B is the Boltzmann constant, T is the temperature, m is the mass of the particle and v is the velocity of the particle. The maximum velocity a particle can have and still remain trapped can be calculated by equating the kinetic energy of a particle with the potential energy of the particle at the trap edge, i.e. the Zeeman shift at a field equivalent to the trap depth.

This maximum trapping velocity is then given by

$$v_{trap} = \sqrt{2\mu B_{depth}/m}, \tag{3.2}$$

where μ is the magnetic moment of the particle and B_{depth} is the trap depth. For Li atoms with a magnetic moment of 1 μ_B in this trap with trap depth of 0.25 T, the maximum trap velocity is 20 m/s. Integrating the Maxwell-Boltzmann distribution from equation 3.1 from 0 up to v_{trap}, with $T = 4$ K, we find the expected fraction of trapped atoms to be 0.63%. This is slightly higher than the 0.5% trapped fraction calculated from the numerical simulations above. The discrepancy is explained by the fact that, in the numerical simulations, the atoms are distributed throughout the trap, with some starting part of the way up the potential hill, instead of them all starting at the bottom of the potential as assumed in the analytical calculation. Both estimates of the trapped fraction are very small, however since laser ablation typically produces in the region of $10^{12}-10^{13}$ atoms [69, 83], only trapping $\lesssim$ 0.5% of the particles is likely to be sufficient.

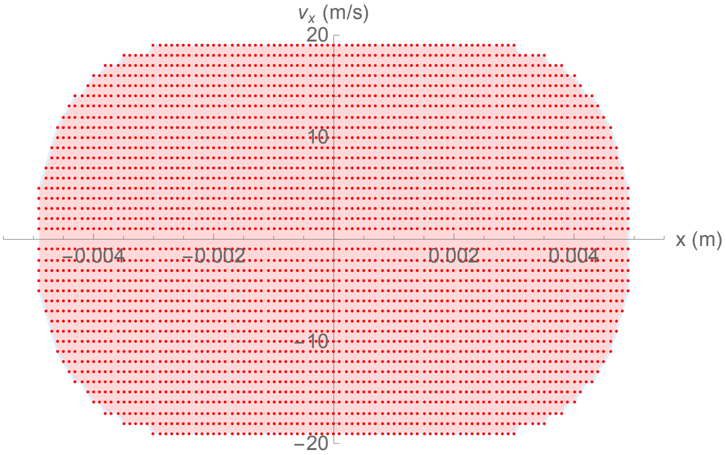

Figure 3.3: Phase space acceptance of the six-magnet trap in the x direction. The plot shows the initial x position against initial velocity in the x direction (v_x) of the particles that can be trapped.

A good way to determine which particles are trapped is to plot the phase space acceptance of the trap. This determines which initial positions and velocities of the Li atoms will remain in the trap. A grid of starting positions against velocities was generated and subsequently the trajectories of particles with each of these initial conditions were calculated. This was then used to determine which initial conditions led to a trapped particle. Figure 3.3 shows the phase space acceptance of the six-magnet trap in the x direction which, due to the symmetry of the trap, will be the same in the y and z directions. The phase space acceptance is almost rectangular because the trap has a flat bottom and steep walls. The trapping region only extends to 5 mm from the centre as this is where the magnets are located. The sharp cut-off in the velocity is a result of the steep gradient of the magnetic field at the trap edges. The numerical value for the one-dimensional phase space acceptance (calculated by integrating over the area of the plot in figure 3.3) of this trap is $350\,\mathrm{mm\,ms^{-1}}$.

3.1.2 Trap loss through Majorana transitions

Even though spin-flip collisions between Li and the He buffer gas are deemed negligible, Majorana transitions to a high field-seeking state are still a possibility. These Majorana transitions are driven by the motion of the atoms through regions where the magnetic field is close to zero. To estimate the rate of Majorana transitions for the Li atoms, typical values of $1/B\ dB/dt$ and the Zeeman shift in frequency units have to be compared [89], where B is the magnitude of the magnetic field and t is time. If the two quantities are comparable then the particle has entered the 'non-adiabatic' region and a Majorana transition to a high field-seeking state becomes likely. To obtain an estimate of the fraction of particles that undergo Majorana transitions, we used the condition,

$$\frac{1}{B}\frac{dB}{dt} > \frac{\mu_B B}{\hbar}. \tag{3.3}$$

If this condition was satisfied at any point during the trajectory of the particle then the particle was deemed to have undergone a Majorana transition. Whilst this is not an exact calculation since the transition still occurs with some non-unity probability, it should give a good order of magnitude estimate of the rate of Majorana transitions.

We started with a cloud of Li atoms at 4 K, and of those atoms not lost over the edge of the trap, the number lost through Majorana transitions was computed using the condition in equation 3.3 and the relation $dB/dt = (\underline{\nabla} B) \cdot \underline{v}$. It was found that after 100 ms, approximately 80% of the atoms were lost through Majorana transitions. This fraction is much too high to achieve a good trapping lifetime of the order of 1 s (see table 2.1) as most of the atoms would have been lost before then. This trap arrangement is therefore unlikely to be useful for Li atoms, however it may still be used in the future for trapping molecules as they are heavier and would therefore be travelling slower than the Li atoms, making the dB/dt term on the left hand side of equation 3.3 smaller, thus reducing the Majorana transition probability.

3.2 Quadrupole trap

To avoid Majorana transitions, a trap design was needed where the field was not so small over such a large area. The dependence of the magnitude of the field with distance from the trap centre is determined by the number of poles making up the trap. This dependence varies as r^{n-1} [102], where r is the radial distance from the trap centre and n is the order of the multipole field, which is equivalent to the number of magnetic dipoles making up the field. This means that for the trap in figure 3.1, with six magnetic dipoles, there is an r^5 dependence of the magnetic field, leading to the 'bucket' shaped trap. To counteract this, the number of poles in the trap must be reduced. This led to the design in figure 3.4(a) using only two cylindrical magnets with their north poles facing each other, giving a quadrupole trap.

From figure 3.4(b) it can be seen that this trap gives a similar field to that from the anti-Helmholtz coil arrangement (figure 2.3), with the field increasing linearly with distance from the centre rather than the previous r^5 dependence. Since only two magnets are used, there is no need to incorporate any holes in the magnets as there is now full access to the trap in the $x-y$ plane for the ablation and detection lasers. This trap arrangement gives a slightly higher trap depth than the trap described in section 3.1, but more importantly there is a very small region where the field is close to zero, which should reduce Majorana transitions as the right hand side of equation 3.3 is now much larger.

Modelling trajectories of Li atoms in the trap as in the previous section, with a distribution of particles starting at 4 K, we found that $\lesssim 0.75\%$ were trapped (i.e. not lost over the edge of the trap), a slight improvement on the six-magnet trap. Calculating the trapped fraction using the Maxwell-Boltzmann distribution from equation 3.1 and the maximum trapped velocity from equation 3.2, gives a value of 0.87%. This is again slightly higher than the fraction from the numerical simulations, for the same reason as explained above in the six-magnet trap section. Figure 3.5 shows a typical trajectory of a particle in this trap. Due to the nature of the trap, the particles follow a much more regular pattern around the trap centre, and if they are in an orbit around the centre, there is an angular momentum barrier that prevents them from reaching the trap minimum. This should significantly reduce the likelihood of Majorana transitions.

Figure 3.6 shows the phase space acceptance of the quadrupole trap in both the x and z directions. The phase space acceptance for this trap is much more oval than that for the six-magnet trap in figure 3.3. This is because the field in this trap increases linearly with distance

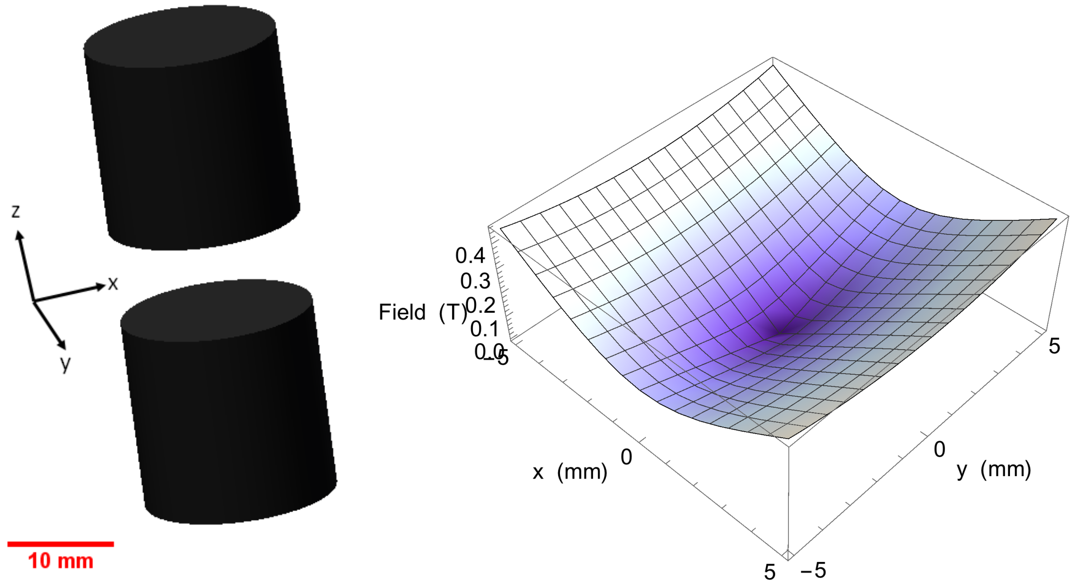

(a) (b)

Figure 3.4: (a) Arrangement of the two NdFeB magnets in the trap with their north poles facing inwards to create the field zero at the centre. The ablation and detection lasers can enter the trap anywhere in the $x-y$ plane. (b) Magnitude of the magnetic field inside the trap with this magnet arrangement. A cross section of the field is shown in the $x-y$ plane at $z=0$ (the centre of the trap). This arrangement leads to a trap depth of approximately 0.31 T

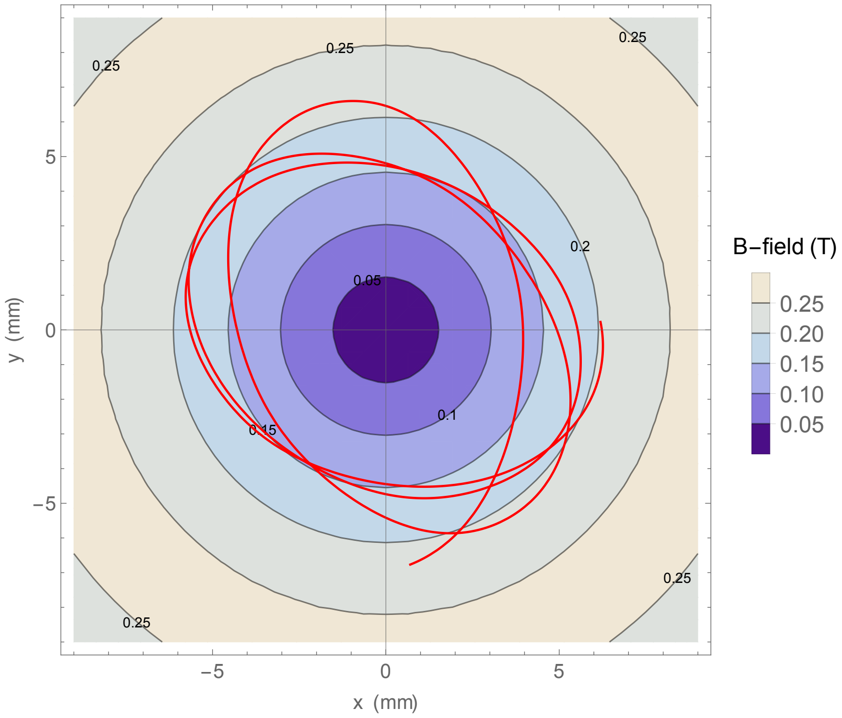

Figure 3.5: 2D projection into the x-y plane of a typical trajectory of a particle in the quadrupole trap (red line). A contour plot of the magnitude of the magnetic field is also shown, with darker areas indicating low magnetic field and lighter areas high magnetic field. Contours are labelled by their magnetic field values. The initial conditions for this trajectory are: $(x,y,z)=(-1.7,0.7,-6.8)$ mm and $(v_x,v_y,v_z)=(-15.2,7.8,3.1)$ m/s.

from the centre, unlike the field for the six-magnet trap which is very flat. As a result, the six-magnet trap is able to accept higher velocities towards the edges of the trap than the quadrupole trap. We see from figure 3.6(a) that the size of the trapping region in the x direction (also in the y direction due to symmetry) is much larger than for the six-magnet trap, extending to just over 9 mm from the centre. It is also a slightly deeper trap, shown by the higher maximum velocity that remains trapped. In the z direction (see figure 3.6(b)) the size of the trapping region only extends to 5 mm from the centre as this is where the magnets are positioned. The numerical values for the one-dimensional phase space acceptance in the x and z directions for this trap are 460 and 300 mm ms^{-1} respectively. In the x and y directions this is significantly higher than for the six-magnet trap, whereas in the z direction it is slightly lower.

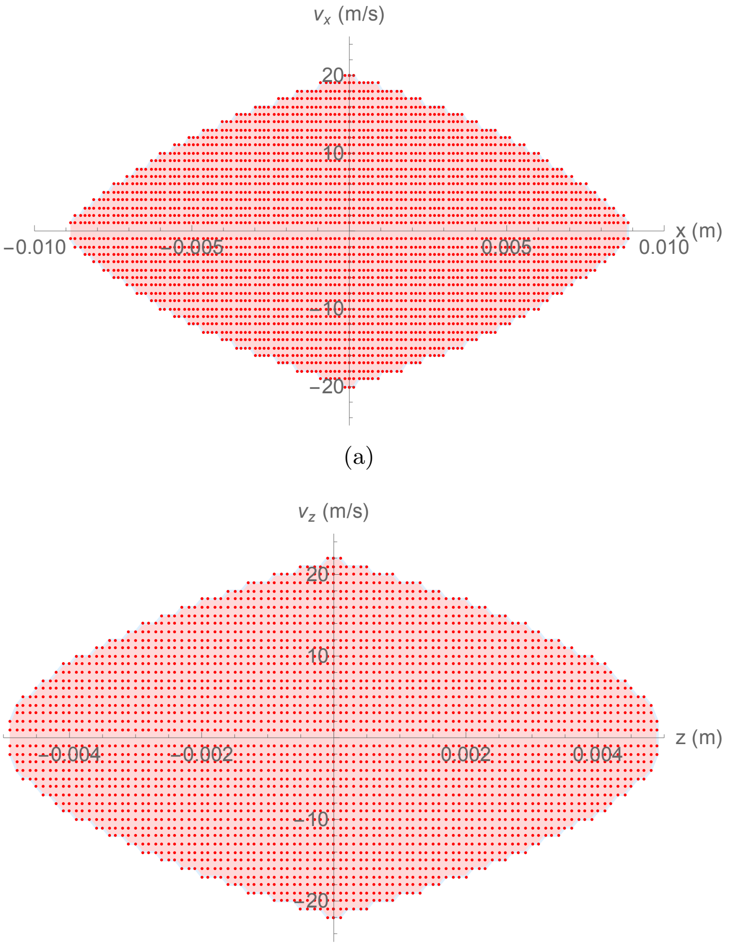

(a)

(b)

Figure 3.6: Phase space acceptance of the quadrupole trap in both the x and z directions, shown in (a) and (b) respectively. Due to the symmetry of the trap, the phase space acceptance in the y direction will be the same as in the x direction.

Going through the same procedure as in section 3.1.2 for calculating Majorana transitions, it was found that none of the Li atoms had undergone a Majorana transition after 100 ms. This is clearly a significant improvement on the six-magnet trap and led to the conclusion that this magnet arrangement is suitable for trapping Li atoms.

When it comes to using this arrangement to trap molecules however, there might be a different problem. Whilst spin-flip collisions between Li and the He buffer gas are negligible, this is not the case for most molecules, as explained in section 2.2.1. Due to the nature of the field in this trap, any inelastic collision between a molecule and He that causes the molecule to spin-flip into a high field-seeking state will result in the molecule being ejected from the trap very rapidly. This is due to the steep gradient in the field leading to a large force on the molecule (see equation 2.9). In the first trap in section 3.1, a spin-flip transition does not result in a rapid ejection of the molecule because the field gradient is small. There is an opportunity for a further spin-flip collision, restoring the molecule to the trapped state before it is lost. The helium density drops with time, and so after a certain time there are no more collisions. In the first trap, molecules are only lost if they are in the high field-seeking state after the collisions cease, while in the second trap they are lost as soon as a collision puts them into the high field-seeking state.

3.3 Conclusion

Typical trajectories, trapped atom fractions and losses were studied for two different trapping arrangements: a six-magnet trap and a simple quadrupole trap. The main findings are summarised in table 3.1. The quadrupole trap has a slightly higher trap depth and phase space acceptance, leading to the higher fraction of 4 K atoms that remain trapped. Numerical simulations of particle trajectories were carried out to get an estimate of Majorana losses for both trapping arrangements. The bucket-like nature of the six-magnet trapping field resulted in severe Majorana losses as the particles spend a large amount of time in regions of near-zero field. This made the six-magnet arrangement unsuitable for trapping Li atoms. In the simple

	Six-magnet trap	Quadrupole trap
Trap depth (T):	0.25	0.31
Fraction of 4 K atoms trapped (numerical):	0.5%	0.75%
Fraction of 4 K atoms trapped (theoretical):	0.63%	0.87%
1D phase space acceptance (mm ms^{-1}):	350	460 (x and y); 300 (z)
Majorana losses after 100 ms:	80%	0

Table 3.1: Summary of the key findings for the two trap arrangements, assuming an initial atomic distribution at 4 K.

quadrupole trap, on the other hand, there are no Majorana losses in the same period. This information, in combination with the higher trap depth of the quadrupole trap, led to the conclusion that the quadrupole arrangement is more suitable to trap atoms and hence was used throughout the experiments in this thesis.

CHAPTER 4

LITHIUM EXPERIMENTS

The ultimate goal of this experiment is to create a permanent magnet trap for molecules. To investigate the feasibility of the trap, the first experiments focus on cooling and trapping lithium atoms as their magnetic moment of $1\mu_B$ is similar to that of most molecules of interest. Additionally, lithium is easy to detect and the inelastic scattering cross section between Li and He is much smaller than the elastic scattering cross section so trap loss through spin-flip collisions should be negligible [69].

4.1 Lithium details

Lithium is detected on the 671 nm D2 $^2S_{1/2}(F=2) \rightarrow {}^2P_{3/2}(F=3)$ transition, shown in figure 4.1. The natural linewidth of the transition is 5.92 MHz so the hyperfine states of the upper

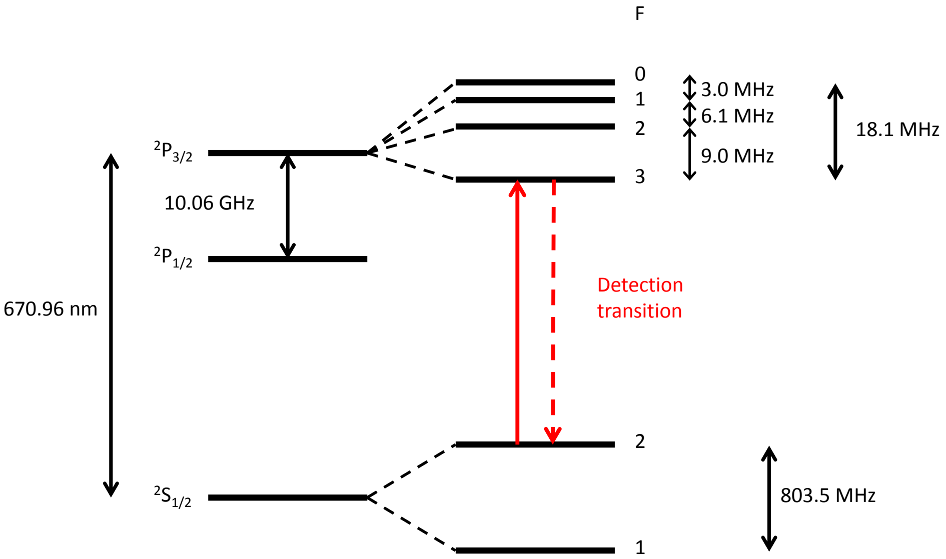

Figure 4.1: Energy level diagram of ^{7}Li showing the D2 transition used for detection.

state are partly unresolved. This can lead to excitations to one of the other hyperfine states and subsequently result in decay down to the ground $F = 1$ state. Once they are optically pumped into the $F = 1$ ground state, the atoms no longer fluoresce since the transition from $F = 1$ is ~800 MHz from resonance. Figure 4.2 shows how the ground state energy levels of Li change with magnetic field, calculated from the Breit-Rabi formula for $J = 1/2$ atoms [103]. As the trap depth in our set-up is about 0.31 T (see figure 3.4(b)), we will mostly be dealing with atoms in the strong field regime. In this regime, all of the F = 1 states are high-field seeking. Of the five F = 2 states, four are low-field seeking and can be trapped. The magnetic moment of Li in the strong-field regime is $1\,\mu_B$, which gives a trap depth in our trap of about 0.14 K.

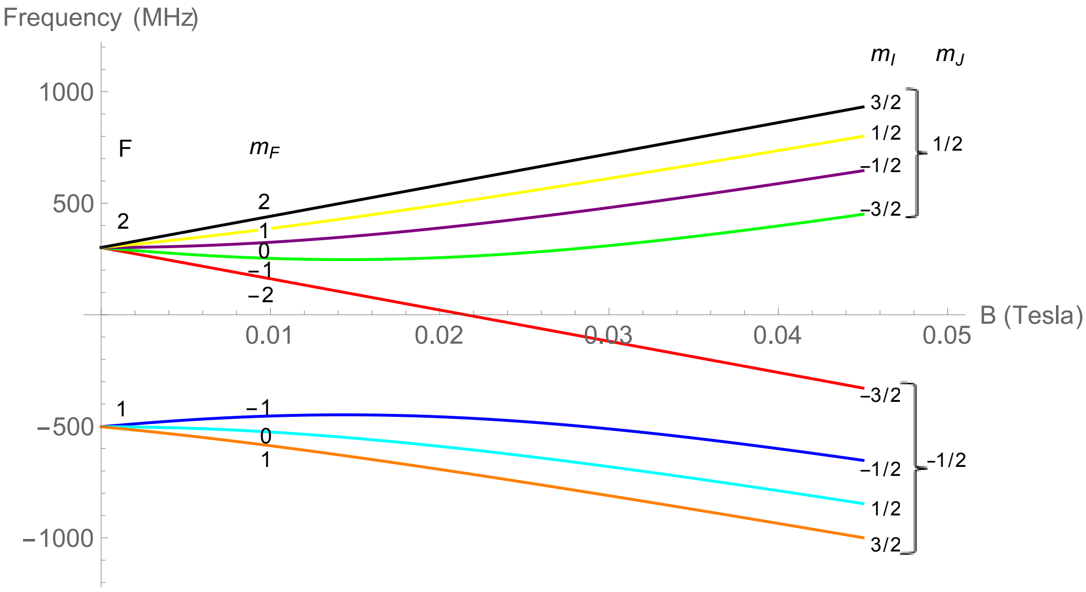

Figure 4.2: Zeeman shift of the $^2S_{1/2}$ ground state of lithium as a function of magnetic field.

The Zeeman splitting of the excited state in the magnetic field looks rather complicated due to the presence of the four closely spaced hyperfine levels. In strong fields, the components separate out into four groups, having $m_J = 3/2$, 1/2, -1/2 and -3/2. Those with $m_J = 3/2$ have a magnetic moment of $2\,\mu_B$, compared to $1\,\mu_B$ for the ground state. With magnetic fields in the trap ranging from 0 up to 0.31 T, this equates to a Zeeman broadening of the transition of up to 4 GHz.

4.2 Experimental set-up

4.2.1 Trap design

The modelling of particles in two different trapping arrangements in chapter 3 led to the conclusion that a simple quadrupole arrangement consisting of two cylindrical NdFeB magnets, as shown in figure 3.4(a), would provide the most effective trapping field. The magnets, supplied by eMagnets UK [104], have a diameter and height of 20 mm and are placed 10 mm apart. To establish whether the field shown in figure 3.4(b), calculated using the Radia plug-in for Mathematica [101], is an accurate representation of the field, measurements of the fields produced by the magnets were made using a Lake Shore Model 425 gaussmeter. Both the field from the magnets individually and the field produced by the magnets in the arrangement from figure 3.4(a) were measured. These were both found to agree with the expected field from the Radia calculations.

A suitable set-up had to be designed to hold the magnets in place in this arrangement. The aim of the trap design was to ensure that it remained as open as possible, rather than the closed cell used in the Doyle group experiments [36], allowing the helium to leave the trapping area without the need to pump it away, thus resulting in fewer spin-flip collisions. This also allows good access for the ablation and detection laser beams. With this in mind, the trap was designed as shown in figure 4.3. An aluminium frame is placed on top of the cold plate to hold the magnets. As the north poles of the two magnets face each other, the repulsion between them keeps the upper magnet in place without the need to fasten it. The helium buffer gas is pulsed into the trap region through a solenoid valve (Parker Series 9 pulsed valve), with a copper tube mounted on top of the valve to guide the pulse towards the trapping region. Typically, the valve delivers on the order of 10^{18} helium atoms per pulse [81]. The lithium is subsequently ablated into the buffer gas from a solid target and allowed to thermalise, with the buffer gas then escaping the trap region, leaving the lithium atoms trapped. The valve is sunk into the copper cold plate to ensure that the buffer gas is as cold as possible before it enters the trap region. We also have the option of placing a copper shield (heatsunk to the cold plate) around the trap region to confine the initial helium pulse and give the atoms/molecules more time to thermalise. Two methods of detection are possible in this set-up: absorption measurements on a photodiode (see figure 4.6) and laser-induced fluorescence detection on a photomultiplier

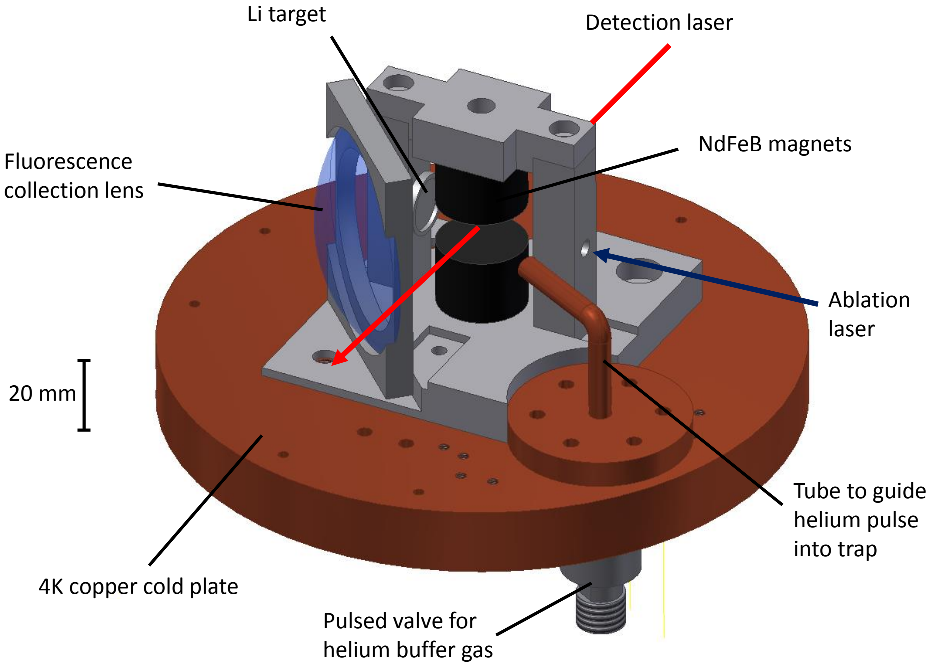

Figure 4.3: The design for the magnetic trap, showing how the magnets are held in place whilst keeping the trap as open as possible. The copper cold plate is attached directly to the 4 K cold plate of the cryocooler, with an aluminium frame placed on top to hold the magnets. The helium gas enters the trap through a pulsed valve, which is sunk into the copper cold plate in order to maximise thermalisation of the helium gas to 4 K. A copper tube is mounted to the top of the valve to guide the pulse into the trap region. The direction of the ablation and detection lasers are also shown (for laser details, see section 4.2.4). A fluorescence collection lens is placed at 45° to the detection beam, with the fluorescence light detected on a photomultiplier tube placed outside the vacuum chamber (details in figure 4.5).

tube (PMT), which is described in more detail in section 4.2.3. Both these detectors are placed outside the vacuum chamber.

4.2.2 The apparatus

The trap from figure 4.3 is attached to the top of the crycooler cold head, as shown in figure 4.4. We use a closed-cycle cryocooler with 1.5 W of cooling power at 4 K (Sumitomo Heavy Industries RDK-415D). It works by pumping cold helium up through the central column of the cryocooler to the cold plate to cool it to 4 K. Several thermometers (made up of ruthenium oxide resistors with a 1.5 kΩ resistance at room temperature) are placed around the cold plate in order to accurately monitor its temperature. The crycooler also has a 50 K stage on which an aluminium radiation shield is placed that surrounds the trap set-up. This acts to prevent any room temperature heat radiation from impinging directly onto the cold plate and raising its temperature. The radiation heat load from the 50 K surfaces onto the 4 K cold head is negligible compared to the 1.5 W of cooling power. A vacuum chamber surrounding the set-up is pumped down using a turbo pump (Leybold Turbovac 361) attached to the bottom of the chamber (see figure 4.4). The base pressure achieved in the chamber when pumped out and cooled down is 1×10^{-7} mbar. The other openings at the bottom of the chamber are necessary to allow access for the gas inlet tubes for the valve, the pressure gauges and the thermometer wiring.

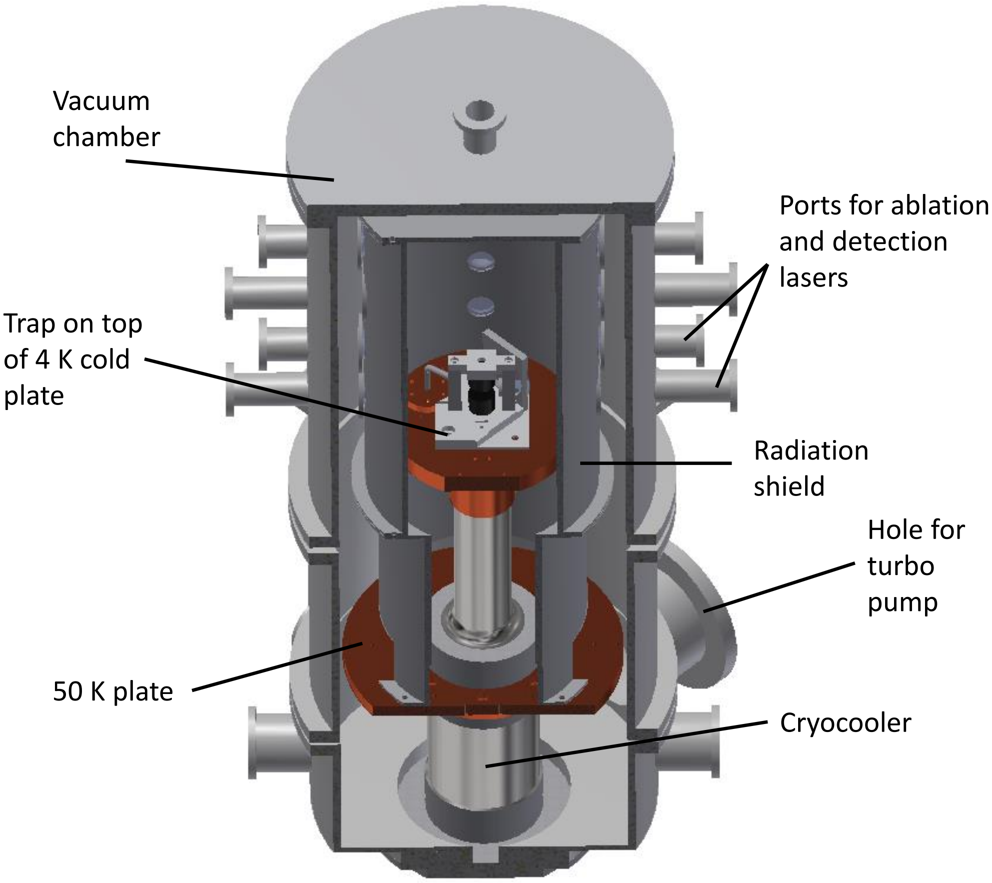

Figure 4.4: Cut-out of the experimental apparatus showing the cryocooler, 4 K plate, 50 K plate, radiation shield and vacuum chamber.

To pump away the helium as quickly as possible, a charcoal sorption pump is placed around the trap which, when cooled to cryogenic temperatures, traps any helium that hits it. Coconut charcoal was used for the sorption pump as it has been identified as one of the best adsorbers of helium at cryogenic temperatures, with a pumping speed of up to 6.7 l/(s cm^2) at 4 K [105]. The coconut charcoal was bonded to the inside of a copper cylinder using a Stycast epoxy, and the cylinder was then placed on top of the cold plate, surrounding the trap, so that it is anchored to the 4 K stage. This provides a total pumping area of approximately 600 cm^2, giving a total pumping speed of up to 4000 l/s for helium. Comparing this to the pumping speed of 340 l/s for helium from the turbo pump (Leybold Turbovac 361), located significantly further away from the trap region, we see that the helium is mainly pumped by the charcoal.

4.2.3 Laser-induced fluorescence detection

The most sensitive detection method in our set-up is the laser-induced fluorescence (LIF) detection on a photomultiplier tube (PMT). As it is an integral part of the experiment, especially when looking for signal of trapped atoms, it is described in more detail here.

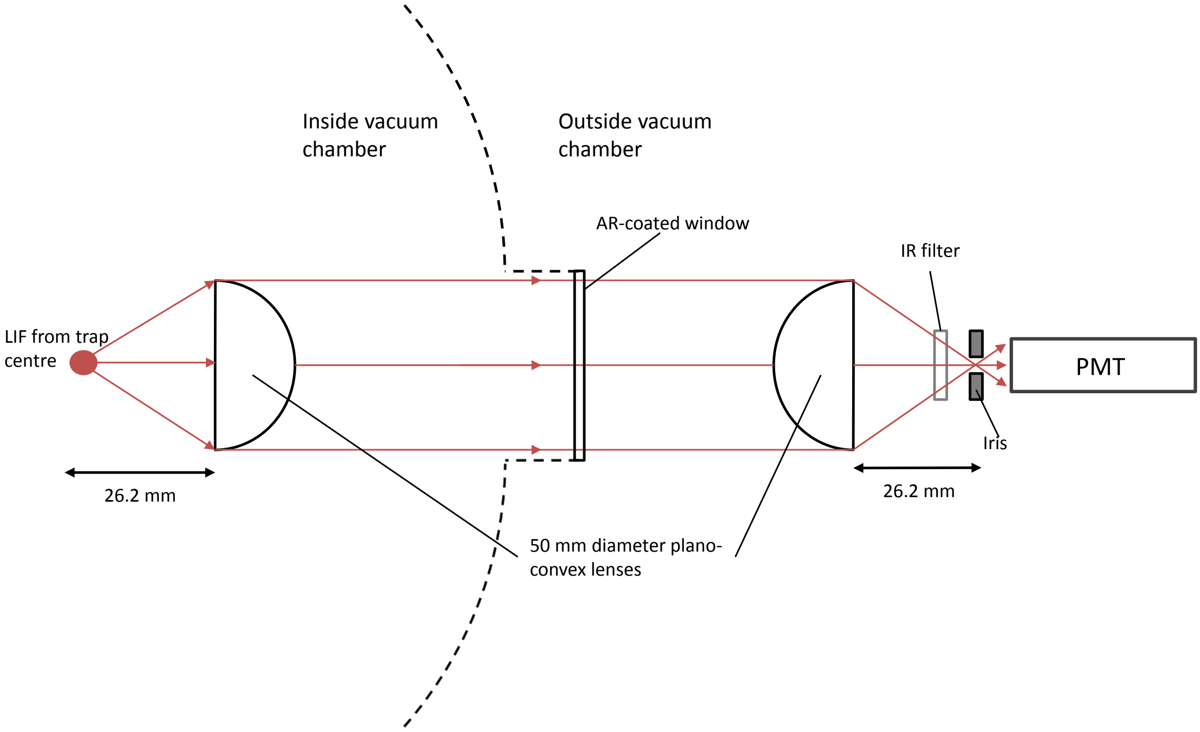

Figure 4.5: Schematic of the LIF detection set-up, showing the positions of the lenses, IR filter and iris. The lenses are 50 mm diameter plano-convex lenses with a back focal length of 26.2 mm. The filter is a coloured glass short-pass filter, transmitting visible light and absorbing IR. The iris has an adjustable diameter from 1 mm to 10 mm.

Figure 4.5 shows the LIF detection set-up for these experiments. A 50 mm diameter, short focal length (26.2 mm) collection lens is situated close to the trap region to capture as much fluorescence as possible. The collimated light is then focused onto the PMT (Hamamatsu Photonics R5070A) with the second lens. An IR filter is placed in front of the PMT to prevent any light from the ablation laser being registered on the PMT. The adjustable iris diameter allows the spatial filtering of light from the trap region to minimise the amount of laser scatter from other regions of the trap reaching the PMT. Additionally, to prevent the probe beam scattering off surfaces in the trap region and being detected on the PMT, the faces of the magnets were coated in black soot.

Taking into account the magnet geometry in the trap and the position of the collection lens, the collection efficiency of this LIF set-up is 6.25%. With a quantum efficiency of the PMT at the relevant wavelength of 7%, the total detection efficiency of the LIF set-up is 0.44%.

Throughout, the PMT was used in ‘current mode’, with a current-voltage amplifier converting the signal into a measurable voltage that was then recorded. However, in order to convert this measured voltage into an actual photon count rate, the PMT had to be calibrated. This was done by placing the PMT inside a blacked out box together with a red LED. A current was passed through the LED and, with the PMT in ‘photon counting mode’, the number of photons registered on the PMT in a 10 μs interval was recorded from an oscilloscope. This was repeated several times to get an average photon count rate for that LED current. The PMT was then plugged into the current-voltage amplifier to measure the voltage output corresponding to that photon count rate. By doing this for several different LED currents, a calibration graph of voltage vs. photon count rate was obtained that could then be used for all subsequent experiments. The calibration was $R = 930\,\mathrm{kHz/V}$.

4.2.4 Laser set-up

Figure 4.6 shows a schematic of the laser set-up. The probe light comes from a 10 mW, 671 nm external cavity diode laser (ECDL), which allows us to address the relevant lithium transitions shown in figure 4.1. The ECDL consists of a 671 nm laser diode, a collimating lens to collimate the diode output and a reflective diffraction grating to provide the optical feedback and form the external cavity. The ECDL is constructed inside a home-built laser housing. The output of the ECDL passes through an optical isolator to prevent back-reflections into the laser from the fibre couplers, and then into a single-mode optical fibre. Going straight into an optical fibre has the advantage that, if anything needs to be adjusted inside the diode laser, the output beam just needs to be realigned with the first fibre, with all subsequent optics remaining unaffected. After the first fibre, several portions of the main beam are split off using non-polarising beam splitters, allowing beams to be sent to the wavemeter, the scanning Fabry-Perot cavity and the two confocal Fabry-Perot cavities. The main beam then enters a second optical fibre to the experiment table. There, a mechanical shutter unblocks or blocks the probe beam so that the laser-induced fluorescence can be distinguished from any other sources of light. Finally, using a neutral density filter wheel (not shown in figure 4.6), the probe power through the trap region can be adjusted to the desired level.

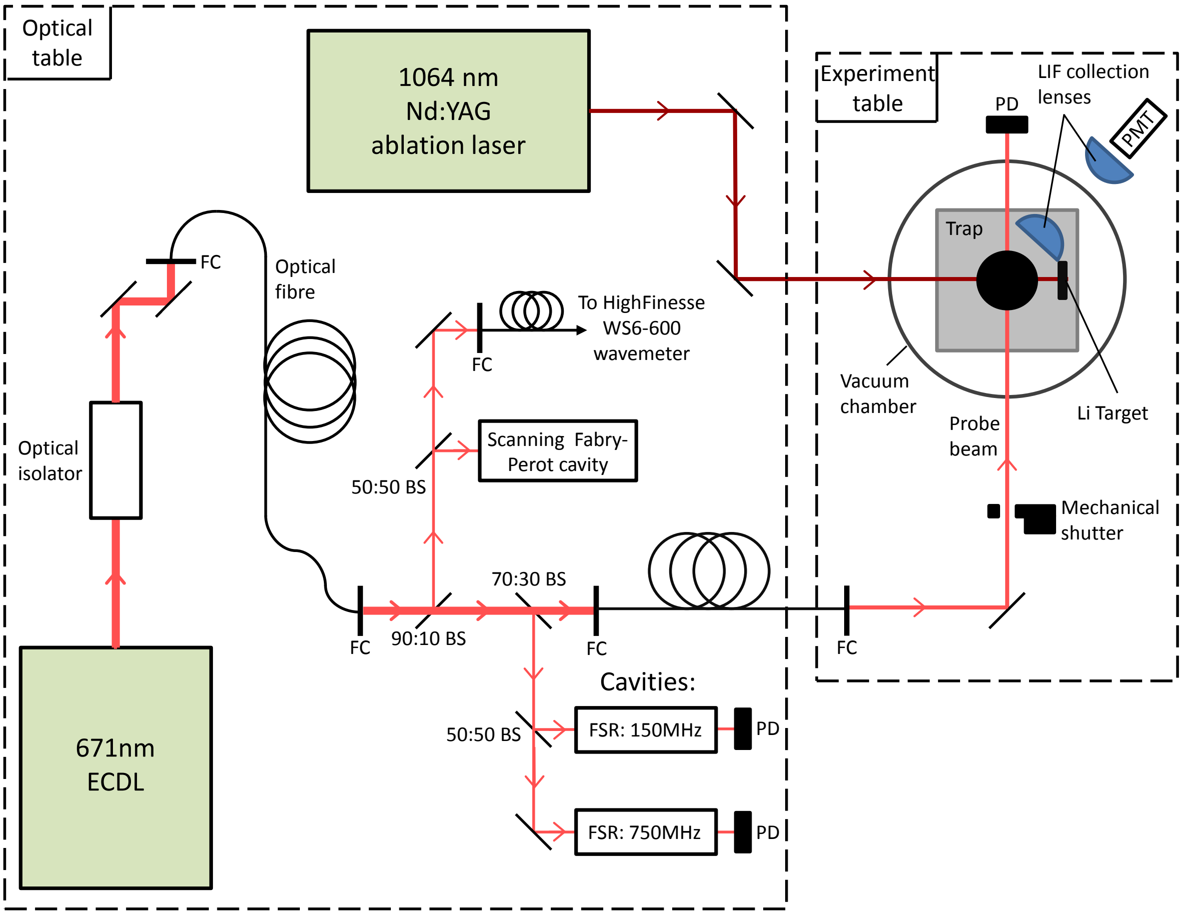

Figure 4.6: Schematic of the laser set-up and how it fits in with the rest of the experimental apparatus. PD refers to Thorlabs DET100A photodiode detectors, BS to beam-splitters with a transmission:reflection ratio, FC to fibre couplers and FSR to the free spectral range of the cavities. The transmission through the cavities is used to monitor the change in laser frequency.

The frequency of the probe light can be scanned by supplying a voltage to a piezo located inside the ECDL to adjust the position of the reflection grating and thereby adjust the length of the external cavity of the diode laser. Typically, a frequency range up to 5 GHz can be scanned in this way before the laser ‘mode hops’ into a different mode. The frequency is monitored using the wavemeter and the transmission through the two cavities, which is measured on the photodiode detectors. The fringes from the scanning Fabry-Perot cavity are monitored on an oscilloscope to ensure the laser stays in single-mode operation.

The Li target is ablated using a Continuum Minilite MLII 1064 nm Q-switched Nd:YAG laser. This has a maximum pulse energy of 50 mJ with a pulse width of 5-7 ns. Throughout all of the data sets in this thesis, unless mentioned otherwise, the pulse power from the ablation laser was kept at 25% of the maximum (12.5 mJ) for the Li experiments, and 20% of the maximum (10 mJ) for the Dy experiments.

4.2.5 Data collection procedure and terminology

The experiment is controlled from a computer, with a National Instruments (NI) PCI-6534 pattern generator used to control the firing of the ablation laser and valve, and an NI PCI-6251 data acquisition card used to record the data. Data collection starts with a He pulse from the valve, fired at a time t_{valve}, with the valve pulse length (VPL) set in advance and typically being on the order of $200\,\mu$s. The Li atoms are subsequently ablated into the gas pulse, with the zero in time taken from the point at which the ablation laser Q-switch is fired. The time at which the valve is fired is typically around $t_{valve} = -500\,\mu$s. The absorption and LIF of Li in the trap region are measured as a function of time on the photodiode detector and PMT respectively, the outputs of which are recorded by the computer. The complete cycle of firing the valve, firing the ablation laser and recording absorption and LIF signals is referred to as a single 'shot'. For all of our data, unless mentioned otherwise, these 'shots' are taken at a rate of 2 Hz, as we found that this gave us the best yield from the target.

When the Li target is hit by the ablation laser, the ablation plume that is released from the target emits a huge amount of fluorescence for the first few $100\,\mu$s, even without the presence of the probe beam. This fluorescence is detected by the PMT and obscures any real LIF signal. In order to eliminate this we take an 'On-shot' with the probe beam on, followed immediately by an 'Off-shot' where the probe beam is blocked using the shutter, but everything else fires as normal. Subtracting the 'Off-shot' signal from the 'On-shot' signal ensures that only the real LIF signal is left. This was done for all of our data, and as a result all plots of LIF signal throughout this thesis refer to the signal after the 'Off-shot' has been subtracted.

Frequency scans are taken by stepping the probe beam frequency between shots for several hundred shots over the desired frequency range. Recording the transmission through the cavities for each shot gives a fringe pattern which calibrates the change in laser frequency. The frequency spectrum is then obtained by averaging the absorption and/or LIF signals between certain time 'gates' for each shot, and plotting the result against the laser frequency. Scans of the valve pulse length (VPL) or t_{valve} can also be taken in the same way by stepping these parameters between one shot and the next, rather than the probe laser frequency. The signal can then be plotted as a function of VPL or t_{valve}.

4.3 Results

4.3.1 Data with 'dummy magnets'

The first stage of the experiment, once the set-up was completed, was to determine the temperature of the lithium atoms after ablation and thermalisation with the buffer gas. The temperature is determined from the Doppler broadening of the spectral lines. This measurement cannot be done with the magnets in place because the Zeeman broadening due to the inhomogeneous magnetic field is far too large. Therefore the permanent magnets in the set-up were replaced by aluminium cylinders of the same shape and size, which are referred to as 'dummy magnets'. This way the physical arrangement of the trap was identical to the one with real magnets, giving us a good indication of what the initial temperature of the lithium atoms will be in the actual set-up.

To measure the temperature, an absorption/LIF profile was obtained by scanning the detection laser frequency (with a power of $12\,\mu$W) across the 671 nm resonance transition. This gives an absorption/fluorescence profile as a function of frequency. For the first few $100\,\mu$s after ablation the signal was sufficient for absorption detection on the photodiode, however at later times fluorescence detection on the photomultiplier tube (PMT) was used. Figure 4.7 shows typical absorption and fluorescence spectra for the Li atoms after thermalisation with a helium pulse from the valve. The absorption profile was recorded $300\,\mu$s after ablation and the fluorescence profile $650\,\mu$s after ablation, with a gate window duration of 50 and $300\,\mu$s respectively [1]. The valve was fired $900\,\mu$s before the ablation laser ($t_{valve} = -900\,\mu$s), as this gave the highest amount of absorption. The He pulse length from the valve (VPL) was $200\,\mu$s. The two peaks in the spectra arise as a result of the two hyperfine ground state levels, with the distance between the two peaks of about 800 MHz corresponding to the splitting between the two levels (see figure 4.1). The most likely candidates for the broadening mechanisms that give rise to the peak shapes are Doppler broadening and pressure broadening, which give rise to a Gaussian profile and a Lorentzian profile respectively. A Voigt profile (a convolution of a Gaussian and Lorentzian profile) fit to the spectra shows that the line shapes are dominated by the Gaussian component and hence Doppler broadening, with pressure broadening negligible. The spectra in

---

[1] This means that the signal was averaged between 275 and $325\,\mu$s for the absorption and averaged between 500 and $800\,\mu$s for the LIF, with the result plotted against the frequency

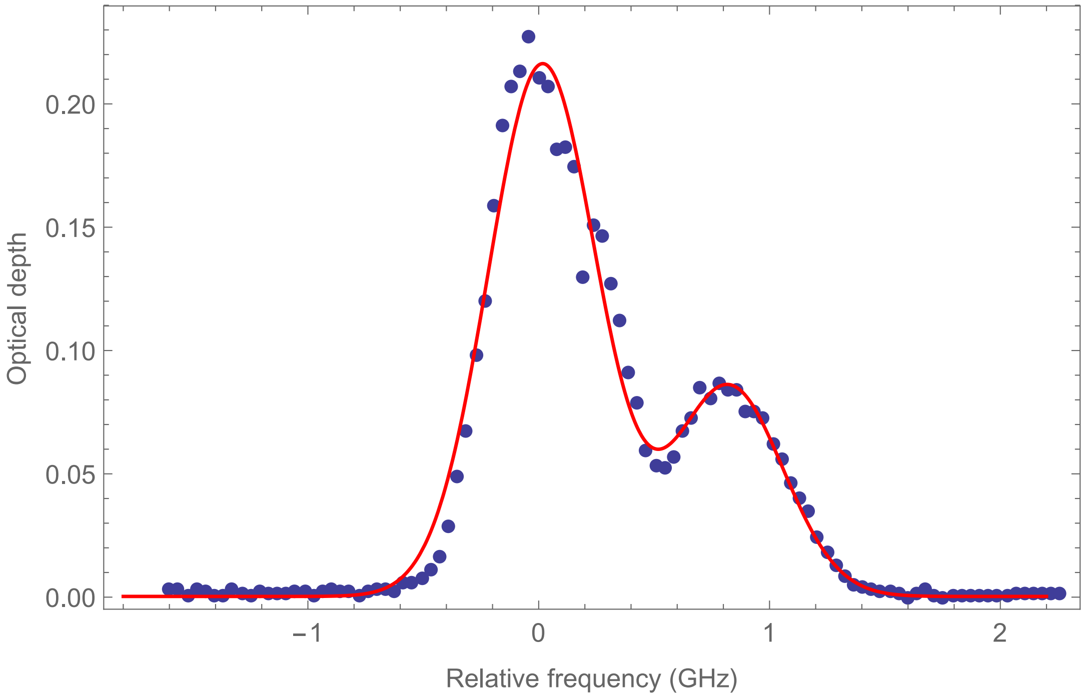

(a) Absorption spectrum of Li taken 300 μs after ablation. Temperature ≈ 22 K.

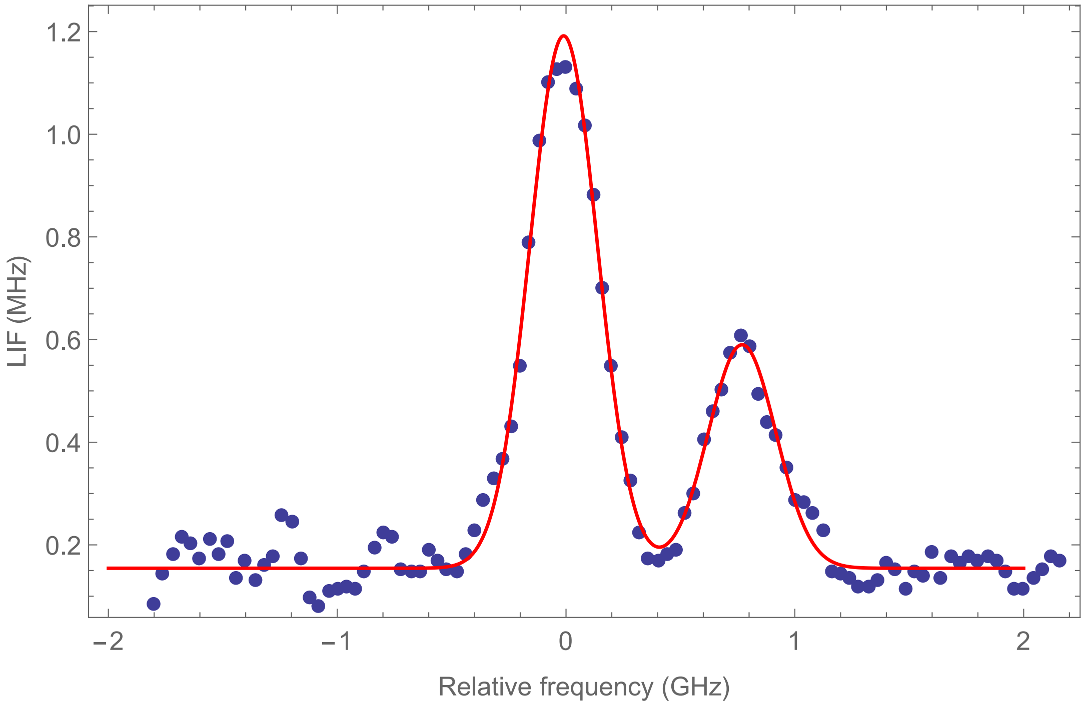

(b) Fluorescence spectrum of Li taken 650 μs after ablation. Temperature ≈ 8 K.

Figure 4.7: Frequency scans over the Li detection transition after thermalisation with the buffer gas, with 'dummy magnets' in place. The frequencies are relative to the transition from the ground $F = 2$ state. Two scans are shown, taken at 300 and 650 μs after firing the ablation laser. Fluorescence detection on a PMT was used at later times as the signal was too low for a good absorption signal. To determine the Doppler width, and hence the temperature of the Li, a Gaussian profile with two peaks was fit to the data points (solid line), and the temperature subsequently calculated from the linewidth of the peaks.

figure 4.7 can therefore be well approximated by fitting a double Gaussian of the form,

$$\text{const.} + A\mathrm{e}^{-\frac{(f-f_{01})^2}{2\sigma^2}} + B\mathrm{e}^{-\frac{(f-f_{02})^2}{2\sigma^2}}, \tag{4.1}$$

where A and B are the heights of the peaks, σ is the width of the peaks, f is the frequency and f_{01} and f_{02} are the centres of each peak. These fits are shown in figure 4.7. The width of the peaks is related to the temperature of the atoms by,

$$T = \frac{m\,(\sigma \times \lambda)^2}{k_B}, \tag{4.2}$$

where λ is the wavelength of the detection, m is the mass of Li and k_B is the Boltzmann constant. For the two plots in figure 4.7, the temperature was calculated to be 22 ± 2 K 300 μs after ablation and 8 ± 1 K 650 μs after ablation. This means that the lithium has successfully thermalised with the buffer gas as the temperature of the lithium coming directly off the ablation target will be more in the region of 10 000 K.

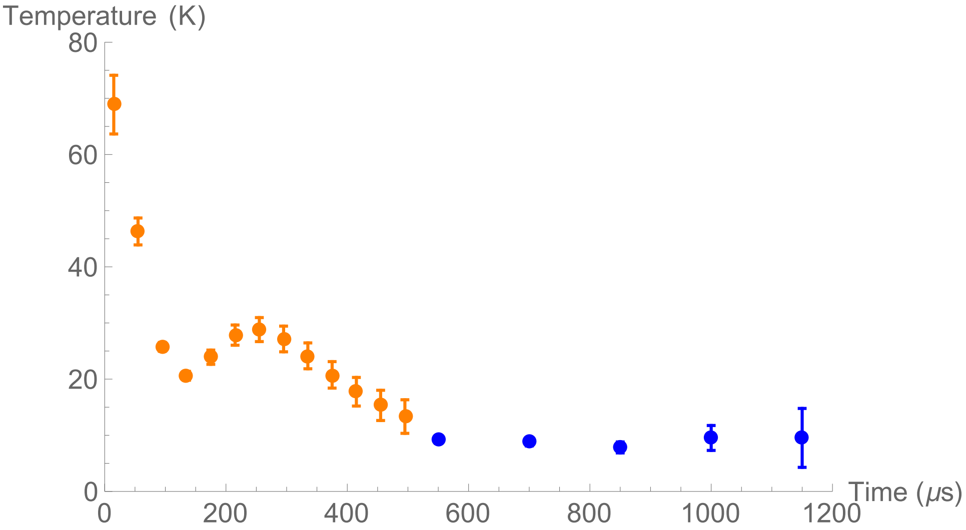

Figure 4.8: Plot of the temperature of the Li atoms against time after ablation, with 'dummy magnets'. Temperatures were obtained by fitting a double Gaussian to the Doppler broadened frequency scans, as in figure 4.7. The orange points are those obtained from absorption measurements and the blue points those obtained from fluorescence measurements.

Figure 4.8 shows how the temperature of the Li atoms changes with time after ablation. As explained above, for early times after ablation, absorption measurements were used to determine the temperature and for later times fluorescence measurements were used. The Li atoms are still at a temperature of about 45 K after 50 μs, however they then thermalise quickly down to 20 K. The slight rise in the temperature 250 μs after ablation appears in all of our data and we do not currently have an explanation for it. The Li subsequently cools further so that, around 500 μs, it has reached its coldest temperature of about 8 K. The data from the fluorescence measurements do not show any further decrease in temperature below 8 K. Even though the Li has clearly thermalised with the buffer gas, it does not reach the 4 K temperature of the cold plate. This slightly higher temperature is likely to be a result of the valve warming up slightly when it fires, meaning that the helium actually comes out of the valve at a temperature slightly above 4 K. This phenomenon has been observed before [41] and is a known problem. Unlike in a standard buffer gas cell, such as illustrated in figure 2.1, the helium here has no walls to collide with to thermalise back to 4 K.

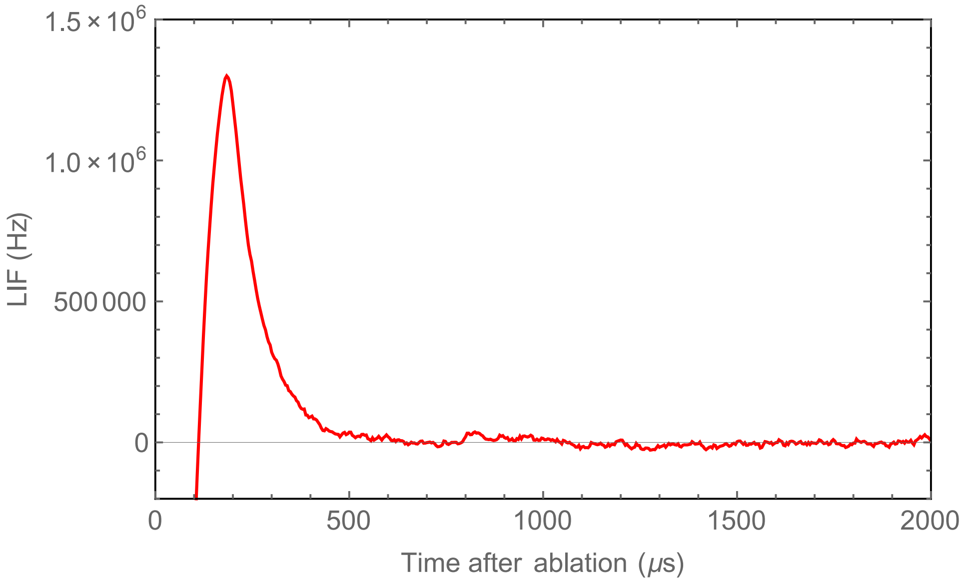

Figure 4.9: Typical laser-induced fluorescence signal from Li in the trap region as a function of time after ablation, with dummy magnets in place. The red line is the recorded signal and the blue line is an exponential fit to the decaying portion of the signal, with a decay lifetime of 76 ± 1 μs.

A typical LIF signal from Li in the trap region is shown figure 4.9. The negative signal at the start, followed by a sharp rise, is a result of the subtraction of the ‘Off-shot’ from the ‘On-shot’, a procedure which is explained in section 4.2.5. The signal saturates the PMT for the first 100 μs for both the ‘On’ and ‘Off-shots’, and the difference between the two results in what appears to be a negative LIF signal. This is an artificial consequence of our data collection procedure. It is only when the signal no longer saturates the PMT that we get an actual measure of the fluorescence from the trap region, hence the sharp rise in signal around 100 μs. Fitting an exponential to the decaying portion of the signal gives a decay time of 76 ± 1 μs. The mean speed of lithium at a temperature of 8 K is about 170 m/s. This means that, in free flight, the

Li atoms would take about 12 μs to leave the probe region of diameter 2 mm. The fact that the decay time is longer than this suggests that the atoms are still colliding with the helium gas. The timescale for atoms to diffuse to the edges of a cube of side-length L is [106],

$$\tau_D = \frac{L^2}{3\pi^2 D}, \tag{4.3}$$

where D is the diffusion coefficient and is given by [106],

$$D = \frac{3}{16}\sqrt{\frac{2\pi k_B T}{\mu}}\frac{1}{n\sigma_D}, \tag{4.4}$$

where T is the temperature, μ is the reduced mass as defined below in equation 4.7, n is the helium density inside the cube and σ_D is the Li-He diffusion cross-section, which is roughly $10^{-18}\,\mathrm{m}^2$ [81]. Using the decay time of 76 μs, this gives an estimate of the He number density in the trap region of $4 \times 10^{22}\,\mathrm{m}^{-3}$.

4.3.2 Data with magnets

The lithium temperature of 8 K after thermalisation with the buffer gas was deemed cold enough to attempt the trapping experiments, so the magnets were placed back inside the set-up. The expected fraction of 8 K atoms that can be trapped is 0.31%, as opposed to 0.87% for 4 K atoms. The presence of the magnetic field makes detection more complicated as the atoms are shifted out of resonance very quickly due to the steep gradient of the field. As explained in section 4.1, the transition is Zeeman broadened by up to 4 GHz- much more than the approximately 100 MHz resulting from Doppler broadening. This means that a single detection frequency can only probe a very small volume of the trapping region, and hence a relatively small number of atoms. As a result, absorption on the photodiode is not sensitive enough and only laser induced fluorescence detection on the PMT was used.

The trapping of Li was optimised by varying the valve pulse length, the time the valve was fired and the helium pressure behind the valve. Figure 4.10 shows fluorescence signal from lithium in the trap region as a function of time, for the settings that provided the best signal. Sometimes we noticed that there was fluorescence from the Li even if the detection beam was off-resonance by more than 100 GHz. In order to make sure that the signal was definitely lithium, both on-resonance and off-resonance data sets were taken for comparison. Both of these are

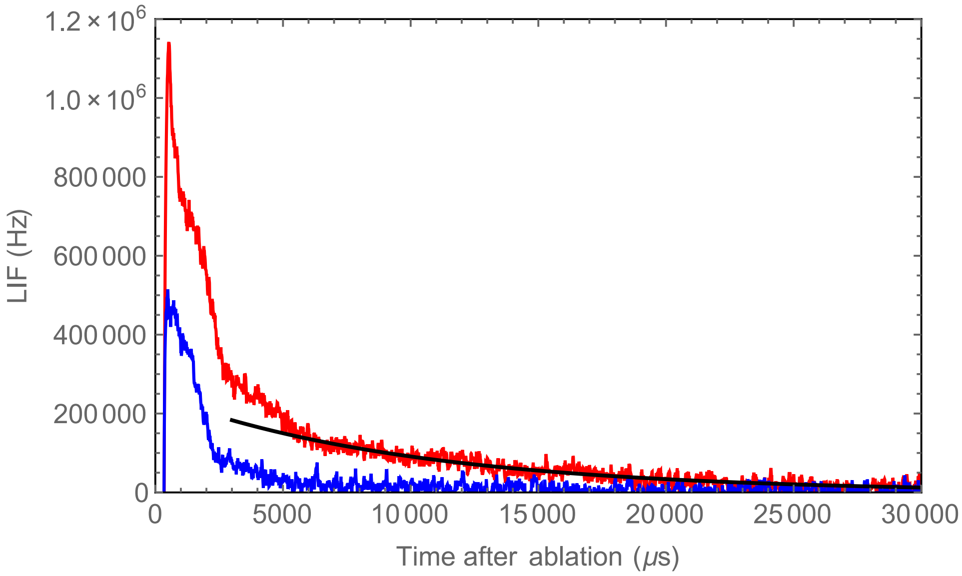

Figure 4.10: Laser-induced fluorescence signal of lithium in the trap region as a function of time, with magnets in place. The probe beam is aligned with the centre of the trap and on-resonance for the red line, and 200 GHz off-resonance for the blue line. The black line is an exponential fit to the tail of the signal, giving a decay time of 10.0 ± 0.2 ms.

plotted in figure 4.10, and it is clear that, whilst there is some signal when the detection beam is off-resonance, there is significantly more when it is on-resonance. Both have a high peak at the start which decays rapidly and is therefore mostly untrapped lithium that is still present after the ablation. After 7 ms however, there is a part of the signal with a longer decay time ($\tau \approx 10$ ms when exponential decay curve was fit to the data) that only appears with the on-resonance detection beam. Comparing this to the typical decay time with 'dummy' magnets of 76 μs, this significantly longer decay time is likely to be a result of lithium trapped for a short period of time.

The period of oscillation in the trap can be estimated by equating the force on the Li atoms due to the magnetic field (equation 2.9) with mv^2/r, the centripetal force at distance r from the trap centre. The oscillation period is then given by,

$$\tau_{osc} = \frac{2\pi r}{v} = 2\pi\sqrt{\frac{m\,r}{\mu \nabla B}}, \tag{4.5}$$

where μ is the magnetic moment of the atom, ∇B is the magnetic field gradient and m is the Li mass. In our trap, the magnetic field gradient in the x and y directions is 33 T/m and in the z direction it is 65 T/m. With the atoms likely residing in orbits no further than 5 mm from the trap centre, the oscillation period can be up to 2.7 ms. This means that the 10 ms decay timescale seen in figure 4.10 corresponds to at least a few oscillations in the trap.

The lifetime could be limited by either optical pumping into an untrapped/undetected state, collisions with background gas in the vacuum chamber, or both. To investigate the first possibility, an acousto-optic modulator (AOM) was used as a fast shutter to switch the probe beam on rapidly at various different times after ablation. This found that the Li signal was not affected by when the probe beam was turned on, which suggests that optical pumping is not a factor in limiting the Li lifetime in the trap. As a result, it is likely that the lifetime is limited by collisions with background gas in the trap region. The average time between collisions with a background gas atom is given by $\tau_{coll} = \lambda / v_{rel}$, where λ is the mean free path and v_{rel} is the relative velocity between He and Li. If it is assumed that every collision with a helium atom knocks a Li atom out of the trap, then τ_{coll} can be used as a rough estimate of the expected lifetime of Li. Using the definitions of λ and v_{rel}, the average time between collisions is,

$$\tau_{coll} = \frac{1}{\sqrt{2}n\sigma}\sqrt{\frac{\pi\mu}{8k_BT}}, \tag{4.6}$$

where σ is the collision cross section, n is the helium density, T is the temperature and μ is the reduced mass,

$$\mu = \frac{m_{Li}m_{He}}{m_{Li}+m_{He}}. \tag{4.7}$$

Converting the helium density to a pressure, we find that a 10 ms trap lifetime corresponds to a background helium pressure of about 3×10^{-7} mbar. With a base pressure in the chamber of 10^{-7} mbar when the valve is not running, a background helium pressure of 3×10^{-7} mbar when the valve is running does not seem unreasonable.

The 10.0 ± 0.2 ms trap lifetime is significantly longer than the 76 ± 1 µs decay time seen with dummy magnets with the same conditions. It is therefore very likely to be the effect of the magnetic field that is trapping the atoms. The background helium pressure that would correspond roughly to a 10 ms trap lifetime also seems very reasonable. In order for the 10 ms decay time to be caused by diffusion through the helium gas rather than magnetic trapping, equation 4.3 shows that the helium pressure in the trap region would have to be > 1 mbar. This is completely inconsistent with the pressure readings on the Penning gauge and it is therefore unlikely that diffusion is causing the long decay time. Unfortunately, the signal in figure 4.10 was not reproducible after this original data set, and as a result we were unable to conduct further tests to determine with absolute certainty that the Li was trapped.

4.4 Improvements to Set-Up

The Li trapping experiments in this chapter brought to light some issues in the set-up that had to be resolved. These issues are outlined in this section, along with the improvements made to solve them.

4.4.1 Photomultiplier tube switching circuit

Constantly exposing the PMT to the bright light from the ablation plume, and thereby saturating it, caused random voltage spikes in the PMT output, even long after the light from the ablation plume had gone. These spikes have a duration of a few tens of μs, an amplitude of about 3-5 V and an average frequency of about 50 Hz. This obscured the signal from Li in the trap and made it very noisy. Attempting to reduce the amount of ablation plume light incident on the PMT by placing a bandpass filter in front of the PMT to only transmit light at our desired detection frequency was unsuccessful as it was found that most of the light from the ablation plume was at a similar frequency. Mechanical shutters that would open and close on a quick enough time scale ($< 100\,\mu$s) to block the initial light from the ablation plume and then open again to transmit the Li fluorescence were not available. As a result, the only way to overcome this problem was to switch off the high voltage supply to the PMT during the initial ablation plume, and then switch it on again a few 100 μs after ablation. This required a fast high voltage switching circuit.

PMT operation

Figure 4.11(a) shows a typical PMT and voltage divider circuit. The PMT consists of a photocathode, a series of dynodes and an anode, with the voltage divider circuit outside the PMT ensuring each of the dynodes is held at the right voltage. The photocathode, typically held at around -1 kV, emits electrons when photons are incident upon it. These electrons are accelerated into the first dynode, producing even more electrons. As each successive dynode is held at a slightly higher voltage than the previous one, electrons are continually accelerated between each dynode, with an ever increasing number of electrons produced at each stage. This cascade eventually results in a large number of electrons hitting the anode, producing a current spike that signals that a photon has been registered on the photocathode.

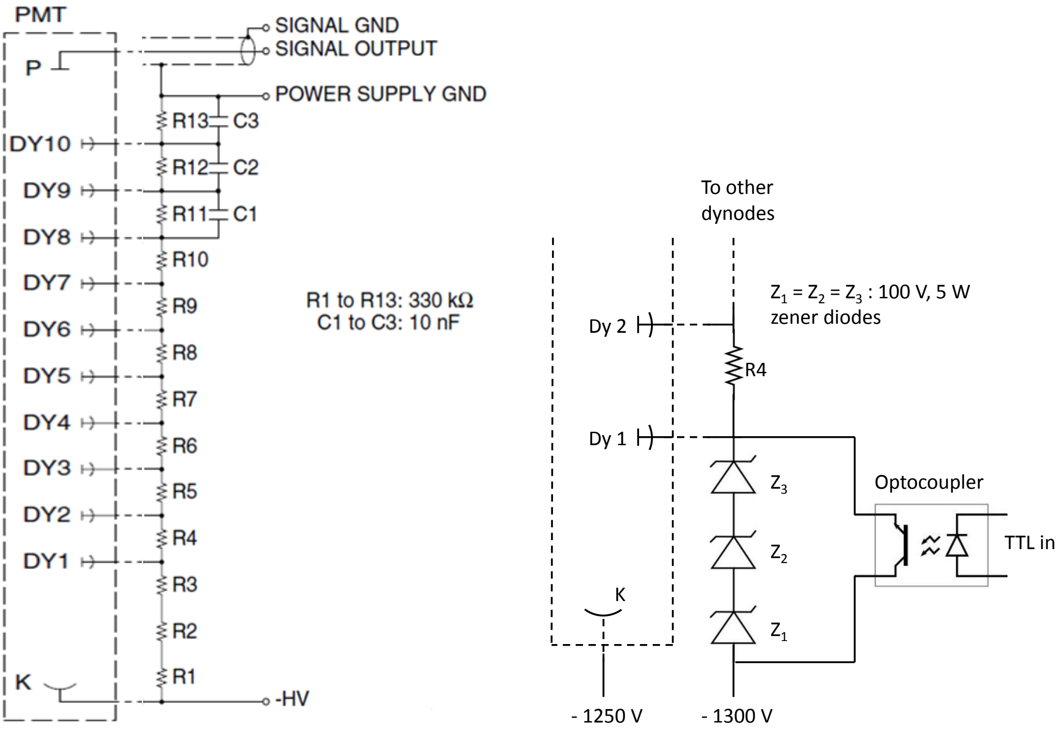

(a) PMT and voltage divider

(b) Switching circuit

Figure 4.11: (a) Schematic of a typical PMT and voltage divider circuit which holds each of the dynodes at the right voltage. K is the photocathode, Dy the dynodes, P the anode and -HV the negative high voltage supplied to the photocathode, typically around $-1\,\text{kV}$. (b) Diagram of the switching circuit. Resistors 1 to 3 in the PMT voltage divider circuit have been replaced by 100 V zener diodes. By supplying the optocoupler with either a TTL 'high' or TTL 'low', dynode 1 can be switched rapidly between -1300 and $-1000\,\text{V}$. At $-1000\,\text{V}$ the PMT will operate as normal, and at $-1300\,\text{V}$ the lower voltage on dynode 1 compared to the photocathode will suppress any gain in the PMT.

Switching circuit

The circuit that was designed to switch the PMT on and off is based on that described in [107], and is shown in figure 4.11(b). The original PMT voltage divider circuit was modified by replacing resistors 1 to 3 with 100 V zener diodes. The photocathode, held at $-1250\,\text{V}$, is supplied from a separate power supply to the one that supplies the dynodes. When the TTL to the optocoupler is 'low', the phototransistor is in the 'open' position and there is a 300 V drop across the zener diodes, holding dynode 1 at $-1000\,\text{V}$. This allows normal operation of the PMT. Switching the TTL to 'high' closes the phototransistor and pulls dynode 1 down to $-1300\,\text{V}$ (the remaining dynodes in the voltage divider circuit are therefore also pulled to a lower voltage). As dynode 1 is now at a lower voltage than the photocathode, there should be no electron acceleration between the photocathode and dynode 1 and hence the PMT gain is

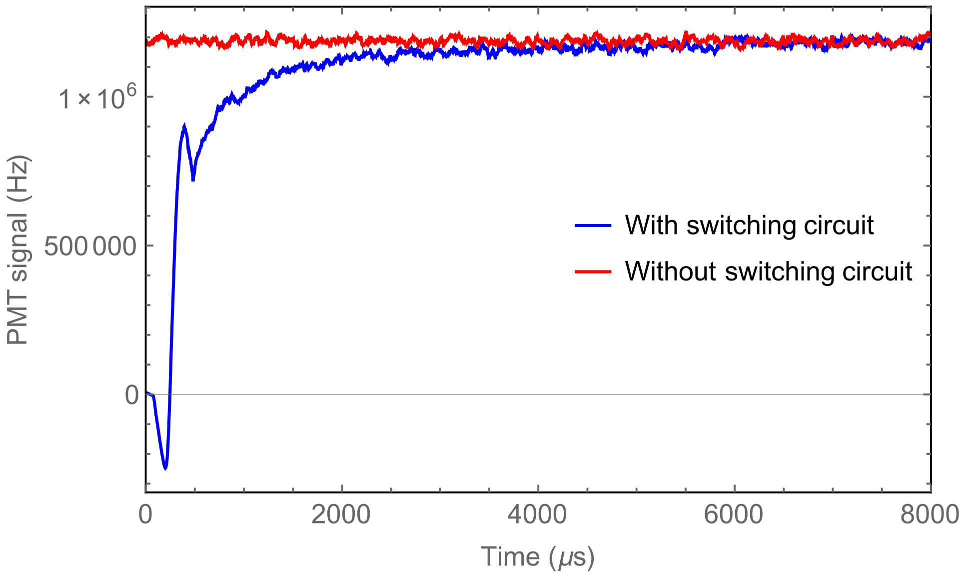

Figure 4.12: Comparison of signal output from the PMT with and without the switching circuit, with a constant background photon rate incident on the PMT. Red line: no switching circuit; blue line: PMT switched on at $t = 0$.

greatly suppressed. Switching the TTL to 'low' again returns the PMT to normal operation.

Figure 4.12 shows a constant background signal recorded on the PMT with and without the switching circuit. When using the switching circuit to switch the PMT on at $t = 0$, it can be seen that, whilst the PMT switches on within a few 100 μs, the gain is initially slightly low, rising up to the full gain within about 5 ms. This is likely to be a result of the other dynodes taking time to reach the right voltage as this timescale is of the same order of magnitude as the time constant for an RC circuit of $\tau = \mathrm{R}_{11}\mathrm{C}_1$. Despite the gain being low for the first few ms, the PMT will still see any signal present, it will just be slightly lower than if the PMT had reached the full gain.

Figure 4.13 shows the signal on the PMT with and without the switching circuit for a Li data set (i.e. with ablation and detection lasers on and the valve firing). Without the switching circuit (red line), the ablation plume saturates the PMT for 600 μs and then decays. The blue line shows that, even with the PMT 'off', the ablation plume still saturates the PMT for the first 200 μs. The switching circuit is then switched 'on' after 200 μs, with the PMT taking a few 100 μs to actually switch on. The LIF signal observed when the switching circuit is used decays much more rapidly than when it is not used. This leads to the conclusion that a large part of the signal recorded without the switching circuit is artificial, and is merely a residue of the huge saturation the PMT experiences due to the ablation plume. This makes the use of the switching circuit crucial in distinguishing real signal from artificial signal. The timescale on

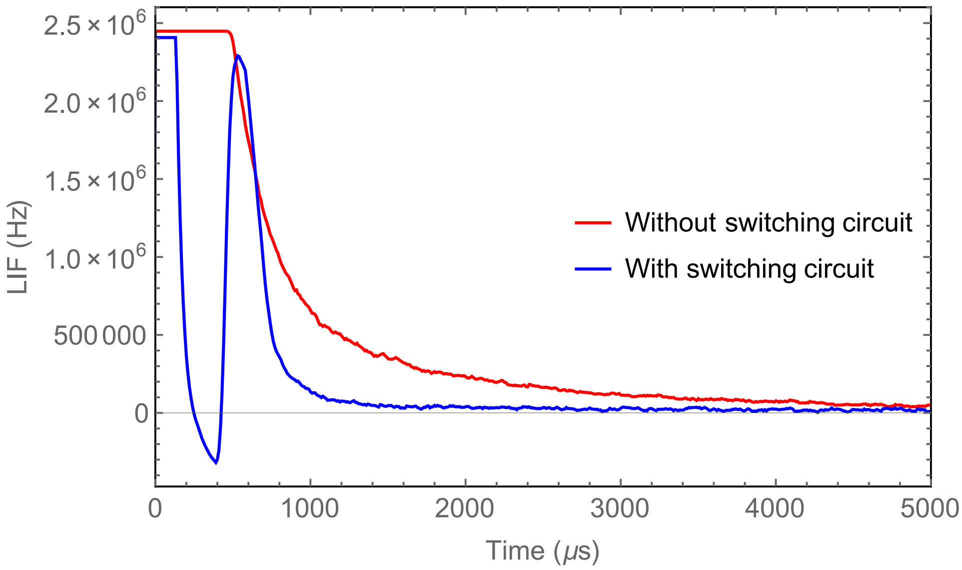

Figure 4.13: Comparison of laser-induced fluorescence signal on the PMT with and without the switching circuit, with ablation and detection lasers on. Red line: no switching circuit; blue line: PMT switched on 200 μs after ablation. To illustrate the effect of the switching circuit, only the 'On-shots' are shown, rather than the usual subtraction of the 'Off-shot' from the 'On-shot'.

which this artificial PMT saturation signal decays is about 2 ms. This is not too far removed from the 10 ms trap lifetime seen in figure 4.10. It is therefore possible that this 10 ms trap lifetime is actually an artefact of the PMT saturation rather than real LIF signal.

4.4.2 Lengths of He gas pulses

One of the requirements for our trapping arrangement to work is that the He buffer gas is introduced as a short pulse into the trapping region, and subsequently disappears quickly once the Li has thermalised with it. To measure the lengths of the He gas pulses coming out of the valve, a fast ionisation gauge (Beam Dynamics Model FIG-1 Fast Ionization Gauge) was mounted at the exit of the tube that guides the pulse into the trap region (see figures 4.3 and 4.14(a)). The length of the gas pulse with various configurations was measured: 1) with the valve tube mounted as in figure 4.3; 2) with a 1 cm straight valve tube mounted on top of the valve; 3) from the valve directly. These three configurations are shown in figure 4.14.

Figure 4.15 shows the results from these measurements, with the total area under each curve (and therefore total amount of gas expelled) normalised to 1. With the original valve tube, the rise time of the gas pulse is quite long and the decay time of 3700 μs is likely to be far too long for the requirements of the experiment. Having a shorter, straight tube makes the pulse shorter and concentrates more of the pulse in the initial peak. It also reduces the decay time

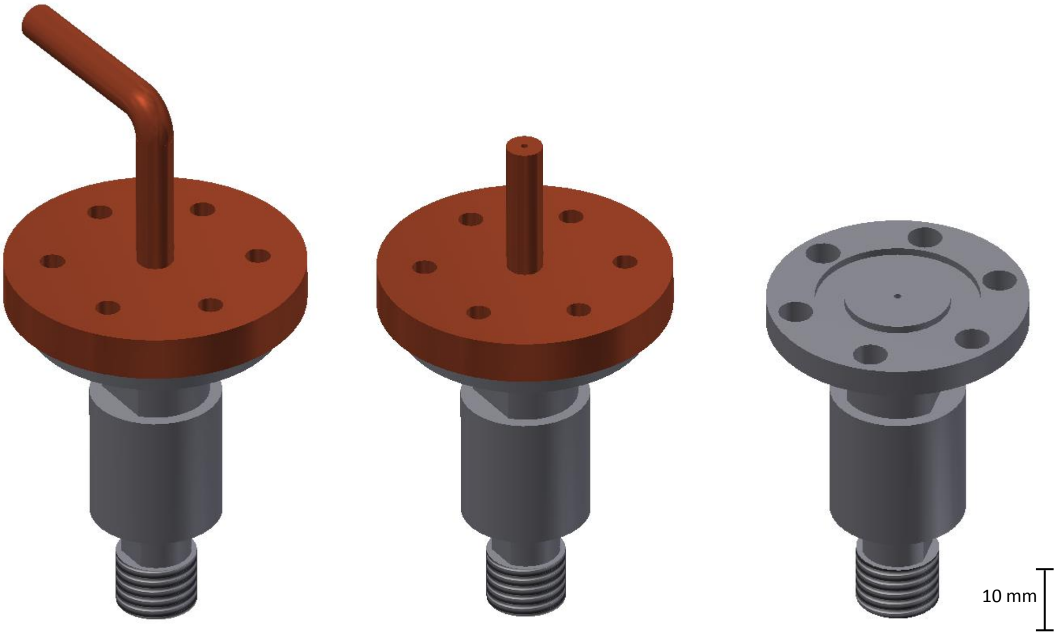

(a) Valve with original tube

(b) Valve with shorter, straight tube

(c) Valve on its own

Figure 4.14: The configurations of the solenoid valve (Parker Series 9 pulsed valve) whose He pulse lengths were tested, with the original valve tube (a), the shorter, straight valve tube (b) and the valve on its own (c).

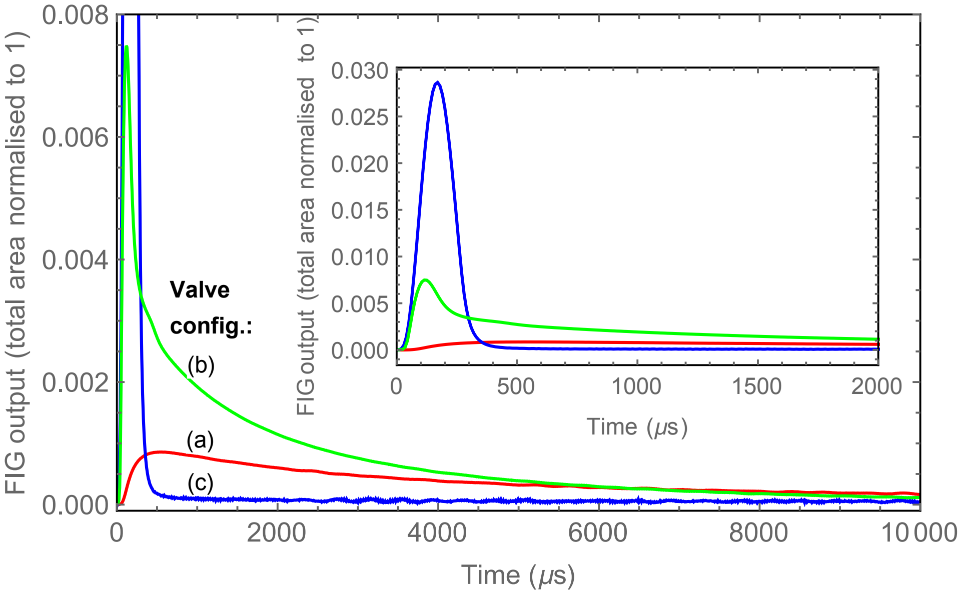

Figure 4.15: Comparison of He pulse lengths from the valve, in the configurations from figure 4.14, measured using a fast ionisation gauge (FIG). By fitting exponential decay curves to the data the decay times, τ, were found. Red line: original valve tube, $\tau = 3700\,\mu$s; green line: 1 cm straight valve tube, $\tau = 1300\,\mu$s; blue line: valve directly, $\tau = 30\,\mu$s. Plots are normalised so that the area under each curve is 1. Inset shows the same plot with different axis scales to show the full height of the blue peak.

to 1300 μs. It is clear however, that the pulse from the valve directly is much shorter than the pulses coming from either of the valve tubes, with all of the gas coming out in one clean pulse without a long tail. The decay tail that emerges when a tube is mounted on top of the valve is likely to be detrimental to the trap lifetimes as helium is still flowing out of the tube, and therefore colliding with trapped Li atoms, for tens of milliseconds after the valve has fired.

Horizontal valve set-up

In light of the data on the lengths of the gas pulses from the valve (see figure 4.15), it was decided that the valve should be mounted in such a way that a valve tube is not necessary anymore. This would avoid any problems due to long gas pulses emerging from the valve tube. Figure 4.16 shows the new design for the trap arrangement, with the valve mounted horizontally, directly facing the trapping region. This avoids the use of any valve tubes, and therefore gives much more control over the length of the gas pulse entering the trap region.

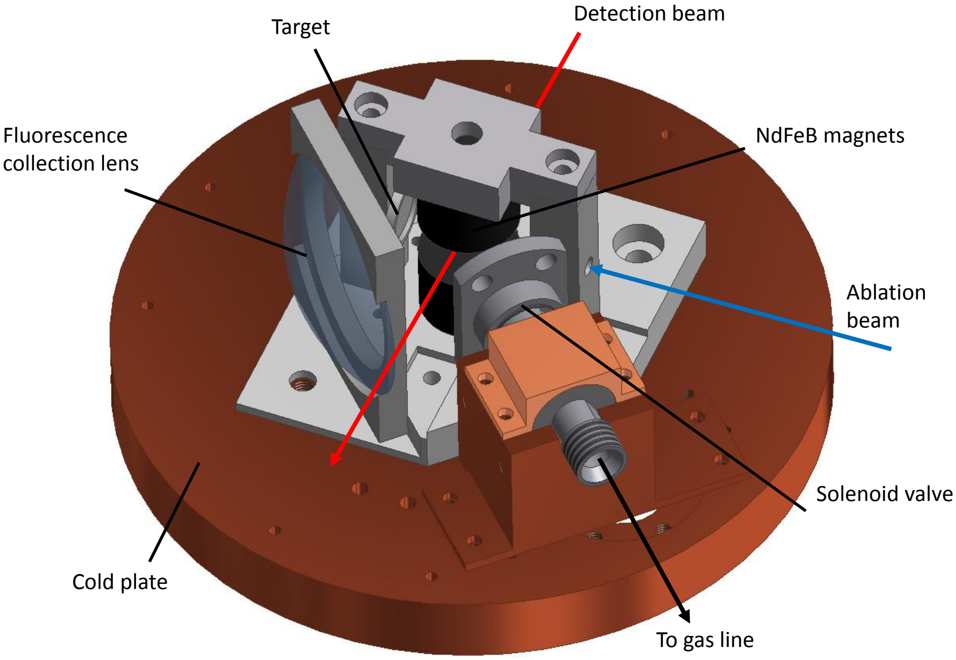

Figure 4.16: New design for the magnetic trap with the valve mounted horizontally, therefore pulsing the helium directly into the trapping region.

4.5 Conclusion

The thermalisation of the Li atoms with the helium buffer gas was studied with aluminium cylinders in place of the magnets. Absorption and LIF spectra of Li were taken, and by measuring the widths of the transition peaks at different times after ablation, the temperature of the Li was calculated as a function of time. It was found that the Li thermalised with the buffer gas within about $500\,\mu$s of the ablation laser firing. However, the final temperature of the Li atoms was in the region of $8\,$K, rather than the $4\,$K temperature of the cold plate. This was attributed to the fact that the valve warms up slightly as it fires the helium pulse, with the helium therefore exiting the valve at a temperature slightly above $4\,$K. The open nature of the trap means that the helium does not then have many walls to collide with to cool back down to $4\,$K.

The magnets were returned to the set-up to attempt the Li trapping experiments. The best lifetime observed for Li in the trap was $10.0 \pm 0.2\,$ms. This is thought to be limited by collisions with background helium gas that is still present in the trap region; an assumption that is consistent with a rough calculation of the helium pressures involved. However, later tests showed that this observed decay time could be an artefact of the saturation of the PMT due to the ablation plume light.

CHAPTER 5

DYSPROSIUM EXPERIMENTS

From the lithium experiments in chapter 4, we were unable to improve upon, or reproduce, the data from figure 4.10, where Li was detected in the trap with a decay time of 10 ms. As a result, we decided to try the experiment with dysprosium instead of lithium. With a magnetic moment in the most favourable state of $10\mu_B$, rather than lithium's $1\mu_B$, the trap depth of dysprosium (Dy) in the same trap is a factor of 10 higher than that of of Li; 1.4 K compared to 0.14 K. Assuming a starting atomic distribution at 4 K, the fraction of Dy atoms that can be trapped in our arrangement is about 20% of those atoms that are in the most favourable trapping state. The corresponding fraction for Li is just under 1%. Changing the set-up to trap Dy was straightforward as it simply required the Li target to be changed for a Dy metal target, with the rest of the set-up remaining identical to that for the Li experiments, described in section 4.2. The detection frequency for Dy at 684 nm (see section 5.1) does not differ hugely from the 671 nm Li transition, and the same laser system could be used as for Li by just replacing the 671 nm diode in the external cavity diode laser for a 684 nm diode.

With the ten-fold increase in trap depth compared to Li, trapping Dy atoms should be significantly easier. However, the higher mass of Dy (162.5 a.m.u. compared to 7 a.m.u for Li) means that significantly more collisions will be necessary in order to cool the Dy atoms down to 4 K after ablation. Using equation 2.6, the number of collisions required to cool Dy atoms to within 30% of the buffer gas temperature of 4 K, assuming they come off the ablation target with a temperature of 10 000 K, is 190. This is compared to the approximately 20 collisions required to cool Li to the same temperature. Nevertheless, the advantages to be gained from the higher trap depth make Dy a worthwhile candidate for trapping experiments.

5.1 Dysprosium details

We detect Dy by exciting the 683.731 nm transition from the ground $4f^{10}6s^2\ {}^5I_8$ state to the excited $4f^9(^6H^0)5d6s^2\ {}^5I_8$ state. The natural linewidth of this transition is 95±13kHz [108]. This is a relatively unstudied transition, with a more commonly used transition being the ${}^5I_8$ to ${}^5K_9$ transition at 421 nm (see for example [109]). However, our transition was chosen because of the ease of changing our detection laser from the Li frequency to this Dy frequency. Additionally, our transition is not Zeeman broadened hugely in the magnetic field (see section 5.1.1), making detection easier.

Isotope	Mass (a.m.u.)	Abundance (%)	Nuclear spin
^{160}Dy	159.93	2.34	0
^{161}Dy	160.93	18.91	5/2
^{162}Dy	161.93	25.51	0
^{163}Dy	162.93	24.90	5/2
^{164}Dy	163.93	28.18	0

Table 5.1: The different isotopes of dysprosium along with their mass, relative abundance and nuclear spin.

Table 5.1 lists the five stable isotopes of dysprosium, along with their mass, relative abundance and nuclear spin. The two odd isotopes have nuclear spin and therefore exhibit ground state hyperfine structure. With a nuclear spin of $I = 5/2$ and angular momentum of $J = 8$, the ground states of these two isotopes are split into six hyperfine levels from $F = 11/2$ to $F = 21/2$. The hyperfine shift of a level with total angular momentum F is given by

$$E_{HFS} = \frac{A}{2}K + \frac{B}{4}\frac{\frac{3}{2}K(K+1) - 2I(I+1)J(J+1)}{I(2I-1)J(2J-1)} \tag{5.1}$$

where $K = F(F+1) - J(J+1) - I(I+1)$ and A and B are the magnetic dipole and electric quadrupole coupling constants [110]. Using the values for the A and B constants from [111], the hyperfine splitting of the ground state can be calculated and is shown for both isotopes in figure 5.1(a). Since the angular momentum of the excited state is also $J = 8$, it is also split into six hyperfine levels, shown in figure 5.1(b) for both isotopes and calculated from the A and B constants taken from [108].

In addition to the hyperfine structure of ^{161}Dy and ^{163}Dy, there is a shift in transition frequency for each isotope. These isotope shifts for the 684 nm transition are shown in table 5.2, where the ^{164}Dy transition is taken to be the centre frequency, with the shifts of the other

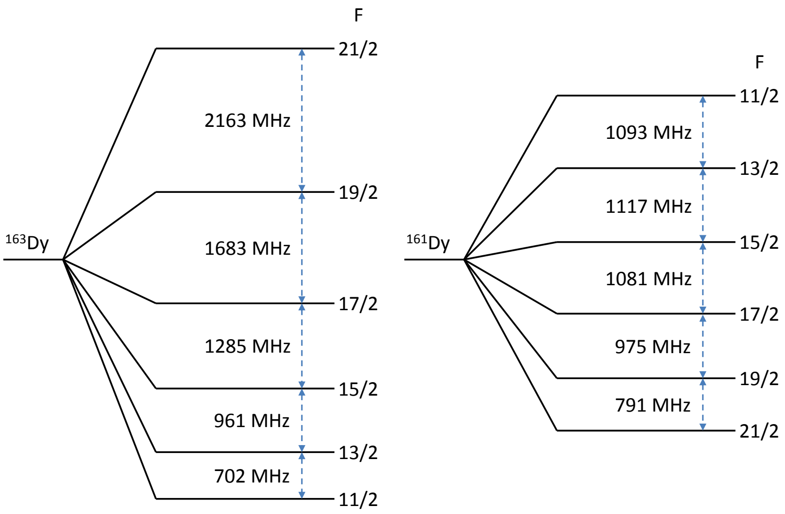

(a) Ground state

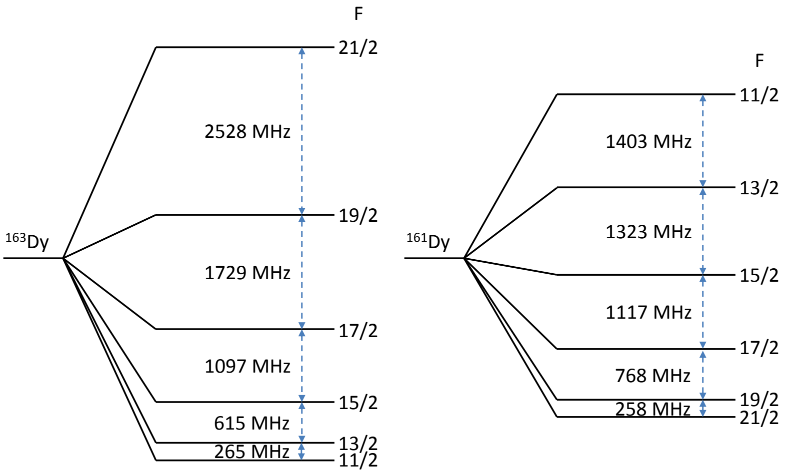

(b) Excited state

Figure 5.1: Hyperfine splitting of the ground $^5I_8$ state and the $4f^9(^6H^0)5d6s^2\ ^5I_8$ excited state of Dy, for both the 161 and 163 isotopes.

isotopes shown with respect to this frequency.

The combination of five different isotopes, with two of them exhibiting hyperfine structure, leads to a relatively complex field-free spectrum of Dy. As our transition is a $\Delta J = 0$ transition with large J, the strongest transitions between the hyperfine levels of the odd isotopes are the $\Delta F = 0$ transitions. This leads to six visible hyperfine components for each of the two odd isotopes, with the other hyperfine components negligibly weak by comparison. These, in addition to the three transitions for the even isotopes, mean a total of 15 transitions will contribute to the full spectrum. Using the values from figure 5.1 and table 5.2, the theoretical spectrum can be calculated, and is shown in figure 5.2. The relative intensity of each line

Isotope	Isotope shift (MHz)
^{164}Dy	0
^{163}Dy	-823
^{162}Dy	-1091
^{161}Dy	-2099
^{160}Dy	-2309

Table 5.2: Isotope shifts of the Dy isotopes with respect to the ^{164}Dy isotope. Values taken from [108].

is estimated from the abundance of each isotope, with the hyperfine components for the odd isotopes each assumed to have a line intensity proportional to their statistical weights of $2F+1$. Whilst this is not exactly correct, it suffices as an approximation to the spectrum, as shown by the fact that the theoretical spectrum is in close agreement with the measured Doppler-free spectrum from [108]. As is evident from figure 5.2, the ^{164}Dy line is the most intense line that is furthest away from other transitions that might artificially broaden it, and as a result this is the line we detect for most of the data in this chapter.

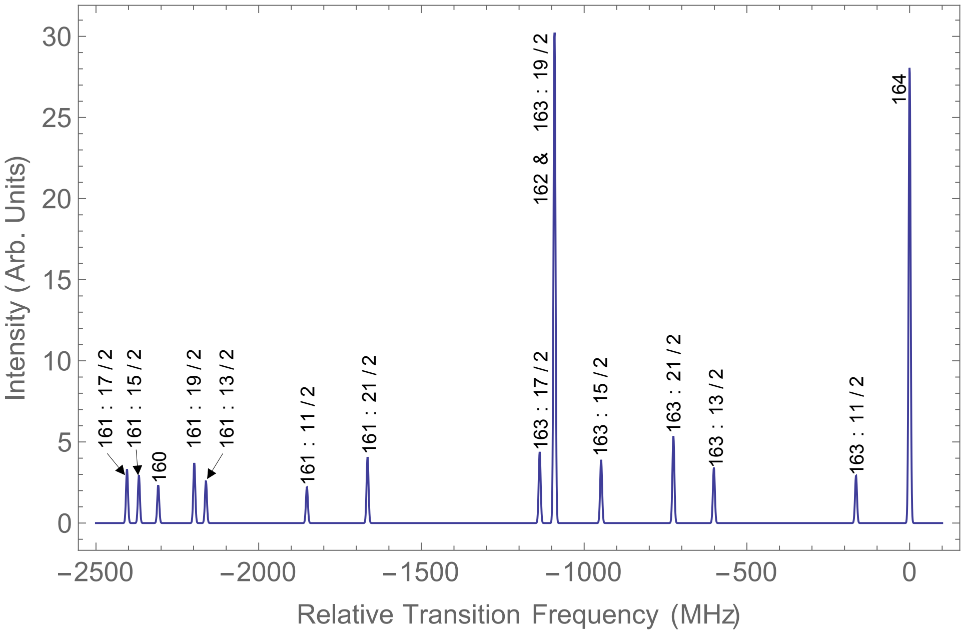

Figure 5.2: Theoretical field-free, Doppler-free Dy spectrum, showing the transitions from all five isotopes. Each transition is labelled in the form 'isotope: F sub-level', except where the isotope does not display hyperfine structure, in which case it is just labelled by the isotope number. Transitions with $\Delta F = \pm 1$ are very weak compared to those with $\Delta F = 0$, and so are not shown.

From the excited state there exists an additional decay channel down to the metastable $4f^{10}6s^2\ ^5I_7$ state [112]. A previous study of the 684 nm transition in [108] measured an upper limit to the branching ratio between these two decay channels of 1:100. The existence of this additional decay channel has to be considered in the detection process which could artificially lower the measured trap lifetime by optically pumping the atoms into an undetected state.

5.1.1 Zeeman shift

In the presence of the magnetic field, the Dy energy levels will be split into their Zeeman sub-levels. For the even isotopes without any hyperfine splitting, the Zeeman shift is given by

$$\Delta E = g_J \, m_J \, \mu_B \, B, \tag{5.2}$$

where g_J is the g-factor corresponding to the total electronic angular momentum, m_J denotes the magnetic sub-level and μ_B is the Bohr magneton. For the ground state of Dy, $g_J = 1.242$ [113] and the state will split into its 17 m_J sub-levels from $m_J = -8$ to $m_J = 8$. Figure 5.3 shows how the ground state of the even Dy isotopes is split into its magnetic sub-levels as a result of the magnetic field. The $m_J = 8$ level is the most favourable for trapping as it has the highest magnetic moment of almost $10\mu_B$, and therefore has a trap depth in our trap of 1.4 K. Note that the fine-structure splitting between $J = 8$ and $J = 7$ is 124 THz, which is more than 3 orders of magnitude larger than the Zeeman shift in all relevant fields. Mixing between J levels is therefore negligible and the Zeeman shift is indeed linear at all fields, as expressed by equation 5.2.

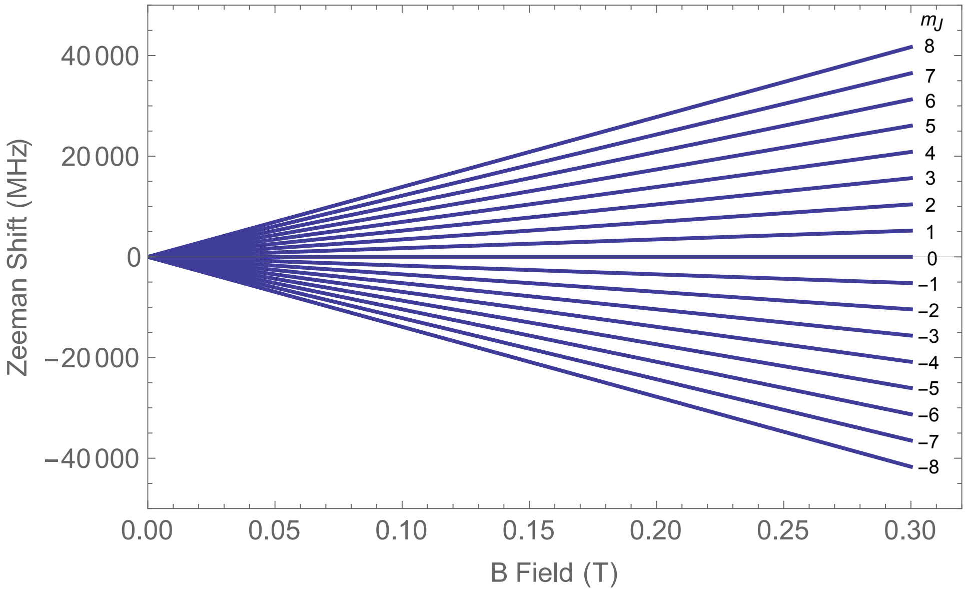

Figure 5.3: Zeeman splitting of the ground state of Dy, for the even isotopes that exhibit no hyperfine structure.

The excited state used to detect the atoms has $g_J = 1.25$ [113], which does not differ significantly from that of the ground state. Given that the excited state is also a $J = 8$ level, the Zeeman shift for the excited state will be almost exactly the same as that in figure 5.3 for the ground state. This makes detection in the magnetic field significantly easier for Dy than for Li, as the $\Delta M = 0$ components will hardly be broadened in the inhomogeneous magnetic field. For the two odd isotopes with hyperfine structure, the Zeeman splitting is significantly more complicated because the magnetic field strongly mixes states of different F that have the same value of m_F. However, since we have chosen to detect primarily on the ^{164}Dy transition, these do not have to be considered in detail.

5.2 Dummy magnet experiments

As with the Li experiments, the first step here was to replace the magnets in the trap with aluminium cylinders of the same shape and size, which we refer to as 'dummy magnets'. This allows the initial parameters required to cool the Dy atoms to be determined without the complicating effects of the magnetic field.

5.2.1 Measured Dy spectrum & temperature

Despite coming to the conclusion in section 4.4.2 that it was necessary to switch to a horizontal valve set-up, the data here were still taken with the original 'valve tube' set-up from figure 4.3 as the new horizontal valve set-up was still being built. Figure 5.4 shows a typical Dy absorption spectrum after thermalisation with the helium buffer gas, taken in exactly the same way as for the Li spectra (see section 4.2.5). The spectrum shows the lines from all of the isotopes and hyperfine levels that are expected from the theoretical spectrum in figure 5.2, but Doppler broadening of the lines means that they are not all individually resolved.

Since the frequency spacing between each of the peaks in the spectrum is known from the theoretical calculations, it is possible to fit Gaussians to all 15 peaks of the full spectrum from figure 5.4 by fixing the spacing between the peaks and having as the free parameters: the height of each peak, the width of the peaks (assumed to be the same for each peak) and an overall constant offset. For later times when the signal is lower, only the two main ^{164}Dy and ^{162}Dy peaks can be resolved, and so only two Gaussian peaks are fit to the spectrum. The function used to fit to the spectra is

$$\text{const.} + \Sigma_i \, A_i \mathrm{e}^{-\frac{(f-f_{0i})^2}{2\sigma^2}}, \tag{5.3}$$

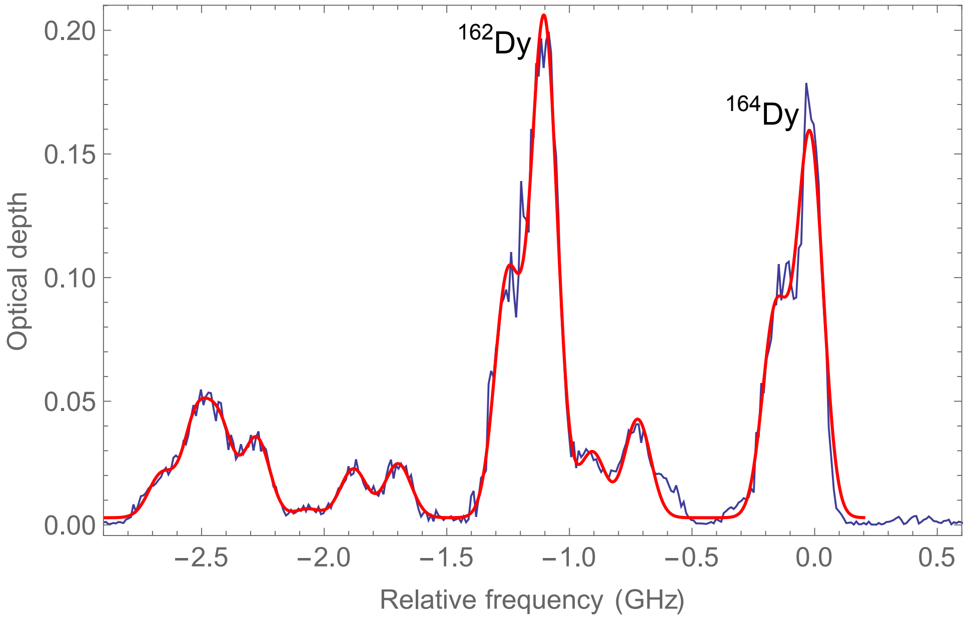

Figure 5.4: Measured absorption spectrum of Dy after thermalisation with the He buffer gas (blue line), with dummy magnets. Frequencies are relative to the ^{164}Dy transition. The two largest peaks result from the ^{164}Dy and ^{162}Dy isotopes. The spectrum was taken from data gated between 80-100 μs after ablation, with a VPL $= 200\,\mu$s and $t_{valve} = -900\,\mu$s. The red line is a 15-Gaussian fit (equation 5.3) to the full spectrum, with $\sigma = 53 \pm 1$ MHz.

where the A_i are the heights of the peaks, σ is the width of the peaks, f is the frequency and the f_{0i} are the centres of each peak and are fixed from the theoretical values. The sum runs from $i = 1$ to 15 when all 15 peaks are resolved and from $i = 1$ to 2 when only the two main peaks are resolved. A typical full fit to all 15 peaks is plotted in figure 5.4, showing good agreement with the measured spectrum. Not all 15 peaks from the theoretical spectrum in figure 5.2 are individually resolved in the fit in figure 5.4 as a result of Doppler broadening. Fitting equation 5.3 to spectra at different times after ablation allows the widths of the peaks to be determined as a function of time. Assuming that the width of the peaks arises purely as a result of Doppler broadening, it is related directly to the temperature of the Dy atoms through equation 4.2. The temperature of the Dy atoms can then be determined as a function of time after ablation, and is shown in figure 5.5. This shows that the Dy thermalises with the buffer gas within 200 μs which, as with the Li experiments (see section 4.3.1), seems to be at a temperature of around 8 K. A two-exponential decay curve fit to the data is also shown, which confirms the final temperature of about 8 K and gives an initial temperature of about 700 K. Despite the increased number of collisions required for thermalisation due to its heavier mass, the time taken for the Dy to thermalise is significantly shorter at 200 μs than the 500 μs required for Li. Given that the valve settings are identical for both the Dy and Li data (VPL $= 200\,\mu$s and $t_{valve} = -900\,\mu$s), the Dy

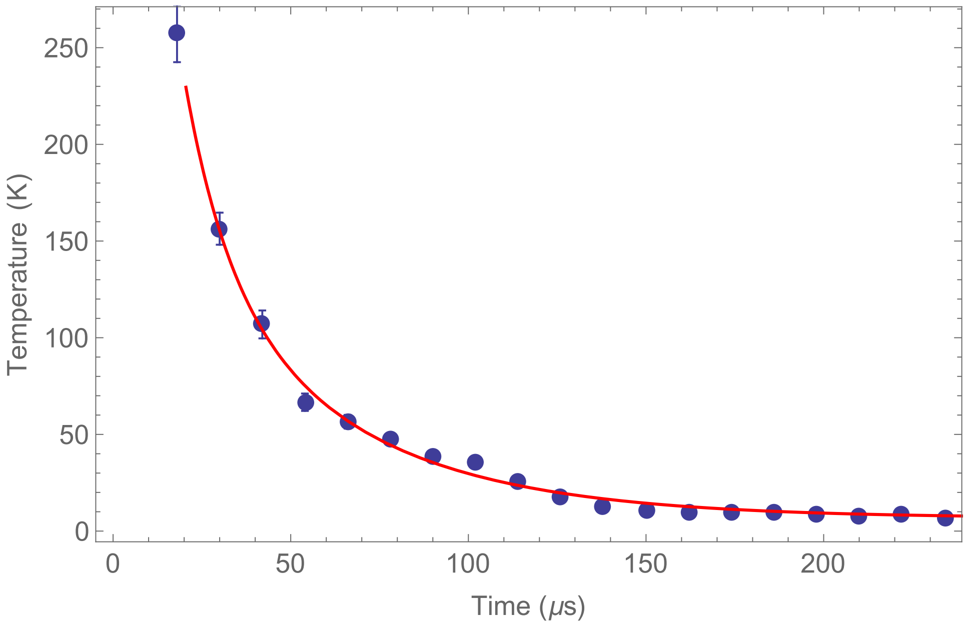

Figure 5.5: Temperature of Dy as a function of time after ablation, showing that it thermalises with the He buffer gas within 200 μs. A two-exponential decay fit to the data is also shown, indicating a final temperature of about 8 K.

thermalisation time would have been expected to be much longer than the Li thermalisation time. This could be explained by significantly different ablation dynamics for the Dy target compared to the Li target, particularly as a slightly higher ablation pulse power of 12.5 mJ had to be used for the Li target compared to the 10 mJ used for the Dy target. This could mean the Dy atoms come off the target at a lower temperature than the Li atoms. It is also possible that the cross-section for collisions with helium is much larger for Dy than for Li.

5.2.2 Velocity measurements of the Dy cloud

In order to ascertain whether the Dy cloud is stationary after thermalisation with the He buffer gas, the centre-of-mass velocity of the Dy cloud was measured. This was done by introducing a counter-propagating probe beam into the set-up, overlapped exactly with the original probe beam. Two absorption spectra were then taken simultaneously, one for each of the probe beams, using two separate photodiodes. If the Dy cloud has no centre-of-mass velocity, the two spectra should overlap perfectly, but if the cloud does have some centre-of-mass velocity in the direction of the probe beams then the two spectra will be Doppler shifted with respect to each other. The frequency difference between the two spectra is then related to the velocity by

$$v_{cloud} = \frac{\Delta f}{2} \frac{c}{f_0} \tag{5.4}$$

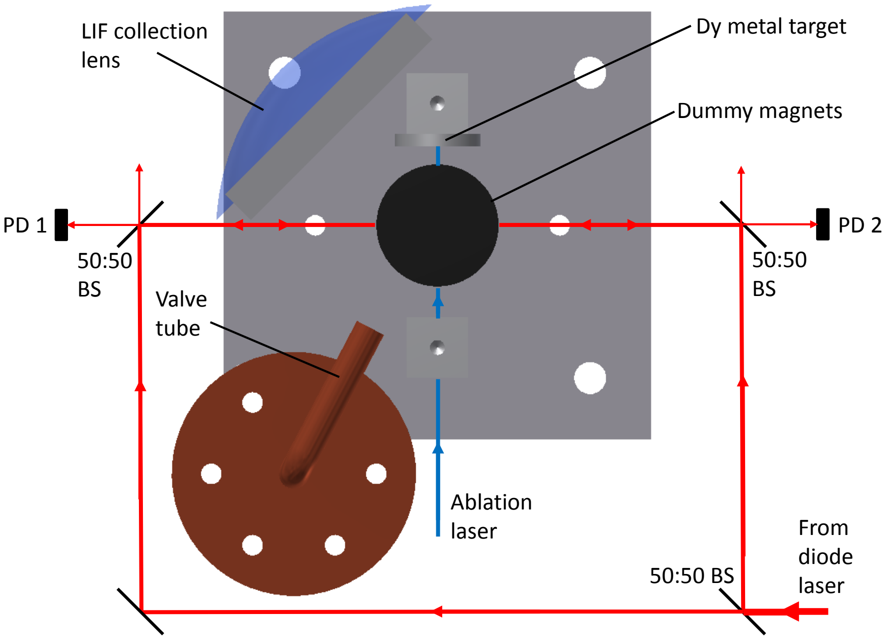

Figure 5.6: Top-down view of the counter-propagating beam arrangement in the original set-up with the valve tube. PDs are photodiodes and BSs beamsplitters.

where v_{cloud} is the cloud velocity in the probe beam direction, f_0 is the actual transition frequency, c is the speed of light and Δf is the frequency difference between the two spectra. The frequency difference between the two spectra is divided by two to account for the fact that one spectrum will be shifted to a higher frequency and one spectrum to a lower frequency, with the actual transition frequency f_0 being exactly halfway between the two spectra. Note that this measurement only determines the centre-of-mass velocity component in the direction of the probe beam.

Original valve tube

Figure 5.6 shows how the counter-propagating probe beam arrangement fits into the original set-up with the valve tube (see figure 4.3) that was used for all previous data in this thesis. The probe beam is split using a 50:50 beamsplitter, with one beam sent into the chamber in the original direction and the other beam in the opposite direction, overlapped exactly with the original probe beam. Two further beamsplitters pick off the appropriate beams to allow the two photodiodes to record the absorption spectra. The probe beam light is delivered to the experiment in the same way as depicted in figure 4.6, with the optical isolator in the set-up preventing any light feeding back into the diode laser.

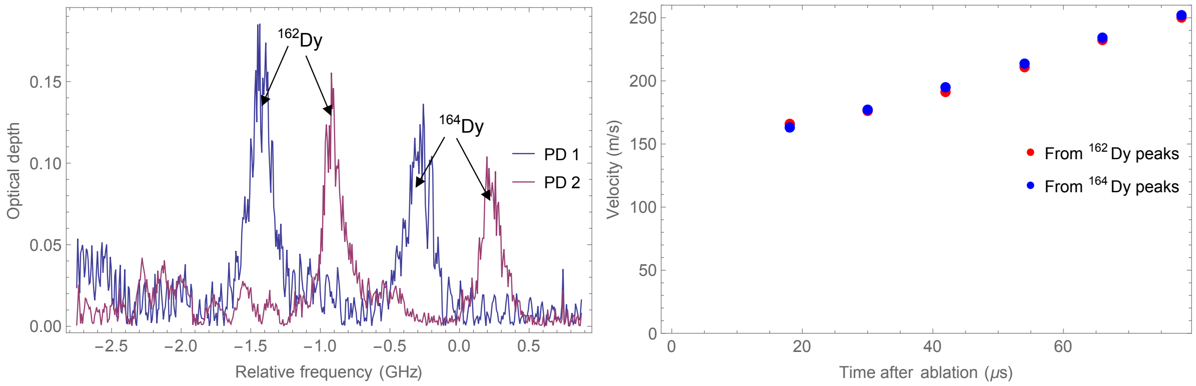

(a) Absorption spectra

(b) Centre-of-mass velocity of Dy cloud against time after ablation

Figure 5.7: Counter-propagating probe beam data with original valve tube. (a) Absorption spectra from the two photodiodes (gated 36-48 μs), showing the frequency shift between the two spectra. Frequencies are relative to the ^{164}Dy transition. (b) Centre-of-mass velocity of the Dy cloud as a function of time after ablation, with the red points calculated from the shift between the ^{162}Dy peaks and the blue points from the shift between the ^{164}Dy peaks. Positive velocities are in the direction away from PD 1. Error bars are on the order of a few m/s and are not visible on this scale. For these data, VPL $= 235\,\mu$s and $t_{valve} = -600\,\mu$s.

Figure 5.7(a) shows the absorption spectra of Dy from the two probe beams, from data gated between 36-48 μs after ablation. There is a clear frequency shift between the spectra from PD 1 and PD 2, signifying a centre of mass velocity of the Dy cloud away from PD 1. The direction that the Dy cloud appears to be travelling in is consistent with the fact that the valve tube points slightly away from PD 1 (see figure 5.6). This suggests that the Dy cloud is entrained in the flow of gas from the valve. By fitting Gaussians to the ^{162}Dy and ^{164}Dy peaks of each spectrum, the centres of those peaks can be determined. The frequency shift between the ^{162}Dy peaks and between the ^{164}Dy peaks of the two spectra can then be calculated and inserted into equation 5.4 to extract a velocity for the Dy cloud. Doing this for several different times after ablation gives a velocity of the Dy cloud as a function of time, shown in figure 5.7(b). The velocities calculated from the 162 and 164 peaks are in very close agreement, as expected, and show that there is a centre-of-mass velocity of the Dy cloud from $t \approx 20\,\mu$s onwards. The cloud starts off at a velocity of about 160 m/s and increases to about 250 m/s at 80 μs, after which time the Dy signal becomes too low to observe the absorption spectra. The large velocity can be explained by the helium pulse expanding supersonically into vacuum and, given that it has almost no walls to collide with after it comes out of the valve tube, it will keep this velocity when the Dy is ablated into it. The Dy then becomes entrained in the helium flow, giving it a large centre-of-mass velocity. The maximum supersonic flow velocity of a beam is given by

[114],

$$u = \sqrt{\frac{\gamma}{\gamma - 1}}\sqrt{\frac{2k_BT}{m}}, \tag{5.5}$$

where T is the temperature, m is the mass of the particles and $\gamma = 5/3$ for a monatomic gas. We know from the Li experiments and figure 5.5 that the He emerges from the valve at a temperature in the region of 8 K. Inserting this, along with the mass of He, into equation 5.5 gives a supersonic velocity of 288 m/s for an 8 K helium pulse. As the valve tube is at an approximately 45° angle with respect to the probe beam, the expected supersonic velocity in the probe beam direction is 204 m/s. Figure 5.7(b) shows the velocity of the Dy cloud increasing over time from 160 m/s to 250 m/s. This increase might be explained by the initial ablation plume arresting the He velocity slightly, with further helium from the valve tube (see section 4.4.2) subsequently pushing this velocity back up again. The final velocity of 250 m/s is greater than the 204 m/s expected from an 8 K supersonic beam. This discrepancy could arise due to the fact that the angle of the valve tube with respect to the probe beam is not exactly known, but is more likely to be indicative of the temperature of the valve increasing slightly, with the helium therefore emerging slightly warmer. A supersonic velocity in the direction of the probe beam of 250 m/s would correspond to a helium temperature of about 12 K. Given that the trap depth in our trap of 1.4 K corresponds to a capture velocity of 14.5 m/s for Dy, it is extremely unlikely that we would have been able to trap Dy in this arrangement as the centre-of-mass velocity of the Dy cloud is far too high.

Valve tube with copper cylinder

In order to contain the helium pulse from the valve tube and to prevent the Dy cloud from having such a large centre-of-mass velocity, a copper cylinder, depicted in figure 5.8(a), was introduced into the set-up around the (dummy) magnets and the trap region. It was screwed onto the cold plate so that it thermalised to 4 K, with the helium pulsed into the copper cylinder through a small hole. The helium can then collide a few times with the 4 K walls of the cylinder to lose its centre-of-mass velocity and become more of a cloud. The Dy can then be ablated into the helium cloud, rather than being entrained in a flow from the valve tube as before. Figure 5.8(b) shows how the copper cylinder fits into the set-up with the counter-propagating beams. The top of the copper cylinder is left open so that the helium can still leave the trap region rapidly once the Dy has thermalised.

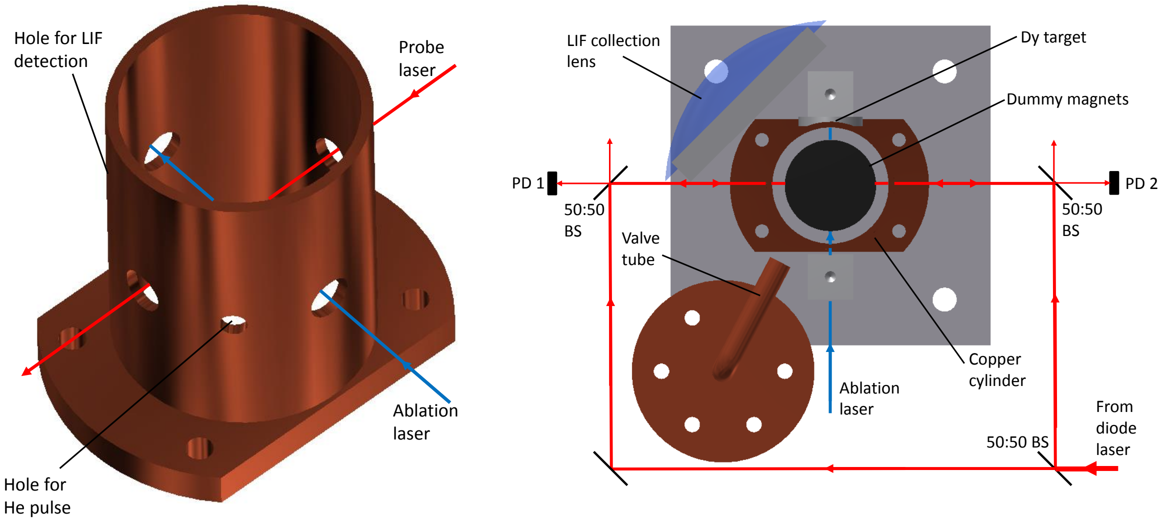

(a) Copper cylinder

(b) Beam arrangement with copper cylinder

Figure 5.8: (a) Copper cylinder used to contain the initial helium pulse, showing the holes for the probe and ablation lasers, as well as for the helium pulse and LIF detection. (b) Top-down view of the counter-propagating beam arrangement with the copper cylinder in place. PDs are photodiodes and BSs beamsplitters.

To determine whether the copper cylinder reduces the centre-of-mass velocity of the Dy cloud, the same counter-propagating probe beam measurements were taken as above for the original set-up. The VPL of 245 μs and t_{valve} of $-4000\,\mu$s used in this set-up are different to those used in the set-up without the copper cylinder. The VPL is higher because the helium now has to enter the trap region through a small hole in the copper cylinder (see figure 5.8(a)), resulting in a fraction of the helium pulse not reaching the trap region. Having a lower helium density in the trap region results in a lower Dy signal after ablation. Therefore, in order to have roughly the same amount of helium in the trap centre as before, the VPL was increased until the Dy signal was comparable to that without the copper cylinder. The helium also needs to be given time to collide with the cylinder walls and lose its centre-of-mass velocity before the Dy is ablated into it, hence the earlier t_{valve} of $-4000\,\mu$s rather than the $-600\,\mu$s used for the data in figure 5.7.

Comparing the frequency shift between the spectra from the two probe beams again allowed the velocity of the Dy cloud to be determined as a function of time after ablation, and is shown in figure 5.9. With the copper cylinder in place, the Dy cloud initially has almost no centre-of-mass velocity, however the velocity then increases up to 120 m/s at about 370 μs, after which the signal becomes too low to observe a spectrum and hence determine the velocity. The first thing to note is that the copper cylinder causes the Dy to stay in the trap region for much longer than

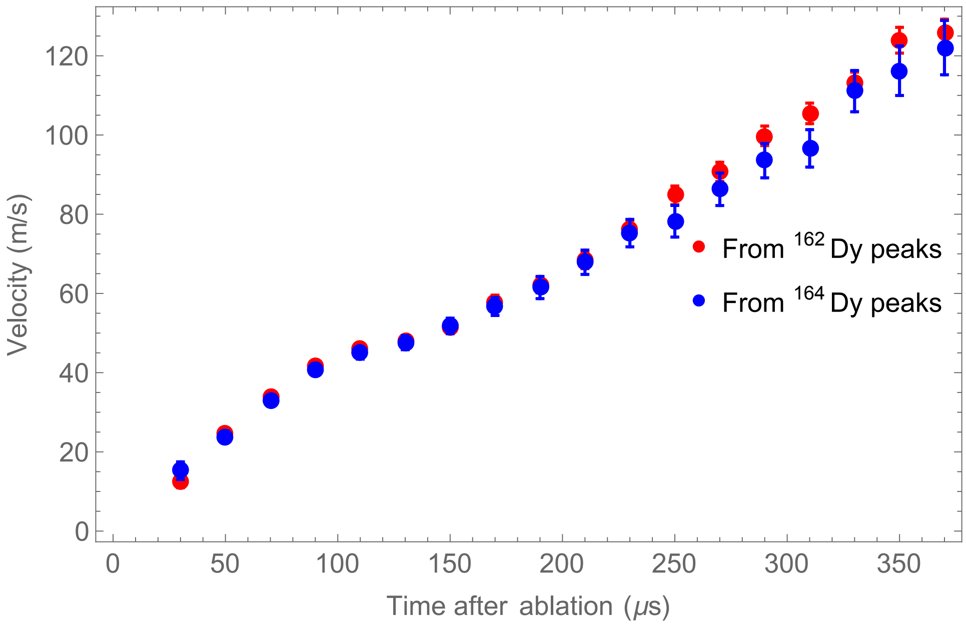

Figure 5.9: Centre-of-mass velocity of the Dy cloud as a function of time after ablation, with the valve tube and copper cylinder. The red points are calculated from the shift between the ^{162}Dy peaks and the blue points from the shift between the ^{164}Dy peaks. Positive velocities are in the direction away from PD 1. For these data, VPL $= 245\,\mu$s and $t_{valve} = -4000\,\mu$s.

without it, given that the signal in figure 5.7(b) had already disappeared after 80 μs, compared to 370 μs here. This is explained by the copper cylinder containing the initial helium pulse and preventing it from acquiring the large centre-of-mass velocities that caused the signal to disappear so quickly in figure 5.7(b). Unfortunately however, even though the Dy cloud initially has a low velocity, it still acquires a large velocity over time. The copper cylinder initially results in a 'cloud' of helium with no centre-of-mass velocity that the Dy is ablated into, but the large tail of the helium pulse coming out of the valve tube (see section 4.4.2) then proceeds to collide with the Dy, slowly giving it a centre-of-mass velocity. This further emphasises the need for a horizontal valve set-up, as in figure 4.16, and confirms that this trapping arrangement will not work with the valve tube in place as the long-lasting helium tail will cause all Dy atoms to be lost from the trap very quickly.

Horizontal valve with copper cylinder

As a result of the above conclusions, the switch was made to the horizontal valve set-up shown in figure 4.16. The copper cylinder was still kept in place to contain the initial helium pulse from the valve. To determine whether this arrangement leads to a stationary cloud of Dy in the trap region, the same counter-propagating probe beam experiments were conducted, with

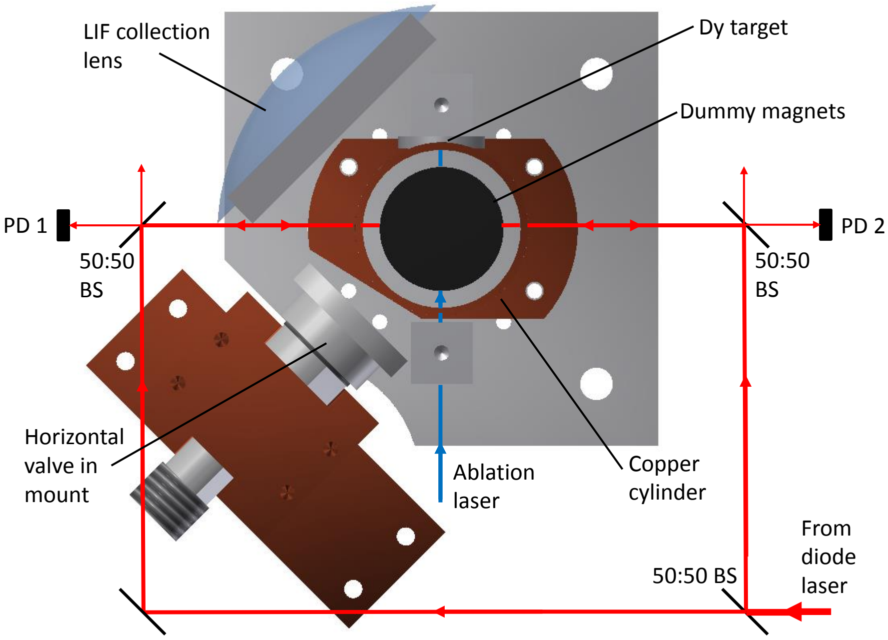

Figure 5.10: Top-down view of the counter-propagating beam arrangement with the horizontal valve and copper cylinder. PDs are photodiodes and BSs beamsplitters.

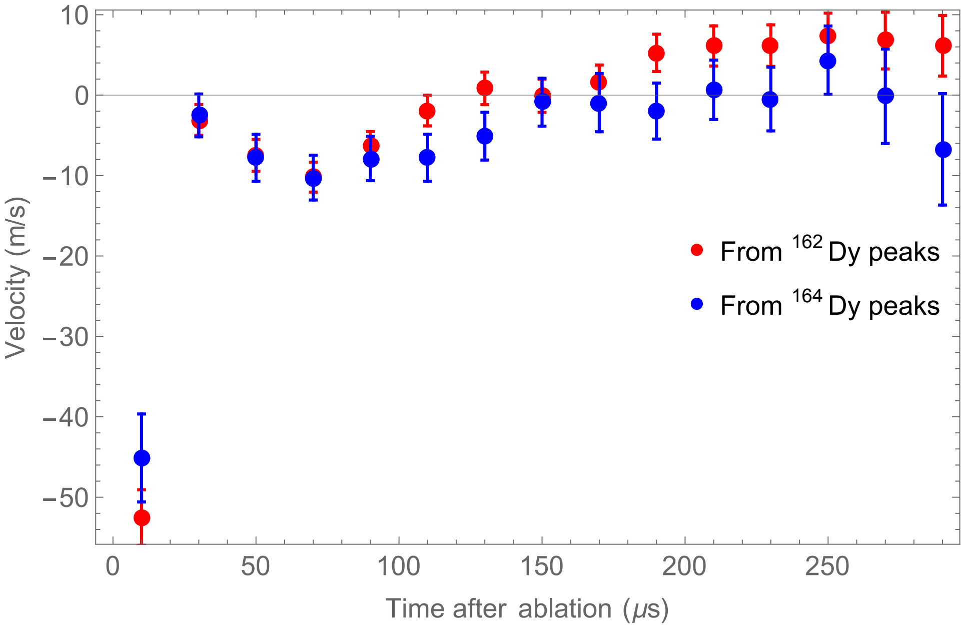

Figure 5.11: Centre-of-mass velocity of the Dy cloud as a function of time after ablation, with the horizontal valve and copper cylinder. The red points are calculated from the shift between the ^{162}Dy peaks and the blue points from the shift between the ^{164}Dy peaks. Positive velocities are in the direction away from PD 1. For these data, VPL $= 200\,\mu$s and $t_{valve} = -2000\,\mu$s.

a diagram of the set-up shown in figure 5.10.

Again, comparing the frequency shift between the spectra from the two probe beams at different times after ablation allows the velocity of the Dy cloud to be determined as a function of time, shown in figure 5.11. The initial, negative centre-of-mass velocity could arise as a result of the Dy coming off the target at a slight angle before the ablation plume is arrested. After 25 µs, despite a few fluctuations in velocity between 10 and -10 m/s, the Dy cloud does not acquire any large centre-of-mass velocity. This shows that switching to the horizontal valve set-up has been successful in eliminating the problems with centre-of-mass velocity that were encountered with the valve tube. The ±10 m/s threshold that the velocity does not seem to surpass is below the capture velocity for this trap arrangement of 14.5 m/s.

5.3 Magnet experiments

The 'dummy magnet' experiments showed that we have managed to create a stationary cloud of Dy atoms in the trap centre that thermalises with the He buffer gas. The next step was to return the magnets to the set-up and attempt the trapping of Dy atoms.

5.3.1 Voltage on magnet

Despite the use of the PMT switching circuit (see section 4.4.1), it was noticed that there was still a significant background fluorescence signal that did not depend on the presence of the probe laser. It can therefore not be due to LIF from Dy in the trap region. This is a problem because it gives a large background that masks the real signal we are searching for. The black line in figure 5.12 shows a typical fluorescence signal from the trap region with magnets in place, which is very comparable to the signal from figure 4.13 where the switching circuit was first used. The black line in figure 5.12 however, is very similar irrespective of whether the probe beam is on or off. This indicates that the slowly-decaying tail is likely due to something other than Dy atoms fluorescing in the trap region for a significant period of time after ablation. Given that laser ablation is a complex process (see section 2.1.2), different ions and clusters are likely to be released from the target during ablation. These could give rise to the observed fluorescence. To remove the ions, a wire was attached to the top magnet in the trap set-up so that a voltage could be applied to the magnet surface. To electrically isolate the magnet from the rest of the set up, the parts that were touching the magnet holder (see figure 4.3) were

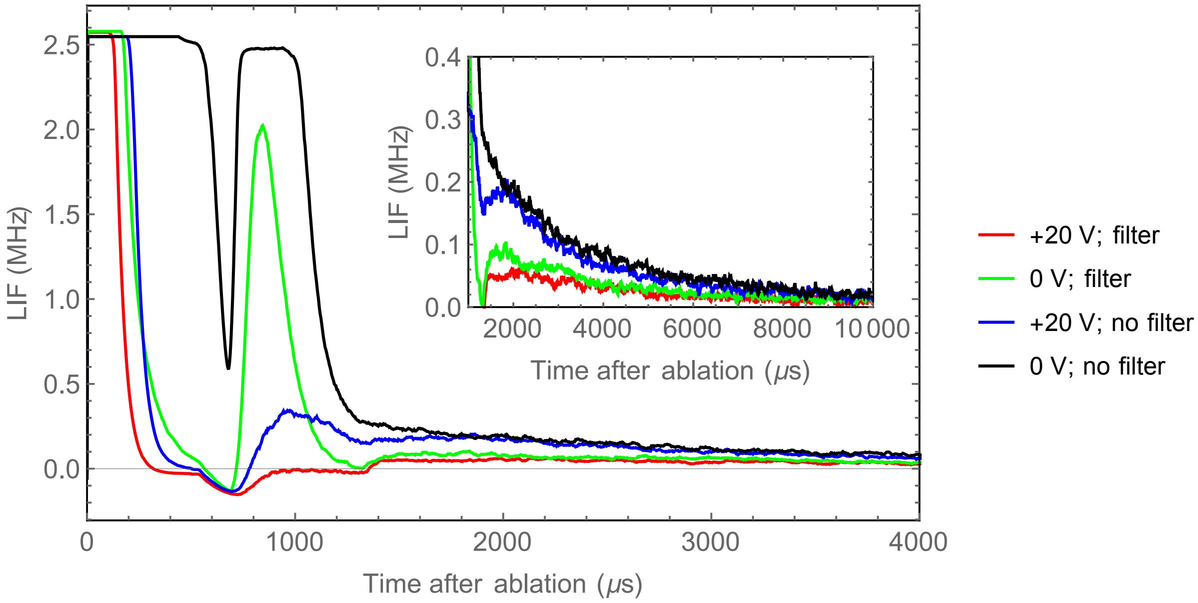

Figure 5.12: Fluorescence signal from the trap region with magnets in place, showing the effect of the magnet voltage and red filter on the signal. The probe laser is off, i.e. only Off-shots are shown. Black line: no voltage, no filter; green line: no voltage, filter; blue line: 20 V on magnet, no filter; red line: 20 V on magnet, filter. The PMT is switched on at 500 μs. Inset shows the same data with the tail signal magnified.

coated in boron nitride. This is a good electrical insulator but a thermal conductor, allowing the magnet to still cool down to 4 K. The bottom magnet remained grounded to the cold plate, creating a voltage difference between the top and bottom magnets. With 20 V applied to the top magnet (blue line in figure 5.12), a significant part of both the first and second peak in the signal that was previously present has disappeared. This part of the signal must have been due to ions in the trap region, and removing the ions in this way ensures that the signal that is left is only due to neutral particles. The negative part of the signal occurs because the PMT takes some time to switch on fully (see section 4.4.1). Removing the ions does not seem to eliminate the slow decaying tail that is present regardless of whether the probe beam is on or off. In order to reduce this tail, a red coloured glass filter (610 nm long pass) was placed in front of the PMT so that only light in the desired frequency range is transmitted. The green line in figure 5.12 shows the signal with no voltage on the magnet, but the filter in place, showing that the tail of the signal is reduced significantly. This shows that a large part of the fluorescence is at wavelengths shorter than 610 nm. We also know from placing a visible light band-pass filter in front of the PMT that the fluorescence is at wavelengths longer than 300 nm, as this filter did not reduce the tail signal. The red line in figure 5.12 shows the signal with both 20 V on the magnet and the red filter in place, clearly showing that both the peaks and the slowly-decaying tail have been reduced significantly. The background fluorescence tail has not been eliminated entirely, but it is now small enough that the rest of it is easily eliminated by subtracting the ‘Off-shot’ from the ‘On-shot’– a procedure that is explained in section 4.2.5. Now when we turn on the probe laser, we only see resonant Dy signal. All subsequent data in the experiment were therefore taken with 20 V on the magnet and the filter in place in front of the PMT.

5.3.2 Spectrum of Dy in the magnetic field

With the background fluorescence signal eliminated as a result of the changes described in section 5.3.1, it was possible to take LIF spectra of Dy in the magnetic field. Figure 5.13 shows a typical spectrum of Dy in the magnetic field, with both the ^{164}Dy and ^{162}Dy lines present. The probe beam is vertically polarised, making it approximately parallel to the magnetic field. This suppresses the broad background signal from the Zeeman broadened $\Delta m = \pm 1$ transitions, making the $\Delta m = 0$ transitions more easily observable. The ^{164}Dy and ^{162}Dy spectral lines are split into two separate peaks that were not present in the field free spectrum from figure 5.4, where the 164 and 162 lines were just single peaks. The splitting of the lines into these two separate peaks is due to the presence of both low-field seeking (LFS) and high-field seeking (HFS) Dy atoms i.e. those with positive and negative values of m_J respectively. In section 5.1.1, it was noted that, for the ground state $g_{J_g} = 1.242$, whereas for the excited state $g_{J_e} = 1.25$. This leads to slightly different Zeeman shifts for the ground and excited states. The $\Delta m = 0$ transitions will therefore be shifted from the field-free resonance by,

$$\Delta E = m_J \, \mu_B \, B(g_{J_e} - g_{J_g}), \tag{5.6}$$

where μ_B is the Bohr Magneton and B is the magnitude of the magnetic field. The relative intensities of the $\Delta m = 0$ transitions are given by the square of the Clebsch-Gordan coefficients, and hence the square of the 3-j symbol,

$$\begin{pmatrix} J & 1 & J \\ m_J & 0 & -m_J \end{pmatrix}, \tag{5.7}$$

where $J = 8$ and m_J ranges from -8 to 8. The highest intensities are for the largest m_J values, i.e. -8 and 8, with the intensity going to zero for $m_J = 0$. It is this property of the relative intensities that results in a splitting of the peaks in the magnetic field, rather than a broadening,

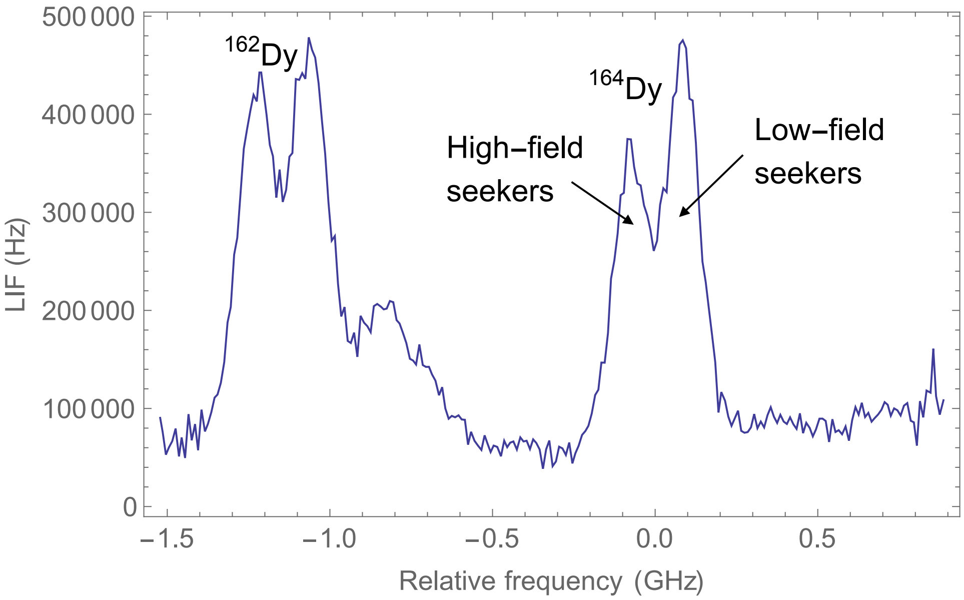

Figure 5.13: LIF spectrum of Dy in the magnetic field, showing the split peak of the ^{164}Dy and ^{162}Dy lines. Frequencies are relative to the field-free ^{164}Dy transition. The probe beam is vertically polarised, i.e. approximately parallel to the magnetic field.

as those highest m_J transitions also have the largest shift from field free resonance (see equation 5.6). By adding together Doppler broadened Gaussian peaks for each of the $\Delta m = 0$ transitions, with the relative frequencies given by equation 5.6 and the relative intensities given by equation 5.7, the split peaks can be modelled. This model can be fit to the data from figure 5.13 to extract a value for the magnetic field of 0.11 T. This magnetic field value is well within the range of fields present in the trap (the trap depth is 0.31 T), and suggests that the probe beam is not quite going through the trap centre, where the magnetic field would be closer to zero. The above model is somewhat simplified as the probe beam has a FWHM diameter of 1.15 mm, and therefore the PMT collects fluorescence from a relatively large area with a spread of magnetic field values. This means the detected atoms do not necessarily reside at a single magnetic field value. However, the calculation does serve to reassure that the splitting of the lines into two peaks is indeed a result of the presence of LFS and HFS states.

5.3.3 Decay of low and high-field seeking peaks

The low-field seeking atoms can be trapped in the magnetic trap, whereas the high-field seeking ones are ejected. This makes it possible to study whether atoms are successfully trapped by measuring the relative heights of the LFS and HFS peaks in the spectrum. The ^{164}Dy peak was chosen for this analysis as it is furthest away from any other transitions (see figure 5.2). A

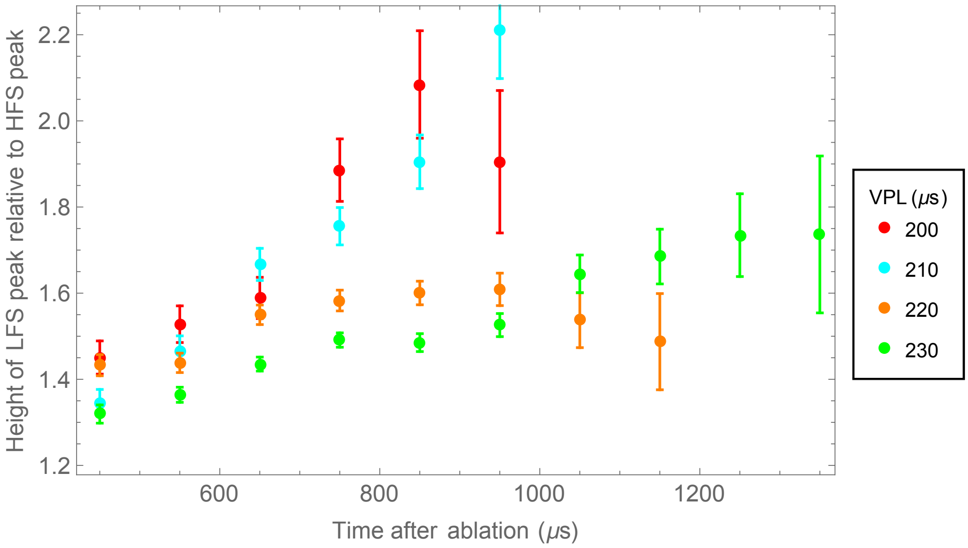

Figure 5.14: Height of the LFS peak relative to the height of the HFS peak as a function of time after ablation. Data points are shown for four different valve pulse lengths (VPLs) from 200 to 230 μs.

double Gaussian was fit to the split 164 peak in the spectrum, and the heights of the two peaks determined from the fit. This was done for several different times after ablation in order to get the heights of the LFS and HFS peaks as a function of time. Given that both peaks will decay over time as the overall signal decays, the height of the LFS peak was calculated relative to the height of the HFS peak to look for any trapping effects. Figure 5.14 shows the relative height of the LFS peak as a function of time for four different valve pulse lengths (VPLs). For each of the VPLs, the relative height of the LFS peak seems to increase with time, which is expected given that the atoms in HFS states are ejected from the trap region much more quickly than those in LFS states. This is clear evidence that the Dy atoms are indeed trapped. We also see that the relative height of the LFS peak increases much more rapidly for the two lowest VPLs of 200 and 210 μs, than for the two higher VPLs of 220 and 230 μs. There are three factors that can influence the decay time and hence the ratio of the peaks: diffusion through the helium gas; magnetic trapping/ejection for the LFS/HFS atoms; and spin-flip collisions and momentum kicks from background helium gas that cause trap losses. For the highest VPLs, where the highest amount of helium gas is pulsed into the trap region, the HFS atoms will take some time to diffuse through the helium and they will therefore be ejected from the trap relatively slowly. This explains why the relative height of the LFS peaks does not increase as rapidly, because the decay of both the LFS and HFS signal is dominated by diffusion and therefore occurs at almost the same rate. For the lower VPLs on the other hand, the amount of helium pulsed into the trap region is lower, which allows the HFS atoms to leave the trap region more rapidly, with the LFS atoms remaining trapped for some time. This results in the sharper increase in relative height of the LFS peak in figure 5.14. Estimates of the helium densities involved in each case are calculated in subsequent paragraphs. The lower VPLs seem to lead to more favourable trapping conditions as the helium density in the trap region is lower and therefore trap losses through spin-flip collisions will be minimised. However, there is a balance to be struck, as reducing the VPL below 200 μs reduces the overall signal so much that it was not possible to observe the spectra. This is because the plume of ablated Dy atoms has to be stopped inside the trap region, requiring a sufficiently high helium density. The difference in amount of helium gas pulsed into the trap region for the different VPLs also accounts for the fact that for the highest VPL of 230 μs, the spectrum can be observed for longer as the higher helium density leads to a longer diffusion time for the Dy atoms (see equation 4.3).

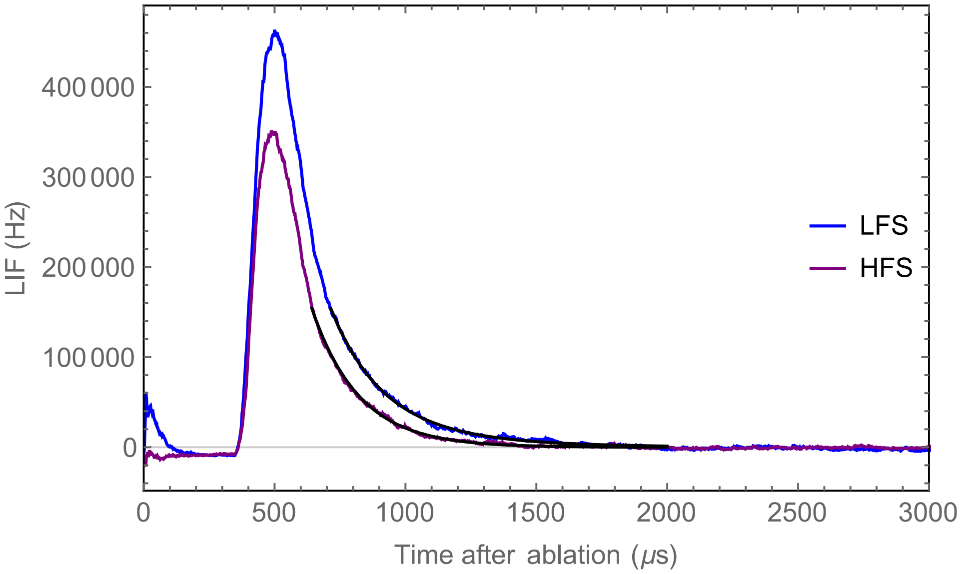

Figure 5.15: Laser-induced fluorescence signal of Dy in the trap for both the LFS and HFS resonance, with a VPL = 210 μs. The black lines are the exponential fits to the decaying part of the curves. The HFS curve is best approximated by a single exponential with a decay time of 180 ± 1 µs, and the LFS curve by a two-exponential decay (equation 5.8) with decay times of 180 ± 9 µs and 360 ± 30 µs.

Figure 5.15 shows time-resolved fluorescence signals from Dy in the trap for the LFS and HFS resonance separately. The settings were the same as those which provided the best ratio of LFS to HFS peaks in figure 5.14. The signal from the LFS atoms is greater than that from the HFS atoms almost from the start. The decaying part of the HFS signal is well approximated by a single exponential decay curve with a decay time of 180 ± 1 µs. The typical timescale for 4 K Dy atoms to leave the probe beam region in free flight is about 80 µs, suggesting that this longer decay time is due to diffusion through the helium gas. The LFS signal on the other hand, is better approximated by a two-exponential decay curve of the form,

$$A\mathrm{e}^{-t/\tau_1} + B\mathrm{e}^{-t/\tau_2}, \tag{5.8}$$

where A and B are constants and τ_1 and τ_2 are the decay times. Fitting this function to the decaying part of the LFS signal, the first decay time is the same as for the HFS signal at $180 \pm 9\,\mu\mathrm{s}$, but the second decay time is significantly longer at $360 \pm 3\,\mu\mathrm{s}$ and is likely to be a result of the trapping field. Using equation 4.5, the oscillation period for Dy in the trap can be up to 4 ms for the widest orbits. This means that the $360\,\mu$s decay time is not likely to correspond to any full oscillations of Dy in the trap. Despite the relatively short trap time-scales involved of only a few $100\,\mu$s, the presence of an additional, longer second decay time for the LFS atoms shows clear evidence that there is trapping of Dy atoms in this set-up.

If we assume, as for the Li experiments, that every collision with a helium atom kicks a Dy atom out of the trap, then using equation 4.6 with a Dy-He collision cross-section of $10^{-18}\,\mathrm{m}^2$ [115], a trap lifetime of $360 \pm 30\,\mu\mathrm{s}$ would correspond to a background helium density of approximately $2 \times 10^{19}\,\mathrm{m}^{-3}$ (pressure of 10^{-5} mbar). In reality, as Dy is much heavier than He, not every collision will cause a Dy atom to be lost from the trap and the actual helium density is therefore likely to be higher than this. As they are not trapped, the decay time of the HFS atoms of $180 \pm 1\,\mu\mathrm{s}$ is likely dominated by diffusion through the helium gas. Using equation 4.3, a diffusion time of $180\,\mu\mathrm{s}$ corresponds to a helium density of approximately $6 \times 10^{22}\,\mathrm{m}^{-3}$. This is probably a more accurate estimate of the He density in the trap region. Past measurements have shown that on the order of 10^{18} atoms are released per valve pulse [116]. This corresponds to a density of about $3 \times 10^{22}\,\mathrm{m}^{-3}$ if these atoms fill the copper cylinder surrounding the trap region. This is fairly consistent with the density of $6 \times 10^{22}\,\mathrm{m}^{-3}$ calculated from the diffusion time of the HFS atoms. Whilst this helium density will decay over time, the fact that it is relatively high means that the $360 \pm 30\,\mu\mathrm{s}$ decay time for the LFS atoms is caused by a combination of both diffusion through the helium and magnetic trapping.

The same time-resolved LIF signals as in figure 5.15 can be plotted for the data from figure 5.14 with the highest VPL of $230\,\mu$s. Both the HFS and LFS signals are well approximated by single exponential decays in this case, with a LFS decay time of $287 \pm 1\,\mu\mathrm{s}$ and a HFS decay time of $281 \pm 2\,\mu\mathrm{s}$. The LFS decay time is slightly longer than the HFS decay time, but the difference is much less than for the VPL of 210 μs, which explains the much slower increase in the relative height of the LFS peak in figure 5.14. Using equation 4.3, the HFS decay time of $281 \pm 2\,\mu$s corresponds to a helium density of approximately $9 \times 10^{22}\,\mathrm{m}^{-3}$, assuming the decay is dominated by diffusion. The decay time of $287 \pm 1\,\mu$s for the LFS atoms is actually less than the $360 \pm 30\,\mu$s decay time with the lower VPL of 210 μs. This suggests that the higher helium density is increasing trap losses through either spin-flip collisions or momentum kicks, and therefore confirms that using the lower VPLs leads to more favourable trapping conditions.

5.3.4 Mapping out the field

From equation 5.6 and the model in section 5.3.2, we saw that the splitting between the double peak of the transition lines in the spectra depends on the magnetic field being experienced by the Dy atoms. Given that the magnetic field, in turn, depends on the position of the Dy atoms in the magnetic trap, moving the probe beam to a different position to detect a different part of the Dy cloud should give rise to a different splitting in the double peak. By taking spectra of the ^{164}Dy line with the probe beam in different positions, the spacing between the double peaks can be mapped out as a function of position. This is shown in the 3-dimensional plot in figure 5.16, where the spacing between the peaks is plotted as a function of position of the probe beam in the vertical and horizontal directions. The probe beam position was measured on a CCD camera placed at the exit of the vacuum chamber. The apertures in the

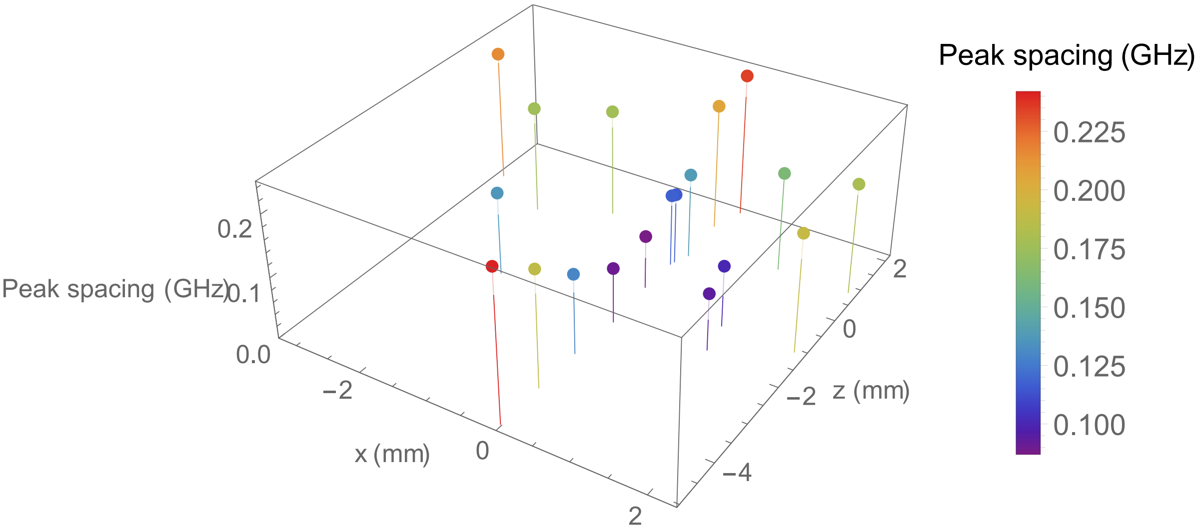

Figure 5.16: Spacing between the double peak of the ^{164}Dy line for different probe beam positions. The z direction is the vertical direction perpendicular to the faces of the magnets, with x being in the horizontal plane.

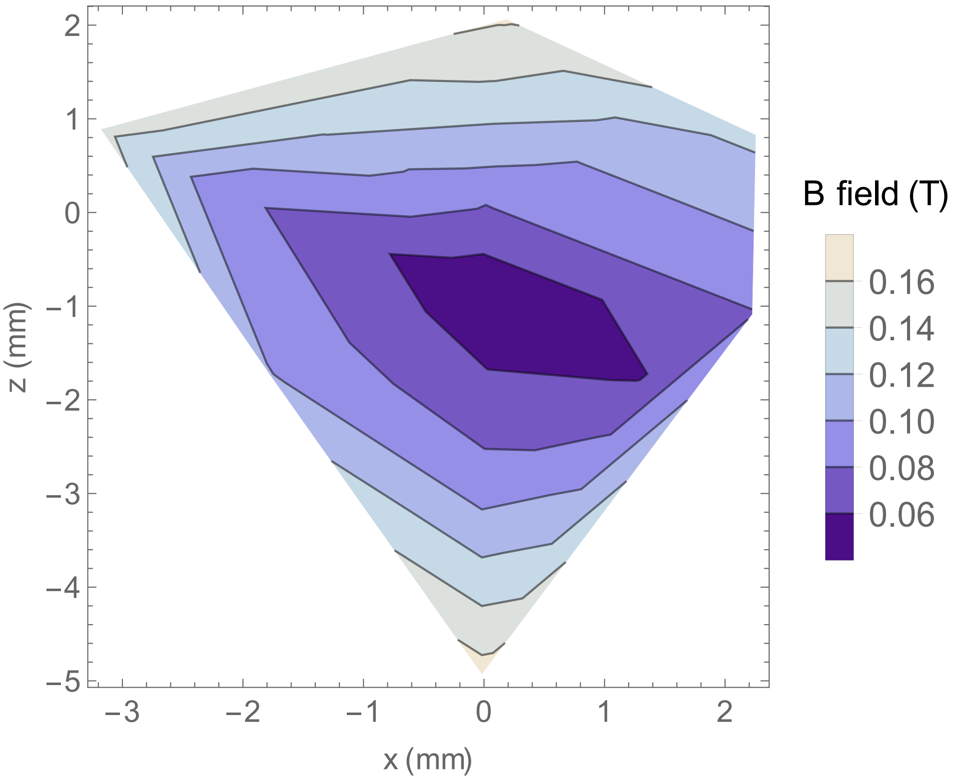

Figure 5.17: Contour plot of the magnetic field corresponding to the spacing between the double peak of the ^{164}Dy line, as a function of probe beam position. This is drawn from the same data as figure 5.16, but with the peak spacing converted to a magnetic field using the model from section 5.3.2.

copper cylinder surrounding the trap region (see figure 5.8(a)) that the probe beam has to pass through constrained the range of positions we were able to probe to the ones shown in figure 5.16. A contour plot derived from the same points is shown in figure 5.17, with the spacing between the double peaks converted to a magnetic field using the model in section 5.3.2. The area where the magnetic field is lowest corresponds to the probe beam going through the trap centre. The magnetic field never goes down to zero because the probe beam has a finite size (FWHM diameter of 1.15 mm), and therefore, even with the probe beam going through the trap centre, regions of non-zero field are still probed. One thing to note from figure 5.17 is that the trap centre appears to be at $z = -1.5$ mm, rather than at $z = 0$ mm. This is caused by a combination of the incorrect determination of the trap centre when positioning the probe beam and the thermal contraction of the set-up when it cools down to 4 K, which lowers the trap centre. The shape of the contour plot resembles that of the magnetic field produced by the magnets (see figure 3.4(b)), which confirms the assumption that the splitting of the lines into two peaks is caused by LFS and HFS states.

5.4 Buffer gas in the trap region

In section 5.3.3, evidence of the trapping of Dy in the magnetic trap was shown. However, the trap lifetimes were never longer than a few 100 μs. The likely cause of this short trap lifetime is the presence of excess helium buffer gas in the trap region causing trap losses. There is no way to measure the helium density in the trap region directly, and estimating it from the amount of helium released in the valve pulse is difficult as the trap still has a relatively open geometry, making it difficult to determine how quickly the helium will leave the trap region. The presence of helium can however be indirectly estimated by looking at the Dy signal. To obtain an indication of the length of time that helium remains in the trap region, the Dy signal was measured for different valve firing times, t_{valve}. Recall that the valve is pulsed first, to inject helium gas, then the ablation laser fires. We also know that the Dy signal measured shortly after ablation depends strongly on the helium density in the trap region, because the helium has to stop the ablation plume. Therefore, by measuring the Dy signal as a function of time delay between the valve and the ablation laser firing, we can determine how the helium density in the trap region changes with time. Figure 5.18 shows such a t_{valve} scan for the set-up with magnets and copper cylinder in place. The Dy signal with the probe beam on the ^{164}Dy resonance is measured as a function of t_{valve}, with t_{valve} changed between each shot (see section 4.2.5). From the plot it is clear that the largest signal appears at around $t_{valve} = -500\,\mu$s, which is similar

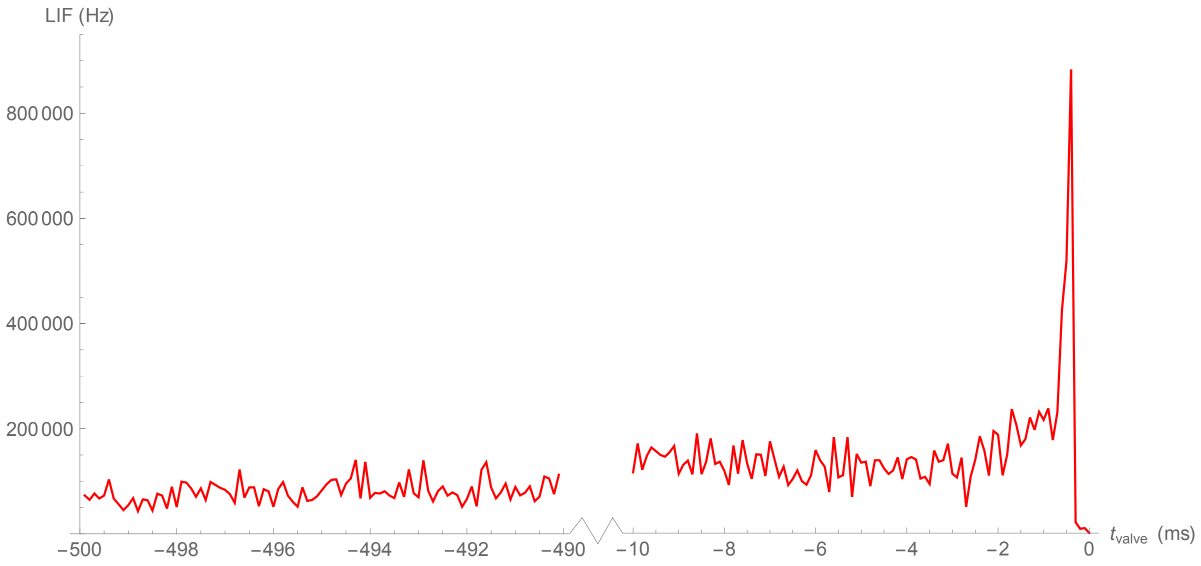

Figure 5.18: Dy LIF signal as a function of the valve firing time, t_{valve}, with copper cylinder. The ablation laser is fired at $t = 0$, with the LIF signal recorded 450 μs after ablation.

to the t_{valve} used to take the data in section 5.3. This is not surprising since that t_{valve} was chosen as it provided the highest Dy signal. However, figure 5.18 also shows that the signal does not go down to zero even if the valve is fired at $t_{valve} = -500\,\mathrm{ms}$. This means that, even if the Dy target is ablated 500 ms after the valve has fired, we still see a significant amount of Dy signal. If the valve is not fired at all, the signal is truly zero, so figure 5.18 suggests that there must still be a significant amount of helium in the trap region even 500 ms after the valve has fired. This is detrimental to trap lifetimes, and suggests that the geometry of the trap is not open enough and does not allow the helium to dissipate quickly enough. In section 5.4.2 a calibration procedure is explained which allows figure 5.18 to be converted into a plot of the helium density as a function of time. First though, we introduce a strategy for allowing the helium to leave the trap region more rapidly.

5.4.1 Mesh cylinder

To allow the helium buffer gas to leave the trap region more quickly, the copper cylinder surrounding the trap region was replaced with a more open cylinder, shown in figure 5.19. A 70% reflective, 30% transparent copper mesh was placed around the rectangular holes so that the

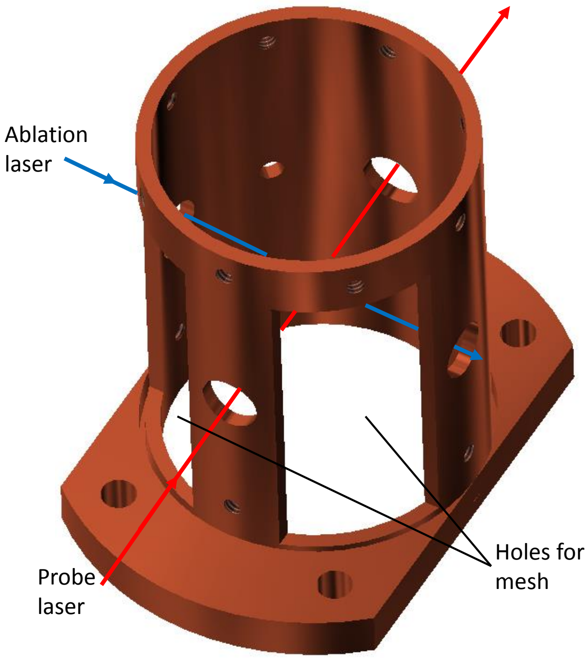

Figure 5.19: The new copper cylinder placed around the trap region, with a transparent copper mesh used to cover up the rectangular holes. This should allow the helium to leave the trap region more quickly.

trap region was still partially enclosed, otherwise the same problem with centre-of-mass velocity of the Dy cloud would have been encountered as in section 5.2.2. Counter-propagating probe beam measurements with the mesh cylinder confirmed that the Dy cloud is indeed stationary in this case. The partially transparent mesh should however allow the helium to escape the trap region more quickly.

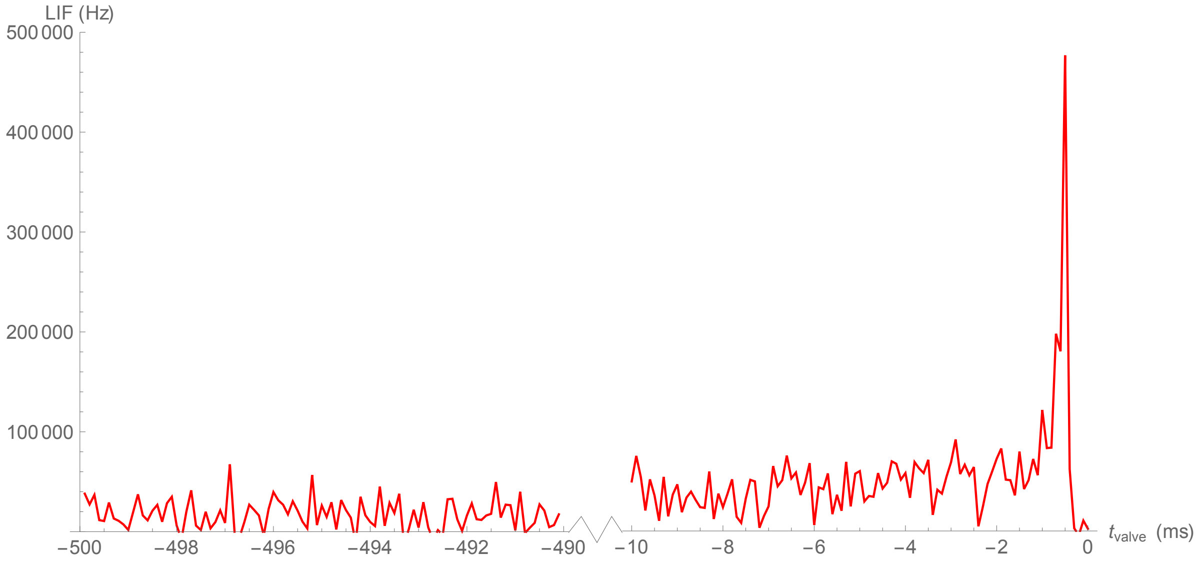

Figure 5.20: Dy LIF signal as a function of the valve firing time, t_{valve}, with mesh cylinder. The ablation laser is fired at $t = 0$, with the LIF signal recorded 450 μs after ablation.

Figure 5.20 shows a scan of the LIF signal versus valve firing time, t_{valve}, with the mesh cylinder in place. As with the full copper cylinder, the peak in the plot is at around $t_{valve} = -500\,\mu$s. However, for earlier t_{valve} times the signal is much lower than it was in figure 5.20, implying that there is less helium gas remaining in the trap region than before. Especially approaching a $t_{valve} = -500$ ms the signal is almost at zero. The signal is not completely zero however, which suggests that even though the amount of helium remaining in the trap region for long periods of time is lower, the mesh cylinder has not allowed it to completely disappear within the time scales measured. The reason for this is discussed in section 5.4.2.

Figure 5.21 shows the LFS and HFS signal from Dy in the trap with the mesh cylinder instead of the copper cylinder. The plot shown is for the settings that gave the longest decay time of the LFS atoms. Fitting a single exponential to the HFS signal gives a decay time of $118 \pm 1\,\mu$s, which is less than the $180 \pm 1\,\mu$s decay time for the HFS atoms with copper cylinder from figure 5.15. As the helium density is expected to be lower for the mesh cylinder, it follows that the decay time for the HFS atoms is also lower as it is expected to be dominated by

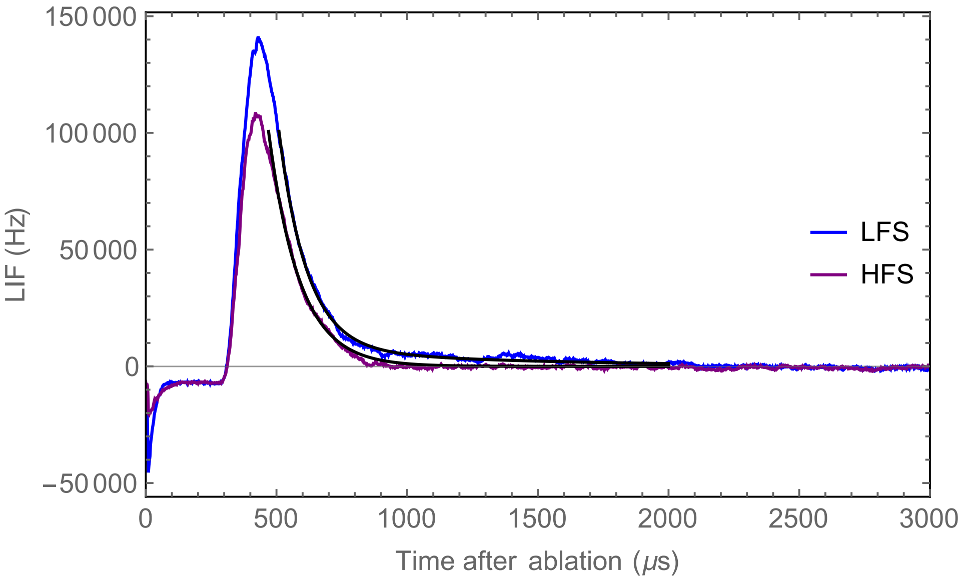

Figure 5.21: Laser-induced fluorescence signal of Dy in the trap for both the LFS and HFS resonance, with a VPL = 210 μs. The copper cylinder has been replaced by the mesh cylinder. The black lines are the exponential fits to the decaying part of the curves. The HFS curve is best approximated by a single exponential with a decay time of 118 ± 1 µs, and the LFS curve by a two-exponential decay (equation 5.8) with decay times of 108 ± 1 µs and 800 ± 30 µs.

diffusion through the helium. The LFS signal in figure 5.21 clearly has a second, longer decay tail that the HFS signal does not have. Therefore, as with the copper cylinder data, the LFS signal is better modelled by the two-exponential decay from equation 5.8. This gives a second decay time of 800 ± 30 µs, compared to 360 ± 30 µs for the copper cylinder data. Introducing the mesh cylinder has therefore improved the trap lifetime of the LFS atoms compared to that with the copper cylinder. This is likely to be a result of a lower buffer gas density in the trap region after the Dy has thermalised.

5.4.2 Calibrating LIF signal with helium pressure

The plots in figures 5.18 and 5.20 show that helium is present in the trap region long after the valve has fired. As the plots only look at the Dy signal, there is no measure of what helium pressure this corresponds to. To convert the Dy signal to a helium pressure, it is necessary to calibrate the Dy signal against helium pressure. This was achieved by closing off the turbo pump attached to the vacuum chamber, and filling the chamber with a constant background pressure of helium gas. To ensure that the background pressure remained constant throughout the chamber, the charcoal sorption pumps that normally sit around the trap region were removed as they would otherwise pump away helium near the trap region. The pressure was measured on the Penning gauge attached to the outside of the chamber. The Dy target was then ablated without firing the valve, and the LIF signal on resonance measured for different background pressures of He gas.

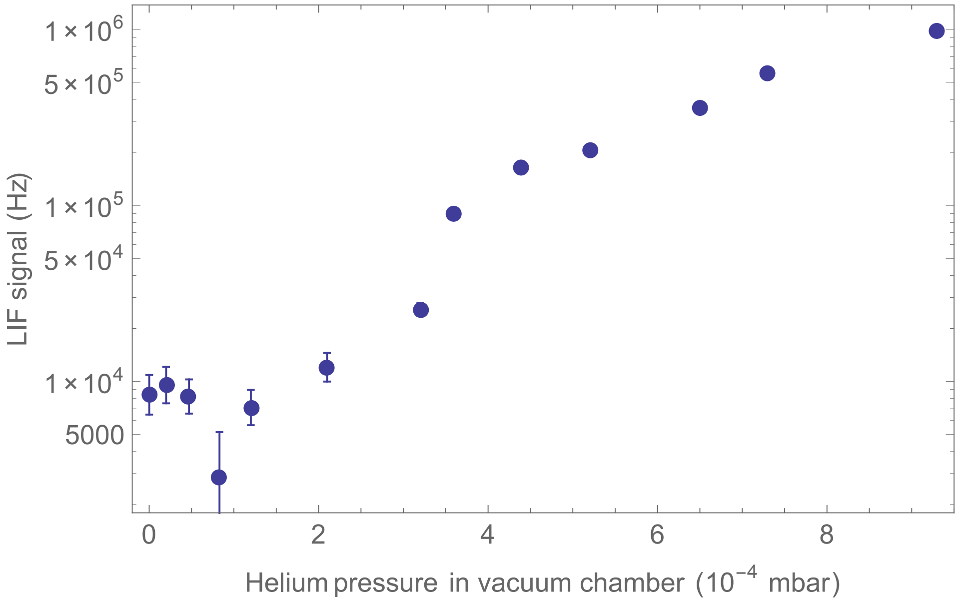

Figure 5.22: Dy LIF signal on resonance for different constant background pressures in the chamber. The point that appears at a pressure of 0 mbar is taken at 10^{-7} mbar, the base pressure in the chamber.

Figure 5.22 shows the LIF signal on resonance against helium pressure in the chamber. For pressures from 10^{-7} mbar (the base pressure in the chamber) up to 10^{-4} mbar, the Dy signal hovers between a few thousand up to 10^4 Hz, but there is no obvious change in signal. This means that for pressures below 10^{-4} mbar, there is not enough helium present to affect the Dy signal. When the background pressure is increased above 10^{-4} mbar the Dy signal starts to increase exponentially with it (note the plot is logarithmic), with the Dy signal reaching 10^6 Hz once the background pressure approaches 10^{-3} mbar. The trend from the plot allows for an estimation of the helium pressures in the trap region.

Returning to the signal against t_{valve} plot in figure 5.18, where the full copper cylinder is still in place, the average signal for the points from $t_{valve} = -500$ to -490 ms is 8×10^4 Hz. Comparing this to the calibration plot in figure 5.22, it suggests that there is a helium pressure of at least 3×10^{-4} mbar (density of $5 \times 10^{20}\,\mathrm{m}^{-3}$) in the trap region 500 ms after the valve has fired. Given that the base pressure in the chamber when the valve does not fire is 10^{-7} mbar, there must be residual helium from the valve pulse in the trap region for this length of time. Even when the copper cylinder was replaced by the mesh cylinder (see figure 5.20), the average signal for a t_{valve} from -500 to $-490\,$ms is $1.7 \times 10^4\,$Hz, which from the calibration plot corresponds to a pressure in the trap region of at least $10^{-4}\,$mbar (density of $2 \times 10^{20}\,\mathrm{m}^{-3}$). These are unlikely to be accurate measures of the helium pressure in the trap region, but even as an order of magnitude estimate, it is clear that the pressure is high even 500 ms after the valve has fired. Using equation 4.6 to estimate the time between collisions at a pressure of $10^{-4}\,$mbar gives a value of $40\,\mu$s. If we pessimistically assume that a single collision is enough to remove a Dy atom from the trap, the expected trap lifetime is also $40\,\mu$s. This confirms that the trap lifetime in figure 5.21 was likely limited by collisions with helium in the trap region.

The time it should take for the helium from the valve pulse to leave the trap region is much shorter than the 500 ms time-scales on which we still see a significant fraction of helium present. The theoretical atom flux from a supersonic valve is given by [114],

$$\phi = A\sqrt{\frac{2k_BT}{m}}n_0\left(\left(\frac{\gamma}{\gamma+1}\right)^{1/2}\left(\frac{2}{\gamma+1}\right)^{1/(\gamma-1)}\right), \tag{5.9}$$

where A is the area of the valve nozzle, T is the temperature of the atoms, m is the mass of the atoms, n_0 is the atom density behind the valve and $\gamma = 5/3$ for a monatomic gas. The valve exit aperture has a diameter of 1 mm and the helium pressure behind the valve is 2 bar, which corresponds to a helium density of $10^{27}\,\mathrm{m}^{-3}$ at 4 K. Inserting those values into equation 5.9 and multiplying by a typical valve pulse length (VPL) of $200\,\mu$s gives 10^{19} atoms expelled from the valve in a single pulse. Assuming all of the helium from the valve enters the copper cylinder surrounding the trap region, this gives an initial helium density in the trap region of $6 \times 10^{23}\,\mathrm{m}^{-3}$. In fact, it is well established that the solenoid valve does not fully open on these short timescales and the actual quantity of helium delivered is smaller. This is confirmed by experimental observations of a similar solenoid valve, which suggest that on the order of 10^{18} atoms are delivered per valve pulse [116]. Numerical simulations of helium in our trap region estimate that, even with the full copper cylinder in place, the helium should dissipate from the trap region on a timescale of at most 10 ms, and therefore the helium density should be below $10^{20}\,\mathrm{m}^{-3}$ after 90 ms and several orders of magnitude lower still after 500 ms. The fact that the helium density is still $5 \times 10^{20}\,\mathrm{m}^{-3}$ 500 ms after the valve has fired suggests that there is another 'source' of helium keeping the helium density in the trap region high. The obvious candidate for this would be a leak from the valve, even when it is closed. However, this would mean that there should be Dy signal when there is helium pressure behind the valve but the

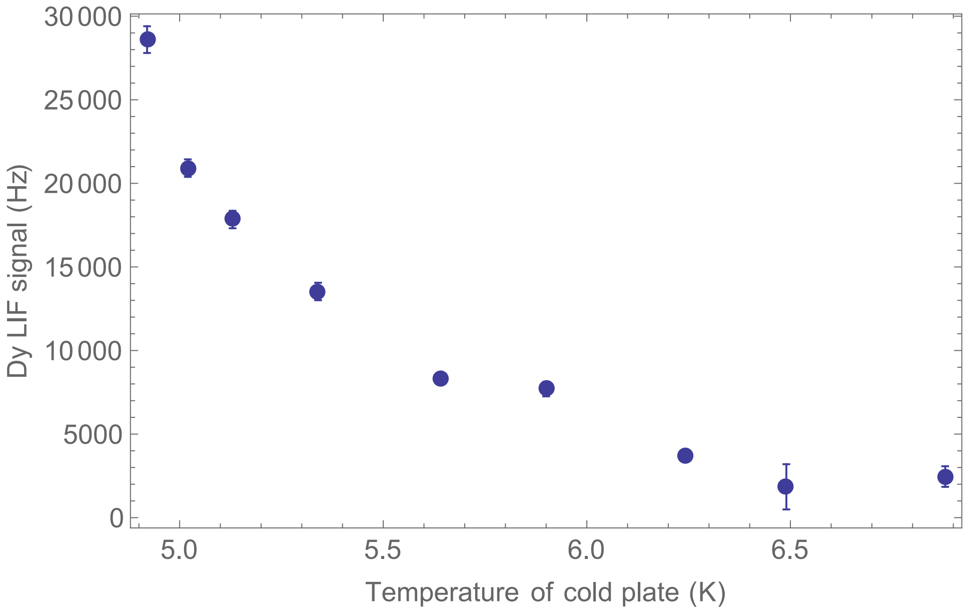

Figure 5.23: LIF signal of Dy on resonance for different temperatures of the cold plate, with the valve firing 50 ms before the ablation laser.

valve is not fired. This is not the case and therefore rules out a leaking valve. The other option that remains is that some of the helium is adsorbed onto a cold surface near the trap region, and slowly released over time. This would be alleviated by raising the temperature of the cold plate to above 4 K as it is unlikely that this adsorption would take place at temperatures significantly above the condensation temperature of helium. This is explored further in the next section.

5.4.3 Data at different temperatures

The temperature of the trap set-up was adjusted by placing a small heater on the cold plate. The delay between firing the valve and firing the ablation laser was set to 50 ms (i.e. $t_{valve} = -50\,\text{ms}$), and the corresponding Dy signal measured for different temperatures of the cold plate. This is therefore a measure of how much helium is still in the trap region 50 ms after the valve has fired. Figure 5.23 shows the Dy signal as a function of the cold plate temperature. It is immediately clear that the amount of signal reduces with increasing temperature, with the signal appearing to level off above 6.5 K. The value to which it levels off is consistent with the region in figure 5.22 that corresponds to helium pressures between 10^{-7} and 10^{-4} mbar. Unfortunately the calibration does not allow for a more accurate determination of the helium pressure, but it is clear from figure 5.23 that the higher temperatures result in the helium leaving the trap region much more quickly. This is consistent with the assertion that the helium is somehow being adsorbed by a cold surface and slowly released over time, and that this does not now occur at the slightly higher cold plate temperature. We already know from section 4.2.2, as well as [105] and [117], that carbon materials, in particular coconut charcoal, are good adsorbers of helium at 4 K. Perhaps this means that the soot used to coat the magnet surfaces and inside of the copper cylinder to minimise laser scatter (see section 4.2.3) adsorbs some of the helium at 4 K and slowly releases it over time, keeping the helium pressure in the trap region high.

The high helium pressure in the trap region when at a temperature of 4 K, even hundreds of ms after the valve has fired, is clearly going to severely limit trap lifetimes. Running the experiment at 6.5 K to limit the amount of helium, and comparing the lifetimes of the LFS and HFS atoms, did not improve upon the trap lifetime from figure 5.21. This leads to the wider conclusion that it is extremely difficult to get long trap lifetimes in this set-up because of the close proximity of the ablation to the trap region. The unpredictable and explosive nature of the ablation process means that a high helium density is required initially in order to contain the ablation plume. This already results in the presence of significant amounts of helium in the trap region, which then has to be removed. With the ablation happening so close to the trap region, the trap lifetimes will also be extremely sensitive to exactly how the ablation plume comes off the target, and any slight changes will probably result in no trapping at all. With so many effects having to come together successfully in a relatively small trap region, it makes it very difficult to achieve longer trap lifetimes in this trap.

5.5 Conclusion

Spectra of Dy taken with the ‘dummy magnets’ in place showed all of the transitions expected from the theoretical calculations. Fitting Gaussians to the spectrum at different times after ablation allowed the widths of the lines to be extracted, and hence the temperature of Dy to be measured as a function of time. This showed that the Dy thermalised with the helium buffer gas within 200 μs.

By introducing a counter-propagating probe beam to the ‘dummy magnet’ set-up, the centre-of-mass velocity of the Dy cloud was measured. This showed that for the original set-up, the Dy cloud had a centre-of-mass velocity of at least 180 m/s right from the start. Introducing a copper cylinder to contain the initial helium pulse limited the centre-of-mass velocity immediately after ablation, however the long tail of the helium pulse emanating from the valve tube resulted in the Dy cloud still acquiring a velocity of over 100 m/s at later times. This problem was finally solved by switching to the horizontal valve set-up with copper cylinder, which showed that the Dy cloud acquired very little, or no centre-of-mass velocity. This was the arrangement that was then used for subsequent experiments with magnets.

When switching to the real magnets, it was noticed that there was still a significant background fluorescence signal that did not depend on the probe laser. This background was not present with the dummy magnets and so is related to the magnetic field. By applying a voltage to the top magnet to remove ions created after ablation, and by placing a red filter in front of the PMT, this background fluorescence was removed, leaving only resonant Dy signal. This allowed the spectrum of Dy to be taken in the magnetic field, which showed the splitting of the resonance lines into high and low-field seeking components. By measuring the heights of the HFS and LFS peaks, evidence of the trapping of Dy was seen in the magnetic trap when the valve pulse length was set to an optimum value. The trap lifetime of $360 \pm 30\,\mu$s is suspected to be limited by excess helium in the trap region.

This was confirmed in subsequent experiments, where the helium density in the trap region was measured as a function of time. There was still a helium density of about $2 \times 10^{20}\,\mathrm{m}^{-3}$ in the trap region 500 ms after the valve was fired. Replacing the copper cylinder with a mesh cylinder so that the helium can leave the trap region more quickly improved the trap lifetime to $800 \pm 30\,\mu$s but did not solve the problem entirely. Finally, it was discovered that this large background helium density could be greatly reduced by raising the temperature of the cold plate to 6.5 K. This suggests that the helium is being adsorbed onto a cold surface near the trap region and slowly released over time, keeping the helium pressure in the trap region very high. Unfortunately, running the experiment at 6.5 K did not improve the trap lifetimes of Dy in the trap.

Whilst there has been clear evidence of trapping in this arrangement, it is also clear from the data in this chapter that it will be difficult to improve upon the trap lifetime of $800 \pm 30\,\mu$s. There are many different factors that need to come together in order for this trap to be more successful. Given the explosive and unpredictable nature of laser ablation, the close proximity of the ablation target to the trap region makes this very difficult as it would require a very precise set of circumstances to stop and cool the ablation plume within the trap region, but with the helium then dissipating quickly enough to leave the Dy trapped. The solution to this problem would seem to be to separate the ablation process from the trapping process, and this is explored in more detail in the next chapter.

CHAPTER 6

TRAPPING FROM A DYSPROSIUM BEAM

Motivated by the conclusions drawn from the Dy experiments in the previous chapter, we decided to try trapping atoms from a beam rather than directly from the ablation source. The close proximity of the ablation source to the trap region made it difficult to achieve long trap lifetimes due to the high densities of helium required in the trap region to stop and cool the ablation plume. By separating the ablation and cooling process from the trapping process, this problem is avoided. A high helium density is used inside a buffer gas cell to create a beam of Dy atoms. This beam then propagates to the trap region where it is stopped by a second helium pulse with a much lower density, leaving the Dy atoms trapped whilst the helium atoms dissipate.

6.1 Buffer gas beams

Buffer gas beams are typically extracted through an exit aperture in one side of a buffer gas cell, as explained in section 2.1.3. This allows a beam of the atoms/molecules of interest to escape from the cell entrained in the buffer gas flow. The properties of the beam are usually dominated by the properties of the buffer gas, as the buffer gas density is typically several orders of magnitude higher than that of the atoms/molecules of interest. A summary of the flow properties and types of buffer gas beams is presented in [37]. There are three different types of flow regime for a buffer gas beam: effusive, intermediate/partially hydrodynamic and supersonic (fully hydrodynamic). In effusive sources, there are very few or no collisions near the exit aperture of the cell and therefore the velocity distribution of the beam that emerges from the cell is simply a sampling of the thermal distribution inside the cell, weighted by the speed since the probability of finding the aperture is proportional to the speed. Such sources are typically made by ensuring that the aperture thickness and diameter are smaller than the mean free path of the gas, ensuring that no collisions occur near the aperture. In the supersonic regime there are many collisions near the aperture and the buffer gas behaves more like a fluid, with the beam properties similar to those of a supersonically cooled beam. In the intermediate regime there are enough collisions near the aperture to alter the beam properties from those of the thermal distribution inside the cell, but not enough for the beam to be fully supersonic.

Buffer gas beams of molecules are now created on a regular basis and are relatively easy to produce, such as for ThO [118] and YbF [119]. Most of these have velocities between 100 and 200 m/s. Creating a purely effusive source with a high flux at cryogenic temperatures is difficult, and in reality most buffer gas beams are in the partially hydrodynamic regime [37]. However, a new two-stage cell technique [62] has been developed to produce beams with near effusive velocity distributions. Such a two-stage cell has been used to directly load CaF molecules into a magnetic trap without any further laser cooling of the beam [61]. These two-stage buffer gas cells may become the norm to produce slow, intense sources of cold molecules, but in the experiments in this chapter a single stage buffer gas cell was used.

6.2 The experiment

The design for our set-up is shown in figures 6.1 and 6.2. This new set-up is placed on a cold plate into the same experimental apparatus described in section 4.2.2. The cell consists of a vertical cylindrical column with 2 mm access holes for the helium pulse and ablation laser; a simplified version of the final design in chapter 6 of [120]. The Dy target is clamped to the outside of one of the cell walls, with the face of the target exposed to the central column, and hence the ablation laser (see figure 6.3(a)). The helium is then pulsed into the cell through the valve, and propagates through the access hole into the central column. At this point the Dy target is ablated and the plume released from the target thermalises with the helium buffer gas inside the central column. Two cross sections through the cell are depicted in figure 6.3, showing how the helium pulse propagates into the central column and out of the cell. A pulsed beam of cold Dy and helium atoms then emerges from the top of the cell and propagates in the direction of the magnets.

This beam travels through the 8 mm hole in the centre of the bottom magnet and into the trap region. The top valve is pulsed so that the helium pulse reaches the trap region at the

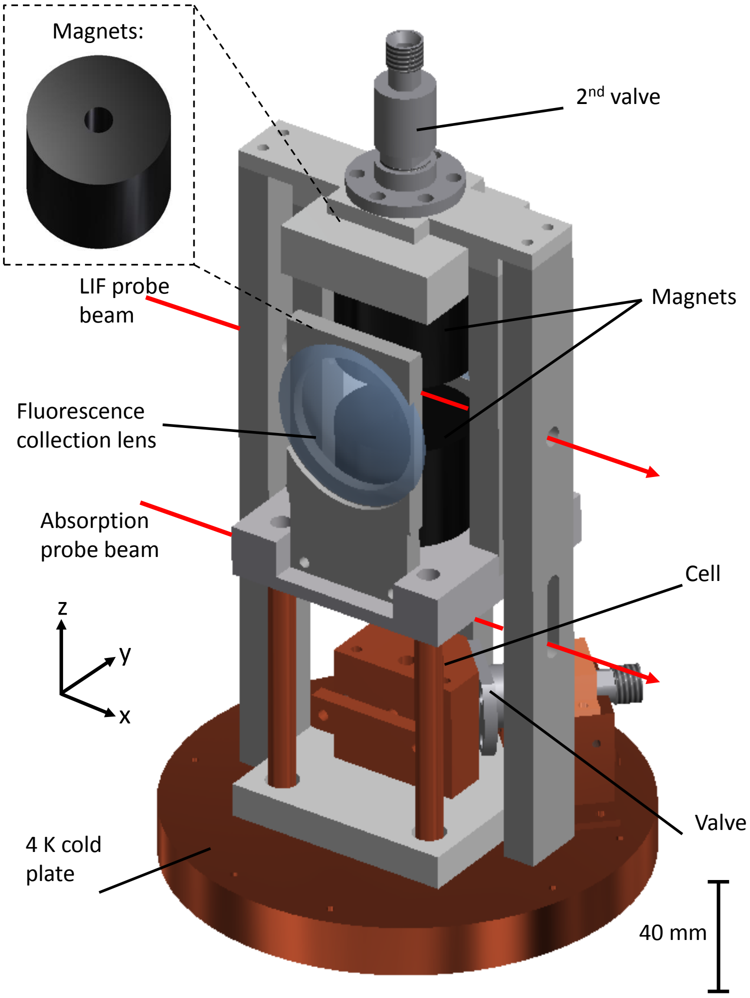

Figure 6.1: Experimental set-up of the two-valve configuration, showing the positions of the cell and second valve relative to the trap region. One of the NdFeB magnets is depicted, showing the 8 mm hole through the centre that allows the Dy beam and second helium pulse to reach the trap region. The positions of the two probe beams are also depicted.

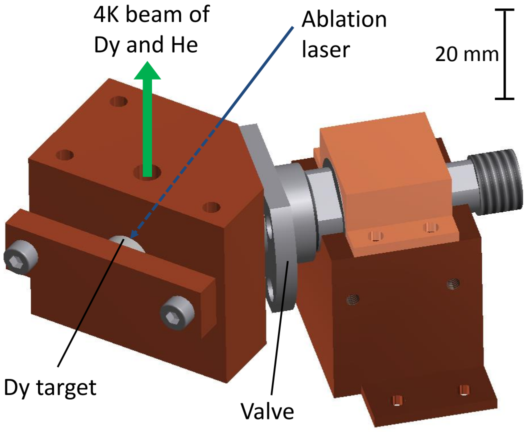

Figure 6.2: Close-up depiction of the cell, showing how the valve pulses the helium into the main central column of the cell through a 2 mm hole. The target is clamped to one side of the cell, with the ablation laser passing through a 2 mm hole from the opposite side to ablate the Dy into the helium pulse. The beam then emerges from the hole in the top of the cell.

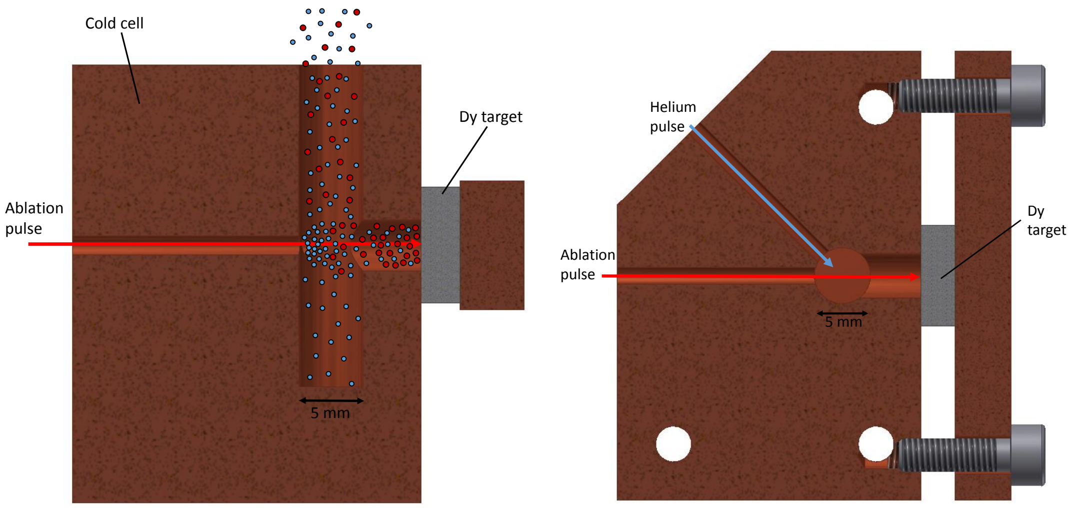

(a) Side view cross section (b) Top view cross section

Figure 6.3: Cross sections of the cell from figure 6.2, showing the channels for the helium pulse and ablation laser, and the central column from which the beam emerges.

same time as the Dy beam. Collisions between the beam and the helium pulse then result in a fraction of the beam remaining stationary in the trap centre. As the trap has a very open geometry, the helium will then dissipate quickly, leaving the Dy trapped. Since there only need to be enough collisions in the trap region to stop a Dy beam that is already cold, rather than having to stop and cool a 10 000 K ablation plume as in the original set-up, the density of helium required in the trap region is much lower. With the second valve positioned so close to the top magnet, the magnetic field from the magnet actually forces the valve to be permanently open. To avoid this, a steel plate is placed between the top magnet and the second valve to shield the second valve from the magnetic field.

The atoms in the trap are detected by collecting laser-induced fluorescence (LIF) on a photomultiplier tube (PMT). The probe laser beam passes through the trap centre, and fluorescence is collected by a large lens placed perpendicular to the probe beam, as shown in figure 6.1. The arrangement of the lenses, iris and PMT is the same as in figure 4.5. In addition, in order to maximise the collection efficiency, a curved mirror is placed directly opposite the collection lens on the other side of the trap region to reflect and focus fluorescence back towards the trap centre, so that it can also be collected by the lens and detected. This gives a collection efficiency of 10%. The quantum efficiency of the PMT at the relevant wavelength is 7%, and so the total detection efficiency of this LIF set-up is 0.7%.

The two NdFeB magnets forming the trap region are cylindrical, with their north poles facing each other to form a quadrupole trap with a magnetic field zero at the centre. This is a

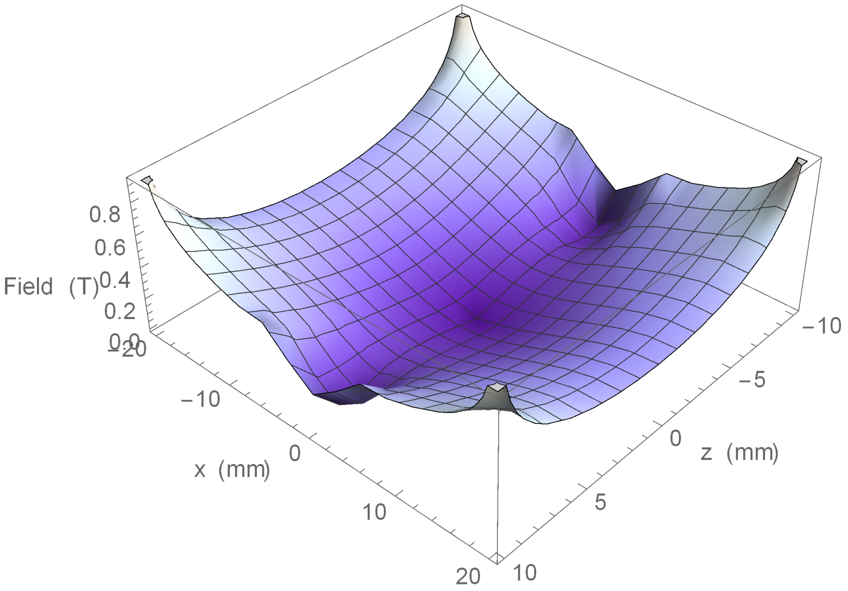

Figure 6.4: Magnitude of the magnetic field formed by the two magnets in the set-up. The x direction is parallel to the magnet faces, with the z direction perpendicular to the magnet faces. The trap depth is 0.23 T.

similar arrangement to the previous trapping experiments for Dy and Li, shown in figure 3.4(a), with the key difference being an 8 mm cylindrical hole through the middle of the magnets (see figure 6.1) to allow the beam of Dy atoms to propagate into the trap region. This magnet geometry was required to ensure access for the probe beam and fluorescence collection optics. The hole in the top magnet is necessary not just to make the trap symmetrical, but also to allow the helium pulse from the top valve into the trap region. The magnets have been scaled up compared to the magnets used in the previous experiments, with a diameter and height of 40 mm, spaced 20 mm apart. This leads to a trap shape similar to the previous set-up (see figure 3.4(b)), where the magnets had a 20 mm diameter and height, and were placed 10 mm apart, but gives a much larger trapping area, making it easier to capture the Dy atoms in the trap. Figure 6.4 shows the magnitude of the magnetic field in this magnet arrangement, calculated using the Radia plug-in [101] for Mathematica. In the plane parallel to the faces of the magnets the field essentially increases linearly away from the centre, similar to the field from the full cylindrical magnets shown in figure 3.4(b). It is only in the direction perpendicular to the magnet faces that the effect of the hole becomes apparent in a reduction in the trap depth. The trap depth in this arrangement is 0.23 T, compared to the 0.31 T from the previous set-up. Whilst this is a significant reduction, the trap depth of 0.23 T, corresponding to 1 K for Dy in the most favourable state, is still high enough to trap 14% of 4 K Dy atoms (see equation 3.1).

The capture velocity of Dy in this trap is calculated using equation 3.2 to be 12.5 m/s.

6.3 Calculations

Calculations were carried out to determine the feasibility of this new two-valve trap arrangement. We would like to know the helium density needed to stop Dy atoms from the beam, and how this compares to the density needed to stop and cool the ablation plume in the previous set-up.

From the hard sphere model in equation 2.6, it was calculated that roughly 190 collisions are required to cool the Dy atoms from a temperature of 10 000 K immediately after ablation to within 30% of the buffer gas temperature of 4 K. If this is to occur inside the cell from figure 6.2, these 190 collisions have to take place before the Dy atoms have a chance to diffuse to the cell walls, where they will be lost. The mean free path of the atoms is given by,

$$\lambda = \frac{1}{\sqrt{2}n\sigma}, \tag{6.1}$$

where n is the helium density and σ is the elastic collision cross section between Dy and helium and is well estimated to be $10^{-18}\,\mathrm{m}^2$ [115]. For the required 190 collisions to occur before the Dy atoms diffuse to the cell walls, the mean free path can be no bigger than the diameter of the cell divided by the number of collisions. This means that, with a cell diameter of 5 mm (see figure 6.3(a)), a helium number density of at least $10^{22}\,\mathrm{m}^{-3}$ is needed inside the cell. A bigger cell diameter would not require as high a helium density to ensure the Dy atoms thermalised before hitting the cell walls. However, with such a wide cell diameter, an aperture would have to be placed at the exit of the cell to ensure a collimated beam. Having an aperture that is smaller than the cell diameter leads to recirculation regions and hence a loss of flux [119]. To avoid this, the cell diameter was kept at 5 mm.

It is important that not all of the helium that is pulsed into the cell ends up propagating through to the trap region along with the Dy beam, as this would hugely increase the helium density inside the trap region. As the beam leaves the cell, the centre-line helium density decreases as [121],

$$\rho = \frac{\rho_0}{\left(1 + \frac{\gamma-1}{2}M^2\right)^{1/(\gamma-1)}}, \tag{6.2}$$

where ρ_0 is the initial helium density, $\gamma = 5/3$ for a monatomic gas and M is the Mach number. The Mach number is defined as the ratio of the local value of flow speed, u, to the speed of sound, a, in the given gas,

$$M = \frac{u}{a}. \tag{6.3}$$

As the pressure in the cell is significantly greater than the background pressure in the chamber, we are operating in the supersonic flow regime. In this regime, and for an axisymmetric expansion (which is the case in our cell as the exit aperture is circular), the Mach number is defined as a function of distance from the cell as [121],

$$M(X) = X^{\gamma-1}\left(C_1(\gamma) + \frac{C_2(\gamma)}{X} + \frac{C_3(\gamma)}{X^2} + \frac{C_4(\gamma)}{X^3}\right) \text{ for } X > 0.5, \tag{6.4}$$

$$M(X) = 1 + A(\gamma)X^2 + B(\gamma)X^3 \text{ for } 0 < X < 1, \tag{6.5}$$

where X is the reduced distance from the cell exit aperture and is defined as $X = x/D$, with x the actual distance from the cell and D the diameter of the exit aperture. The coefficients C_1, C_2, C_3, C_4, A and B for a monatomic gas (i.e. $\gamma = 5/3$) are known [121].

It is imperative to ensure that the helium density has decreased sufficiently by the time the beam reaches the bottom magnet, after which it is guided through the hole in the magnet into the trap region. Inserting the definition of the Mach number, M, from equations 6.4 and 6.5 into equation 6.2 for the beam density allows the helium density in the beam to be calculated as a function of reduced distance from the cell. The diameter of our cell aperture is 5 mm, and the distance from the top of the cell to the bottom of the first magnet is 32.6 mm, leading to a reduced distance of 6.52 to the bottom of the first magnet. With a starting density in the cell of $10^{22}\,\mathrm{m}^{-3}$, the helium density of the beam at the entry point to the first magnet will be $4 \times 10^{19}\,\mathrm{m}^{-3}$. This is almost two orders of magnitude lower than the density required from the second valve to stop the Dy in the trap region (see calculation in following paragraph) and therefore the helium density from the first beam will not contribute significantly to the total helium density inside the trap region.

For this trap set-up to be successful, the number of collisions required to stop the beam of Dy atoms inside the trap region has to be realistically achievable. Using equation 5.5, a 4 K supersonic beam would emerge from the cell at 204 m/s. Using a 1-dimensional hard sphere model, and assuming the collisions between the Dy and helium atoms to be elastic, energy and momentum conservation can be used to estimate that 14 collisions with helium atoms are required to stop a Dy atom from 200 m/s to rest. This is significantly less than the 190 collisions required to cool the Dy atoms after ablation. These 14 collisions need to take place within the trap region, otherwise the Dy atoms will be lost from the trap. Even though the full distance between the magnet faces is 2 cm, the cooling collisions are likely to have to take place within a 1 cm region at the trap centre to have a realistic chance of trapping the atoms. Using the mean free path equation 6.1, and assuming that 14 collisions have to take place within a distance of 1 cm, then $14\lambda < 0.01$, which leads to a minimum helium density in the trap region of 10^{21} m^{-3}. This is an order of magnitude lower than the helium density required to stop and cool the Dy atoms from the ablation plume, and therefore greatly improves the trapping prospects compared to the set-up in chapter 5, where the Dy atoms were trapped directly after ablation.

6.4 Results

The calculations above have shown that there is a significant helium density advantage to this trap set-up over the set-up from figure 4.3. As with the Li (chapter 4) and Dy (chapter 5) experiments, aluminium 'dummy magnet' cylinders were placed inside the set-up in place of the magnets to test various trapping parameters before reverting back to real magnets for the trapping experiments.

6.4.1 Dummy magnet results

In addition to the probe beam through the centre of the trap region, a second probe beam 28 mm above the cell exit aperture was used to record the absorption signal of the Dy beam as it emerges from the cell (see figure 6.1). Figure 6.5 shows a typical absorption signal above the cell and the corresponding LIF signal in the trap region, in the case where dummy magnets are used and the top valve is not fired. The probe frequency was tuned to the ^{164}Dy transition as for the data in the previous chapter. Taking the time between the maxima of the absorption and LIF peaks, along with the distance between the two probe beams, a rough estimate of the velocity of the beam is 170 ± 10 m/s. From equation 5.5, the velocity of a supersonic beam entrained in 4 K helium should be 204 m/s. This shows that the number of collisions near the cell exit aperture is not quite high enough to lead to a fully supersonic beam (see section 6.1). The maximum absorption measured 28 mm above the exit aperture of the cell is 2.4%. The

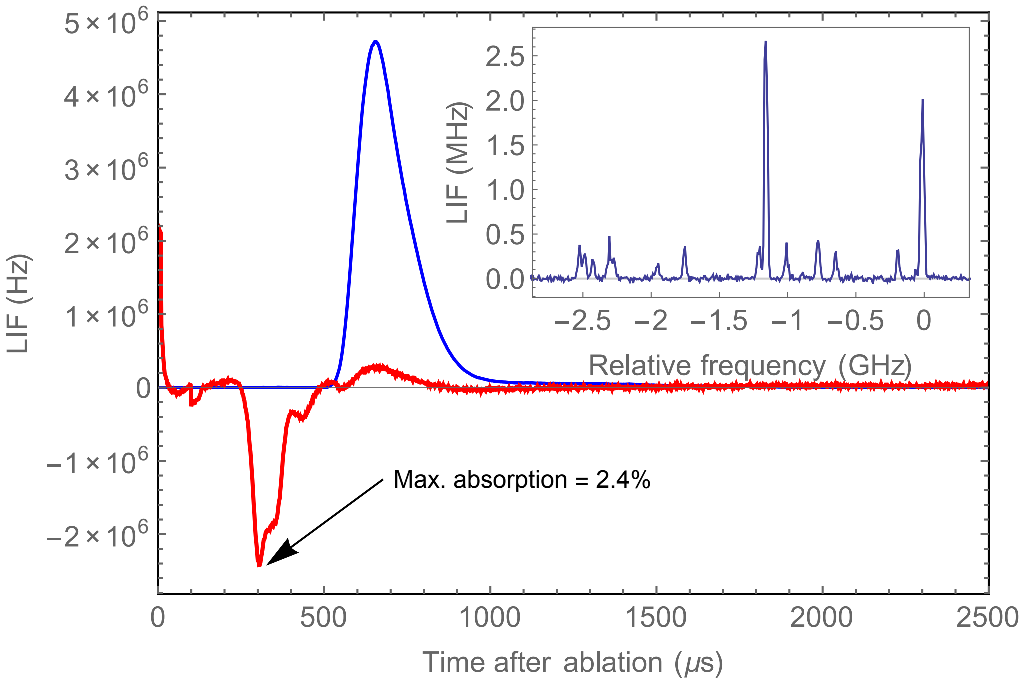

Figure 6.5: Laser-induced fluorescence (LIF) signal of Dy in the trap region (blue line). Absorption signal just above the cell (red line) scaled up and plotted on the same axes. The y-axis values only apply to the blue LIF line, with the red absorption peak maximum being at 2.4% absorption. Inset is a LIF spectrum of Dy in the trap region, taken from average LIF signal gated between $t = 800$ and $900\,\mu$s, with the frequencies relative to the ^{164}Dy transition. All data taken with 'dummy magnets' and without the top valve.

fractional absorption is given by,

$$\frac{I_{in} - I_{out}}{I_{in}} = 1 - e^{-n\sigma z}, \tag{6.6}$$

where n is the atom number density, z is the width of the Dy beam and $\sigma = 3\lambda^2/(2\pi)$ is the resonant absorption cross-section for a two-level atom. Of course, Dy is not really a two-level atom, and as a result the absorption cross-section for the transition used for detection is calculated to be a factor of 6 smaller. Additionally, the resonance line is Doppler broadened, with a temperature of 4 K. At any one frequency, we are only resonant with a particular velocity sub-set of the Doppler broadened line. To take into account all atoms in the Doppler profile, the absorption cross section is reduced by a factor $\sqrt{2\pi}\Gamma/(12\sigma_D)$, where Γ is the natural linewidth of the transition and $\sigma_D = \sqrt{k_B T/(m\lambda^2)}$ is the Doppler broadened linewidth. Using this adjusted absorption cross section, a value for nz can be found using equation 6.6. Integrating this over the beam area (FWHM diameter of 1.15 mm), the instantaneous ^{164}Dy atom number at peak absorption is calculated to be approximately 2×10^8. Integrated over the full length of the absorption signal, this gives a total number of ^{164}Dy atoms in the pulse of about 3×10^9.

The inset in figure 6.5 shows a typical Dy LIF spectrum in the trap region with dummy magnets. Each of the theoretical transitions from figure 5.2 are much more clearly resolved than they were in figure 5.4, which shows the Dy spectrum in the original trap. This is because the Dy in the original trap was in an 8 K thermal distribution, and so the lines were Doppler broadened. In the beam set-up on the other hand, the spatial filtering of the LIF detection due to the iris in front of the PMT means that only the central part of the beam is detected, where atoms with low transverse velocities reside. This reduces the Doppler broadening of the spectral features.

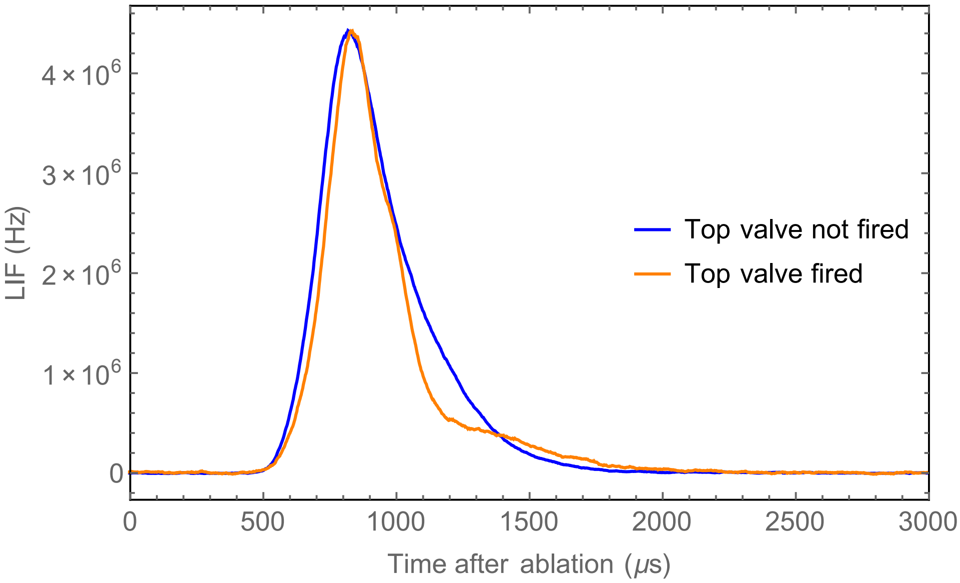

Figure 6.6: LIF signal from the trap region with dummy magnets, with (orange) and without (blue) the top valve. Fitting exponential decay curves to the part of the signal where $t > 1400\,\mu$s, gives a decay time of $156 \pm 1\,\mu$s without firing the top valve and $240 \pm 2\,\mu$s when the top valve is fired.

Figure 6.6 shows the effect on the LIF signal when firing the top valve. The valve pulse length of the top valve, VPL$_{top}$, was 230 μs; the top valve firing time, $t_{top\,valve}$, was 300 μs; and the He pressure behind the top valve was 100 mbar. As before, $t = 0$ is taken from the time at which the ablation laser is fired. Firing the top valve with these settings arrests part of the Dy signal from about 1050 μs after ablation, and leaves a second tail in the signal with a longer decay time than the signal without the top valve. The decay time of this second tail is $240 \pm 2\,\mu$s, compared to the $156 \pm 1\,\mu$s decay time when the top valve is not fired. This clearly shows that the helium pulse from the top valve affects the passage of the Dy beam through the trap region, and causes a portion of it to have a longer decay time. The increase in decay time

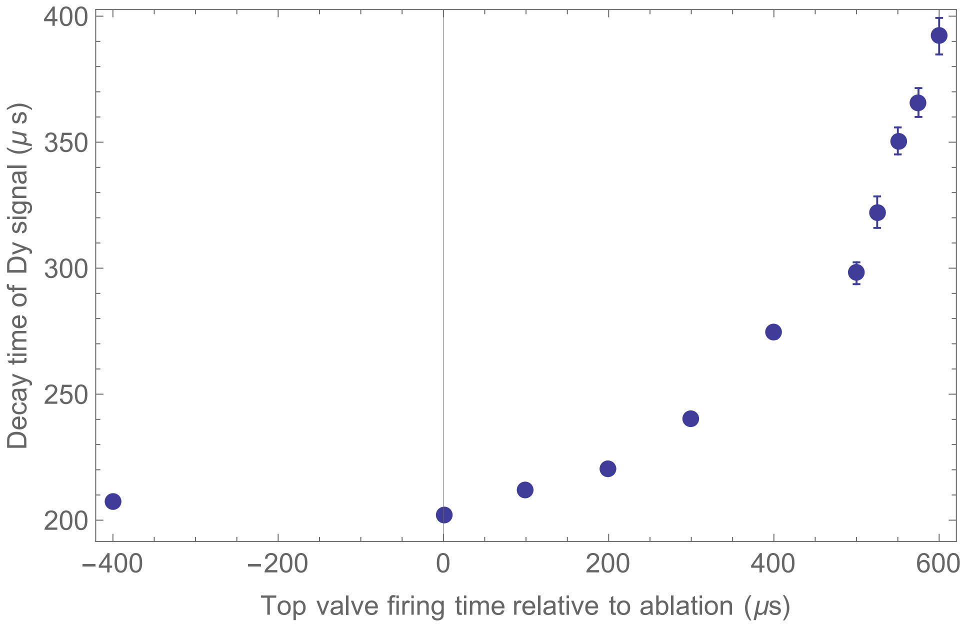

Figure 6.7: Dummy magnet data for the decay time of the Dy signal in the trap region for different top valve firing times, $t_{top\,valve}$. The ablation laser is fired at $t = 0$.

when the top valve is fired can be explained by the fact that the helium pulse from the top valve collides with the Dy beam in the trap region and slows down the tail of the beam, which then appears as a longer decay time.

The signal in figure 6.6 was for a top valve firing time, $t_{top\,valve}$, of 300 μs, and resulted in the signal being arrested from about 1050 μs. The time at which the top valve is fired determines at what point the He pulse starts to affect the Dy pulse, with later $t_{top\,valve}$ times resulting in the signal being arrested at later times. This 'arrest time' is directly related to $t_{top\,valve}$, i.e. increasing $t_{top\,valve}$ by 100 μs would result in the signal also being arrested 100 μs later. Fitting exponential decay curves to the tail of the signal for different top valve firing times gives the plot in figure 6.7 of decay time against $t_{top\,valve}$. The later the top valve firing time, the longer the decay time of the tail of the signal. This increases all the way to a decay time of almost 400 μs with a $t_{top\,valve}$ of 600 μs. Beyond the $t_{top\,valve}$ of 600 μs, the signal gets arrested so late that there is not enough signal left to detect. The longer decay times for later $t_{top\,valve}$ times are speculated to arise because firing the top valve later increases the density of helium in the trap region at the time when the tail of the Dy beam arrives. This increases the diffusion time of the arrested Dy in the trap region.

There are two ways to control the amount of helium introduced into the trap region from the top valve; either by changing the helium pressure behind the valve or by varying the valve

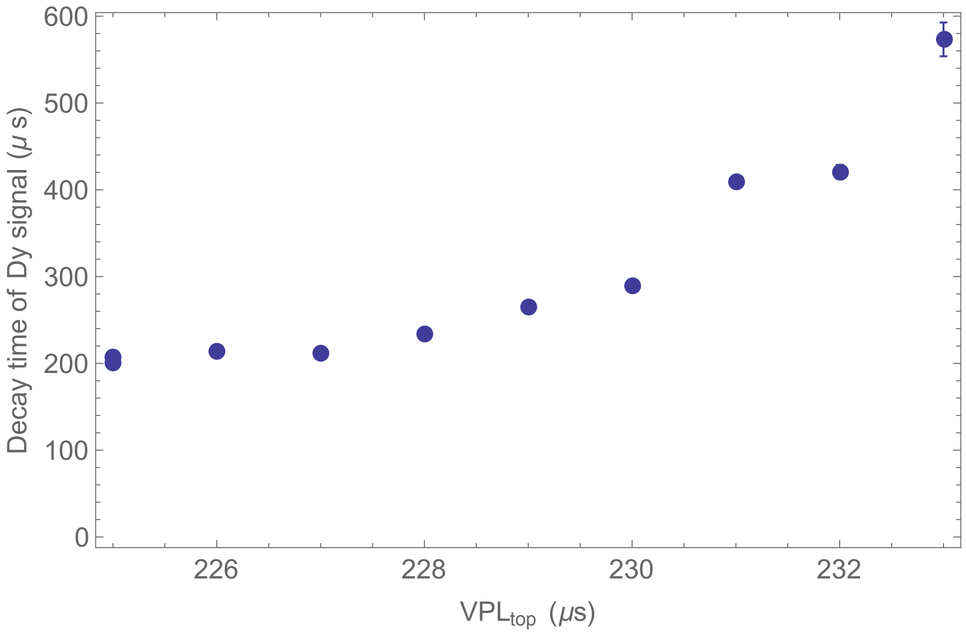

Figure 6.8: Dummy magnet data for the decay time of the Dy signal in the trap region against the valve pulse length of the top valve, VPL$_{top}$.

pulse length of the top valve, VPL$_{top}$. Changing VPL$_{top}$ has proven to give us finer control over the amount of helium and therefore this was varied whilst keeping the pressure behind the valve constant. Fixing $t_{top\,valve}$ at 400 μs and increasing VPL$_{top}$ led to a larger fraction of the Dy signal being arrested. Fitting exponential decay curves to the remaining Dy signal tail shows that the decay time increases for increasing VPL$_{top}$. A plot of the decay time against VPL$_{top}$ is shown in figure 6.8. This shows that the decay time goes up to 570 ± 20 μs for a VPL$_{top}$ of 233 μs, after which the VPL$_{top}$ becomes so high that all of the Dy signal is cut off. The increase in decay time can be explained by a higher helium density in the trap region due to the increased VPL$_{top}$, and therefore the Dy takes longer to diffuse away through the helium (see equation 4.3). The fact that the decay time keeps increasing for increasing VPL$_{top}$ confirms that the decay time is likely dominated by diffusion through the helium gas.

It is clear from figure 6.8 that increasing the amount of helium pulsed into the trap region from the top valve increases the observed decay time of the Dy signal. However, this is not necessarily the best strategy when trapping Dy in the magnetic field as high background helium densities can lead to trap loss through spin-flip collisions or momentum kicks. For the best trapping prospects, we want the helium pulse to convert as high a fraction of the Dy beam as possible into a cloud of 4 K atoms with no centre of mass velocity. It is possible to get an indication of whether the Dy is in a cloud or a beam by looking at the widths of the peaks in the frequency spectra.

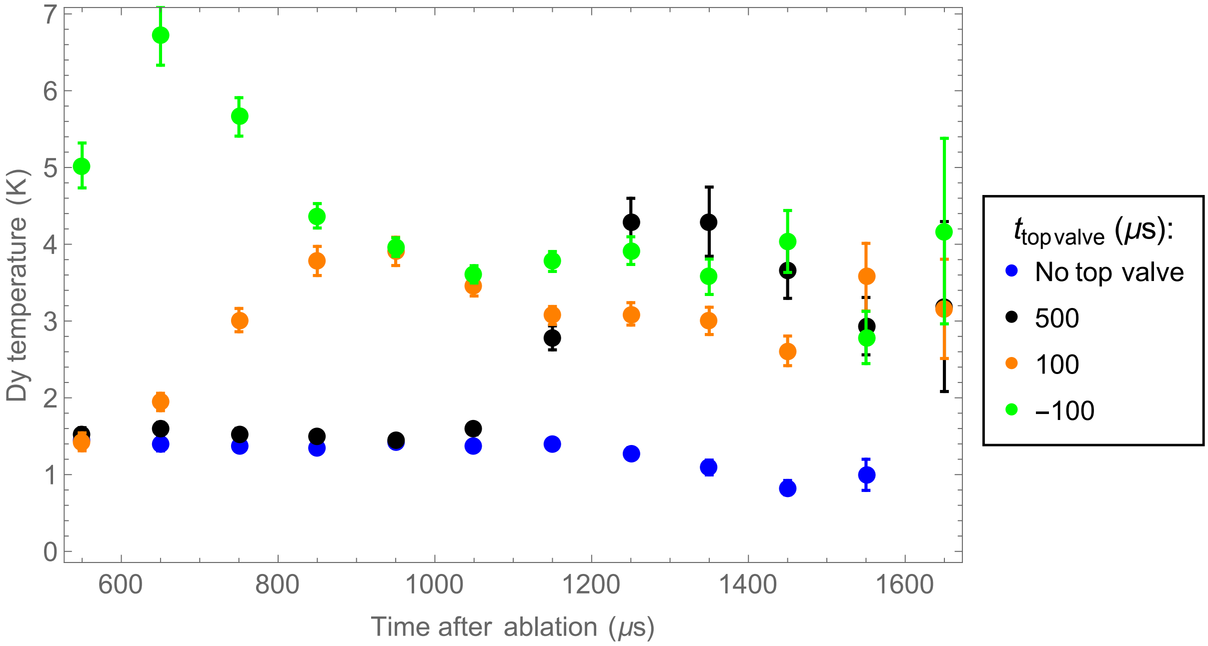

Figure 6.9: Temperature of Dy in the trap region against time after ablation, with dummy magnets. Data were taken for several different top valve firing times, with: $t_{top\,valve} = 500\,\mu$s (black dots); $t_{top\,valve} = 100\,\mu$s (orange dots); $t_{top\,valve} = -100\,\mu$s (green dots). The temperature when the top valve is not fired is also shown (blue dots).

In a beam, there are very few collisions between Dy atoms and the helium, or between the Dy atoms themselves. This means that as the beam propagates, the atoms with lower transverse velocities remain in the central part of the beam, and those with higher transverse velocities move to the outer edges of the beam. With the limited field of view of the LIF detection, only those atoms with the lowest transverse spreads are detected, which leads to the narrow spectral lines seen in the inset in figure 6.5. Fitting Gaussian peaks to those spectral lines shows a ‘temperature’ of about 1.4 K. The iris in front of the PMT has a diameter of about 10 mm which means that the field of view of the PMT is also 10 mm in diameter as the LIF set-up has one-to-one imaging (see figure 4.5). By taking the beam to be at a temperature of 4 K, and assuming it to expand freely in the transverse direction from the moment it exits the cell, only those atoms with a temperature below 0.6 K are expected to be within the field of view of the PMT. However, the beam is not actually a point source, and atoms starting in ‘off-axis’ positions in the exit aperture can still reach the LIF detection region, even with higher transverse velocities. Taking this into account, we can expect to detect atoms with temperatures up to 1.4 K. This is in agreement with experimental observations.

In a cloud of atoms, collisions between the atoms ensure that the whole cloud is thermalised to the same temperature. This means that, despite the limited field of view of the LIF detection, if an atomic cloud is present in the trap region, the temperature measured from the line widths in the spectrum will correspond to the actual temperature of the cloud. By taking spectra of the ^{164}Dy line at different times after ablation and fitting a Gaussian to the peak, the temperature of the Dy can be determined as a function of time. Figure 6.9 shows this temperature variation for three different $t_{top\,valve}$ times, as well as for no top valve. When the top valve is not fired, the 'temperature' of the Dy remains at about 1.4 K as the Dy just passes through the trap region as a beam. However, if the top valve is fired, the temperature of the Dy appears to rise and levels off between 3-4 K. This temperature increase can be attributed to part of the Dy beam being turned into a cloud as a result of the top valve helium pulse. The later $t_{top\,valve}$, the later this increase in temperature occurs. Figure 6.9 shows that for a $t_{top\,valve}$ of 100 μs, the fastest part of the Dy pulse passes through as a beam, but then the rest of the pulse is turned into a cloud at a temperature between 3 and 4 K. Firing the top valve later, at $t_{top\,valve} = 500\,\mu$s, results in the cloud being formed much later and therefore a lot of the Dy will have passed through the trap region as a beam. Firing the top valve earlier, at $t_{top\,valve} = -100\,\mu$s, means that none of the pulse passes through the trap region as a beam, but the temperature is initially significantly higher at almost 7 K. It has been observed before that the early part of a pulsed, supersonic helium buffer gas beam emerges from the cell hotter than the rest of the pulse [120]. This would explain why the temperature is observed to be at almost 7 K initially, for $t_{top\,valve} = -100\,\mu$s. It would seem desirable to let the fastest atoms pass through the trap region as a beam as they are unlikely to be trapped anyway, and therefore a $t_{top\,valve}$ of 100 μs was used for the magnet results below.

6.4.2 Magnet results

The dummy magnet results have shown that the top valve pulse has a clear effect on the Dy beam and that it turns at least part of the Dy beam into a cloud at the trap centre. The magnets were therefore returned to the set-up to see if these atoms could be trapped. The strength of the magnetic field from the two magnets affected the firing of both the solenoid valves. If the same valve pulse length, on the order of 200 μs, was used as with the dummy magnets then no helium was released during the pulse as the field from the magnets slowed the valve opening process and therefore prevented the valve from opening in that time. In order to release the same amount of helium during the valve pulses (determined by measuring the

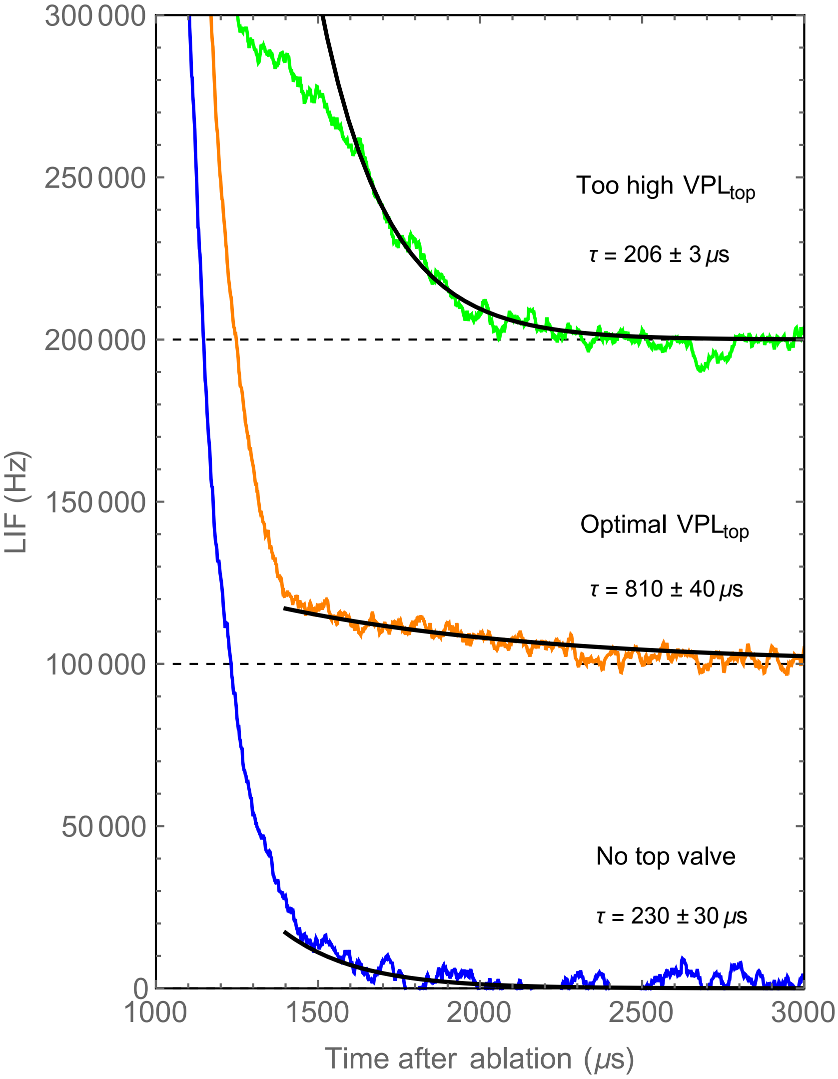

Figure 6.10: Plot of LIF signal from the trap region with magnets in place, for different top valve pulse lengths, VPL$_{top}$. Blue line: no top valve; orange line: VPL$_{top}$ = 315 μs; green line: VPL$_{top}$ = 325 μs. Plots are offset by 100 kHz for comparison.

background pressure in the chamber on the Penning gauge whilst firing the valve), the valve pulse lengths had to be increased to values on the order of 300 μs.

As with the dummy magnets, firing the top valve caused a second decay tail to appear in the signal. Figure 6.10 shows a comparison of the LIF signal from the trap region for two different top valve pulse lengths, VPL$_{top}$, as well as for no top valve. Without the top valve, the signal passes through the trap region with a 230 ± 30 µs decay time. This is similar to the lowest decay times seen for dummy magnets in figures 6.7 and 6.8. When the top valve is fired with an optimum VPL$_{top}$ of 315 μs, there is a very clear tail in the signal with a longer decay time of 810 ± 40 µs. If VPL$_{top}$ is increased even further to 325 μs, then a large second bump appears in the signal. When fitting an exponential to the decaying part of the bump though, the decay

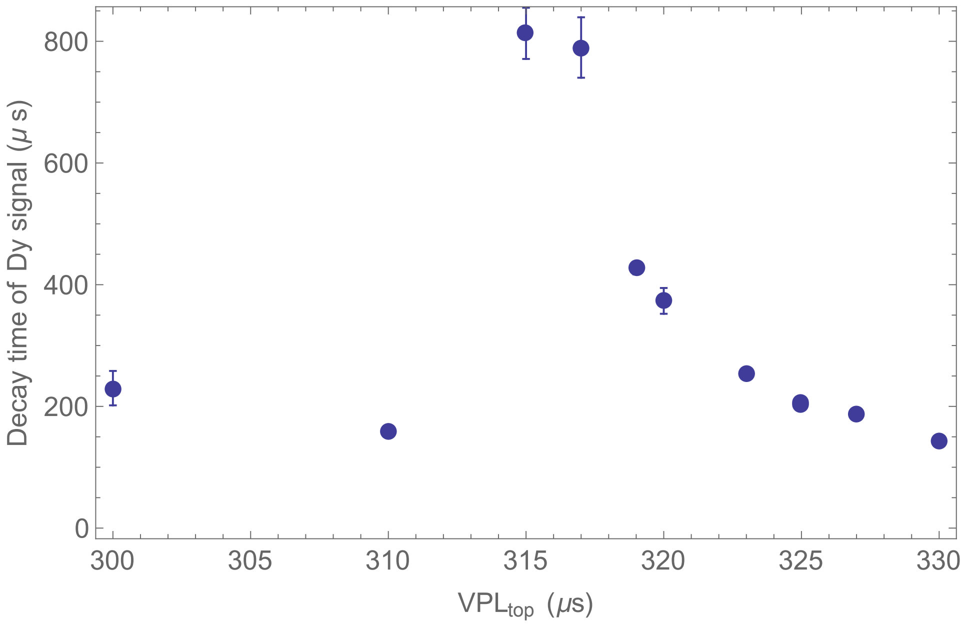

Figure 6.11: Magnet data for the decay time of the Dy signal in the trap region against the valve pulse length of the top valve, VPL$_{top}$.

time is only 206 ± 3 μs– similar to the decay timescale without a top valve pulse.

Figure 6.11 shows the decay time against VPL$_{top}$. There is a clear peak in the decay time for a VPL$_{top}$ of 315 μs, after which the decay time decreases again down to similar levels seen without the top valve pulse. This decrease in decay times is not seen in figure 6.8 for dummy magnets, where the decay time keeps increasing for increasing VPL$_{top}$. This means that, with the magnets in place, the increase in decay time is unlikely to be due to diffusion through higher helium densities, as it was for dummy magnets. For an optimum VPL$_{top}$, the helium pulse slows the Dy just enough to be in the capture range of the magnetic trap. Increasing VPL$_{top}$ further then adds more background helium to the trap region which increases trap losses and therefore decreases the decay time. When VPL$_{top}$ reaches 325 μs, the decay time is no longer than the timescale on which we see the original beam pulse decay.

The VPL$_{top}$ was 315 μs for the maximum decay time with magnets, and 233 μs for the maximum decay time with dummy magnets. Unfortunately, these VPL$_{top}$ values are not a good comparison of the amount of helium pulsed into the trap region in each case as the magnetic field affects the firing of the valve. Instead, we can look at the background pressure in the chamber measured on the Penning gauge during each data run. This pressure was 3×10^{-6} mbar for the magnet data with the longest decay time of 810 ± 40 μs. The decay time with dummy magnets for a VPL$_{top}$ that gave a similar background pressure was much lower at around 200 ± 2 μs. The maximum decay time seen with dummy magnets of 570 ± 20 µs is for a much higher background pressure of 9×10^{-6} mbar, and even then this decay time is less than that seen with magnets. It is therefore clear that, to achieve the longest dummy magnet decay time, a lot more helium was pulsed into the trap region and it is therefore very unlikely that the long decay time with magnets could have been caused by diffusion through the helium. This, in combination with the fact that there is a clear peak in decay time against VPL_{top} for magnets (figure 6.11), as opposed to the continual increase in decay times with VPL_{top} for dummy magnets (figure 6.8), provides strong evidence that there is magnetic trapping in this set-up. Using equation 4.5, the oscillation period in this trap can be up to 7-8 ms, and therefore the trap lifetime of 810 ± 40 µs is not likely to correspond to many oscillations of Dy in the trap. The trap lifetime is likely limited by collisions with background helium gas, as the decrease in lifetime for higher VPL_{top} in figure 6.11 confirms.

In the direct trapping experiments in chapter 5, we were able to distinguish between LFS and HFS atoms by a splitting of the peaks in the spectrum (see figure 5.13). The difference between the decay times of the LFS and HFS atoms were then used to determine whether the atoms were trapped. In the magnetic trap used for the beam trapping experiments in this chapter, the magnets were scaled up, increasing the trap area and decreasing the magnetic field gradient. This decreased magnetic field gradient means that, with the probe laser passing through the trap centre, the range of magnetic field values probed with the LIF detection is smaller. This reduces the spacing between the peaks, as this is determined by the strength of the magnetic field (see equation 5.6). It is therefore difficult to accurately distinguish between the LFS and HFS peaks in the spectrum, which is why this method is not used here.

6.5 Conclusion

In these experiments, Dy atoms were trapped from a buffer gas cooled beam, rather than directly after ablation. Trapping in this way was motivated by results from the previous Dy trapping experiments in chapter 5, which indicated that the helium density required to stop the ablation plume resulted in high helium densities in the trap region which limited trap lifetimes. By separating the ablation process from the trapping process, these high helium densities were avoided.

A simple copper cell with a circular exit aperture was used to create a pulsed beam of Dy atoms with a forward velocity of about 170 m/s and a ^{164}Dy atom number of approximately 3×10^9. This beam propagated vertically into the trap region, where it was stopped by a counter-propagating helium pulse from a valve placed just above the trap region.

The first experiments with dummy magnets showed that the helium pulse from the top valve has a clear effect on the Dy beam. The helium pulse arrested the tail of the Dy beam, leaving a new tail with a significantly longer decay time. Increasing the top valve pulse length, VPL_{top}, to introduce more helium into the trap region increased the decay timescale for this second tail, until VPL_{top} was so long that the whole tail of the Dy beam was arrested, with no signal left to detect. The longest decay timescale seen was about $570 \pm 20\,\mu\mathrm{s}$. The increase in the decay timescale with increasing helium in the trap region is likely a result of diffusion through the higher helium gas density.

Conducting the same experiments with the magnets in place did not lead to a continual increase in decay time with increasing VPL_{top}, but instead showed a peak in decay time for a VPL_{top} of $315\,\mu$s, after which the decay time decreased again. The longest decay time observed in this case was $810 \pm 40\,\mu\mathrm{s}$. This decay time was observed for a much lower helium density in the trap region than the longest decay time with dummy magnets, and can therefore be attributed to magnetic trapping of the Dy atoms. The trap lifetime here is comparable to the longest lifetimes seen in the direct Dy trapping experiments in chapter 5, however in this case the trapping process was much more repeatable and controllable due to the separation of the ablation from the trap region. The fact that the trap lifetime here is not longer than that seen in the direct trapping experiments, despite the helium density being lower, could be explained by the slightly lower trap depth, which means that fewer collision are required to lose atoms. A potential way to increase the trap depth is discussed in section 7.3.

CHAPTER 7

CONCLUSION & OUTLOOK

Two different arrangements for magnetically trapping buffer gas cooled atoms are presented in this thesis: direct trapping from an ablation plume and trapping from a cryogenic buffer gas cooled beam. A summary of the main findings is presented in this chapter, as well as some potential improvements to increase trap lifetimes.

7.1 Direct trapping after ablation

The first experiments focused on trapping Li and Dy atoms directly from a buffer gas cooled ablation plume. The magnetic trap consisted of two cylindrical NdFeB magnets, aligned with their north poles facing each other to form a quadrupole trapping field at the centre. A 4 K helium pulse was delivered to the trap region through a solenoid valve, and the Li/Dy atoms were subsequently ablated into the helium pulse from a solid precursor target. This allowed the atoms to thermalise with the helium in the trap region. The trap geometry was kept as open as possible to then allow the non-magnetic helium atoms to disappear, leaving the Li/Dy atoms trapped in the magnetic field. Detection of the atoms in the trap was either through absorption, detected on a photodiode, or through laser-induced fluorescence (LIF) detected on a photomultiplier tube (PMT). The whole trap set-up was kept at 4 K by a closed-cycle cryocooler.

For the initial Li experiments, the helium pulse was guided into the trap region through a copper tube attached to the solenoid valve. With aluminium cylinders (‘dummy magnets’) in place of the magnets, non-Zeeman broadened frequency spectra of Li were taken. The widths of the transition peaks were then used to determine the temperature of the Li as a function of time, showing that the Li thermalised with the helium within about 500 μs. With the magnets back in place, it appeared that Li was detected in the trap region with a lifetime of 10 ms, thought to be limited by collisions with background helium gas. However, as later tests showed, it is also possible that this long decay time was actually a result of artificial signal caused by saturation of the PMT due to the bright ablation plume light.

To avoid this problem in subsequent experiments, a fast switching circuit was designed that allowed the PMT to be switched off during the ablation pulse and then rapidly switched on again to capture the LIF signal. Additionally, delivering the helium to the trap region with the copper tube attached to the valve resulted in a long helium tail emerging from the tube for many tens of milliseconds, when the helium pulse should actually be around 200 μs long. This is likely to have severely reduced trap lifetimes. To ensure that a short helium pulse was delivered as required, the trap was redesigned so that the valve pulses the helium directly into the trap region, without the need for a copper tube. This gave much greater control over the length and timing of the helium pulses.

With these modifications implemented, we decided to try the experiment with Dy atoms instead of Li. With a magnetic moment of $10\mu_B$, Dy has a tenfold higher trap depth than Li in the same trap arrangement. A favourable detection transition also means that the Zeeman broadening of the resonance lines is very small, making detection in the magnetic field significantly easier. With 'dummy magnets' in place, the thermalisation of Dy with the helium buffer gas was observed to occur within about 200 μs. Absorption data taken with a counter-propagating probe beam in place showed a centre-of-mass velocity of the Dy cloud up to 250 m/s in the direction of the probe beam. This was assumed to result from the Dy being entrained in a supersonic helium pulse from the valve; and was solved by placing a copper cylinder around the trap region to contain the helium pulse, and ablating the Dy into the resulting, stationary helium cloud.

Having achieved a cold, stationary Dy cloud in the trap region, the magnets were placed back into the set-up. In the magnetic field, it was possible to distinguish between low-field seeking (LFS) and high-field seeking (HFS) atoms due to a splitting of the spectral peaks. This showed that, for an optimal valve pulse length, the LFS signal decayed with a lifetime of $360 \pm 30\,\mu$s, compared to the decay time of the HFS atoms of $180 \pm 1\,\mu$s. This is clear evidence that there is trapping of the Dy in the magnetic field. The LFS decay time is assumed to be limited by collisions with excess helium in the trap region. Subsequent experiments confirmed that there was indeed a high helium density in the trap region of about $2 \times 10^{20}\,\mathrm{m}^{-3}$, even 500 ms after the valve had fired. Replacing the copper cylinder with a mesh cylinder increased the lifetime of the LFS atoms to 800 ± 30 µs, as the helium was able to leave the trap region more rapidly, but a significant background helium density still remained. By increasing the temperature of the set-up to 6.5 K, this background helium density was reduced significantly, suggesting that the helium was being adsorbed onto a cold surface near the trap region and slowly released over time. Running the experiment at these slightly higher temperatures did not appear to improve the trap lifetimes though.

There has been clear evidence, particularly from the Dy experiments, of magnetic trapping directly from the buffer gas cooled ablation plume. However, the data from both the Dy and Li experiments highlighted the difficulty of improving upon the trap lifetimes achieved in this set-up. With the ablation happening so close to a relatively small trap region, a high helium density is required initially in order to stop and cool the ablation plume. Removing such a high helium density quickly enough to achieve longer trap lifetimes has proven difficult. The relatively unpredictable nature of laser ablation demands that a very precise set of circumstances have to come together in a small area in order for the atoms to be trapped, and whilst we have seen evidence of this, achieving repeatable, longer lifetimes has proven difficult.

7.2 Trapping from a Dy beam

To reduce the helium density inside the trap region, we decided to separate the ablation process from the trapping process. In doing so, it allowed higher helium densities to be used to stop and cool the ablation plume, with lower helium densities then required in the trap region. This separation was realised by having a cell on top of the cold plate, in which the Dy was ablated and thermalised with a helium pulse from a solenoid valve. An aperture at the top of the cell allows a buffer gas cooled beam of Dy atoms to emerge, and travel vertically up towards the trap region, which was again constructed from cylindrical NdFeB magnets and placed a few cm above the cell. Once the beam had reached the trap region, a second helium pulse from another valve was pulsed into the trap region from above, arresting some of the Dy beam and leaving it trapped. Holes through the centre of the magnets allowed the Dy beam and helium pulse to enter the trap region. The helium densities required in the trap region to stop the Dy beam were calculated to be significantly lower than those required to stop and cool the ablation plume.

As with the direct trapping experiments, ‘dummy magnets’ were first placed into the set-up to test the effect of the second valve pulse on the Dy beam. These experiments showed that the Dy beam emerged from the cell with a velocity of about 170 m/s and a ^{164}Dy atom number of approximately 3×10^9. When pulsing the second valve, a clear effect on the Dy beam was observed, with part of the beam being arrested by the helium pulse, leaving a tail with a longer decay time. This decay time increased the more helium was pulsed into the trap region, and was therefore attributed to diffusion through the helium gas. With magnets in place, the decay time did not keep increasing with increasing helium density, but instead peaked at a much lower helium density in the trap region, with the decay time decreasing again if the helium density was increased further. With an optimum amount of helium pulsed into the trap region, the Dy pulse is slowed just enough for the atoms to be trapped, with any further increase in helium resulting in trap losses, which lowers the observed trap lifetime. The decay time achieved of $810 \pm 40\,\mu$s was significantly longer than the $200 \pm 2\,\mu$s decay time observed with ‘dummy magnets’ for similar helium pressures, providing clear evidence of magnetic trapping.

Even though the decay time observed in this set-up was no longer than that from the direct trapping experiments, this method was much more repeatable and reliable. Having the atoms enter the trap region from a buffer gas cooled beam ensured a more consistent delivery than directly from an ablation plume. As buffer gas beams of molecules are now readily produced (see for example [84]), this trapping method would, just like the direct trapping, not be limited to atomic species.

7.3 Possible improvements

There are a number of possible improvements that could be made to the set-up to increase the trap lifetimes. The main limiting factor to the trap lifetimes appears to be collisions with background helium gas. The most drastic improvements would therefore be achieved by reducing the background buffer gas density.

One main problem that was identified, particularly in the Dy experiments in chapter 5, is the adsorption of helium onto cold surfaces near the trap region. This helium is then slowly released over time, keeping the background helium density in the trap region relatively high, and reducing trap lifetimes. In the case of the beam trapping experiments in chapter 6, we can estimate how quickly the helium should dissipate from the trap region. The diffusion equation

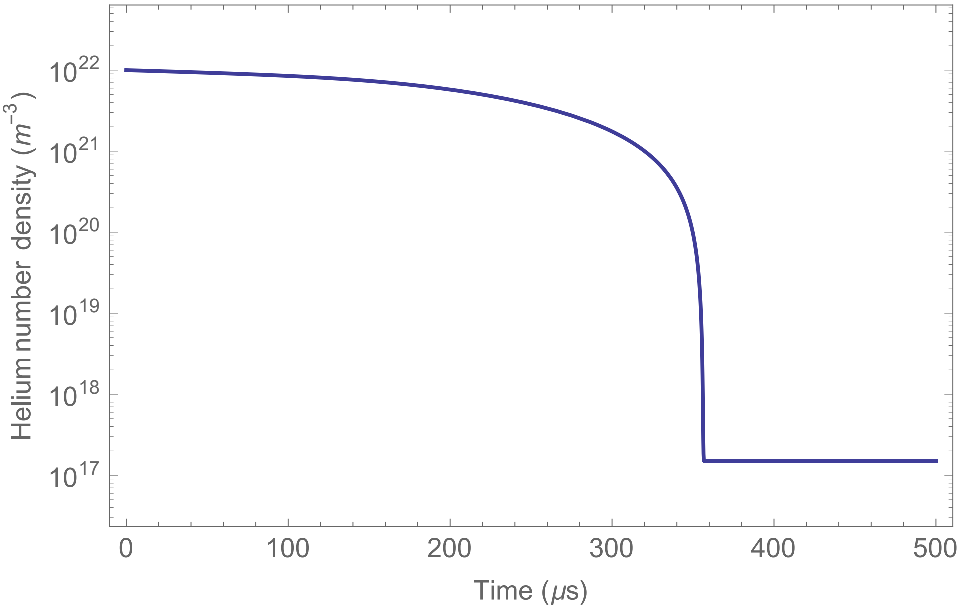

Figure 7.1: Expected helium density as a function of time at the centre of the trap region of the beam trapping set-up. Density is calculated by solving the diffusion equation. A density of $10^{17}\,\mathrm{m}^{-3}$ corresponds to the background pressure in the chamber.

is given by,

$$\frac{\partial n(r,t)}{\partial t} = \nabla \cdot [D(n(r,t))\nabla n(r,t)], \tag{7.1}$$

where $n(r,t)$ is the helium number density as a function of position and time and D is the diffusion coefficient as a function of density, given by equation 4.4. The collision cross section for He-He collisions is approximately $10^{-18}\,\mathrm{m}^2$ [122]. To estimate the helium density as a function of time, equation 7.1 can be solved in one dimension, assuming a trap region 2 cm in diameter, with a starting Gaussian distribution of helium atoms with a peak density of $10^{22}\,\mathrm{m}^{-3}$. The result is plotted in figure 7.1. This shows that the helium density at the centre of the trap region initially decays fairly slowly as the helium density is still high, resulting in a short mean-free-path. Then as the density drops, the mean-free-path also becomes much longer, resulting in a rapid decay at around 350 μs down to the background level in the chamber. The plot shows that the helium should dissipate completely within a few hundred microseconds, however the Dy trap lifetimes observed are no longer than for the direct trapping experiments. This suggests that in the beam trapping set-up there is also some adsorption of helium onto surfaces, which is slowly released over time. To reduce background helium densities and increase trap lifetimes, this adsorption has to be avoided.

One possible route for exploration is to use a neon buffer gas, rather than helium, such as in [76]. As neon has a higher boiling point than helium, it would immediately stick to any 4 K surfaces it comes into contact with. This effectively means that all surfaces act as sorption pumps, rather than just the charcoal, as is the case for helium. Neon would therefore be pumped away far more efficiently than helium. The disadvantage is that, in order to deliver a sufficiently high density of neon to the trap region, the gas lines and valve (and the cell in the beam experiments) would have to be kept at around 15 K [76] to keep the neon in the gas phase. This brings with it some added experimental complications as these components would have to be thermally isolated from the 4 K parts of the set-up. Additionally, the atoms/molecules to be trapped would start at 15 K rather than 4 K, reducing the fraction that can be trapped.

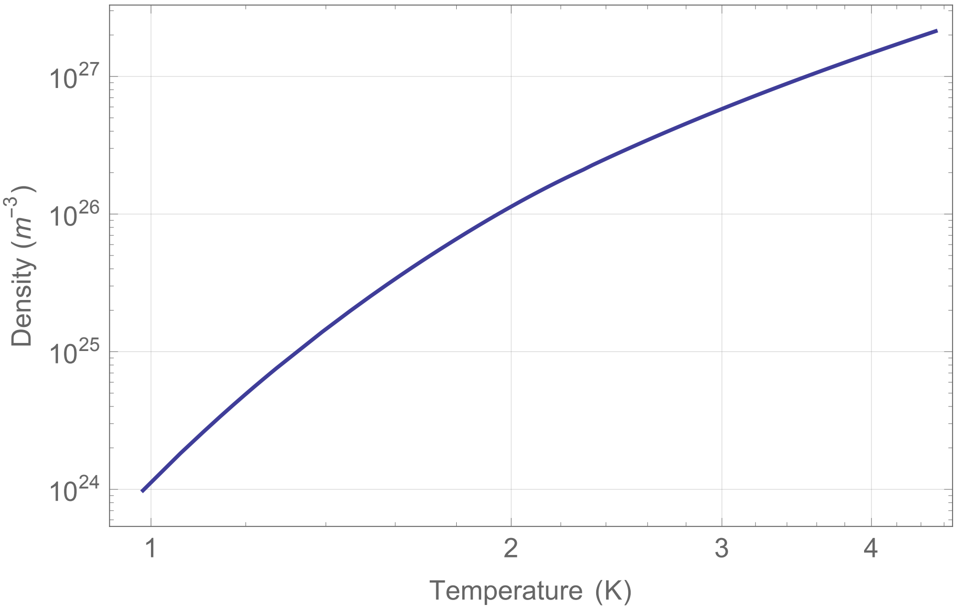

Figure 7.2: Saturated He vapour density as a function of the temperature of the gas [123].

Lowering the temperature of the trap, to 1 K for example, would allow the helium to remain frozen onto the surfaces rather than being released over time. Figure 7.2 shows the saturated vapour density of helium as a function of temperature. This shows that, at 1 K, the saturated vapour density is about 3 orders of magnitude lower than at 4 K. This means that any helium on 1 K surfaces would be frozen out. Temperatures below 4 K can be achieved by using a dilution refrigerator [123], or more practically by installing a small reservoir of helium cooled by the cryocooler and then pumping on the vapour above the liquid in the reservoir in order to lower its temperature to 1 or 2 K. The trap could initially be held at 4 K to create the cloud of helium gas in the trap region for the Dy atoms to thermalise with. Then, once the atoms are trapped, the temperature could quickly be reduced down to 1 K to freeze out any helium on the surfaces and prevent it from being released over time.

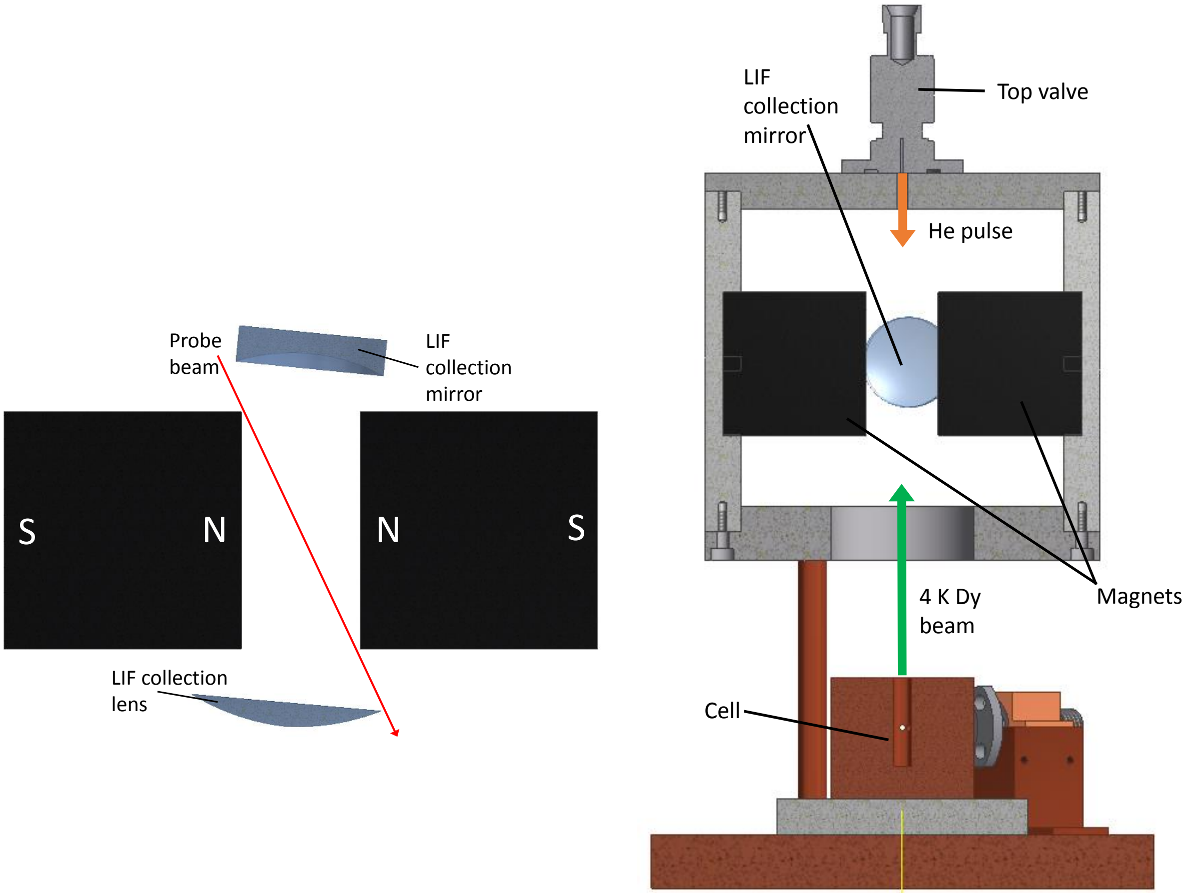

(a) Top-down cross section

(b) Side view cross section

Figure 7.3: Potential horizontal magnet arrangement for the beam trapping experiments. (a) Top-down cross section of the cylindrical magnet and LIF arrangement, showing the LIF collection optics at a slight angle so that the holes in the magnets are no longer needed. (b) Side view cross section showing how the new magnet arrangement would fit into the same beam trapping set-up as in figure 6.1.

In the case of the beam trapping experiments, a significant improvement could be made by placing the magnets horizontally rather than vertically, as shown in figure 7.3. In this arrangement, the Dy beam would travel in free space all the way to the trap region, rather than being guided through the hole in the bottom magnet as in figure 6.1. This would allow the helium density in the beam to decrease by an additional order of magnitude before reaching the trap region (see equation 6.2), reducing the amount of background helium. Additionally, if the LIF set-up is arranged as shown in figure 7.3, there would be no need for any holes in the magnets for the probe beam, increasing the trap depth from 0.23 T to 0.36 T. Even if the LIF detection could not be arranged in this way, and the probe beam had to pass through the magnets, the hole diameter could easily be reduced to 2 mm, giving a trap depth of 0.3 T. In both cases, a higher fraction of the 4 K Dy could be trapped. Trap lifetimes should also increase, as collisions with background helium gas would be less likely to eject an atom from the trap through momentum kicks.

BIBLIOGRAPHY

[1] I. Bloch, Nature **453**, 1016 (2008).

[2] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, Phys. Rev. Lett. **75**, 3969 (1995).

[3] C. Chin, V. V. Flambaum, and M. G. Kozlov, New. J. Phys. **11**, 055048 (2009).

[4] J. P. Uzan, Rev. Mod. Phys. **75**, 403 (2003).

[5] E. R. Hudson, H. J. Lewandowski, B. C. Sawyer, and J. Ye, Phys. Rev. Lett. **96**, 143004 (2006).

[6] S. Truppe, R. J. Hendricks, S. K. Tokunaga, H. J. Lewandowski, M. G. Kozlov, C. Henkel, E. A. Hinds, and M. R. Tarbutt, Nature Communications **4**, 2600 (2003).

[7] K. Beloy, M. G. Kozlov, A. Borschevsky, A. W. Hauser, V. V. Flambaum, and P. Schwerdtfeger, Phys. Rev. A **83**, 062514 (2011).

[8] E. J. Salumbides, M. L. Niu, J. Bagdonaite, N. de Oliveira, D. Joyeux, L. Nahon, and W. Ubachs, Phys. Rev. A **86**, 022510 (2012).

[9] M. Kajita, G. Gopakumar, M. Abe, M. Hada, and M. Keller, Phys. Rev. A **89**, 032509 (2014).

[10] E. Reinhold, R. Buning, U. Hollenstein, A. Ivanchik, P. Petitjean, and W. Ubachs, Phys. Rev. Lett. **96**, 151101 (2006).

[11] J. Bagdonaite, P. Jansen, C. Henkel, H. L. Bethlem, K. M. Menten, and W. Ubachs, Science **339**, 46 (2013).

[12] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, Phys. Rev. Lett. **89**, 023003 (2002).

[13] J. J. Hudson, D. M. Kara, I. J. Smallman, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, Nature **473**, 493 (2011).

[14] The ACME Collaboration, J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, et al., Science **343**, 269 (2014).

[15] M. R. Tarbutt, B. E. Sauer, J. J. Hudson, and E. A. Hinds, New. J. Phys. **15**, 053034 (2013).

[16] M. Pospelov and A. Ritz, Ann. Phys. **318**, 119 (2005).

[17] A. Micheli, G. K. Brenner, and P. Zoller, Nature **2**, 341 (2006).

[18] A. André, D. DeMille, J. M. Doyle, M. D. Lukin, S. E. Maxwell, P. Rabl, R. J. Schoelkopf, and P. Zoller, Nat. Phys. **2**, 636 (2006).

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

[20] G. Pupillo, A. Micheli, H. P. Büchler, and P. Zoller, in *Cold molecules: Theory, Experiment, Applications*, edited by R. V. Krems, B. Friedrich, and W. C. Stalley (CRC press, 2009), pp. 421–472.

[21] S. F. Yelin, D. DeMille, and R. Côté, in *Cold Molecules: Theory, Experiment, Applications*, edited by R. V. Krems, B. Friedrich, and W. C. Stalley (CRC press, 2009), pp. 629–648.

[22] A. Chotia, B. Neyenhuis, S. A. Moses, B. Yan, J. P. Covey, M. Foss-Feig, A. M. Rey, D. S. Jin, and J. Ye, Phys. Rev. Lett. **108**, 080405 (2012).

[23] R. V. Krems, Phys. Chem. Chem. Phys. **10**, 4079 (2008).

[24] J. J. McClelland, R. E. Scholten, E. C. Palm, and R. J. Celotta, Science **262**, 877 (1993).

[25] C. O'Dwyer, G. Gay, B. V. de Lesegno, J. Weiner, A. Camposeo, F. Tantussi, F. Fuso, M. Allegrini, and E. Arimondo, Nanotechnology **16**, 1536 (2005).

[26] L. D. Carr, D. DeMille, R. V. Krems, and J. Ye, New. J. Phys. **11**, 055049 (2009).

[27] H. L. Bethlem, G. Berden, F. M. H. Crompvoets, R. T. Jongma, A. J. A. van Roij, and G. Meijer, Nature **406**, 491 (2000).

[28] H. Sabbah, L. Biennier, I. R. Sims, Y. Georgievskii, S. J. Klippenstein, and I. W. M. Smith, Science **317**, 102 (2007).

[29] K. M. Jones, E. Tiesinga, P. D. Lett, and P. S. Julienne, Rev. Mod. Phys. **78**, 483 (2006).

[30] P. Pellegrini, M. Gacesa, and R. Côté, Phys. Rev. Lett. **101**, 053201 (2008).

[31] J. Deiglmayr, A. Grochola, M. Repp, K. Mörtlbauer, C. Glück, J. Lange, O. Dulieu, R. Wester, and M. Weidemüller, Phys. Rev. Lett. **101**, 133004 (2008).

[32] T. Shimasaki, M. Bellos, C. D. Bruzewicz, Z. Lasner, and D. DeMille (2014), arXiv:1407.7512v1 [physics.atom-ph].

[33] K.-K. Ni, S. Ospelkaus, M. H. G. de Miranda, A. Pe'er, B. Neyenhuis, J. J. Zirbel, S. Kotochigova, P. S. Julienne, D. S. Jin, and J. Ye, Science **317**, 102 (2007).

[34] T. Takekoshi, L. Reichsöllner, A. Schindewolf, J. M. Hutson, C. R. Le Sueur, O. Dulieu, F. Ferlaino, R. Grimm, and H.-C. Nägerl, Phys. Rev. Lett. **113**, 205301 (2014).

[35] P. K. Molony, P. D. Gregory, Z. Ji, B. Lu, M. P. Köppinger, C. R. L. Sueur, C. L. Blackley, J. M. Hutson, and S. L. Cornish (2014), arXiv:1409.1485v3 [physics.atom-ph].

[36] J. D. Weinstein, R. deCarvalho, T. Guillet, B. Friedrich, and J. M. Doyle, Nature **395**, 148 (1998).

[37] N. R. Hutzler, H. I. Lu, and J. M. Doyle, Chemical Reviews **112**, 4803 (2012).

[38] H. L. Bethlem, G. Berden, and G. Meijer, Phys. Rev. Lett. **83**, 1558 (1999).

[39] E. R. Hudson, J. R. Bochinksi, H. J. Lewandowski, and J. Ye, Eur. Phys. J. D **31**, 351 (2004).

[40] M. R. Tarbutt, H. L. Bethlem, J. J. Hudson, V. L. Ryabov, V. A. Ryzhov, B. E. Sauer, G. Meijer, and E. A. Hinds, Phys. Rev. Lett. **92**, 173002 (2004).

[41] N. E. Bulleid, R. J. Hendricks, E. A. Hinds, S. A. Meek, G. Meijer, A. Osterwalder, and M. R. Tarbutt, Phys. Rev. A **86**, 021404(R) (2012).

[42] M. I. Fabrikant, T. Li, N. J. Fitch, N. Farrow, J. D. Weinstein, and H. J. Lewandowski, Phys. Rev. A **90**, 033418 (2014).

[43] E. S. Shuman, J. F. Barry, and D. DeMille, Nature **467**, 820 (2010).

[44] J. F. Barry, E. S. Shuman, E. B. Norrgard, and D. DeMille, Phys. Rev. Lett. **108**, 103002 (2012).

[45] V. Zhelyazkova, A. Cournol, T. E. Wall, A. Matsushima, J. J. Hudson, E. A. Hinds, M. R. Tarbutt, and B. E. Sauer, Phys. Rev. A **89**, 053416 (2014).

[46] M. T. Hummon, M. Yeo, B. K. Stuhl, A. L. Collopy, Y. Xia, and J. Ye, Phys. Rev. Lett. **110**, 143001 (2013).

[47] J. F. Barry, D. J. McCarron, E. B. Norrgard, M. H. Steinecker, and D. DeMille, Nature **512**, 286 (2014).

[48] I. Smallman, F. Wang, T. Steimle, M. Tarbutt, and E. Hinds, Journal of Molecular Spectroscopy **300**, 3 (2014).

[49] R. J. Hendricks, D. A. Holland, S. Truppe, B. E. Sauer, and M. R. Tarbutt, Frontiers in Physics **2**, 51 (2014).

[50] B. K. Stuhl, M. T. Hummon, M. Yeo, G. Quéméner, J. L. Bohn, and J. Ye, Nature **492**, 396 (2012).

[51] S. K. Tokunaga, W. Skomorowski, P. S. Zuchowski, R. Moszynski, J. M. Hutson, E. A. Hinds, and M. R. Tarbutt, Eur. Phys. J. D **65**, 141 (2011).

[52] T. V. Tscherbul, H.-G. Yu, and A. Dalgarno, Phys. Rev. Lett. **106**, 073201 (2011).

[53] J. van Veldhoven, H. L. Bethlem, and G. Meijer, Phys. Rev. Lett. **94**, 083001 (2005).

[54] D. P. Dunseith, S. Truppe, R. J. Hendricks, B. E. Sauer, E. A. Hinds, and M. R. Tarbutt, J. Phys. B: At. Mol. Opt. Phys. **48**, 045001 (2015).

[55] B. C. Sawyer, B. K. Stuhl, D. Wang, M. Yeo, and J. Ye, Phys. Rev. Lett. **101**, 203203 (2008).

[56] D. J. McCarron, E. B. Norrgard, M. H. Steinecker, and D. DeMille (2014), arXiv:1412.8220v1 [physics.atom-ph].

[57] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, Phys. Rev. Lett. **59**, 2631 (1987).

[58] S. Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable, Phys. Rev. Lett. **57**, 314 (1986).

[59] C. Salomon, J. Dalibard, A. Aspect, H. Metcalf, and C. Cohen-Tannoudji, Phys. Rev. Lett. **59**, 1659 (1987).

[60] M. H. G. de Miranda, A. Chotia, B. Neyenhuis, D. Wang, G. Quéméner, S. Ospelkaus, J. L. Bohn, J. Ye, and D. S. Jin, Nat. Phys. **7**, 502 (2011).

[61] H. I. Lu, I. Kozyryev, B. Hemmerling, J. Piskorski, and J. M. Doyle, Phys. Rev. Lett. **112**, 113006 (2014).

[62] H. I. Lu, J. Rasmussen, M. J. Wright, D. Patterson, and J. M. Doyle, Phys. Chem. Chem. Phys. **13**, 18986 (2011).

[63] S. Chervenkov, X. Wu, J. Bayerl, A. Rohlfes, T. Gantner, M. Zeppenfeld, and G. Rempe, Phys. Rev. Lett. **112**, 013001 (2014).

[64] A. Trottier, D. Carty, and E. Wrede, Molecular Physics **109**, 725 (2011).

[65] A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, Phys. Rev. Lett. **54**, 2596 (1985).

[66] J. Kim, B. Friedrich, D. P. Katz, D. Patterson, J. D. Weinstein, R. DeCarvalho, and J. M. Doyle, Phys. Rev. Lett. **78**, 3665 (1997).

[67] J. J. Tollett, C. C. Bradley, C. A. Sackett, and R. G. Hulet, Phys. Rev. A **51**, R22 (1995).

[68] C. J. Rennick, J. Lam, W. G. Doherty, and T. P. Softley, Phys. Rev. Lett. **112**, 023002 (2014).

[69] N. C. Brahms, Ph.D. thesis, Harvard University (2008).

[70] C. I. Hancox, S. C. Doret, M. T. Hummon, L. Luo, and J. M. Doyle, Nature **431**, 281 (2004).

[71] J. M. Doyle, B. Friedrich, J. Kim, and D. Patterson, Phys. Rev. A **52**, R2515 (1995).

[72] W. C. Campbell, Ph.D. thesis, Harvard University (2008).

[73] W. C. Campbell, E. Tsikata, H. I. Lu, L. D. van Buuren, and J. M. Doyle, Phys. Rev. Lett. **98**, 213001 (2007).

[74] J. K. Messer and F. C. D. Lucia, Phys. Rev. Lett. **53(27)**, 2555 (1984).

[75] D. Patterson and J. M. Doyle, J. Chem. Phys. **126**, 154307 (2007).

[76] D. Patterson, J. Rasmussen, and J. M. Doyle, New. J. Phys. **11**, 055018 (2009).

[77] R. F. Haglund, *Laser ablation and desorption* (Academic Press, 1998), pp. 15–126.

[78] J. Wilson and J. F. B. Hawkes, *Lasers: principles and applications* (Prentice Hall International, 1987).

[79] J. M. Hutson, in *Cold molecules: Theory, Experiment, Applications*, edited by R. V. Krems, B. Friedrich, and W. C. Stalley (CRC press, 2009), pp. 3–38.

[80] W. C. Campbell and J. M. Doyle, in *Cold molecules: Theory, Experiment, Applications*, edited by R. V. Krems, B. Friedrich, and W. C. Stalley (CRC press, 2009), pp. 473–508.

[81] S. M. Skoff, R. J. Hendricks, C. D. J. Sinclair, J. J. Hudson, D. M. Segal, B. E. Sauer, E. A. Hinds, and M. R. Tarbutt, Phys. Rev. A **83**, 023418 (2011).

[82] M. J. Lu and J. D. Weinstein, New. J. Phys. **11**, 055015 (2009).

[83] A. O. Sushkov and D.Budker, Phys. Rev. A **77**, 042707 (2008).

[84] S. E. Maxwell, N. Brahms, R. deCarvalho, S. V. Nguyen, D. Patterson, J. M. Doyle, J. Helton, D. R. Glenn, J. Petricka, and D. DeMille, Phys. Rev. Lett. **95**, 173201 (2005).

[85] T. Junglen, T. Rieger, S. A. Rangwala, P. W. H. Pinkse, and G. Rempe, Eur. Phys. J. D **31**, 365 (2004).

[86] L. D. van Buuren, C. Sommer, M. Motsch, S. Pohle, M. Schenk, J. Bayerl, P. W. H. Pinkse, and G. Rempe, Phys. Rev. Lett. **102**, 033001 (2009).

[87] B. H. Bransden and C. J. Joachain, *Physics of atoms and molecules* (Longman, 1983).

[88] T. Bergeman, G. Erez, and H. J. Metcalf, Phys. Rev. A **35**, 1535 (1987).

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

[90] J. L. Bohn, Phys. Rev. A **61**, 040702 (2000).

[91] R. V. Krems, A. Dalgarno, N. Balakrishnan, and G. C. Groenenboom, Phys. Rev. A **67**, 060703 (2003).

[92] R. V. Krems and A. Dalgarno, J. Chem. Phys. **120(5)**, 2296 (2004).

[93] J. G. E. Harris, R. A. Michniak, S. V. Nguyen, N. Brahms, W. Ketterle, and J. M. Doyle, Europhys. Lett. **67**, 198 (2004).

[94] T. H. Bergeman, P. McNicholl, J. Kycia, H. Metcalf, and N. L. Balazs, J. Opt. Soc. Am. B **6**, 2249 (1989).

[95] C. V. Sukumar and D. M. Brink, Phys. Rev. A **56**, 2451 (1997).

[96] V. S. Bagnato, G. P. Lafyatis, A. G. Martin, E. L. Raab, R. N. Ahmad-Bitar, and D. E. Pritchard, Phys. Rev. Lett. **58**, 2194 (1987).

[97] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, Phys. Rev. Lett. **74**, 3352 (1995).

[98] M. Stoll, J. M. Bakker, T. C. Steimle, G. Meijer, and A. Peters, Phys. Rev. A **78**, 032707 (2008).

[99] E. Tsikata, W. C. Campbell, M. T. Hummon, H.-I. Lu, and J. M. Doyle, New J. Phys. **12**, 065028 (2010).

[100] N. Brahms, B. Newman, C. Johnson, T. Greytak, D. Kleppner, and J. Doyle, Phys. Rev. Lett. **101**, 103002 (2008).

[101] O. Chubar, P. Elleaume, and J. Chavanne, *RADIA* (1997), URL `http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/Radia`.

[102] A. Wolski (2011), arXiv:1103.0713.

[103] G. K. Woodgate, *Elementary Atomic Structure, 2nd Ed.* (Oxford University Press, 1983).

[104] URL `http://e-magnetsuk.com/`.

[105] D. W. Sedgley, A. G. Tobin, T. H. Batzer, and W. R. Call, J. Vac. Sci. Technol., A **5**, 2572 (1987).

[106] Hasted, *Physics of Atomic Collisions, 2nd edition* (Butterworths, 1972).

[107] K. P. Carney and J. M. Goldberg, Appl. Spectrosc. **41**, 308 (1987).

[108] M. Schmitt, E. A. L. Henn, J. Billy, H. Kadau, T. Maier, A. Griesmaier, and T. Pfau, Optics Letters **38**, 637 (2013).

[109] M. Lu, S. H. Youn, and B. L. Lev, Phys. Rev. Lett. **104**, 063001 (2010).

[110] E. R. Eliel, W. Hogervorst, G. J. Zaal, K. A. H. van Leeuwen, and J. Blok, J. Phys. B: Atom. Molec. Phys. **13**, 2195 (1980).

[111] W. J. Childs, Phys. Rev. A **2**, 1692 (1970).

[112] J. G. Conway and E. F. Worden, J. Opt. Soc. Am. **61**, 704 (1971).

[113] W. C. Martin, R. Zalubas, and L. Hagan, in *Nat. Stand. Ref. Data Ser., NSRDS-NBS 60* (Nat. Bur. Stand., U.S., 1978).

[114] G. Sanna and G. Tomassetti, *Introduction to molecular beams gas dynamics* (Imperial College Press, 2005).

[115] C. I. Hancox, Ph.D. thesis, Harvard University (2005).

[116] S. M. Skoff, Ph.D. thesis, Imperial College London (2011).

[117] K. S. W. Sing, in *Adsorption by Powders and Porous Solids: Principles, Methodology and Applications, 2nd Edition*, edited by F. Rouquerol, J. Rouquerol, K. S. W. Sing, P. Llewellyn, and G. Maurin (Academic Press, 2014), pp. 321–392.

[118] N. R. Hutzler, M. F. Parsons, Y. V. Gurevich, P. W. Hess, E. Petrik, B. Spaun, A. C. Vutha, and D. DeMille, Phys. Chem. Chem. Phys. **13**, 18976 (2011).

[119] N. E. Bulleid, S. M. Skoff, R. J. Hendricks, B. E. Sauer, E. A. Hinds, and M. R. Tarbutt, Phys. Chem. Chem. Phys. **15**, 12299 (2013).

[120] N. E. Bulleid, Ph.D. thesis, Imperial College London (2013).

[121] D. R. Miller, in *Atomic and Molecular Beam Methods*, edited by G. Scoles, D. Bassi, U. Buck, and D. Lainé (Oxford University Press, 1988), p. 14.

[122] M. G. Dondi, G. Scoles, F. Torello, and H. Pauly, J. Chem. Phys. **51**, 392 (1969).

[123] F. Pobell, *Matter and Methods at Low Temperatures, 3rd Ed.* (Springer, 2007).