UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING

Physics

Novel techniques for the trapping and manipulation of ultracold atoms

by

Nathan L. Cooper

Thesis for the degree of Doctor of Philosophy

May 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
Physics

Doctor of Philosophy

NOVEL TECHNIQUES FOR THE TRAPPING AND MANIPULATION OF ULTRACOLD ATOMS

by Nathan L. Cooper

We describe several novel techniques for the optical trapping of ultracold atoms and for the production of wavelength-stabilised, coherent light at the frequencies required for use in atomic physics experiments.

The greater part of this thesis deals with work towards the creation of regular arrays of microscopic optical dipole traps formed at the foci of truncated spherical cavities in a metallic film, in which the inter-site spacing can be set anywhere between one and several hundred micrometers. Arrays of such cavities are synthesised and structurally characterised via optical and electron microscopy, and numerical simulations of the light intensity distribution near the foci of such cavities under normal illumination are used to confirm their suitability for dipole trap production. A method for the construction of arrays of magneto-optical traps based on such structures is proposed and theoretically examined, and some preliminary experimental work towards the synthesis of the required microstructures is also described. Possible approaches to the loading of such traps and the imaging of the atoms contained therein are discussed — experimental work towards ballistic atom transfer from a specialised form of magneto-optical trap that can be formed close to a microstructured surface is carried out, and the efficacy of wavelength-selective fluorescence imaging as a means of reducing the effects of background scatter from the surface is experimentally demonstrated.

Further work described herein includes the proposal and experimental demonstration of two novel techniques for the removal of the carrier wave from a phase-modulated laser beam, one of which is based on a fiber-optic Mach–Zehnder interferometer that is shown to be an effective device for splitting or combining beams of nearly equal frequencies. A spontaneous-force based atom trapping mechanism that does not rely on the use of a magnetic field, but rather on spatially-dependent optical pumping between different metastable atomic states, is also proposed, and a proof of principle experiment is carried out to demonstrate the validity of the suggested mechanism. We find that this trapping scheme allows the spatial dependence of the trapping force to be tailored with a greater degree of flexibility than is usually possible with magneto-optical trapping, and that it is also capable of producing traps with stronger spring constants than are typically achievable with magneto-optical trapping under realistic experimental constraints.

Contents

List of Figures

Declaration of Authorship

I, Nathan L. Cooper , declare that the thesis entitled *Novel techniques for the trapping and manipulation of ultracold atoms* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as: [1],[2],[3]

Signed:..

Date:..

Acknowledgements

To begin on a formal note, some of the experimental work described in this thesis was carried out jointly with other researchers, and I should therefore make clear where this is the case. During the microfabrication work described in chapter 3, I was advised and guided by Dr Hamid Ohadi, who also carried out a small amount of the experimental work, although this was mainly left to me. Hamid is also responsible for setting up many of the experimental systems in our lab, which I have made use of throughout the work described in this thesis. Jonathan Woods carried out more than half of the experimental work relating to §8.2.2.3, and is also responsible for producing figures 8.9, 8.10 and 8.11. Several sections of this thesis also refer to or describe in detail the work reported in [4], and it should be made clear that this work was not carried out by the author.

There are many people to whom I am grateful for help or encouragement, but those who most deserve mention are my research group. My supervisor, Tim Freegarde, and postdocs James Bateman, Hamid Ohadi and Mathew Himsworth, have all been very generous with their time and have provided me with invaluable advice and guidance over the course of my PhD. Andre Xuereb provided some useful suggestions and brought my attention to a range of relevant papers that I might otherwise have missed and, with the exceptions of the occasions when he would come into the lab and misalign my optics, was generally a helpful person to have around. Alex Dunning, Joseph Rushton, Alan Forrester, Rachel Gregory and Mathew Aldous have all been fun to work with and generally quite good about letting me borrow their equipment. I should also both thank Jonathan Woods for his work on the fiber-optic MZI, and apologise to him for inflicting such a device on a new and unsuspecting PhD student.

Finally, I would like to thank Sophie, who has put up with me being tired and irritable during the last few weeks of writing and, despite submitting her thesis considerably before me, has only engaged in a bare minimum of gloating during the intervening period.

Chapter 1

Introduction

Cold, trapped atoms have already found several direct applications — primarily in the field of metrology. Cold atom systems now form the world's most accurate clocks [5], and are showing great promise as devices for measuring acceleration [6] as well as magnetic and gravitational fields [7, 8]. However, the importance of cold atom systems goes beyond such existing, direct applications. They have been widely used to facilitate other areas of research, such as precision measurement of fundamental constants [9, 10, 11] and the experimental verification of quantum mechanical predictions via experiments involving Bose-Einstein condensates (which are formed by trapping and cooling bosonic atoms) [12, 13]. Importantly, cold atom systems are also suitable candidates for the experimental realisation of quantum simulation or quantum information processing [14, 15].

However, all of these exciting and useful possibilities depend on our ability to trap, cool and manipulate a range of atomic species. The aim of the work described in this thesis is to enhance the toolkit of experimental techniques used to cool, trap and manipulate atoms, in order to facilitate experiments that would be either difficult or impossible using only existing techniques. To this end, we have carried out a theoretical investigation of and substantial experimental work towards the creation of arrays of microscopic optical dipole and magneto-optical traps for neutral atoms, based on the use of microscopic truncated spherical cavities in a metallic surface. Techniques for loading atoms into such microtraps and imaging them once trapped, on which research was already in progress, have been developed further during the course of this work. In addition, we propose a novel atom trapping scheme based on optical pumping between metastable (or fully stable) atomic states and conduct a proof of principle experiment to demonstrate the validity of the technique. It is anticipated that this approach to atom trapping may prove a valuable alternative to magneto-optical trapping in certain applications, and could be particularly useful as a loading mechanism for our microscopic dipole trap arrays. Finally, we have studied light generation techniques for atomic physics experiments and experimentally demonstrated two novel approaches to the separation of selected frequency components from phase-modulated laser beams, a function that would most likely be required in any potential quantum information experiment based on our microscopic dipole trap arrays.

In this chapter we give a brief overview of how cold atom systems are employed in each of the major applications described above, followed by a discussion of which existing trapping, cooling and manipulation techniques represent the most relevant enabling technologies. We then explain how each of the elements of the research described in this thesis differs from existing methodology, and in what ways this might impact on research involving related systems.

1.1 Applications of cold, trapped atoms

Here we give a brief overview of the most common applications of cold, trapped atoms across three broad areas: metrology, quantum information and tests of fundamental physics. For each set of related applications, we consider the most relevant trapping or manipulation techniques for cold atoms and discuss how these applications could be facilitated by changes in the properties of such techniques.

1.1.1 Metrology

Cold atoms have several distinct applications within the area of metrology — they can be used for timekeeping, to measure gravitational fields and acceleration, or to measure electric or magnetic fields. As these rely on somewhat different techniques, we shall consider each separately.

Timekeeping based on cold atoms usually involves the stabilisation of an external oscillator to the frequency of an atomic transition. This has taken several forms, with the most accurate approach currently being the 'optical lattice clock' [5, 16]. There is even research being conducted into a nuclear clock based on ^{229}Th [17], which may have enhanced stability due to the reduced effect of many unwanted external influences on the nucleus, owing to its small size and high degree of shielding by the atom's electrons. One feature that all of these clocks have in common is that the accuracy of the clock is improved if either the number of atoms used or the interrogation time of those atoms is increased. Such techniques are therefore dependent on our ability to create a high density of cold atoms, for which spontaneous-force based traps such as the magneto-optical trap are essential. Clocks such as optical lattice clocks, in which the atoms remain trapped during interrogation, will also depend on the length of time for which those atoms can be trapped without the trapping process significantly influencing their internal state — in the case of dipole trapping this time will be determined by the rate at which the atoms scatter light from the trapping beams.

Most field or acceleration measurements with cold atoms rely on atom interferometry — see for example [6, 7, 8, 18, 19, 20]. Here the number and temperature of the atoms can both be important, so once again cooling and trapping techniques are vital enabling tools for such work. Indeed, some results have shown that use of a Bose-Einstein condensate for atom interferometry can improve the performance of such devices [19], making ground state cooling of trapped atoms a worthwhile enterprise. Furthermore, many of these experiments (such as [7, 20]) rely on velocity sensitive Raman transitions, and the ability to generate the appropriate optical fields to drive these transitions is therefore also important. Finally, in order to put these measurement devices into use beyond the laboratory, research on miniaturization and enhanced portability is being undertaken [18, 21]. As such, any technique that allows existing technology to be reduced in size or complexity, or decreases its susceptibility to unwanted influences from the local environment, is of considerable interest.

1.1.2 Quantum information and simulation

The criteria that must be satisfied to make a system suitable for the practical implementation of a generalised quantum information processor were established by Vincenzo [22]. They include scalability of the system to large numbers of qubits, the ability to set all the qubits to some known initial state, a decoherence timescale that greatly exceeds the time required for qubit operations, the ability to implement a universal set of qubit operations and the ability to individually address specific qubits for both gate operations and measurement. Systems employing cold, trapped, neutral atoms can potentially possess all of these attributes.

Scalability is inherent in cold atom based systems, as very large numbers of cold atoms can be simultaneously trapped via existing techniques, and, due to the neutrality of the atoms, scaling is not impeded by long range interactions in the same way as it is for systems based on trapped ions. Initialisation of atomic qubits is also straightforward, as many atomic species have electronic transitions with large (MHz and above) spontaneous decay rates and sufficiently large transition energies that the population of the excited state is negligible even at room temperature. This allows optical pumping to be used to set all atoms into a known initial state.

Decoherence can be extremely slow in cold atom systems. For cold atoms suspended (in a suitable optical or magnetic trap) in ultrahigh vacuum, the main sources of decoherence are spontaneous decay and decoherence induced by the trapping process. These factors are both dependent on the trapping and manipulation techniques used to control the atoms. Spontaneous decay rates can be extremely low if appropriate energy levels are chosen, such as the hyperfine ‘ground’ states of an alkali metal atom, for which there is no permitted electric dipole transition between the upper and lower states. However, in order to drive coherent transitions between such states, one typically requires a method for producing two closely-spaced optical frequency components with a known and highly stable frequency difference between them — see §1.2.3. Trap-induced decoherence in an optical dipole trap can be reduced by increasing the detuning of the trapping laser (see §2.5), a change which generally necessitates a corresponding increase in the optical power applied to form the trap. Therefore, in real experimental situations in which unlimited optical power is not available, any technique that increases the efficiency with which that optical power can be applied to the formation of dipole traps has the potential to reduce trap-induced decoherence. The neutrality of the atoms is also beneficial with regard to decoherence — coupling of the atoms to any stray electric or magnetic fields is of much less significance than in ion based systems.

Construction of a universal set of quantum gates is not entirely trivial with cold atoms. Single-qubit gates can easily be implemented via coherent laser pulses, but multi-qubit operations are not so straightforward. However, it has recently been shown that two-qubit operations can be performed via Rydberg blockade [23], and by combining suitable two- and one-qubit gates it is then (in principle) possible to carry out any multi-qubit operation [24].

Addressing individual qubits is also something that has proved experimentally challenging with cold atoms, as the preferred mechanism for generating an array of trapped, cold atoms is the optical lattice [25], which inherently produces an inter-site spacing that is of the order of one half of the wavelength of the trapping light. This small inter-site spacing has made it difficult to resolve individual atoms. However, single-site resolution with a high degree of fidelity has recently been demonstrated in an optical lattice [26], and arrays of atom traps based on microlens arrays (in which the site spacing is larger and each site is easily addressed via its own dedicated microlens) have also been demonstrated [27]. Ease of single-site addressability remains an important criterion to consider when designing a candidate quantum information system.

1.1.3 Tests of fundamental physics

We have already mentioned that cold atom systems have been used for sensitive measurements of fundamental constants [9, 10, 11], as well as various other tests of fundamental physics [12, 13]. In many cases these rely on atom interferometry, and are thus dependent on our ability to produce a dense cloud of cold atoms on demand. Some, such as [28], employ cold atoms confined in an optical lattice, and are therefore also dependent on the properties of the lattice and its ability to contain the atoms without excessive scattering of the trap light.

There is also interest in carrying out some of this work in space-based experiments [29], owing to the improved performance of some of these experiments in low gravity and/or low noise environments. This would require significant miniaturization and simplification of existing experimental equipment, as well as (in some cases) a reduction in power requirements for the experiment. This is therefore another reason why improvements in these areas are relevant to the applicability of cold atom systems.

It has even been suggested that some atom interferometry experiments constitute a sensitive measurement of gravitational redshift that can be used to test the equivalence principle [30], although this has been heavily disputed [31, 32].

1.2 The contents of this thesis

Here we give an overview of the work described in this thesis, and discuss its relevance to the various applications considered above.

1.2.1 Production of microscopic atom trap arrays based on truncated spherical cavities in a reflective surface

Work towards the creation of arrays of both dipole [33] and magneto-optical traps [34] based on truncated spherical cavities in a metallic film is described, with the proposed magneto-optical trap arrays not only being of interest in their own right but also representing a convenient mechanism by which to facilitate loading of the dipole trap arrays. As will be discussed in more detail below, there has been much previous work on the production of arrays of microscopic dipole traps [25, 26, 27, 35, 36, 37] and some work towards the creation of arrays of microscopic magneto-optical traps [38, 39]. However, the author is not aware of any previous work based on the use of arrays of microscopic concave reflectors, and we will go on to describe several key ways in which the properties of dipole trap arrays produced via this approach could be expected to differ from those of previously demonstrated dipole trap arrays. Previous attempts to create microscopic magneto-optical traps have so far not succeeded in producing any magneto-optical trap significantly below a millimetre in size [38, 40], and so are not comparable to the aims of the work described herein. Microscopic magnetic trap arrays have been studied previously and do present a viable alternative to microscopic dipole trap arrays [41, 42]. However, magnetic traps are inherently dependent on the spin state of the trapped particle, and so are not appropriate for a wide range of applications in which the spin state of the trapped particles must be modified while the particle remains trapped. Furthermore magnetic trap arrays, like most existing approaches to the creation of dipole trap arrays, also lack some of the advantages specific to an implementation based on arrays of concave reflectors, which will be discussed in more detail below.

By employing a templated electrodeposition process [43], regular arrays of truncated spherical cavities can be created in a metallic film with inter-site spacing equal to twice the radius of curvature of the cavities, which can be set anywhere between hundreds of micrometres and a few tens of nanometres. In our experimental prototype, a radius of curvature of 50 μm is used. Each such cavity can be used as a shaped mirror to produce either a region of high light intensity suitable for optical dipole trapping, or (with some modifications described in chapter 3) a set of illuminating beams appropriate for magneto-optical trapping about a point on the cavity axis.

Previous demonstrations of arrays of optical dipole traps have been based on holography using spatial light modulators (SLMs) [35, 36], on microlens arrays [27] and on optical lattices [25, 26]. All of these techniques have been demonstrated experimentally and shown to be effective methods for creating arrays of trapped atoms. However, they each come with some associated disadvantages: in optical lattices the inter-site spacing in the array is compelled to be of a similar scale to the wavelength of the light used to generate the lattice, while both optical lattices and holographic trap arrays make rather inefficient use of the available optical power. In the case of optical lattices this is due to the absence of any focusing of the input light, and in holographic traps both the imperfect reflectivity (quoted as 90% in [35]) and the diffraction efficiency of the SLM cause substantial loss of optical power. Meanwhile, systems based on microlens arrays operate on a similar principle to those we aim to produce but with a lower numerical aperture optic addressing each trap site (see below), and are therefore likely to produce weaker traps (for a given laser power) and allow less effective collection of light emitted by the trapped atoms for imaging purposes.

Our approach would allow the creation of an array of microscopic atom traps with a flexible inter-site spacing, that is not tied to the wavelength scale as it is in optical lattices. This could make it easier to selectively address individual sites within the array, and is an important advantage of our technique over dipole trap arrays based on optical lattices. The very tight focusing provided by such cavities also reduces the laser power requirements of such systems (when compared to optical lattices and holographic or microlens-array based systems) and provides the option of using collisional blockade mechanisms (such as that described in [44]) to prevent multiple occupation of a single trap site. These features make this a promising system for quantum simulation or computing applications, as they provide a straightforward approach to generating an array of single trapped atoms that are individually addressable and can be held at a suitable distance from their nearest neighbours for the Rydberg blockade effect to be used to enact two-qubit operations. An additional benefit of the tight focusing capabilities of such structures is that they will allow strong coupling of the trapped atoms to incident light to be achieved without the need to place the atoms into a conventional optical cavity, which would be of use in quantum information experiments [45].

There are also further ways in which microcavity arrays of this kind make attractive experimental systems: each trap site has its own dedicated, large solid angle optic for imaging or optical control — for example, the solid angle taken up by the trapping cavity (from the position of the dipole trap) in our prototype experimental array is approximately $3\pi/4$ steradians, while the solid angle taken up by the microlenses used in [27] is only $\pi/100$ steradians. It is acknowledged in [27] that the comparatively low NA of the microlenses limits the minimum spot size they are able to achieve at each dipole trap site to > 1.3 μm (indeed, the spot size actually achieved at each trap site in [27] was rather larger than this theoretical limit at 3.7 μm) — a limitation not encountered in our prototype system. Furthermore, modification of the properties of our truncated spherical cavities may allow the interactions of the trapped atoms with the vacuum electromagnetic field to be tailored to make them more suited to a particular experiment [46], and the proximity of the atom traps to a structured metallic surface opens up the possibility of using surface-plasmon resonances to enhance the trapping potential [47]. It should be noted that the fact that an enhancement of the trapping potential (via either plasmon resonances or tight focusing of the trapping light) could just as easily be achieved by increasing the total laser power does not render these benefits insignificant. Many of the applications described above employ cold atoms that are confined in an optical lattice or other dipole trap array, and the efficacy of these cold-atom devices is therefore dependent on an ability to trap cold atoms for as long a time as possible with as little scattering of the trapping light as can reasonably be achieved. The limiting factor on how favourable these properties can be made experimentally is often the total laser power available at the appropriate wavelength, and any system that increases the efficiency with which the available laser power can be applied to atom trapping is therefore worth investigating.

The primary disadvantage of our scheme for creating an array of dipole traps is that the close proximity of the dipole trap array to a microstructured surface makes trap loading extremely difficult. While a number of viable schemes by which this could be achieved are presented in this thesis, we are yet to produce an experimental demonstration of any sufficiently effective technique, and at present trap loading remains a major obstacle to the use of such dipole trap arrays. However, in the event of this difficulty being successfully resolved, the experimental complexity of our trapping scheme would not be dissimilar from that of the various techniques described above. While our approach requires advanced loading techniques, optical lattices are produced by the interference of counter-propagating laser beams and so require wavelength-scale path length stabilisation, while holographic intensity-pattern generation using spatial light modulators is non-trivial and requires the use of costly equipment and numerically intensive feedback algorithms [48, 49]. The use of microlens arrays is perhaps the least experimentally challenging approach, as such arrays are now commercially available[1], but this technique offers neither the dynamic configurability of systems based on optical lattices or holography nor the extremely tight focusing of the trapping light and efficient collection of fluorescence photons provided by our approach.

---

[1] See for example Thorlabs product MLA150.

At the time of writing, the author is not yet aware of any convincing experimental demonstration of an array of microscopic magneto-optical traps. While efforts have been made to miniaturise such traps [40, 38, 39], the length scales involved have yet to drop significantly below 1 mm. In addition, the only technique yet demonstrated experimentally that could in principle be scaled to a 2D array involves traps in which the trapped atoms must necessarily be recessed into a surface [38, 39], and employing such a system would therefore severely restrict optical access to the trapped atoms and rule out any study or utilisation of the interactions between atoms in adjacent trap sites.

We also investigate methods by which cold atoms could be loaded into such microtraps, with the primary approach taken in this thesis being the ballistic transfer of atoms from an atom cloud formed in a magneto-optical trap (MOT) centered a few mm below a (downward-facing) microcavity array. An unconventional beam geometry is employed in the MOT used to trap atoms in this location, which has been dubbed the 'Λ-MOT' by its inventors [4]. We undertook significant work towards the improvement of this prototype device, and a characterisation of the enhanced Λ-MOT resulting from this work is given in chapter 5. A launching process, by which atoms can be propelled from the position of the MOT towards a nearby microstructured surface, was also investigated during the course of this work, along with techniques to allow atoms close to such a surface to be imaged optically. These loading and imaging methods are not only a necessary prerequisite for the demonstration of the atom trap arrays described above, but are also of some interest in their own right, as there is much topical research involving the interactions between atoms and microstructured surfaces [47, 50, 51].

1.2.2 Dissipative atom trap based on optical pumping between metastable states

We propose and demonstrate a novel spontaneous-force based atom trapping scheme that is not only of interest in its own right but could also prove to be a highly suitable system for the loading of our microscopic dipole trap arrays. Almost all experimental work with ultracold atoms has involved the use of a magneto-optical trap [34] to produce the dense samples of ultracold atoms required (experiments involving atoms in optical dipole or magnetic traps generally use a MOT to facilitate loading). The reason for this is that traps based on a purely conservative potential (such as optical dipole or magnetic traps) have no associated cooling mechanisms and tend to have much smaller trap depths than a MOT. A few alternatives to the MOT have been proposed [52, 53, 54], with the intention of increasing the versatility of spontaneous-force traps by allowing them to be created in the absence of a magnetic field [53], or permitting the captured atoms to be spin-polarised while remaining trapped [54], but none of these have been widely used by experimentalists and those demonstrated experimentally have typically captured fewer atoms than equivalent magneto-optical traps — see [53] for example.

Increasing the maximum number and/or density of trapped atoms that can be produced in spontaneous-force based atom traps, or allowing the distribution of the atoms within such traps to be adjusted more flexibly, would both aid the loading efficiency of an array of microscopic dipole traps, as well as being of significant use for many of the other applications mentioned in §1.1. Chapter 7 describes a mechanism capable of doing both — trapping via optical pumping between metastable states [3]. Along with a number of further advantages that are discussed in more detail in chapter 7, this trapping mechanism is not only able to operate in the absence of a magnetic field, but because it does not depend on the trapping light being in any particular polarisation state it can also operate normally in the presence of unwanted magnetic fields, in contrast to both [53] and [54]. This makes it more readily compatible with miniaturisation efforts involving surface-based magnetic trapping and/or manipulation [55], as it can be directly overlaid onto a region containing magnetic traps without significant disruption to the trapping process, making it an ideal atom-loading solution for magnetic microtrap arrays as well as those based on dipole trapping. Furthermore, it means that trapping based on optical pumping between metastable states can be applied in combination with magneto-optical trapping, thus allowing the form and magnitude of the trapping force to be modified independently of parameters that also affect cooling and loading rates, such as the laser detuning and magnetic field gradient. This could allow for a substantial increase in the density of atoms that can be generated around the trap centre — especially as, under any realistic set of experimental constraints, traps based on optical pumping between metastable states can produce substantially larger spring constants than magneto-optical traps.

1.2.3 Light generation for atomic physics experiments

Almost all of the applications described above require, at some point, that the atoms involved be addressed optically via near-resonant laser light. As the natural linewidths of atomic transitions are often very narrow — of the order of one part in a billion of the total transition frequency — wavelength-stabilisation of the necessary laser sources can be both costly and time-consuming. Techniques such as electro-optic modulation or laser current modulation, which generate two or more closely spaced wavelengths of illumination from a single laser source, are therefore extremely useful as they allow many laser sources to be replaced by just one. Furthermore, many experiments employ two-photon transitions between initial and final states that are very close to each other in energy (in particular, the potential quantum information applications of the microscopic dipole trap arrays described in this thesis are strongly dependent on an ability to address such transitions, as the proposed qubit states of the trapped atoms would be the two hyperfine 'ground' states of Rb atoms). In such cases, the experiment is very sensitive to the frequency difference between the two relevant wavelength components, but not to their absolute frequencies. The aforementioned modulation techniques are therefore of particular importance in such experiments, as they allow the frequency difference between the resulting wavelength components to be set with a high degree of accuracy and stability.

Both of the above matters are addressed in this thesis. Chapter 8 describes a simple and robust method by which the unmodified wavelength component (alternatively known as the 'carrier') can be removed from a phase-modulated laser beam [1], thus increasing the utility of the modulation techniques mentioned above by allowing this one component to be separated from the other, co-propagating frequency components in the resulting beam — a function that is often required in atomic physics experiments, such as [56], for example. Chapter 8 also describes how the same end can be achieved via the construction and stabilisation of a fibre-optic Mach-Zehnder interferometer [2], which allows the splitting or combining of nearly equal optical frequency components, and is therefore more versatile than the polarisation based technique described earlier in this paragraph as its use is not restricted to carrier removal. Finally, §8.3 describes preliminary investigations of the use of an acousto-optic modulator (AOM) in a ring cavity containing a gain medium to create an optical 'frequency comb' (a term used to refer to a device that generates a beam containing a large number of uniformly spaced optical frequency components) with a tooth spacing of 80 MHz and a span of several GHz, which has potential applications in allowing many slave lasers to be stabilised with well defined frequency differences relative to a single, highly stable master laser.

Chapter 2

Techniques for cooling and trapping atoms in vacuum

In this chapter we describe some of the standard techniques used to trap and cool atoms, as well as discussing some more sophisticated enhancements of these standard techniques that have been demonstrated and giving an indication of the limits of the performance of existing approaches to atom trapping. Though there are a great number of techniques that are frequently employed within the field of atomic physics, there are four specific techniques that are so widely used and of such central relevance to this thesis that they warrant individual, in-depth discussion: Doppler cooling, magneto-optical trapping, magnetic trapping and optical dipole trapping. With the exception of magnetic trapping, all of these techniques depend strongly on the interaction of atoms with laser radiation.

2.1 The interaction of atoms with laser radiation

Continuous wave lasers provide a source of coherent radiation with a very narrow frequency distribution — sufficiently narrow that in most cases laser radiation can be treated as though it were completely monochromatic. We shall therefore proceed by considering the effect of light of a single frequency on an atom. Furthermore, we shall restrict our discussion to the atomic species that are most frequently cooled and trapped via these techniques: the alkali metals. In such atoms, only a single electron interacts significantly with the applied light.

2.1.1 Scattering

By considering the effect of an oscillating electric field as a perturbation to a two-level atom, it can be shown (for example in [57]) that the scattering rate (i.e. the rate at which the atom absorbs incident photons and spontaneously re-emits them) from an incident beam is given by:

$$\mathrm{R}_{\mathrm{scatt}} = \frac{\Gamma\Omega^2}{\delta^2 + \Omega^2/2 + \Gamma^2/4} \tag{2.1}$$

where Γ is the spontaneous decay rate of the upper state (about $2\pi \times 6$ MHz for the D2 line of ^{85}Rb), δ is the angular frequency difference between the atomic transition and the laser light (and consequently should be negative if cooling is desired) and Ω is the Rabi frequency. The Rabi frequency is itself given by:

$$\Omega = \frac{\langle 1|U(\mathbf{r})|2\rangle}{\hbar}. \tag{2.2}$$

Here, $\mathbf{r}$ is a vectorial position relative to the atomic nucleus and $U(\mathbf{r})$ is the potential energy of an electron with position vector $\mathbf{r}$ in the electric field of the applied light, with $|1\rangle$ and $|2\rangle$ being the wavefunctions of the upper and lower energy states of the atom's outermost electron (for that particular transition). We do not include the time-dependent phase factor in the electric field of the light as this has already been factored out of the expression for the Rabi frequency (see [57]). Defining $\mathbf{E}(\mathbf{r})$ as the electric field of the incident radiation (without the time-dependent phase factor), $U(\mathbf{r})$ is then given by:

$$U(\mathbf{r}) = \int_{\mathbf{0}}^{\mathbf{r}} \mathbf{E}(\mathbf{r}') \cdot \mathrm{d}\mathbf{r}', \tag{2.3}$$

where the integral can be taken along any path between the two points, and $\mathrm{d}\mathbf{r}'$ is an element of that path. We are able to define $U(\mathbf{r})$ in this way as the orthogonality of $|1\rangle$ and $|2\rangle$ means that the addition of a position-independent scalar to $U(\mathbf{r})$ leaves the value of $\langle 1|U(\mathbf{r})|2\rangle$ unchanged, and the zero of potential energy can therefore be defined to be at any convenient position. Theoretical calculation of the Rabi frequency is non-trivial for any atoms other than hydrogen, as it requires a knowledge of the electronic wavefunctions $|1\rangle$ and $|2\rangle$ for the lower and upper energy levels involved in the transition. However, both the lifetime of the upper state against spontaneous decay and the scattering rate of an illuminated atom can be measured experimentally. Using equation (2.1) and the fact that the spontaneous decay rate Γ is the reciprocal of the upper state's lifetime, it is then possible to determine the values of Ω and Γ. Furthermore, because the wavelength of the radiation being used is almost always several orders of magnitude greater than the length scales associated with the electronic wavefunctions, the electric field of the light can be accurately approximated as being independent of position in equation (2.2). This is usually referred to as the electric dipole approximation, and it allows us to take the electric field $\mathbf{E}$ to be independent of position and depend only on the polarisation state of the incoming light, and we therefore obtain:

$$\Omega = \frac{e\langle 1|\mathbf{E}\cdot\mathbf{r}|2\rangle}{\hbar}. \tag{2.4}$$

The values of $\langle 1|\mathbf{E}\cdot\mathbf{r}|2\rangle$ for particular transitions can hence be derived from experimental results and can then be used in conjunction with (2.1) to calculate the scattering rate for an atom illuminated with a known intensity and wavelength of radiation. For transitions commonly used to trap or cool atomic species, these values (known as the 'dipole matrix elements' for the relevant transitions) have been accurately calculated from experimental data and tabulated in documents such as [58].

2.1.1.1 The scattering force

When a photon is absorbed by an object the momentum carried by the photon ($\hbar\mathbf{k}$, where $\mathbf{k}$ is the photon's wavevector) is transferred to that object. If that object spontaneously re-emits at a later time, the resulting photon will propagate in a random direction and the (vector) mean momentum transferred to the object by the emission event will be $\mathbf{0}$. The mean net momentum transferred to the object during the whole process is therefore equal to the momentum carried by the initial photon. If a unidirectional beam of photons (propagating along the z axis, for argument's sake) is incident on an object that absorbs and spontaneously re-emits them at a rate of n per second, that object will therefore be subject to a force equal to $n\hbar k\hat{\mathbf{z}}$. This force (and its modification via the Doppler and Zeeman shifts) is the key element of both Doppler cooling and magneto-optical trapping.

2.1.1.2 Scattering in depth: multi-level atoms

The description of scattering given above is somewhat incomplete, as it is only sufficient to describe the behaviour of a two-level atom. For a full description of multi-level atoms, one must begin by considering the expansion of the atom's wavefunction $|\psi\rangle$ in the stationary states of the unperturbed atom:

$$|\psi\rangle = \sum_k c_k |\phi_k\rangle e^{-i\omega_k t}, \tag{2.5}$$

where ω_k is equal to the energy of state $|k\rangle$ divided by $\hbar$. By substituting this into the time-dependent Schrödinger equation and adding a perturbation term $\hat{H}'$ to the Hamiltonian to account for the interaction of the atom's electron with the electric field of the applied light, then left multiplying by $\langle\phi_j|$, it is possible to derive the rate of change of each of the coefficients c_j that determines the admixture of state $|j\rangle$ of the unperturbed atom in the illuminated atom's wavefunction. The result is the set of coupled differential equations given by:

$$i\hbar\frac{\mathrm{d}}{\mathrm{d}t}c_j = \sum_k c_k \langle j|\hat{H}'|k\rangle e^{i(\omega_j-\omega_k)t}. \tag{2.6}$$

The probability of the atom being found to be in state $|j\rangle$ when measured is then equal to $|c_j|^2$. This formula is sufficiently general to deal with the majority of real, multi-level atoms. Unfortunately it is not usually possible to solve these equations analytically without making some simplifying approximations (with an appropriate choice of approximations being dependent on the properties of the system under consideration), but they do allow the response of a multi-level atom to radiation to be computed numerically without further assumptions. A clear derivation of this result, upon which the above is based, is given in the first few pages of [59].

The above does not include a description of spontaneous emission, but this can be taken account of by including additional decay terms for each state — see [59], for example. It should also be borne in mind that results very similar to those obtained with this approach can often be reached via simple, rate-equation based models that ignore atomic coherences, and are much easier to solve for the relevant atomic populations [60].

2.1.2 The optical dipole force

The physical origin of the dipole force is the polarisation of an object by the electric field of the light. Once polarised, this object then has an electric dipole moment, which interacts with the electric field of the light, resulting in a force being exerted on the object. Because the only requirement here is that the object is polarisable, almost any microscopic particle can in principle be trapped and manipulated via the dipole force.

The magnitude and direction of the dipole force on an atom is a standard result, and derivations can be found in many texts on the subject — see for example [57, 33]. The most common approach is to consider only a single transition and express the electronic wavefunction as a superposition of the upper and lower states of the transition. Adding a perturbation term to the atomic Hamiltonian to account for the interaction between the outermost electron and the applied electric field, then substituting the modified wavefunction and Hamiltonian into the time dependent Schrödinger equation allows the derivation of the polarisability of the atom as a function of frequency. It is then straightforward to calculate the force generated by the light-atom interaction and hence derive the form of the dipole force, $\mathbf{F}_{dip}$:

$$\mathbf{F}_{dip} = -\frac{\hbar\delta\Omega}{2\delta^2 + \Omega^2 + \Gamma^2/2}\nabla\Omega, \tag{2.7}$$

where the symbols have the same meanings as in §2.1.1. In the limit of $\delta \gg$ both Γ and Ω, this expression becomes

$$\mathbf{F}_{dip} = -\frac{\hbar\Omega}{2\delta}\nabla\Omega, \tag{2.8}$$

and the dipole force can therefore be defined as the gradient of a dipole potential

$$U_{\mathrm{dipole}} = \frac{\hbar\Omega^2}{4\delta}. \tag{2.9}$$

Using (2.2), it can be seen that Ω^2 is proportional to the light intensity. The dipole potential is also therefore directly proportional to the light intensity, with the sign of the potential being given by the sign of the detuning. Light detuned slightly below an atomic resonance frequency (referred to as 'red detuned') attracts atoms to regions of high light intensity, while that tuned slightly above a resonance ('blue detuned') has the opposite effect.

2.2 Doppler cooling

Doppler cooling is a technique that allows atoms (and more recently certain other species, such as dielectric spheres [61] or simple molecules [62]) to be cooled from temperatures $\sim$300 K and above down to the level of a few tens of μK. It relies on the radiation pressure force resulting from the absorption and subsequent re-emission of radiation.

In order to use the scattering force to cool the thermal motion of a set of objects, it must first be brought to depend upon the velocity of the object it acts upon. In the case of an object for which the rate of absorption and re-emission is dependent upon the wavelength of the radiation, this can be done via the Doppler shift. For an object moving against (with) the direction of propagation of the light, the frequency of the electromagnetic field oscillations will be increased (decreased) by v/λ, where v is the speed of the object. In the non-relativistic limit, Doppler shifts caused by motion transverse to the direction of light propagation can be ignored, as these are second order in v/c.

Now consider an object with a narrow absorption resonance (such as an atom) that is illuminated by laser beams of equal intensity propagating in the positive and negative z directions. If the frequency of these laser beams is slightly below that of the absorption resonance, then a positive Doppler frequency shift will result in an increased scattering rate from the relevant beam, while a negative frequency shift will have the opposite effect. This will cause the object to scatter photons more rapidly from the beam that is propagating against its direction of travel (and vice versa), thus resulting in a net force that opposes the motion of the object — a slowing force. The resulting reduction in the velocity of individual particles corresponds to a decrease in temperature. By illuminating a group of objects with three such pairs of laser beams that are not coplanar, it is possible to cool all three translational degrees of freedom simultaneously.

We can derive the magnitude of this slowing force by considering the portion of the total scattering force on an atom that results from its motion. This can be done by substituting a Doppler-shifted laser frequency into (2.1), then subtracting from this the scattering rate for a stationary atom and multiplying the whole expression by the photon momentum:

$$\text{Slowing Force} = \hbar\mathbf{k}\left(\frac{\Gamma I/2I_{sat}}{1+I/I_{sat}+4(\delta-2\pi v/\lambda)^2/\Gamma^2}-\frac{\Gamma I/2I_{sat}}{1+I/I_{sat}+4\delta^2/\Gamma^2}\right), \quad (2.10)$$

where v is the component of the atom's velocity along the direction of propagation of the incident light. In practice two balanced beams propagating in opposite directions would be used, such that the forces exerted by each on a stationary object exactly counteract each other and only the velocity dependent force elements remain.

The lowest possible temperature that can be reached with Doppler cooling occurs when the rate of this cooling is balanced by the heating resulting from the random momentum transfers associated with spontaneous re-emission. These depend not only on the directions in which photons are spontaneously emitted, but also (in the case of illumination by balanced, counter-propagating beams) on the fact that as absorption is a probabilistic process, there may be an inequality between the number of photons scattered from each of the illuminating beams, and hence a resultant net momentum transfer. A derivation of this limiting temperature is not given here but can be found in many standard texts, such as [57], for example. The minimum temperature, T_D, that can be achieved via Doppler cooling based on a transition with spontaneous decay rate Γ is given by:

$$T_D = \frac{\hbar\Gamma}{2k_B}, \tag{2.11}$$

where k_B is the Boltzmann constant.

2.3 Magnetic trapping

The orbit of an electron around an atomic nucleus is, in effect, a current loop, and therefore generates a magnetic dipole. Furthermore, both electrons and nucleons have an intrinsic spin magnetic dipole moment. If the interactions between these magnetic dipoles within an atom are significantly stronger than any interactions that exist with externally applied magnetic fields, then it is clear that the internal state of the atom must be determined primarily by the form of these internal interactions. Therefore, if only a weak external field is applied to an atom, it is a legitimate approximation to first determine the internal state of the atom (by assuming no external field) and then determine the orientation of the atom's overall magnetic dipole moment within this external field.

Coupling the overall magnetic dipole moment $\vec{\mu}_{atom}$ to an external field $\mathbf{B}$ yields an interaction energy $U = -\vec{\mu}_{atom} \cdot \mathbf{B}$. Furthermore, quantisation of this interaction yields a set of discrete states of the atom in the magnetic field, which represent the eigenstates of the overall Hamiltonian. Because these are stationary states of the system (for a constant magnetic field), atoms that move sufficiently slowly within a spatially varying field will 'follow' the applied field — meaning that the alignment of their magnetic moment *relative* to the direction of the local magnetic field vector will be unchanged.

This means that for any individual atom the energy of interaction between the atom and the external field will simply be proportional to the strength of the applied field, as neither the magnitude of the magnetic dipole moment nor its orientation relative to the direction of the applied field will change as a result of the atom's motion. Atoms with a positive interaction energy can therefore be trapped about a minimum of magnetic field strength via this effect, as movement into a region of higher field strength increases their potential energy and a restoring force equal to $-\nabla U$ is therefore exerted on them (in a similar way, atoms for which the interaction energy is negative will experience anti-trapping about a magnetic field minimum).

The simplest form of magnetic trap employs a pair of anti-Helmholtz coils (as seen in figure 2.2) to produce a magnetic field zero that constitutes a local minimum of magnetic field strength about which atoms can be trapped. In this form of magnetic trap there is a relatively high loss rate of atoms from the trap, owing to the fact that the atoms are trapped about the zero of magnetic field: in very low field regions, the different m_f states of the atoms are nearly degenerate and there is a high rate of spin-flips between the trapped and anti-trapped states, which results in loss of atoms that enter these regions of the trap [63] (the atom's magnetic quantum number, m_f, reflects the alignment of the atom's magnetic moment relative to some chosen axis, which is frequently set to be the direction of the local magnetic field vector). Magnetic traps capable of containing atoms for longer time periods have been constructed via the use of time-orbiting potentials to trap atoms in a location where the time-averaged magnetic field strength is a local minimum but the instantaneous field strength is never zero [64], and by 'plugging' the central region of a standard magnetic trap with blue-detuned laser beam that produces a repulsive dipole force, thus preventing the atoms from reaching the zero of magnetic field and forming a ring-shaped atom trap with minimal spin-flip induced losses [65]. Such advanced forms of magnetic trap, while less straightforward to construct, have been shown to be capable of confining large numbers of atoms for long periods — for example [64] reports trap lifetimes in excess of 100 s and simultaneous confinement of more than 10^6 atoms.

As the trapping potential in a magnetic trap is conservative, atoms can not be continuously loaded into it from a background gas. In order to load a magnetic trap (or, for that matter, an optical dipole trap), a magneto-optical trap is therefore usually employed.

2.4 Magneto-optical trapping

An examination of the values of $\langle 1|\mathbf{E}\cdot\mathbf{r}|2\rangle$ mentioned in the discussion of scattering in §2.1.1 shows that, in the presence of a magnetic field, beams with different polarisation states stimulate different atomic transitions — specifically that linearly polarised light stimulates transitions in which there is no change in the magnetic quantum number of the atom. Meanwhile, depending on their direction of propagation and helicity, beams of circularly polarised light will stimulate transitions in which the atom's magnetic quantum number either increases or decreases by one.

This is the core principle behind magneto-optical trapping. In the presence of an external magnetic field, the energy of the atom will be dependent (as seen above) on its m_f state, and the wavelength at which resonance with a particular transition is achieved will therefore be modified in response to any change in the atom's m_f state that occurs during the transition. This means that the resonant frequency of the transitions stimulated by circularly polarised light will be either increased or decreased with respect to the transition frequency for linearly polarised light, by an amount that is proportional to the strength of the external field. The sign of this frequency shift is dependent both on the helicity and the direction of propagation of the light — swapping either of these will change the sign of the shift.

Figure 2.1 shows how this can be used to create an atom trap (at this stage only in one dimension). If red-detuned laser light is incident from both directions on a magnetic field zero as shown, then correct polarisation choices for the incident beams (both beams being circularly polarised and of the same helicity but opposite direction of travel) will result in a restoring force on any atom displaced from the trap centre. This is because the frequency of the transition induced by the beam propagating towards the trap centre will be reduced by the atom's interaction with the magnetic field, thus bringing the laser closer to resonance with the atomic transition and increasing the scattering rate from that beam. The opposite effect will reduce scattering from the other beam, and the net scattering force on the atom will then point towards the trap centre.

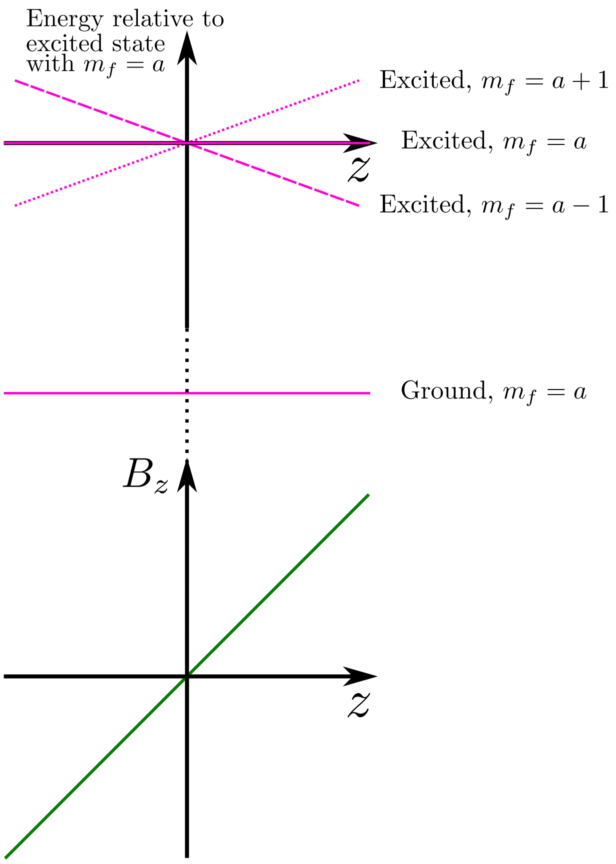

Figure 2.1: The magnetic field structure used for magneto-optical trapping in one dimension, and its effect on the relative energies of the different m_f states. As the polarisation and direction of travel of a beam determine the change in atomic m_f state on absorption of a photon from that beam, the position-dependent energy shift of different m_f states relative to one another can be used (with the correct choice of incoming polarisation states) to produce a position-dependent bias in favour of absorption from one laser beam over another. If this bias favours the laser beam that is propagating towards the trap centre, then a restoring force will be generated and an atom trap created.

Extension to a multi-dimensional trap requires the use of three pairs of (non-coplanar) counter-propagating laser beams and a magnetic field with a suitable structure along the axes of all three beam pairs. Such a magnetic field structure can be produced by a pair of anti-Helmholtz coils, as shown in figure 2.2. This can either be shown by applying the Biot–Savart law and integrating around the relevant current loops or more simply by considering the on-axis magnetic field: when nearer to one coil than another, the field component along the axis must be in the direction associated with that coil, while the field components perpendicular to the axis must be zero due to the symmetry of the system. At a point exactly in between the two coils, the field must therefore be zero. The

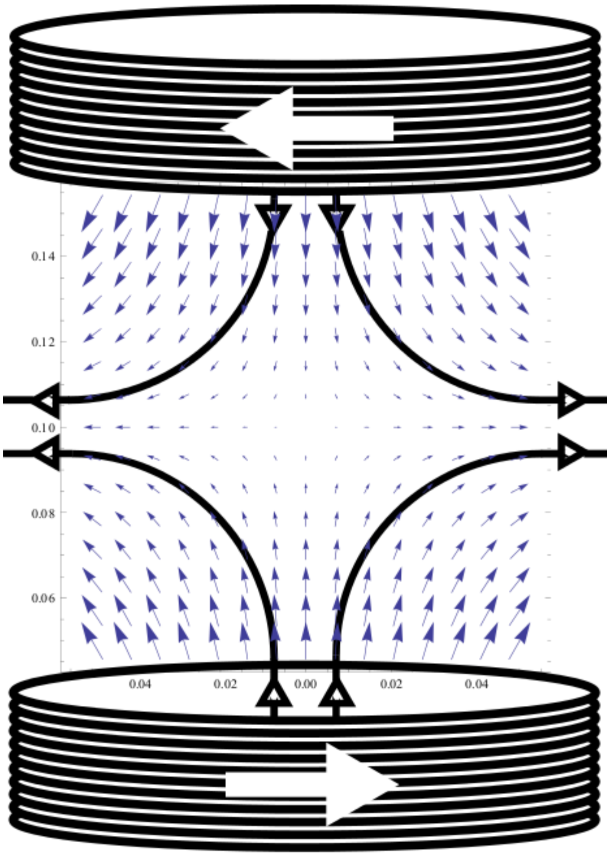

Figure 2.2: A pair of anti-Helmholtz coils and the resulting magnetic field structure — the white arrows show the direction of current flow in the coils. The bold, black arrows and lines are a schematic representation of the form of the field, while the blue arrows are a vector plot (arrow indicates direction and size indicates magnitude) of the magnetic field between two anti-Helmholtz coils in a plane that contains the axis of the coils (derived via numerical integration of the Biot-Savart law).

divergence theorem can then be used to show that the magnetic field vector's orientation must be of the opposite sense when moving away from the field zero along the axis as when moving away perpendicular to the axis — i.e. if, at a point slightly displaced from the field zero along the axis, the magnetic field vector points back towards the field zero, then the opposite will be true for a point displaced slightly from the zero in a direction perpendicular to the axis, and vice versa. An appropriate choice of polarisations for magneto-optical trapping (assuming the use of three orthogonal beam pairs) is then to have four beams of identical helicity traveling in the plane perpendicular to the axis of the coils and one beam pair of the opposite helicity traveling along the axis of the coils.

It is important to note that magneto-optical traps (and other traps based on the scattering force) typically have much greater trap depths than dipole or magnetic traps (usually of the order of a few K, compared with a few mK for dipole traps, for example). Furthermore, because they involve a resonant interaction with red-detuned laser radiation, magneto-optical traps simultaneously trap and Doppler cool atoms, allowing them to load themselves directly from a room temperature background gas. Neither magnetic nor dipole traps are capable of self-loading as they both produce conservative trapping potentials, in which there is no mechanism by which the kinetic energy gained by a particle on entering the trap can be dissipated before the particle escapes from the trapping region.

Magneto-optical traps also possess several enhanced cooling mechanisms, some of which allow atoms to be cooled to below the Doppler limit. Firstly, cooling to the Doppler limit is aided by the magnetic field structure around the trapping region: the position-dependent Zeeman shifts experienced by the different magnetic sublevels within each atomic energy state allow a wider velocity class of atoms to be addressed and Doppler cooled by the illuminating laser beams than would otherwise be possible. The effect is somewhat analogous to a Zeeman slower [66], in which resonance between the atoms in an atomic beam and a laser (used to decelerate the atoms via the scattering force) is maintained as the atoms are decelerated across a wide range of velocities via the application of a magnetic field of appropriate strength to induce a Zeeman shift that counteracts the Doppler shift. Secondly, as most magneto-optical traps make use of retro-reflected laser beams to achieve trapping, there are a number of cooling mechanisms relating to the polarisation gradients that occur in the standing waves resulting from this retro-reflection that are capable of reducing the temperature of the atom cloud in a magneto-optical trap to well below the Doppler limit. These mechanisms are analysed in detail in [67].

Experimentally, magneto-optical traps typically produce a cloud of between 10^5 and 10^8 atoms in a volume that ranges from a cubic mm to a cubic cm [40, 34]. The temperature of the atom cloud is usually somewhere between 20 and 140 μK [34, 68, 69]. The trapped atoms are typically alkali metals, with Rubidium being the most common followed by sodium, potassium and cesium.

2.4.1 Enhancements of the magneto-optical trap

Magneto-optical traps encounter a fundamental limitation on the number and density of atoms they are able to trap as a result of atoms near the outside of the cloud being expelled from the trap by the force exerted by re-radiated light from atoms close to the trap centre. This mechanism generally limits the total number of atoms in a conventional MOT to a maximum of a few times 10^9 atoms, with the maximum densities that have been obtained being of the order of 10^{11} atoms per cubic centimetre — see [70], for example.

At the cost of only a slight increase in experimental complexity, the achievable densities can be significantly increased using the 'dark' or 'compressed' MOT techniques described in [71] and [72]. The principle behind the dark MOT is that by allowing the atoms near the trap centre to be optically pumped into a 'dark state', in which they do not interact strongly with the applied laser light, it is possible to reduce the amount of light re-radiated from these atoms and so lower the expulsive radiation pressure force produced on atoms nearer the extremities of the trapping region. Compressed MOTs involve a sudden increase in magnetic field gradient which brings about a corresponding increase in the strength of the trapping force, which then causes the trapping force to outweigh the repulsive forces caused by re-radiation of the trapping light and so compresses the atom cloud. The increase in magnetic field gradient must be applied transiently as it lowers the loading rate of the MOT, and so continuous application would result in a reduced number of trapped atoms. These approaches can produce a significant increase in the number and/or density of atoms that can be trapped — for example, [72] quotes a transient density increase of more than a factor of 10 while [71] reports achieving an atom density of 10^{12} atoms per cubic centimetre.

2.5 Optical dipole trapping

In §2.1.2 an expression for the optical dipole force was given which showed that, depending on the frequency difference between the illuminating laser light and an atomic transition, a force could be generated that would either attract atoms to or repel them from regions of high light intensity. A local maximum of light intensity can therefore be used to create an atom trap, provided that the laser frequency is tuned slightly below that of the relevant atomic transition ('red-detuned'). Light tuned slightly above an atomic transition frequency ('blue-detuned') can be used to create potential barriers that confine or guide atoms. The number of atoms held in such traps typically varies between 1 and 10^4, while the trap depths are generally sufficient to confine atoms at temperatures of a few mK or below. However, as these features are constrained by practical rather than fundamental considerations, there is a great deal of variability between individual experiments and no small set of examples can legitimately be considered representative of the overall capabilities of such traps.

Equation (2.1) shows that, for large laser detunings ($\delta \gg \Gamma$), the scattering rate is approximately proportional to δ^{-2}. However, (2.7) shows that in the same limit the strength of the dipole force is proportional only to δ^{-1}. This means that in the limit of very large detunings, the dipole force dominates over the scattering force, and with a sufficient level of detuning dipole traps can be constructed in which there is an extremely low scattering rate and the scattering force can be considered negligible. There is of course a balance to be struck, as in practice the available laser power may be limited, and detuning further from resonance will reduce the absolute magnitude of the dipole force, even as it increases its relative strength when compared to the scattering force.

There are two primary mechanisms by which atoms can be lost from dipole traps. Firstly, in an imperfect vacuum an untrapped atom from the background gas (which will typically carry kinetic energy that greatly exceeds the trap depth) can collide with a trapped atom and expel it. The rate at which this occurs is determined almost entirely by the properties of the vacuum system, and shows very little dependence on the parameters of the dipole trap. The second loss mechanism results from heating of the trap contents by unwanted scattering of light from the trapping beams. This can severely limit the lifetime of trapped atoms if the detuning is small, but its effects can be reduced by increasing the detuning, as the ratio of trap depth to scattering rate is proportional to δ. In real experimental situations, one must also allow for some loss and/or heating of the trapped atoms by variations in the power and orientation of the laser beam used to form the traps. These can be minimised by using an appropriate laser power supply (with minimal current fluctuations), by physically enclosing free-space optics to reduce the influence of air currents, by running the experiment in a temperature-stabilised environment and by ensuring that the entire optical system (including the source laser) is mechanically isolated from its surroundings to reduce the influence of unwanted vibrations.

Traps based on blue-detuned light are generally more complicated to construct than those based on red-detuned light, as a closed surface surrounding the trapping region must be illuminated at high intensity rather than just the trapping region itself. However, because the centre of the trap is now a region of low, not high, intensity, the atoms spend less time in regions of high light intensity than in a red detuned trap. This results in a lower scattering rate and hence less heating and a longer trap lifetime. Blue detuned dipole traps have been shown to confine atoms for many seconds [73], and there is no fundamental reason why confinement could not be achieved for much longer times if desirable.

One particularly important application of dipole trapping is the creation of ‘optical lattices’. These are obtained by retro-reflecting a laser beam in order to produce a standing wave, which will contain both nodes and anti-nodes at regular intervals. These regions of high and low intensity can then be used to trap cold atoms using the dipole force. Additional beams can be used along other axes to extend this scheme to two or three dimensions, and it is common to use optical lattices to construct regular, two-dimensional arrays of dipole traps for cold atoms [25].

Chapter 3

Microtraps — fabrication

In this chapter we describe the templated electrodeposition process that we used to synthesise regular arrays of truncated spherical cavities in a metallic film, and display the results of this process via images produced by both optical and scanning electron microscopy. These structures can be used as arrays of microscopic, concave mirrors that, when exposed to normal illumination by a laser beam of an appropriate frequency, will produce optical dipole traps at their foci. We chose to base our work around the use of ^{85}Rb as the atomic species to be trapped, due to its well understood optical properties and the ready availability of diode lasers operating at appropriate wavelengths to trap, image and manipulate atoms thereof. We also propose enhancements to the synthesis process that might allow for the creation of microstructures capable of simultaneously producing both magneto-optical and dipole traps with coincident trap centres, meaning that the magneto-optical traps could be used to facilitate loading of the dipole traps. The results of preliminary work on this enhanced fabrication process, in which cavities with a central hole are synthesised, are also given.

Widely used microfabrication techniques such as photolithography, electron beam lithography and silicon etching [38] are not generally convenient methods for the production of smoothly curved three-dimensional structures, with lithographic techniques typically using only a two-dimensional mask to define the required structures and silicon etching producing angular structures with sharp corners[1]. As a consequence, it would not be straightforward to use such techniques to produce the desired concave mirror arrays for our application. However, electroplating around a template object (followed by removal thereof) allows the creation of a void of identical shape to that of the object within an electroplated film, and monodisperse polymer and silica microspheres are commercially available in a range of sizes from 20 nm to 200 μm [75]. By using these spheres as template objects, it is possible to create smooth, truncated, spherical cavities within a metallic film.

We have created regular arrays of such cavities via the process described in [43], in which gold is electroplated around a self-assembled array of polymer microspheres that are then dissoved with tetrahydrofuran (THF) to leave behind a microstructured gold film. We used a somewhat larger size of microsphere than was employed in [43] — our initial prototype system is based around spheres with a 100 μm diameter. This relatively large cavity size was chosen for our prototype in order to facilitate the observation and loading of the dipole traps we ultimately aim to create using these structures. However, [43] and a number of similar papers have reported that the process is effective with sphere diameters anywhere between 50 μm and a few tens of nm, so we do not anticipate any difficulties with the fabrication of similar structures on smaller length scales.

3.1 Truncated spherical cavities

The principles of the fabrication method for these cavities are illustrated in figure 3.1. A clean glass substrate is placed in an evaporation chamber at a pressure of $\sim 10^{-6}$ mbar. Evaporative deposition of a uniform, 10 nm layer of chromium is then carried out — this serves as an adhesion layer that helps to prevent subsequent layers from peeling away from the substrate. 200 nm of gold is then deposited on top of this via the same technique.

---

[1] Some recent work by Hinds et al. has produced such structures via silicon etching [74], but at present the process appears to be somewhat more involved than the one we employ while not offering any significant advantages for our application.

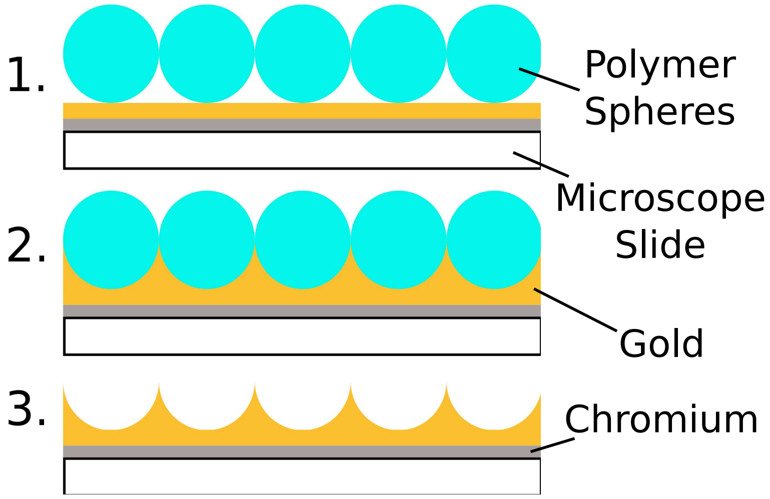

Figure 3.1: The process used to synthesise regular arrays of truncated spherical cavities.

The sample is then placed in a 0.02 M solution of cysteamine hydrochloride in water, where it is left for a period of several hours. The cysteamine hydrochloride molecules adhere both to the gold surface and to the polymer microspheres, which subsequently act as our templates, and are thus useful to prevent the spheres from leaving the surface during the subsequent electroplating process. Following this, a droplet of a suspension of polymer microspheres in water is placed onto the gold surface. This is left under a loose cover to allow gradual evaporation of the water over a period of around 24 hours. During this time, as the water droplet contracts, capillary forces cause the spheres to be drawn into an optimally packed monolayer, as seen in figure 3.2.

Electroplating is then carried out around these spheres. All parts of the conductive surface onto which gold should not be deposited (generally those not covered by the microspheres) are masked with an insulating layer of nail varnish. Gold plating is performed using a commercially available gold plating solution (Metalor ECF60) with an associated brightener (Metalor E3). The brightener is intended to improve surface quality by promoting nucleation of the electrodeposited gold, thus preventing an uneven distribution of the metal arising from heavy deposition around only a small number of naturally occurring nucleation sites. The current and voltage are regulated using an Autolab PGSTAT30 potentiostat/galvanostat. Regulation of the current density is important, as excessive current densities result in poor quality deposition and high surface

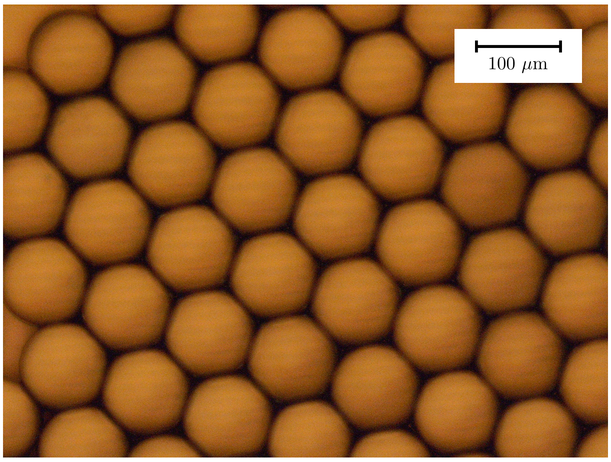

Figure 3.2: Optical microscope image of hexagonally packed microspheres, taken after evaporation of the water in which they were initially suspended but before any electrodeposition. The gaps on the upper and left hand edges of the image are the boundaries of the microsphere layer, not internal defects.

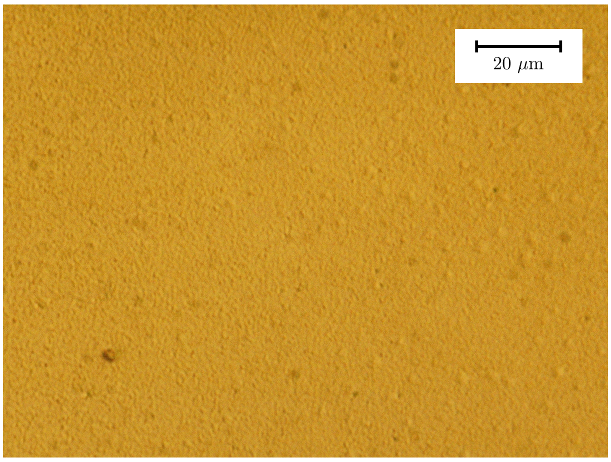

Figure 3.3: Optical microscope image of a gold surface produced by electroplating with a current density of 50 A m^{-2}.

roughness. For example, figure 3.3 shows an image (taken via an optical microscope) of a gold surface produced by electroplating at a current density of 50 A m^{-2} while figure 3.4 shows one for which the current density was set at the higher value of 300 A m^{-2}. A few cursory experiments were undertaken to determine the effect of the temperature of the electroplating solution on surface quality, but as no significant correlation was observed between the two during these early experiments a more thorough investigation was not deemed worthwhile.

The depth of the electroplated gold is controlled by regulating the charge passed. By

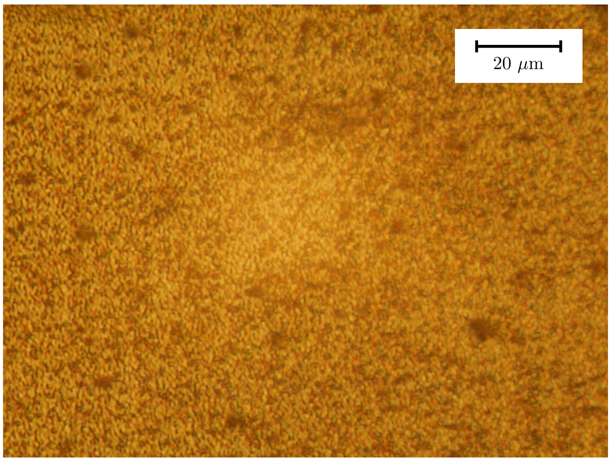

Figure 3.4: Optical microscope image of a gold surface produced by electroplating with a current density of 300 A m^{-2}.

making use of the fact that optimally packed circles on a plane occupy 90.7% of the area of the plane [76], it is possible to carry out a geometric calculation of the total volume of gold deposited in plating to a given depth. Figure 3.5 illustrates this calculation. In the 9.3% of the area not covered by the spheres at any height, the volume deposited is simply equal to the exposed area times the plating depth. Underneath the spheres however, we must refer to figure 3.5. Consider the reduction in the volume V_s of gold plated as a consequence of the presence of a single template sphere of radius R. By considering the geometry shown in figure 3.5, it can be seen that this is given by:

$$V_s = \int_0^h \mathrm{d}a\, \pi\, (x(a))^2 = \int_0^h \mathrm{d}a\, \pi \left(R^2 - (R-a)^2\right) = \pi \left[R^2 a + \frac{1}{3}(R-a)^3\right]_0^h = \pi h^2 R - \frac{\pi}{3} h^3. \tag{3.1}$$

As we now know the volume lost due to the presence of each sphere, and that this much volume is lost out of every $\pi R^2 h$ units of volume across 90.7% of the surface, we can then deduce the total volume of gold plated onto an area A as a function of the plating height h:

$$V = 0.093Ah + 0.907Ah\frac{\pi R^2 h - V_s}{\pi R^2 h} = 0.093Ah + 0.907Ah\left(\frac{R^2 - hR + h^2/3}{R^2}\right). \tag{3.2}$$

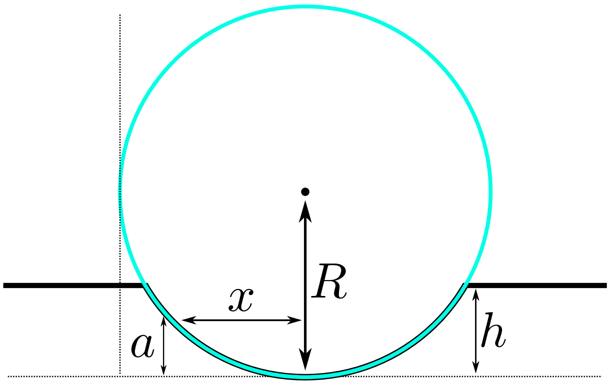

Figure 3.5: Diagram illustrating the calculation of the total volume of gold deposited in plating to a given depth around an array of spherical templates.

Combining this with a knowledge of the density and atomic mass of gold and the fact that, as gold is in the monovalent state in this particular plating solution, one atom of gold is deposited for each electron charge passed, it is possible to derive the following formula for the charge, Q Coulombs, required to plate to a depth of h micrometers:

$$Q = 0.941\left(0.093Ah + 0.907Ah\left(\frac{R^2 - hR + h^2/3}{R^2}\right)\right), \tag{3.3}$$

where A is the area onto which one is plating (in cm^2) and R is the radius of the microspheres (in μm).

Once the electroplating process is complete, the template spheres and nail varnish are dissolved away using tetrahydrofuran, leaving behind the desired structures. Figure 3.6 shows images of these structures taken with both optical and scanning electron microscopy. Several important factors can clearly be seen from these images. Firstly, images (b) and (d) show the regular size and formation of the cavities. Secondly, image (b) demonstrates that each cavity produces a tight focus on the axis of the cavity for reflected light that was incident normal to the surface. Finally, images (a) and (c) show that while surface roughness outside the cavities is substantial (and determined primarily by the properties of the electroplating process), surface roughness inside the cavities is greatly reduced. We surmise that this results from the continuation of gold plating at all points underneath the template spheres until the gold surface becomes flush with the surface of the microsphere, meaning that surface roughness inside the cavities is

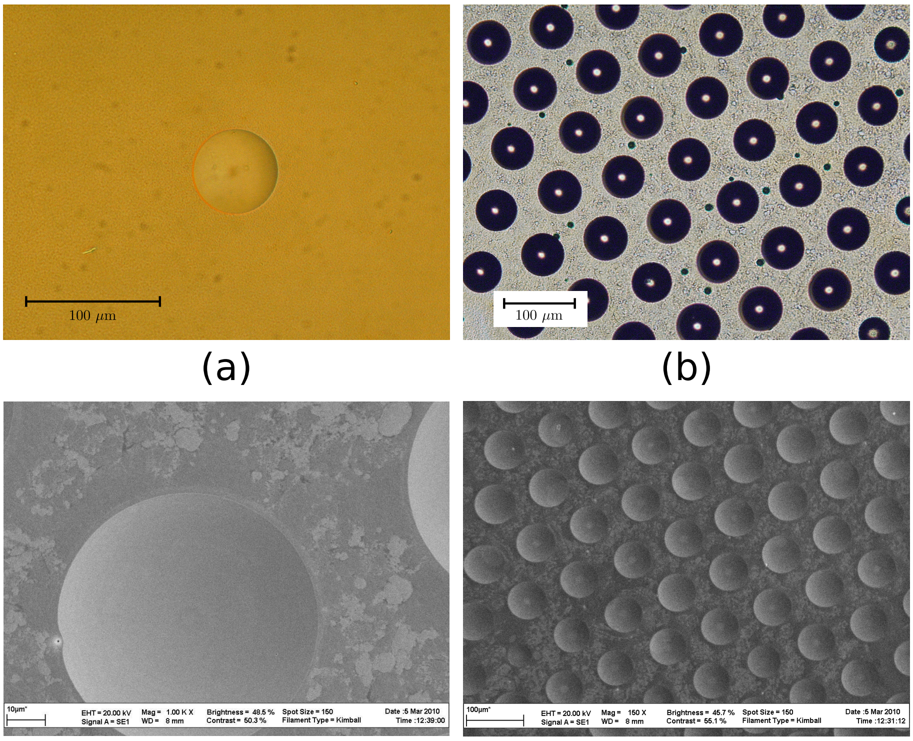

(a) (b)

(c) (d)

Figure 3.6: Images of individual truncated spherical cavities (a and c) and parts of the cavity array (b and d) taken using optical (a and b) or scanning electron (c and d) microscopy. These images are illustrative only, and the sample was viewed at a higher zoom level in order to estimate the surface roughness inside the cavities.

determined primarily by the roughness of the template surface around which the gold was plated. In principle, an accurate value for the RMS surface roughness could have been obtained using an atomic force microscope, but as even very approximate observations using a scanning electron microscope (with an adjustable viewing angle) showed that roughness inside the cavities was well below the required level this was not considered necessary (these measurements suggested roughness levels of the order of 1 μm outside the cavities and 1 nm inside). Similar structures have previously been synthesised via the same technique and used in optical physics experiments — see for example [77, 78].

3.2 Cavities with a central hole

As has been discussed previously, we aim to create not only microscopic dipole trap arrays but also arrays of microscopic magneto-optical traps. The illumination received by an atom at the focus of one of the cavities described above, while appropriate for optical dipole trapping, is not suitable for magneto-optical trapping. We have therefore carried out some preliminary work towards the creation of cavities with a central hole, as these structures are necessary for our proposed approach to the production of microscopic magneto-optical traps — see §3.3. In this section we give both our results to date and our proposals for further work on this topic, as our experimental investigations remain unfinished at this time.

The creation of cavities with a central hole would enhance the versatility of the system by allowing optical access from behind the trapped atom sample, which would increase the range of options available for dipole trap loading, while photolithographic techniques could make use of the template thus produced to selectively deposit ferromagnetic materials behind the centre of each cavity, in a pattern appropriate for the creation of microscopic magneto-optical traps. Furthermore, central holes do not prevent the formation of a region of high reflected intensity suitable for dipole trapping — indeed cavities containing a central hole could provide appropriate illumination for the production of dipole and magneto-optical traps with coincident centres, as shall be seen in the next chapter. This would then allow these microscopic magneto-optical traps to be used to load atoms into the microscopic dipole trap array.

Our prototype procedure (based on that used in the preceding section) for the creation of cavities containing central holes is outlined in figure 3.7. Firstly, the chromium layer is omitted during evaporative deposition, and the gold layer is replaced with copper. This results in poorer adhesion of the metallic layer to the glass substrate — in particular exposure to cysteamine hydrochloride in water was found to cause the layer to peel away from the substrate. Cysteamine treatment was therefore carried out using cysteamine hydrochloride dissolved variously in methanol, ethanol or acetone, all of which were found to be effective. It is our hypothesis that the peeling of the film in response to

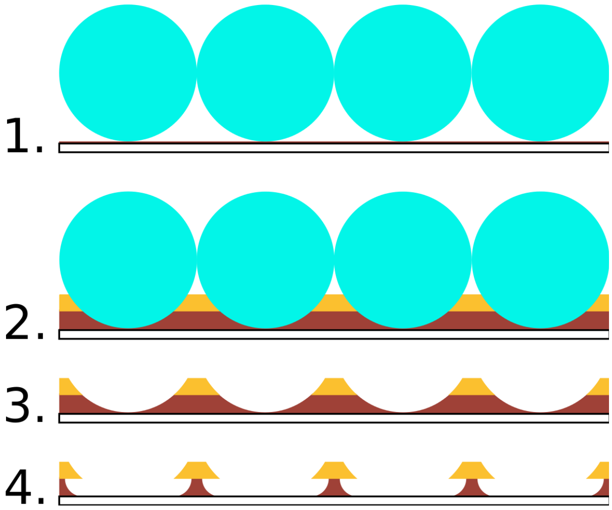

Figure 3.7: The synthesis process for cavities with a central hole. The yellow regions are gold while the brown regions indicate copper.

cysteamine hydrochloride in water resulted from the large permanent electric dipole moment of the water molecules, which causes significant adhesive forces between the water and the copper film, thus making it energetically favourable for the film to peel away from the substrate and be exposed to water on both sides. If this is the case, replacing water with one of the substances listed above (the molecules of which have a smaller permanent electric dipole moment per unit mass) would be expected to have the observed effect of reducing the propensity of the film to peel away from the substrate.

Spheres are then placed onto the surface in the usual way, and once the sample has dried copper is electroplated around them from a solution of copper (II) sulphate in water. Once plating commenced, there was a rapid increase in the thickness of the film which we assume to have caused a corresponding increase in the rigidity of the film and thus prevented peeling, as no significant peeling was encountered during the electroplating process so long as plating commenced promptly once the sample was placed in the plating solution.

Once copper had been plated to the desired depth, gold was then electroplated on top of this before the spheres were removed. After dissolution of the microspheres, the sample was exposed to 1 M nitric acid. Experimental tests with blank copper substrates indicated that exposure to this acid caused all exposed copper surfaces to recede at approximately 1 nm per second. By selecting the correct time interval for the relevant

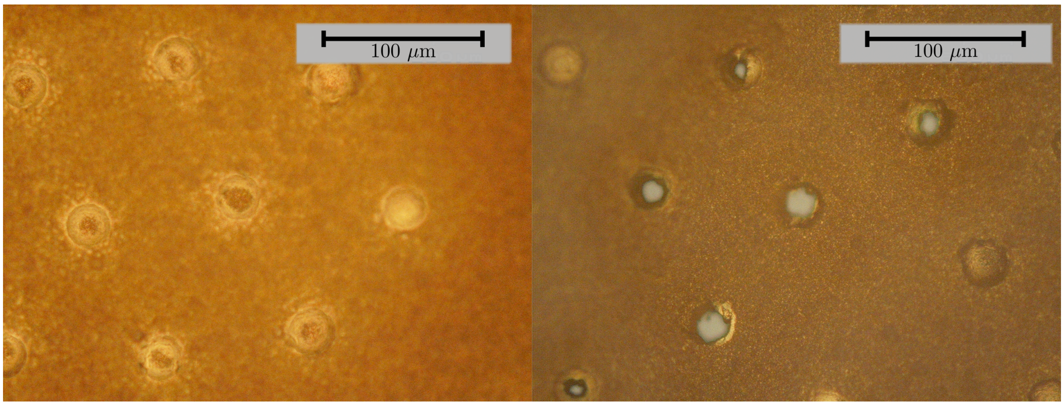

Figure 3.8: Optical microscope images of the structures resulting from the prototype synthesis process for cavities with a central hole. The left hand panel shows a sample after removal of the template spheres but before exposure to nitric acid, while the right hand panel shows a completed sample.

plating depth, it was therefore possible to produce structures of the form shown in the last part of figure 3.7, by dissolving away the exposed region of copper in the centre of each cavity. Experimental results from this prototype process are shown in figure 3.8. It is clear from these images that, while the principles of the process are sound (regions of exposed copper can be seen in the centre of most cavities once the spheres have been removed, and exposure to nitric acid does then result in the formation of a hole at the centre of these cavities), the quality of the samples produced is very low — the deposited copper shows significant variations in depth and a high surface roughness which prevents the formation of regular samples with neat circular holes in every cavity. We attribute this to the rudimentary nature of our copper plating process and the fact that the plating solution used to produce these samples was not acidified.

Most commercial copper plating using copper (II) sulphate is done with acidified plating solutions produced via the addition of around 25 grams per litre of sulphuric acid. In order to improve on the results shown in figure 3.8, we therefore employed a slightly acidified solution (5 grams per litre of sulphuric acid) in order to ascertain whether any improvement in the quality of the deposited copper could be observed. Figure 3.9 shows optical microscope images of copper surfaces that were electroplated with and without acidification of the plating solution, and a very clear improvement can be seen as a result of acidification.

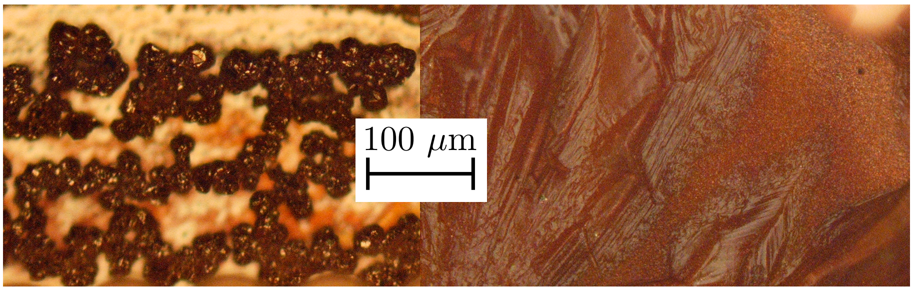

Figure 3.9: Optical microscope images of copper surfaces produced by electroplating with unacidified (left) and acidified (right, 5 g l^{-1} H_2SO_4) copper sulphate solutions. It is clear that in the left hand panel the quality of the deposited copper is extremely poor — grain size is very large and the deposition is heavily islandised, with many regions having no covering of copper at all. In the right hand image quality is much improved. The substance covering much of the surface in the right hand image is $CuSO_4.5H_2O$, which is left on the substrate when it is removed from the plating solution and allowed to dry, but it can easily be removed by rinsing with deionised water.

However, a further problem was then encountered as sulphuric acid was found to dissolve the polystyrene microspheres used in the process, thus preventing the entire templated deposition process from working. In order to avoid this difficulty, further work with acidified plating solutions should be conducted using silica microspheres, which are more resistant to attack by sulphuric acid. We expect that the combination of an acidified plating solution with silica microspheres may yet produce a high quality array of cavities with central holes, but have not yet had time to investigate the matter further. We therefore conclude that it should be possible to fabricate arrays of truncated spherical cavities with a hole of well-defined size at the centre of each cavity, and that such arrays could be of use in optics and atomic physics experiments. A small amount of further work on the technique described above may well yield a succesful experimental demonstration.

3.3 Microscopic magneto-optical traps

Microscopic magneto-optical traps require both the optical and magnetic fields to be tailored on a micrometre length scale. Figure 3.10 shows how, by using the process described above to produce cavities with central holes, and using a thin quarter wave plate as the deposition substrate, it should possible to create an optical field with a

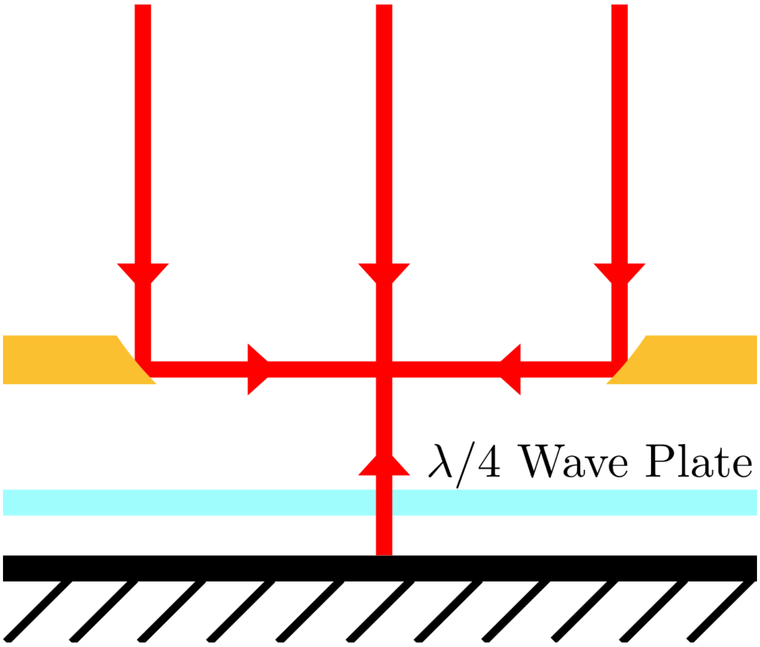

Figure 3.10: Creation of the required optical fields for a micro-MOT by a truncated spherical cavity with a central hole under normal illumination. The substrate onto which the cavity is grown must be a thin quarter wave plate that backs onto a planar reflector as shown, in order that the waveplate can counteract the polarisation change undergone on reflection from the mirror behind the central hole. Combining this with the fact that the light incident on the inclined sides of the cavity will undergo a polarisation change on reflection, it becomes possible to obtain the necessary polarisation states for all light approaching the trapping region while using only a single, circularly-polarised laser beam to illuminate the cavity.

suitable form (including beam polarisations) within each cavity. However, in order to obtain the required magnetic field structure, further modifications are required.

We propose a process in which the existing holes at the centre of each cavity are used as masks for the photolithographic deposition of a ferromagnetic layer on the reverse surface of the substrate. The process is illustrated in figure 3.11. A suitable photoresist (such as PMMA) is deposited on the reverse face of the substrate. This photoresist is then exposed *through* the sample before etching, such that only the areas of photoresist directly behind the cavities' central holes are removed. This is then followed by evaporative or sputter deposition of a reflective layer and a ferromagnetic layer (consisting of a suitable ferromagnetic material such as NdFeB) onto the reverse of the substrate. The remaining photoresist can then be removed to leave circular islands behind the centre

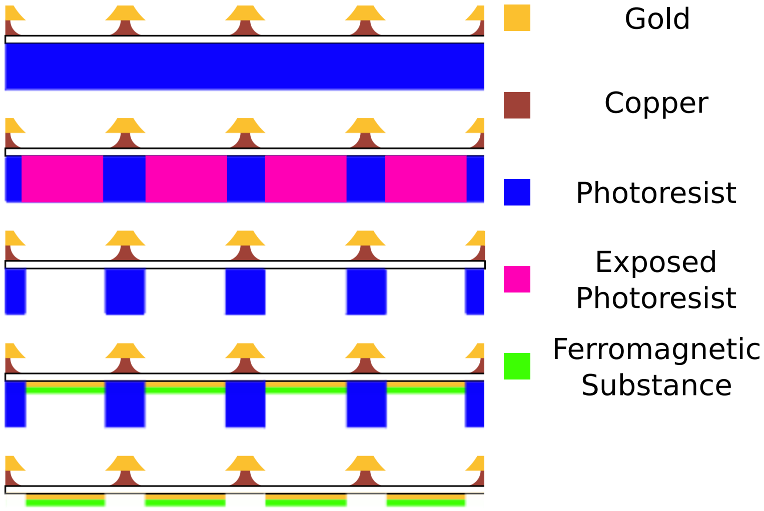

Figure 3.11: Proposed process for the creation of a sample with a ferromagnetic island behind the centre of each cavity. Blue regions are photoresist, pink exposed photoresist and green a suitable ferromagnetic substance such as NdFeB.

of each cavity. Each of these islands consists of an upper reflective layer and a lower ferromagnetic layer.

By exposing the sample to a strong external magnetic field while it is heated to above the Curie temperature, it should then be possible to magnetise these ferromagnetic islands such that their magnetic dipole moments are perpendicular to the plane of the substrate [79]. When subsequently placed in a weaker external bias field of the opposite orientation, these islands will then produce the required magnetic field structure for the creation of micro-MOTs. We can see this without any formal calculation in the following way: firstly, in the limit of very large distances from the structure the external field must dominate, while (for a suitably weak external field) the field generated by the magnetic island must dominate when very close to the ferromagnetic substance. If these fields oppose one another (which they will on the axis of each island), then there must therefore be a field zero somewhere between these two regimes. By applying the divergence theorem in combination with Maxwell's law ($\nabla.\vec{B} = 0$), we can see that the field structure within a sufficiently small region around that field zero must be identical to that shown in figure 2.2.

We therefore expect that the creation of structures suitable for the production of regular arrays of microscopic magneto-optical traps should be possible via this method. The properties of these traps, along with those of dipole traps based on the more standard structures described in the preceding sections, will be discussed and analysed in more detail in the following chapter.

Chapter 4

Microtraps — theory

4.1 Microscopic dipole traps based on truncated spherical cavities

Each cavity within an array such as that shown in figure 3.6(b) and 3.6(d) is a concave, spherical reflector with a focal length approximately equal to half the radius of the spheres used in the synthesis process. Under normal illumination, each cavity will therefore produce a high intensity focal spot on its central axis, at about half a sphere radius above the lowest point in the cavity. If the illuminating light is red-detuned with respect to an atomic transition, then these focal spots might be used to form an array of dipole traps for the relevant atomic species.

The properties of these traps could be approximately determined via ray tracing arguments, which would provide a reasonably accurate result in the case of our 100 μm diameter prototype cavities. However, ray tracing becomes increasingly inaccurate as the size of the cavities is reduced and the importance of diffraction effects increases.

Instead, we therefore use the approach described in [80], in which the relevant Kirchoff diffraction integral is evaluated for points near to the cavity axis. In this treatment, it is assumed that the displacement of all points of interest from the cavity axis is small compared to the radius of the cavity and also that both the incident and doubly reflected light fields make a negligible contribution to the light intensity near the focus,

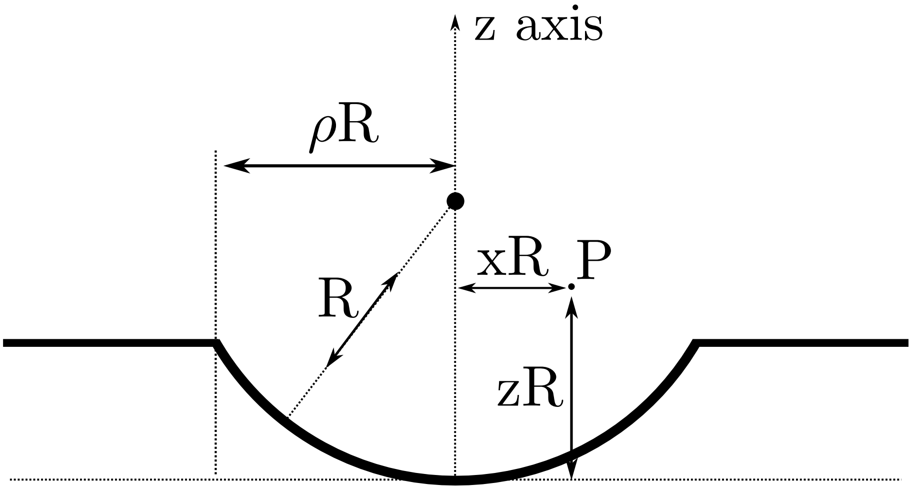

Figure 4.1: Geometry of a micro-cavity, showing the dimensions of the cavity and the position P at which the reflected wavefunction is evaluated in equation (4.1).

on account of being heavily outweighed by the tightly focused light that returns after a single reflection. As all of these conditions are met in our case, it provides an appropriate means by which to determine the shape and depth of our dipole traps.

We take the radius of the spheres used to produce the cavity to be R, the radius of the circle that marks the rim of the cavity to be ρR and the z axis to be along the central axis of the cavity. We also define α as equal to $\sqrt{1-\rho^2}$. The result of such a treatment (as given in [80]) is that the reflected light amplitude ψ_{ref} at position $(xR, 0, zR)$ as shown in figure 4.1 is given by:

$$\psi_{ref} = \frac{k\psi_0}{2i} \int_\alpha^1 \mathrm{d}\chi \frac{e^{ikR(\chi+\sqrt{1+z^2-2\chi z})}}{\sqrt{1+z^2-2\chi z}} \left(\chi + \frac{1-\chi z}{\sqrt{1+z^2-2\chi z}}\right) J_0\left(\frac{kRx\sqrt{1-\chi^2}}{\sqrt{1+z^2-2\chi z}}\right), \tag{4.1}$$

where we have neglected an arbitrary overall phase factor, assumed $x \ll 1$, and taken the illumination to be by a normally incident plane wave with amplitude ψ_0 and wavevector $-k\hat{\mathbf{z}}$. Figure 4.2 shows the light intensity distribution predicted by this formula near the foci of cavities with diameters of 100, 10 and 1 μm and $\rho = 0.6$ when illuminated with 780 nm light at normal incidence (where the intensity is calculated by taking the square of the modulus of the ratio of ψ_{ref} to ψ_0). In the largest cavity the light intensity near the focus reaches about 5000 times the incident light intensity. Even in the smallest of these cavities, the intensity at the focal point can be seen to be ~2.5 × the incident light intensity, showing that some focusing effect still exists (although the intensity plot should no longer be considered accurate as the approximation that the reflected light intensity near the focus greatly exceeds the incident intensity no longer holds).

Equation (2.9) shows that, in the limit of large detunings, the magnitude of the dipole potential is directly proportional to the light intensity. Having determined the intensity distribution within these cavities, it is therefore trivial to derive the dipole potential as a function of position, and hence to obtain the shape and depth of the dipole traps. As an example, suppose that a 100 μm cavity with $\rho = 0.6$ were illuminated with 800 nm light at an intensity of 1 W cm^{-2}, to be used as a dipole trap for ^{85}Rb. The trap depth would then be given by (4.1) and (2.9) as 320 μK, with the approximate FWHM of the trap potential being 0.25 μm in the transverse directions and 4 μm axially.

4.1.1 Cavities with a central hole

The intensity distribution inside a cavity with a central hole can be found in a straightforward manner by applying (4.1) and subtracting the amplitude that would result from a cavity of size equal to the central hole from the amplitude resulting from a cavity of the correct overall size but without a central hole, before once again taking the square of the modulus of the ratio of amplitudes. Mathematically, this is to say that:

$$\text{Intensity in cavity with hole} = \text{Incident intensity} \times \left| \frac{\psi_{ref}(\rho_2) - \psi_{ref}(\rho_1)}{\psi_0} \right|^2, \qquad (4.2)$$

where ρ_1 relates to the width of the hole and ρ_2 to the width of the overall cavity. Figure 4.3 shows the intensity distribution inside a cavity of radius 50 μm with $\rho_1 = 0.6$ and $\rho_2 = 0.8$. A calculation of the depth of a dipole trap formed at its centre, under the same illumination conditions as before, yields a trap depth of 190 μK — clearly showing that such structures retain the ability to form dipole traps of a worthwhile depth when the central region of the cavity is removed. The periodic intensity variation seen in the

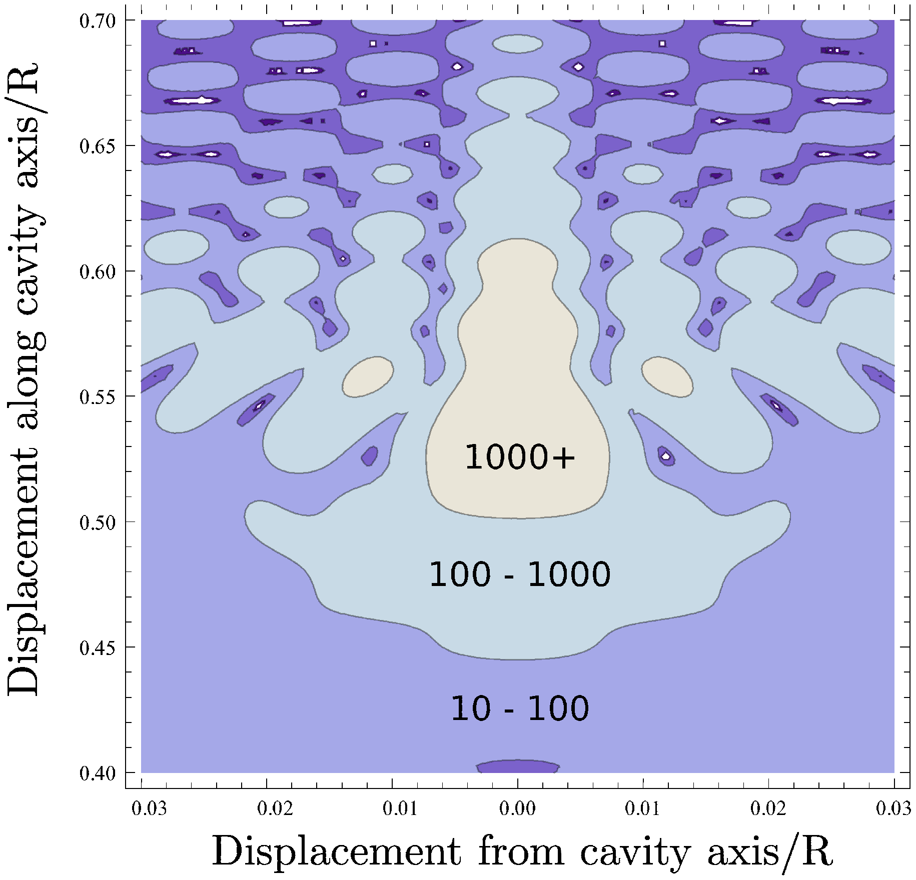
Displacement along cavity axis/R
0.70
0.65
0.60
0.55
0.50
0.45
0.40
1000+
100 - 1000
10 - 100
0.03
0.02
0.01
0.00
0.01
0.02
0.03
Displacement from cavity axis/R

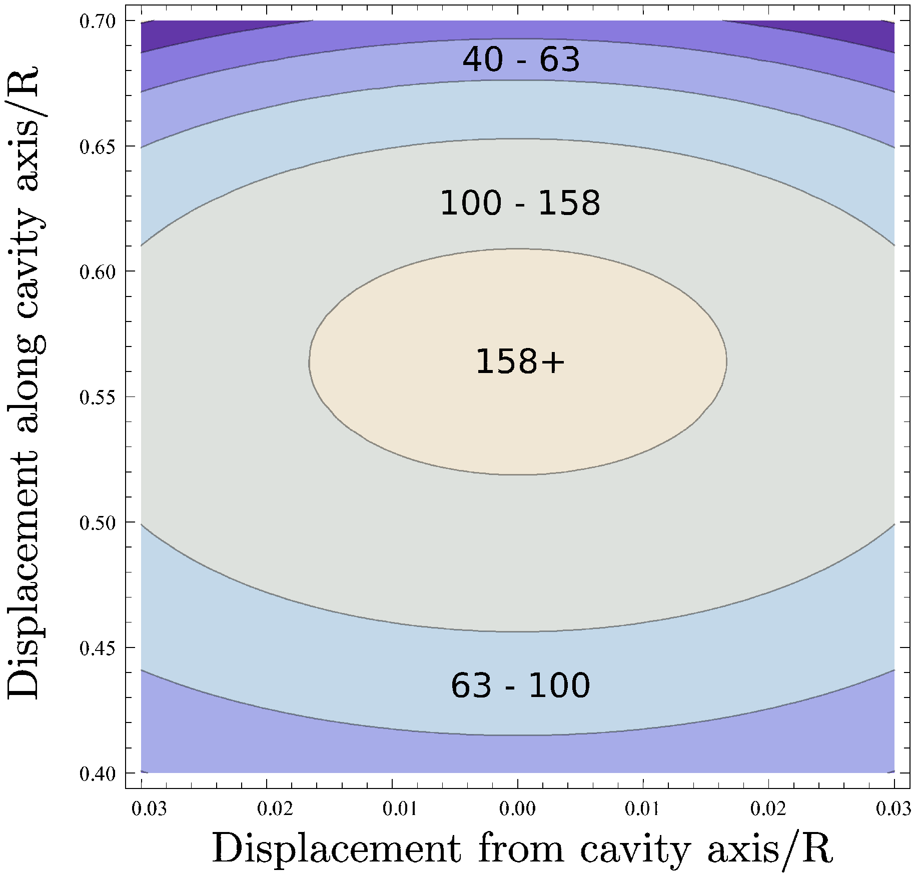
Displacement along cavity axis/R
0.70
0.65
0.60
0.55
0.50
0.45
0.40
40 - 63
100 - 158
158+
63 - 100
0.03
0.02
0.01
0.00
0.01
0.02
0.03
Displacement from cavity axis/R

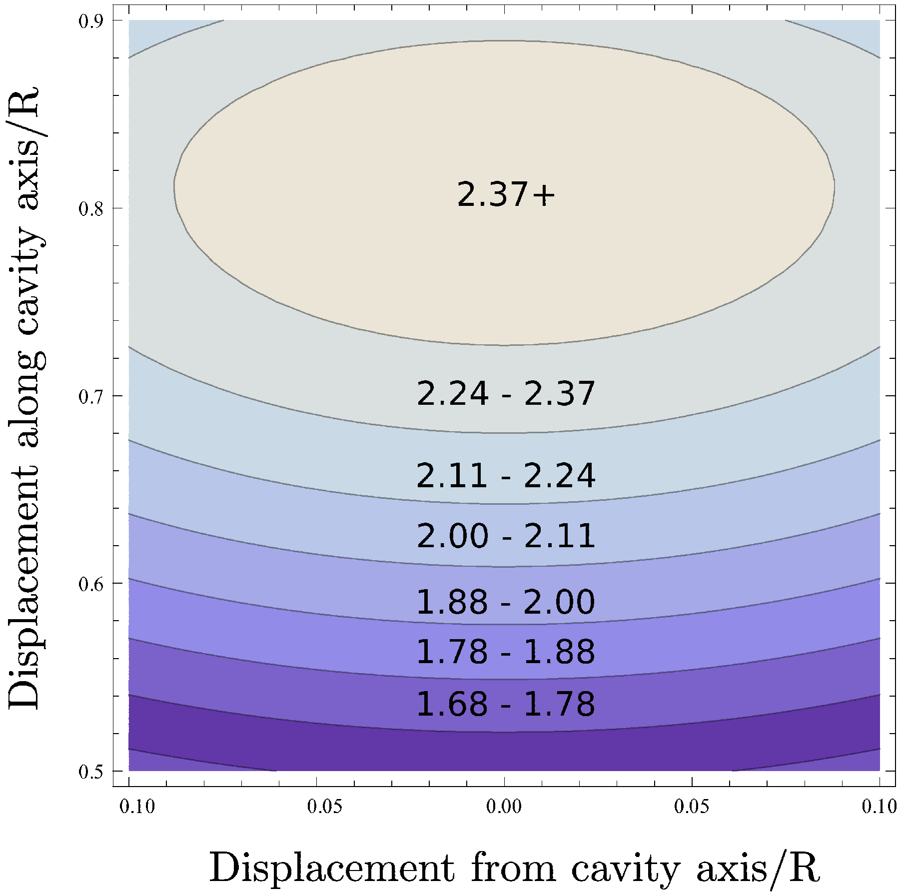

Figure 4.2: Contour plots of simulated light intensity distributions inside cavities with $\rho = 0.6$ based on spheres with R = 50, 5 and 0.5 μm (top to bottom) when illuminated at normal incidence with a 780 nm plane wave. The numbers on the plots show the ratio of the reflected light intensity in the region they are printed in to the incident light intensity. The origin of the vertical axes is set at the position where the cavity axis intersects the reflective surface. The displayed function is the base ten logarithm of the ratio of the reflected to incident light intensity, with the value of the highest contour and contour separation being equal to (top to bottom) 3 & 1, 2.2 & 0.2 and 0.375 & 0.025.

figure is reminiscent of a standing wave or optical lattice, and its origin is the interference of the counterpropagating light rays reflected from opposite sides of the microcavity, as previously illustrated in figure 3.10.

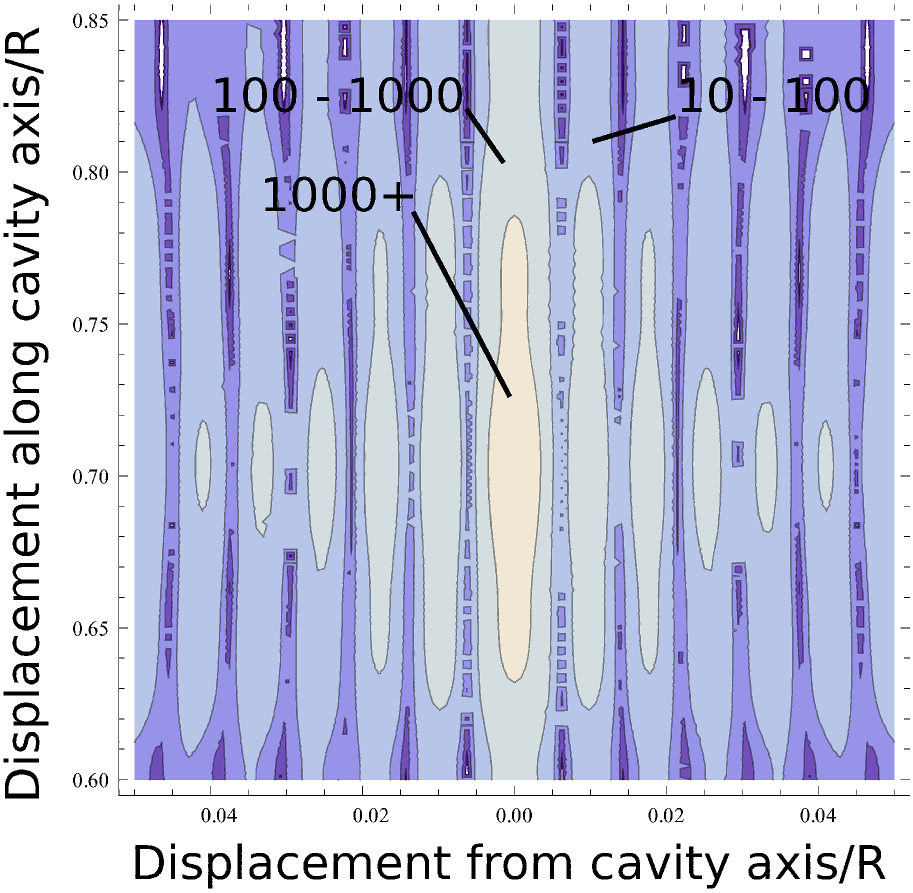

Figure 4.3: Contour plot of simulated light intensity distribution inside a cavity with R = 50 μm and $\rho_2 = 0.8$ containing a hole with $\rho_1 = 0.6$ when illuminated at normal incidence with a 780 nm plane wave. The numbers on the plot show the ratio of the reflected light intensity in the region indicated to the incident light intensity. The origin of the vertical axis is set at the position where the cavity axis intersects the surface of the template sphere. The displayed function is the base ten logarithm of the ratio of the reflected to incident light intensity, with the value of the highest contour and contour separation being equal to 3 and 1.

4.2 Microscopic magneto-optical traps

4.2.1 Magnetic field structure

After magnetisation in a strong external field, provided that any further applied fields are comparatively weak, a bulk ferromagnetic substance has an approximately constant magnetic dipole moment per unit volume. By taking the standard result for the magnetic field produced at position $\mathbf{x}$ relative to a magnetic dipole of moment $\mathbf{m}$:

$$\mathbf{B} = \frac{\mu_0}{4\pi}\left(\frac{3\mathbf{x}(\mathbf{m}\cdot\mathbf{x})}{|\mathbf{x}|^5} - \frac{\mathbf{m}}{|\mathbf{x}|^3}\right), \tag{4.3}$$

which can be found in many standard texts such as [81] or derived by applying the Biot-Savart law to a current loop and allowing the dimensions of the loop to tend to zero, it is possible to numerically integrate over the contribution of all regions of a ferromagnetic substance in order to determine the resultant magnetic field at any given position. This method can be applied to the structures shown in figure 3.11 to derive the form of the expected field.

Fortunately, the fact that the magnetic field strength is inversely proportional to the cube of the distance from a magnetic dipole means that even near a pseudo-infinite plane of identical structures, we can still legitimately ignore the contribution of all structures at large distances from the point in question. This makes it possible to numerically evaluate the magnetic field around a ferromagnetic island in the interior of a very large array of such structures, as contributions to the magnetic field from all islands beyond a certain distance from the point of interest can safely be ignored. In our case the contribution to the overall field strength from islands more than ∼10 trap sites away was found to be small, while (as shown by figures 4.4, 4.5 and 4.6) the contributions made by nearby ferromagnetic islands can be significant.

The results of this are shown in figures 4.4–4.6, which show the form of the magnetic fields resulting from a single ferromagnetic island, an island and its six nearest neighbours and an island and all other islands within ten rows of it in either direction, where these last two assume that the islands are co-centered with optimally packed circles on a plane, as would be the case for structures produced in the way illustrated in figure 3.11. Figure 4.7 shows quantitative results for the axial component of the magnetic field displayed qualitatively in figure 4.6, as a function of on-axis position. In each case, our calculations were based on the use of a 0.7 μm thick sputtered film of $Nd_2Fe_{14}B$, magnetised to saturation before being returned to a low field environment, to form a magnetic island of diameter 60 μm behind the center of each cavity. The dipole moment per unit volume is taken to be 6×10^5 A m^{-1} (this value being the one given in [82] for a sputtered

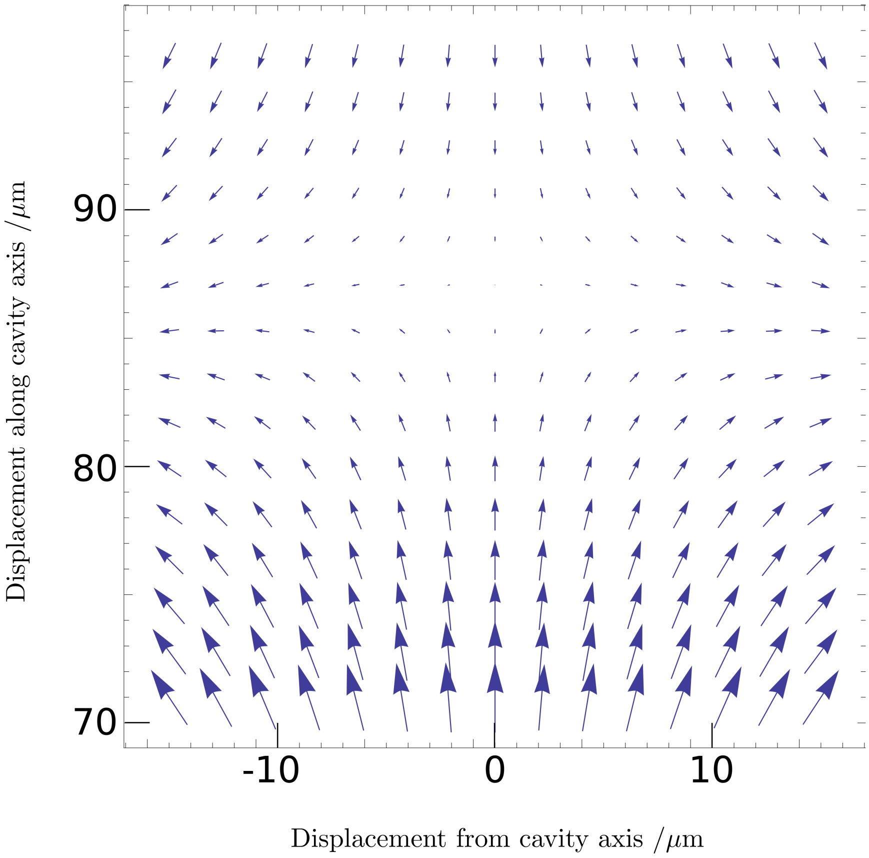

Figure 4.4: Vector plot of the magnetic field structure resulting from a single ferromagnetic island, in a plane containing the cavity axis. The origin of the vertical axis is at the center of the ferromagnetic island and an external bias field of -252 mG is applied along the vertical axis.

film of the same thickness and composition). We also choose to apply an external bias field that places the magnetic field zero about 85 μm above the magnetic island, as the deposition substrate for the process shown in figure 3.11 must be a quarter-wave plate and the thinnest commercially available waveplates are typically of the order of 50 μm in thickness [83] (the remaining 35 μm is the approximate displacement between the top of the substrate and the position where the MOT must be formed in a cavity based on a 100 μm diameter sphere — see figure 4.3).

These figures show that, in the presence of a weak external bias field, a magnetic field

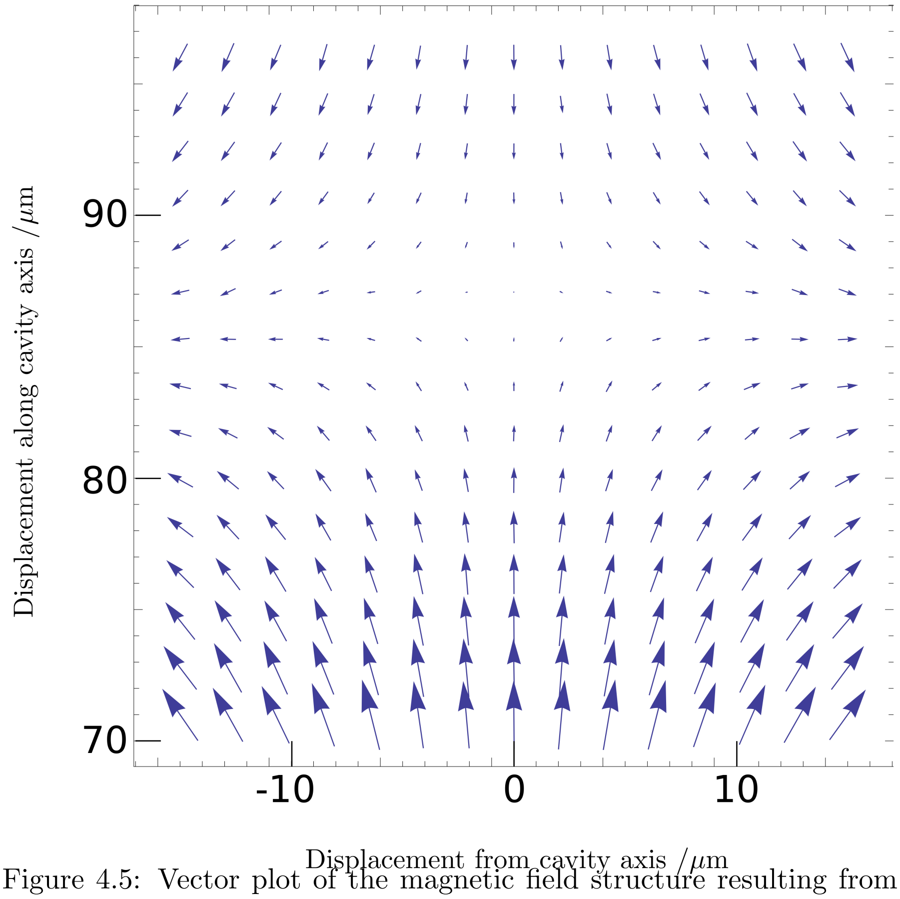

Figure 4.5: Vector plot of the magnetic field structure resulting from a ferromagnetic island and its six nearest neighbours, in a plane containing the cavity axis. The origin of the vertical axis is at the center of the ferromagnetic island and an external bias field of -336 mG is applied along the vertical axis.

of the form required to produce a micro-MOT might be generated by such structures, and that a suitably large field gradient is also achievable with realistic experimental conditions. For example, figure 4.7 shows that the field gradient along the z axis (about the field zero) is roughly 0.25 T m^{-1}, which is comparable with (indeed somewhat greater than) the typical field gradient in a macroscopic MOT.

4.2.2 Loading rates, lifetimes and steady state atom numbers

In order to ascertain what properties we might expect from such microtraps, as well as to determine whether they can be expected to self-load to a reasonable degree or will require loading via atom transfer from a larger MOT, it would be useful to have some idea how the properties of magneto-optical traps scale with trap size. Several

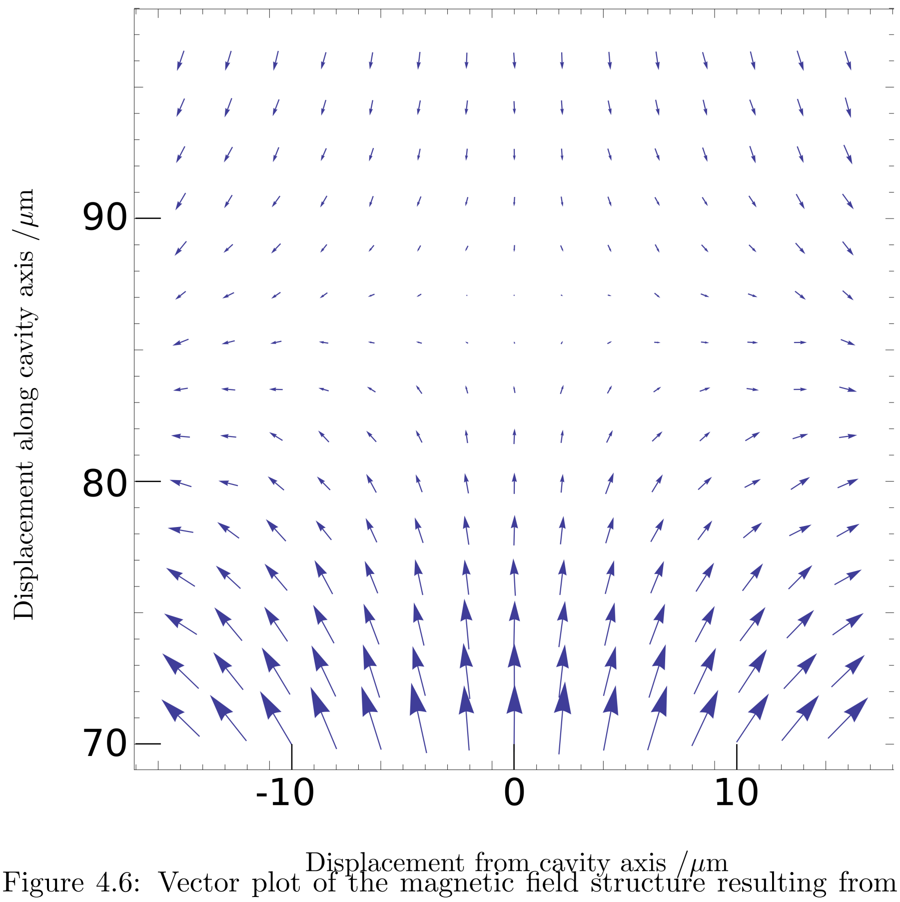

Figure 4.6: Vector plot of the magnetic field structure resulting from a ferromagnetic island and all others within ten rows in either direction, in a plane containing the cavity axis. The origin of the vertical axis is at the center of the ferromagnetic island and an external bias field of -97 mG is applied along the vertical axis.

attempts have been made to model the way in which the loading rates and steady state atom capture numbers of magneto-optical traps vary with the size of the trap. However, existing models often fail to accurately match experimental data without the inclusion of free parameters that are set according to experimentally derived values, and even then are often only valid for a limited range of trap sizes (which our micro-MOTs fall well outside of) — see for example [40]. Furthermore, these models are frequently based on assumptions or approximations for which no rigorous justification is given, such as the assumption used in both [84] and [85] that there is a single 'capture velocity' for all atoms entering the trapping region, which does not appear to be consistent with the fact that the maximum speed at which atoms can be captured will depend on their position and angle of incidence onto the boundary of the trapping region.

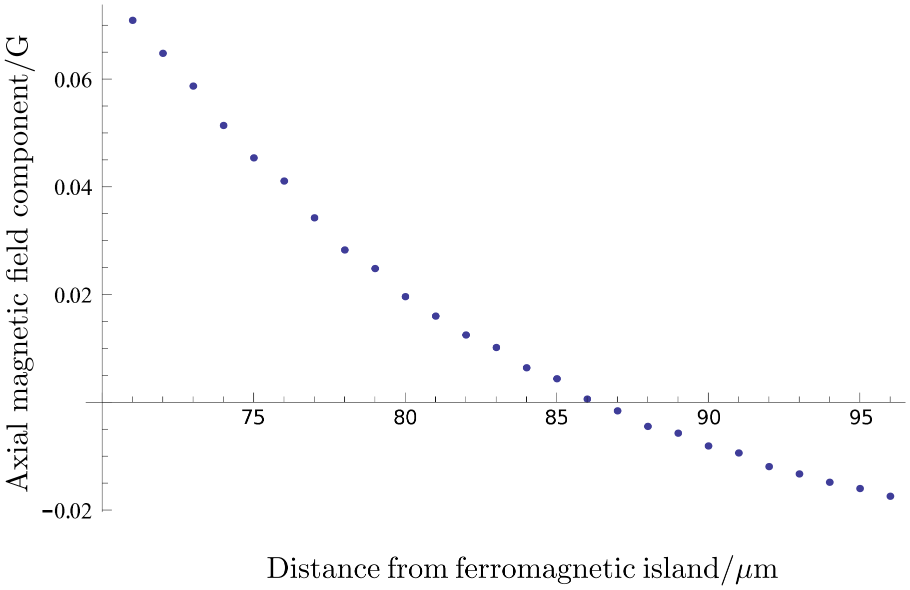

Figure 4.7: Plot of the axial component of the magnetic field (vertical axis) from figure 4.6 as a function of position (horizontal axis) on the cavity axis.

In order to make reasonable estimates of the properties expected from micro-MOTs based on our truncated spherical cavities, we therefore derive our own model that is appropriate to the situation. Fortunately, this can be done quite simply by combining a coordinate substitution with a set of physically justifiable approximations that improve in accuracy as the trap size is reduced. Our model assumes that 'scaling' a MOT consists of modifying the trapping beams' diameters while retaining the same magnetic field gradient about the centre of the MOT.

Almost all models of the loading rate and atom number in magneto-optical traps predict that the atom number scales as a power of the (one-dimensional) trap size that is greater than 3 [40, 85], and this observation is also consistent with experimental evidence [38, 40]. This means that the atom density in a MOT must decrease as the trap becomes smaller. We can therefore legitimately assume that in a very small MOT, effects that increase in importance with larger atom densities, such as many body loss processes and the repulsive inter-atom forces generated by re-radiation of the trapping light, will be of relatively little significance. We therefore expect that, in a very small MOT (with beam diameters ∼1 mm or less), the lifetime of atoms in the trap will be approximately equal to the reciprocal of the loss rate due to collisions with the background gas, and the steady state atom number will be determined by the balance between the loading and loss rates. Therefore, if we can determine how the loading rate of a magneto-optical trap scales with trap size in this regime, we should also be able to derive (to a good approximation) the scaling of the lifetime and steady state atom number of the trap.

In order to determine the scaling of the loading rate, we should consider the equation of motion for an atom in the trapping region. It can be shown theoretically and is generally accepted — see [59], for example — that optical molasses (the term used to describe illumination with balanced, counter-propagating, red-detuned laser beams) induce a slowing force on an atom that is approximately proportional to its velocity. This approximation has been shown to be accurate even for MOTs with beam diameters of several mm, and to improve as the size of the MOT is reduced [38]. Sub-Doppler cooling mechanisms only have a significant effect on atoms that are already confined in the trap, and so do not significantly influence loading rates [86]. Similarly, magneto-optical trapping along any given axis produces a restoring force on an atom that is proportional to its distance from the trap centre[1], and MOTs are often characterised in terms of their 'spring constant' [68].

These approximations are already widely accepted for macroscopic MOTs, and their accuracy improves as the size of the trap is reduced, as smaller traps will not only exhibit a lower maximum capture velocity (atoms moving well above the capture velocity of the MOT will certainly not be captured, and can therefore legitimately be ignored when determining the loading rate) but will also, for a given magnetic field gradient, produce a smaller maximum Zeeman shift on atoms in the trapping region. Since the magnitude of deviations from the two approximations given above can be shown to be at least second order in the atom's velocity or the Zeeman shift experienced by the atom respectively [57, 59], restricting these variables to smaller values clearly improves the accuracy of the approximations.

Combining the forces resulting from optical molasses and magneto-optical trapping, one therefore finds that the one-dimensional equation of motion for an atom near the centre of a MOT is identical to that for a damped harmonic oscillator, with the total force on the atom being the sum of terms proportional to both its displacement and its velocity.

---

[1] For large MOTs with a high atom density, this may no longer be accurate, but it holds well for small traps of the kind we are interested in here [86].

However, in a three-dimensional system the intensities of any beams propagating along a given axis will be functions of the atom's position in the two orthogonal directions, and as a consequence so will the spring constant and damping coefficient along that axis. Therefore, writing an atom's position relative to the trap centre as $\mathbf{x} = (x_1, x_2, x_3)$ and defining the spring constant and damping coefficient along the x_1 axis as $f_1(x_2, x_3)$ and $g_1(x_2, x_3)$ respectively we obtain

$$m\frac{\mathrm{d}^2 x_1}{\mathrm{d}t^2} = -x_1 f_1(x_2, x_3) - \frac{\mathrm{d}x_1}{\mathrm{d}t} g_1(x_2, x_3), \tag{4.4}$$

where m is the mass of the atoms being trapped. Similar expressions can clearly be derived for motion along the two orthogonal axes. Now consider the effect of shrinking the trap by reducing the size of the trapping beams by a factor of ϵ while retaining the same magnetic field gradient. We then have

$$m\frac{\mathrm{d}^2 x_1}{\mathrm{d}t^2} = -x_1 f_1(\epsilon x_2, \epsilon x_3) - \frac{\mathrm{d}x_1}{\mathrm{d}t} g_1(\epsilon x_2, \epsilon x_3). \tag{4.5}$$

Suppose that we express this instead in terms of $\mathbf{y} = \epsilon\mathbf{x}$. We then obtain

$$\frac{m}{\epsilon}\frac{\mathrm{d}^2 y_1}{\mathrm{d}t^2} = -\frac{1}{\epsilon} y_1 f_1(y_2, y_3) - \frac{1}{\epsilon}\frac{\mathrm{d}y_1}{\mathrm{d}t} g_1(y_2, y_3), \tag{4.6}$$

which, when multiplied through by ϵ, is identical to the original equation for the larger trap. This means that $\mathbf{y}(t)$ for an atom entering the smaller trap at time t_0 with position $\mathbf{y}_0 = \mathbf{a}$ and velocity $\dot{\mathbf{y}}_0 = \mathbf{b}$ will be equal to $\mathbf{x}(t)$ for an atom entering the larger trap at time t_0 with position $\mathbf{x}_0 = \mathbf{a}$ and velocity $\dot{\mathbf{x}}_0 = \mathbf{b}$. An atom will become trapped if the magnitude of its position vector remains finite as $t \mapsto \infty$. We can therefore conclude that, if an atom entering the larger trap with position $\mathbf{a}$ and velocity $\mathbf{b}$ would/would not become trapped, then so too will/will not an atom entering the smaller trap with position $\mathbf{a}/\epsilon$ and velocity $\mathbf{b}/\epsilon$.

We now make some approximations relating to the background gas from which the MOT is to be loaded. We wish to assume that the one-dimensional spatial and velocity densities of the atoms in this background gas are uniform. We define the velocity density, or the probability density in velocity space, to be the limit of the probability of finding an atom's velocity to be between v and $v + \delta v$, divided by δv, as δv tends to zero. It is clear from this how to extend the concept to multi-dimensional velocity spaces. The spatial density of the background gas atoms clearly will be uniform, as the thermal background atoms can be very accurately approximated as an ideal gas. The one-dimensional velocity density is not uniform, as this is given by a uniform density of states multiplied by a non-uniform Boltzmann factor. However, provided that the capture velocity of the trap is well below the characteristic thermal velocity of the atoms composing the background gas from which the trap is loaded, so that over the relevant velocity range the Boltzmann factor is approximately equal to unity, the one-dimensional velocity density will be approximately constant. This is true even for macroscopic MOTs (compare, for example, the 20 ms^{-1} capture velocity claimed in [87] with the typical thermal velocities of several hundred metres per second for most atomic species at room temperature), and again the approximation improves for smaller traps. We can therefore conclude that our assumption of uniform one-dimensional spatial and velocity densities for the relevant portion of the background gas is an accurate one.

Consider atoms incident on the larger trap: suppose that all atoms incident at positions between $\mathbf{a}$ and $\mathbf{a} + \delta\mathbf{a}$ with velocities between $\mathbf{b}$ and $\mathbf{b} + \delta\mathbf{b}$ are captured. Given our reasoning above, it follows that all atoms incident on the smaller trap at positions between $\mathbf{a}/\epsilon$ and $(\mathbf{a}+\delta\mathbf{a})/\epsilon$ with velocities between $\mathbf{b}/\epsilon$ and $(\mathbf{b}+\delta\mathbf{b})/\epsilon$ are also captured. As we know that the one-dimensional spatial and velocity densities of the background gas atoms are uniform, we can conclude that the number of atoms incident within such a region per unit time is directly proportional to the size of the region in each of the six applicable dimensions ($x_1, x_2, x_3, \dot{x}_1, \dot{x}_2$ and $\dot{x}_3$). The region corresponding to the smaller trap can be seen to be a factor of ϵ smaller than that corresponding to the larger trap in each applicable dimension, and the total number of atoms incident per unit time will therefore be a factor of ϵ^6 smaller within the region corresponding to the smaller trap than within the region corresponding to the larger trap. As the above reasoning holds for any six-dimensional space–velocity region that can be defined for either trap, the overall loading rate of the smaller trap must be a factor of ϵ^6 lower than that for the larger trap.

We can therefore conclude that, to a good approximation, the loading rate of small (less than about 1 mm in diameter) magneto-optical traps scales as the sixth power of the one-dimensional trap size, or the square of the trap volume. This is consistent with some existing models and experimental results [38, 40], but the above reasoning makes it clear that the validity of this result extends well beyond the region of these existing results, and that this scaling law should continue to be obeyed (with increasing accuracy) as trap size is reduced to the micrometre scale and beyond. The only stipulation here is that the trap size must be reduced by dropping the beam diameters while retaining the same magnetic field gradient.

The trapped atom number should therefore also scale in the same way, as the atom lifetime will be determined primarily by the rate of collisions between trapped atoms and the background gas, and is therefore almost independent of the trap size. There is one caveat with this reasoning: our assumption that the loss rate is nearly independent of the trap size is only valid if the depth of the trap is substantially higher than the characteristic energy of an atom at the cooling limit of the trap. This condition is met in the example below, but for very small traps this may not be an acceptable assumption, especially if the magnetic field gradient is also relatively low. If this condition is not met then significant atom loss may occur, even in the absence of any collisional processes, as the cooling provided by the trapping beams would no longer be sufficient to keep the trapped atoms well below the escape velocity of the trap.

We are now in a position to estimate the properties of a micro-MOT. Let us suppose that, based on a cavity created from a 100 μm diameter template sphere, we form a micro-MOT with a beam diameter of $\sim$10 μm. A typical MOT with 1 mm diameter beams might be expected to capture around 5×10^3 atoms [40], and according to the scaling law derived above we should therefore expect an atom capture number of around 5×10^{-9} atoms. Clearly therefore, these traps are not suitable for independent operation with only self-loading from the background gas, but rather require some enhanced loading process such as the transfer of a cloud of cold, dense atoms from a macroscopic MOT. This fact informs our approach to the experiment in the subsequent chapters.

Chapter 5

The Λ-MOT

As the microtraps we are aiming to create are based on structured, reflective surfaces, our work towards devising a loading method for these traps has necessarily involved seeking to design a system capable of bringing cold atoms close to a structured surface. We have considered several approaches and practical steps have been taken towards one of them — the creation of a cold atom cloud a few mm from the surface via the use of a 'Λ-MOT' [4] followed by ballistic atom transfer via optical launching. This chapter explains the principles of the Λ-MOT, describes an important modification made to it during the course of this work and gives the results of a thorough experimental characterisation of the Λ-MOT in its modified form.

The work on microscopic atom traps described in this thesis is by no means the only research being undertaken on the interactions between atoms and microstructured surfaces. Plasmon enhanced dipole traps have been demonstrated [47], atoms have been used as sensitive probes of surface properties [50, 51] and theoretical work has indicated that proximity to structured surfaces can be used to modify the spontaneous decay rates for some atomic energy states [46]. The ability to bring large numbers of cold atoms close to a microstructured surface is therefore not only vital for the loading of our microscopic atom trap arrays, but also useful in a range of other experiments.

Important incentives therefore exist for the creation of a system in which substantial numbers of ultracold atoms can be brought close to a microstructured surface. However, many conventional methods for bringing atoms close to a surface are not easily applied when the surface is structured on such a length scale. Techniques such as the creation of a 'mirror-MOT' [88, 89, 90, 91, 92] or the loading of atoms into microscopic, surface-based magnetic traps [93] have so far mostly required that the surface in question be a planar reflector. Near-surface dipole potentials created by evanescent optical fields, another common tool for bringing cold atoms close to surfaces [94], require not only a planar surface but also that the substance of which the surface is made does not heavily attenuate transmitted light. Structured surfaces therefore render most existing approaches ineffective. The Λ-MOT represents one element of our proposed solution to this problem, as it allows a dense cloud of cold atoms to be formed a few mm from a microstructured surface. The next chapter shall describe the second stage of this solution, in which atoms from the cloud formed by the Λ-MOT are brought closer to the microstructured surface via ballistic transfer techniques.

5.1 Details of the Λ-MOT

The conventional way to form a MOT close to a reflective surface is to use a mirror-MOT [89]. However, this requires the region of surface closest to the atoms to be an effective mirror, and microstructuring a surface causes light incident on it to be scattered across a range of angles rather than reflected as from a planar mirror. It is therefore not possible to form a trap directly above a microstructured surface region using a mirror-MOT. Instead, the Λ-MOT is based on the geometry shown in figure 5.1, in which beams are reflected from a planar surface either side of a microstructured surface region, with the atom cloud produced directly above the structured area as desired — an image of the resulting atom cloud near a microstructured surface is shown in figure 5.2.

The original implementation of the Λ-MOT, as described in [4], used a pair of anti-Helmholtz coils orientated such that their axis would be along the vertical direction in figure 5.1. The magnetic field structure produced by this coil configuration is shown in figure 2.2. Comparing figures 5.1 and 2.2, it is clear that this results in the direction of propagation of the incident beams (the 'W-beams' from figure 5.1) being almost

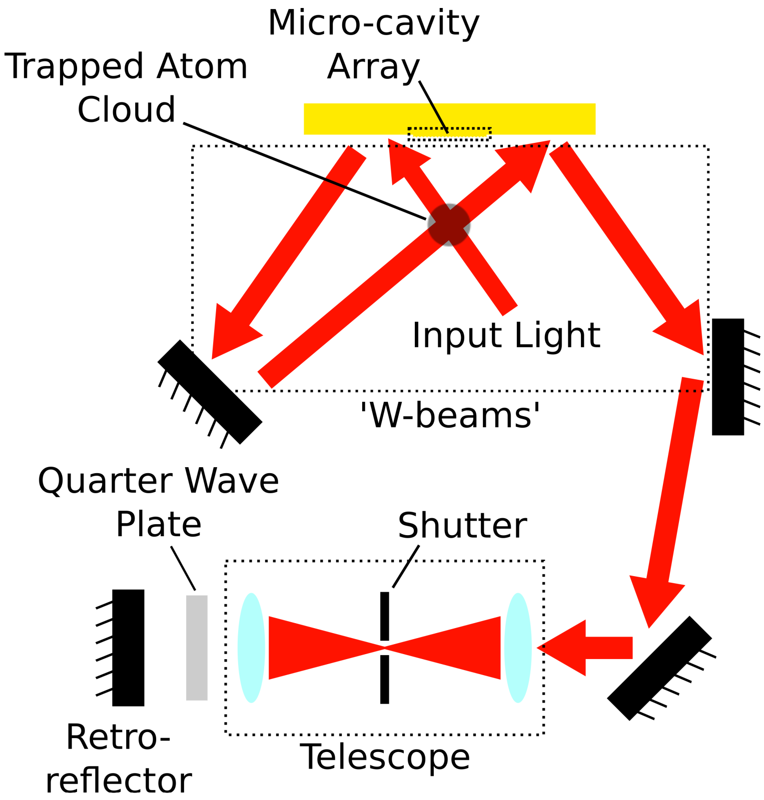

Figure 5.1: Beam geometry used to create a Λ-MOT. The input beam is circularly polarised, with a helicity that depends on the magnetic field structure about the trap centre. In addition to the beams shown, there is one further beam with a direction of propagation that is perpendicular to the plane of the diagram. It passes through the trapping region and then through a quarter wave plate and onto a mirror that reflects it back along its original path, thus allowing this beam to produce cooling and trapping effects along its axis of propagation.

perpendicular to the magnetic field lines around the trapping region, with exact perpendicularity occurring when the beams are incident on the gold mirror at an angle of $\pi/4$. This is not an appropriate magnetic field distribution for the production of a magneto-optical trap, and the efficacy of the magneto-optical trapping is substantially reduced under these conditions, which are believed to be partly responsible for the unusually low atom capture number recorded for the Λ-MOT: with this configuration only around 3×10^4 atoms were captured. This is substantially lower than is typical with magneto-optical trapping — a more normal atom capture number would be between 10^6 and 10^8 atoms, depending on the beam parameters and the density of rubidium atoms in the vacuum chamber.

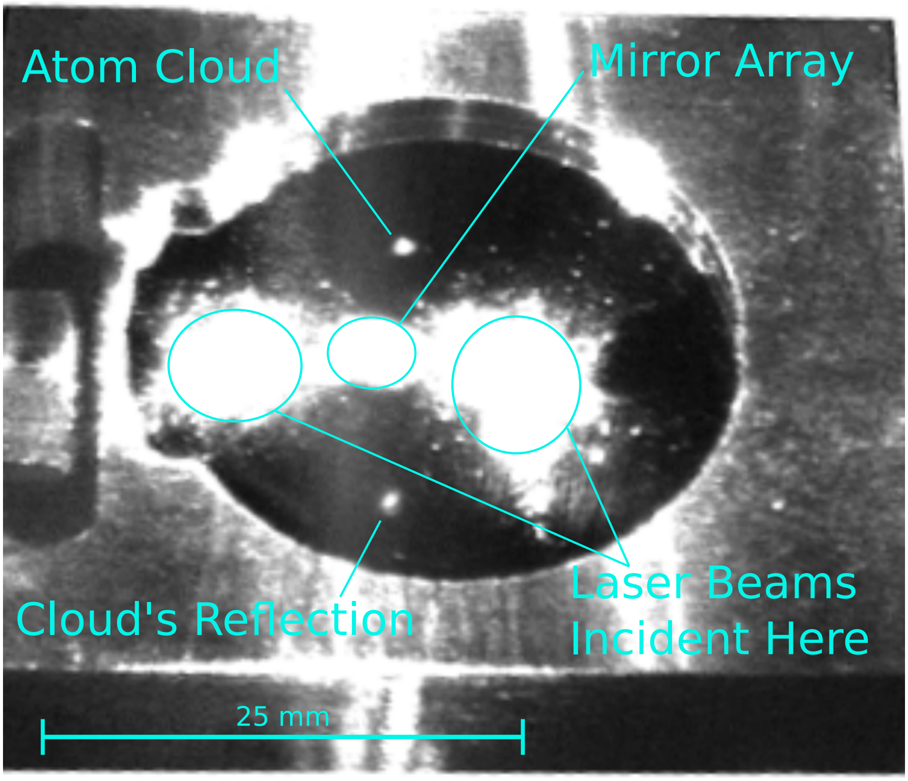

Figure 5.2: Image of the atom cloud formed by the Λ-MOT, showing the position of the atom cloud relative to the microcavity array.

It was decided by the author that this situation should be remedied by rotating the coils so that their axis ran perpendicular to the plane shown in figure 5.1, and this modification was then carried out. This produces a more appropriate field structure for magneto-optical trapping, and using this new arrangement around 2.7×10^5 atoms were ultimately captured (after optimisation), a figure which is substantially greater than that for the original Λ-MOT, though still much lower than would be expected for a trap with the same beam diameters ($\sim$6.6 mm) — see for example the figure of $\sim 4 \times 10^7$ atoms indicated for a MOT with these beam diameters by the results given in [40].

The original coil orientation also created a further difficulty: it appeared to be impossible to produce effective trapping at positions less than $\sim$5 mm from the array. This resulted from the fact that, with the coils in their original positions, the efficacy of the MOT was observed to decrease significantly as the angle of incidence of the W-beams (see figure 5.1) onto the surface containing the cavity array approached $\pi/4$, and (despite several attempts) no trapping effect was observed for angles of incidence $\geq \pi/4$. This can be understood by considering the magnetic field structure about the position of the Λ-MOT, as the projection of the magnetic field vector onto the direction of propagation of the incident beams reduces to zero when the angle of incidence is equal to $\pi/4$ and changes sign as the angle is swept through this value. This implies that a change of polarisation state of the light used in the W-beams would have permitted trapping at angles of incidence greater than $\pi/4$. More importantly, with the coils in their new position, it should be possible to scan the angle of incidence of the W-beams through $\pi/4$ without being unable to achieve trapping at any angle, or needing to alter the polarisation state of the light. Clearly there will be practical limitations on how far this angle can be adjusted: loss of trapped atoms from the MOT due to collisions with the surface, the vastly increased projected area of the beams on the mirror surrounding the cavity array for angles of incidence approaching $\pi/2$ and the reduction in the MOT's spring constant along any axis as the projection of the trapping beams onto that axis is reduced, to name a few. These will place some limit on how close to the cavity array the atom cloud can be formed. However, the minimum displacement of the MOT centre from the surface is nevertheless expected to be significantly less than was thought to be possible with the original Λ-MOT[1].

In order to maximise the number of atoms captured and minimise their temperature it was our initial intention to measure the number and temperature of atoms captured by the MOT as functions of several trap parameters, such as laser detuning, magnetic field gradient, laser power etc. However, it soon became clear that the behaviour of the trap's properties as a function of one trap parameter could not easily be decoupled from the values of any of the other trap parameters, including the alignment of each mirror used to control the relevant laser beams[2]. This rendered such characterisations of the effect of individual parameters on the properties of the MOT almost meaningless, as the same behaviour could only be reproduced if all other conditions were entirely unchanged. The

---

[1] One of our first actions on determining that it would be possible to move the MOT centre closer to the sample would ideally have been to do so. However, the restrictions on optical access to our vacuum chamber imposed by the adjustment of the coil positions (unanticipated at the time at which the chamber was designed) prevented us from doing so, and given the practical time constraints on this work it was not a realistic option to completely re-design and re-construct our vacuum system.

[2] One might predict such a result by considering the physics involved, which produces an expression for the force on an atom with a given position and velocity that is a non-separable function of many trap parameters.

trap was therefore optimised by eye (the fluorescence from the MOT cloud being viewed via a CCD camera) and thorough characterisation was carried out for the specific set of conditions relating to the optimised MOT. It was however considered worthwhile to characterise the variations of the temperature of the atom cloud as a function of laser detuning, since Doppler cooling of the atoms is heavily dependent on the laser detuning but not so drastically affected by other trap parameters as the number of captured atoms or the loading rate, and the variation of the atom cloud's temperature with laser detuning was found to be reasonably consistent across a range of other trap parameters.

5.2 Characterisation methods for the Λ-MOT

Below, we describe the methods by which the number of atoms in our MOT and the temperature of the atom cloud were measured experimentally. In order to allow comparison with other forms of magneto-optical trap on a similar scale, the intensity profile of the trapping beams was also measured in order to allow their width to be accurately quantified. The axial magnetic field gradient in the optimised case was calculated to be $\sim$0.077 T m^{-1}, with small variations about this value having very little effect on the trap's properties.

5.2.1 Trapping beam intensity profiles in the Λ-MOT

The intensity profile of the trapping beams was obtained very simply by attaching a beam blocker containing small circular aperture ($\sim$50 μm in diameter) to a translation stage and moving the aperture through the centre of the beam in a direction perpendicular to its propagation. Recording the optical power that reached a commercial power meter, positioned immediately behind the aperture, as a function of the position of the stage allowed the intensity profile of the beam to be determined. As all our trapping beams are ultimately sourced by splitting a single initial beam, their intensity profiles can be assumed to have approximately the same form on entry into the MOT optics.

The results of this process are shown in figure 5.3. They are close to those expected for a Gaussian beam with a $1/e$ intensity radius of 3.3 mm. Though there are some systematic deviations from a Gaussian intensity profile at large displacements from the beam centre, the structure of these peripheral regions of the beam is not likely to have a significant effect on our experiments. It should however be noted that this is the profile of the laser beam before it enters the optics used for the creation of the MOT, and that substantial clipping of the beams occurs when they are reflected from the limited amount of space that exists between the microcavity array and the edge of the mirror on which the array was fabricated. This reduces the effective beam size, and may be partially responsible for the fact that, even with the coils in a more appropriate orientation, the Λ-MOT still captures many fewer atoms than one would typically expect from a MOT with similar beam diameters. The total time-averaged optical power in the beam was measured to be 12.63 ± 0.04 mW, with the errors in the power measurements resulting from small variations in power that take place on a timescale $\sim$1 s and are thought to be caused by air currents and mechanical vibrations producing changes in the efficiency with which light is coupled into the optical fiber used to convey it to the vicinity of the vacuum chamber. Four measurements of power (each separated by many seconds) were made for each data point and the resulting fractional random error in each power measurement was derived statistically and found to be $\pm0.6\%$. There was additionally a systematic error in all power readings of $\pm3\%$, owing to the uncertainty in the calibration of the photosensor. However, as this effect was consistent between all measurements, it had no significant implications for the determination of the beam profile. The translation stage allowed the relative position of the aperture to be determined to within ±10 μm.

5.2.2 Number of trapped atoms in the Λ-MOT

The number of atoms trapped in a MOT can be determined by measuring the total fluorescence power radiated from the atom cloud when it is illuminated by a high intensity laser beam that is close to resonance with an atomic transition. This illumination must be sufficient to heavily saturate the relevant atomic transition, such that all atoms in the trap can be approximated as scattering photons at a rate equal to half the spontaneous decay rate for the relevant transition. In our case this requirement is met during normal operation of the MOT. Dividing the total scattered optical power by the product of the

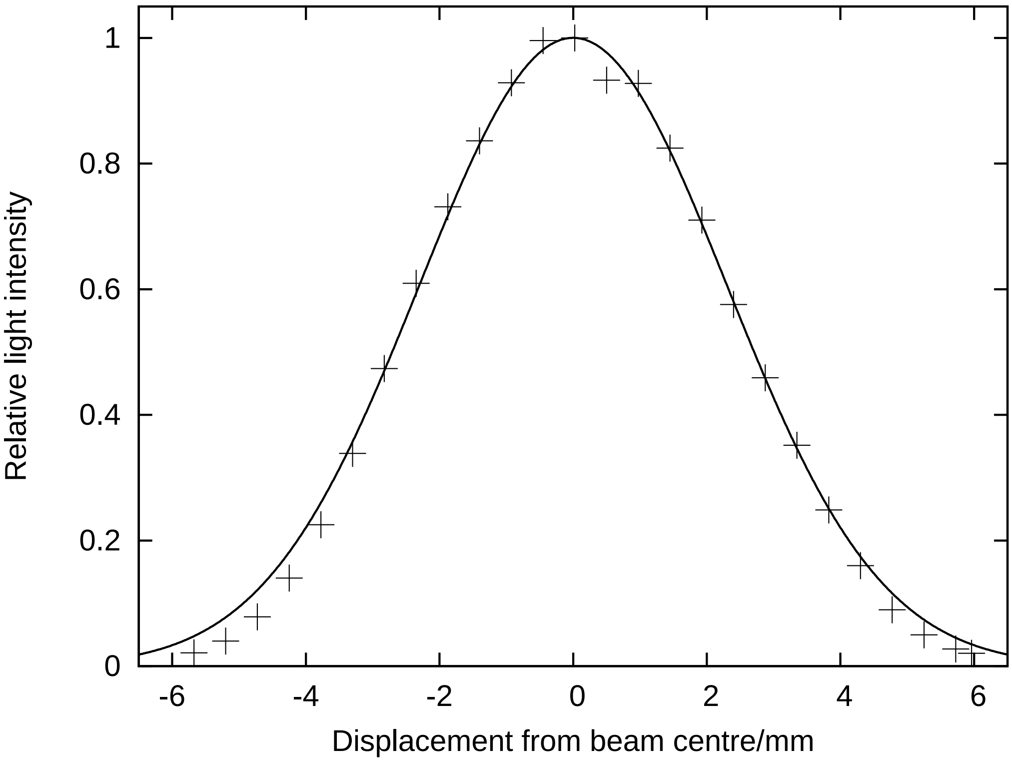

Figure 5.3: Light intensity profile of the MOT beams. The points are experimental data and the solid line is the theoretical result for a Gaussian beam with a $1/e$ intensity radius of 3.3 mm. The position of the beam centre was chosen to fit the data.

energy per photon and the scattering rate per atom then yields the total number of atoms in the MOT.

In our case determination of the total radiated power was carried out by imaging the MOT cloud with an EM CCD camera (Andor iXon 897) and determining the sum of the (background-subtracted) number of counts across all affected pixels. This camera was then calibrated by applying known laser powers at the same wavelength directly to the sensor and measuring the response. The results of this process are displayed in figure 5.4, which shows the background-subtracted count rate against the incident optical power at 780 nm. In order to reduce the effect of background illumination on the measurements, a narrow-band wavelength filter (Thorlabs part FB780–10) was used that heavily attenuated light with wavelengths differing from 780 nm by more than a few nm. In principle the attenuation of 780 nm light by this filter could have been measured and the scaling coefficient between applied power and count rate adjusted accordingly, but as

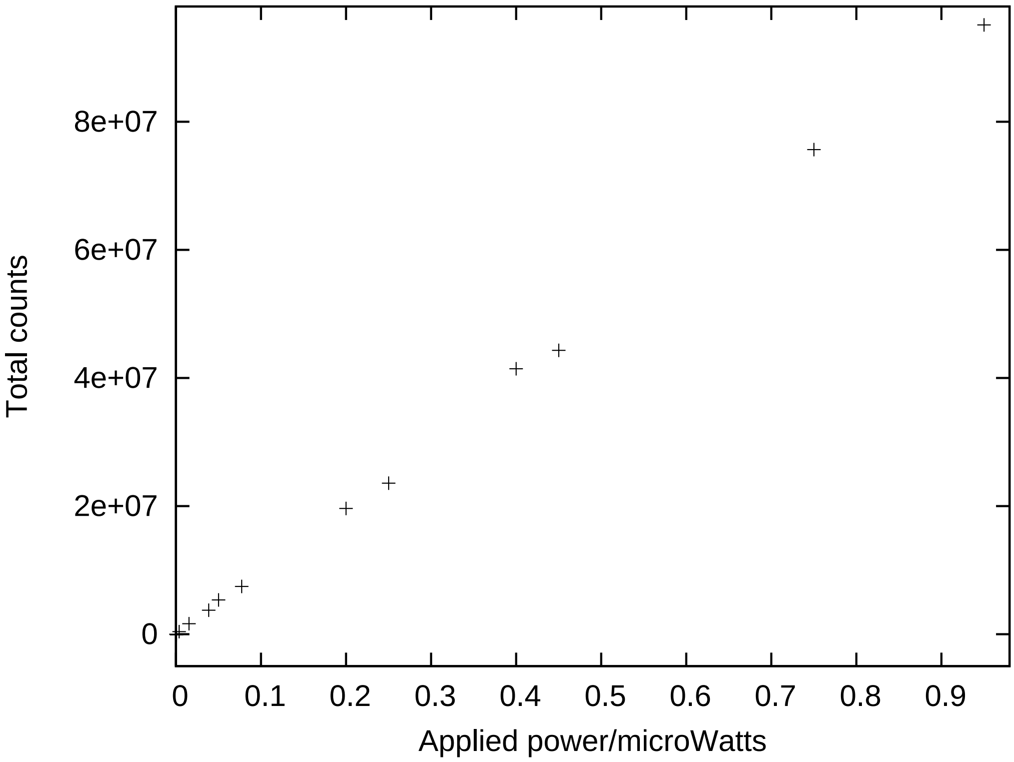

Figure 5.4: Total number of background-subtracted counts recorded across all pixels on the EM CCD camera (Andor iXon 897) against optical power applied to the sensor at 780 nm. The exposure time was 1 ms and the EM gain was set to 10×. It should be borne in mind that the count rates displayed are only valid when the camera is used in conjunction with a 780 nm bandpass filter (Thorlabs part FB780–10).

this same filter was used during observations of the atom cloud it was simpler to ignore its effect in both cases. As no deviations from linearity were found in the response of the camera for a range of illumination intensities that extended well beyond that expected from the atom number experiments, a simple constant scaling factor was determined to convert between count rate and incident optical power; similar checks were carried out to ensure a linear scaling of the count rate with exposure time and EM gain. The power incident on the camera was then divided by the product of the proportion of the total solid angle into which light could be radiated from the atom cloud and reach the sensor with the intensity transmission coefficients of all surfaces between the atom cloud and the sensor, in order to obtain the total optical power radiated by the atom cloud.

The primary source of systematic error in this measurement arose from the determination of the solid angle represented by the detection apparatus from the position of the atom cloud. Though there were several other sources of systematic error, they were small enough to neglect. Figure 5.5 shows how only a small part of the lens used with the CCD camera could be illuminated at any one time, owing to the restrictions placed on optical access to our vacuum chamber by the change of orientation of the MOT coils (the chamber was originally designed to have the coils in the less appropriate orientation described in [4]). The lens-camera system had already been manually aligned as symmetrically as possible into the window of available space, as this maximises the total solid angle available. Measurement of the values of a and b in figure 5.5 was then carried out by fixing the lens and a millimeter scale ruler in place while removing the camera, so that the eye could be directly lined up behind the ruler and used to determine how far along the scale one could move before losing line of sight to the microcavity array that was visible in the centre of the chamber (as the atom cloud in the MOT forms directly below this the same positions will also constitute the limits on line of sight to the MOT). Unfortunately this does not allow quite the same degree of accuracy as reading off a millimeter scale under ordinary conditions, and the measurement error for each reading here was taken to be ±1 mm. The proportional systematic error in atom number $SysErr_{AN}$ is therefore given by:

$$SysErr_{AN} = \frac{1}{\Omega}\left(\left(\delta a \frac{\mathrm{d}\Omega}{\mathrm{d}a}\right)^2 + \left(\delta b \frac{\mathrm{d}\Omega}{\mathrm{d}b}\right)^2\right)^{0.5}, \tag{5.1}$$

where Ω is the solid angle taken up by the exposed region of camera lens when viewed from the position of the atom cloud. Given the large distance from the MOT to the lens relative to the dimensions of the lens, this can be assumed to be approximately proportional to the exposed area of the lens. We have also added the effects of the (independent) measurement errors in a and b in quadrature[3]. The exposed area A in figure 5.5 is given by:

$$A = r^2 \arcsin(a/r) + a\sqrt{r^2 - a^2} + r^2 \arcsin(b/r) + b\sqrt{r^2 - b^2}. \tag{5.2}$$

---

[3] Strictly, the errors in a and b are not completely independent, as some portion of the error in each will result from the error made in positioning the origin of the millimeter scale relative to the centre of the lens. However, this is much smaller than the error made in determining which parts of the lens have line of sight to the atom cloud, and therefore the approximation that these errors can be treated as being independent is a good one.

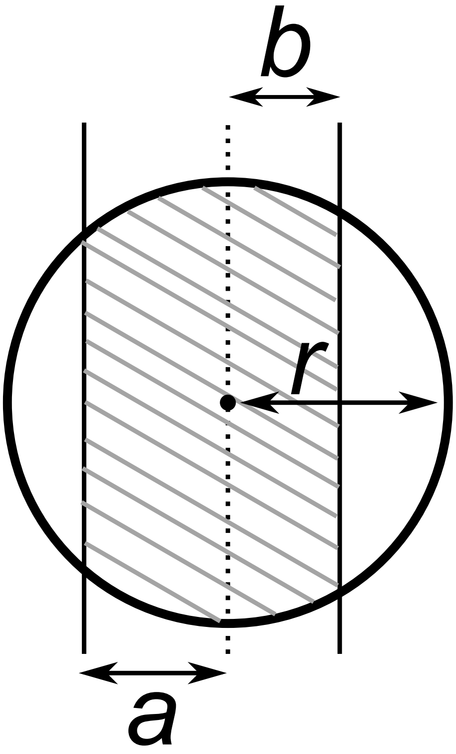

Figure 5.5: Portion of the (25.4 mm diameter) lens from which a direct line of sight exists to the position of the atom cloud (shaded area has line of sight).

Therefore, we can calculate $\frac{\mathrm{d}A}{\mathrm{d}a}$ and $\frac{\mathrm{d}A}{\mathrm{d}b}$ by differentiating this expression. This yields:

$$\frac{\mathrm{d}A}{\mathrm{d}a} = \frac{r}{\sqrt{1-a^2/r^2}} + \sqrt{r^2-a^2} - \frac{a^2}{\sqrt{r^2-a^2}} = 2\sqrt{r^2-a^2} \tag{5.3}$$

and an identical expression for $\frac{\mathrm{d}A}{\mathrm{d}b}$.

By substituting A in place of Ω in (5.1) and using the measured values of a and b of 5.5±1 and 6.5±1 mm, we then obtain the value of the area and its error estimate as 293±32 mm^2. This gives the proportional systematic error in atom number as ±0.11.

5.2.3 Temperature of trapped atoms in the Λ-MOT

Measuring the temperature of the atom cloud formed by a MOT is by no means trivial. There are two primary reasons for this: firstly the actual velocity distribution of the atoms may well be athermal, and so it may be altogether impossible to assign a single-valued 'temperature' to the cloud, and secondly the velocity distribution itself is likely to be a function of position within the cloud. Indeed, some theoretical work has suggested that there are no stable states of a trapped atom cloud in a MOT in which the velocity distribution is independent of position [95]. However, as we are merely concerned with loading as many atoms as possible into our dipole trap arrays, the derivation of an approximate 'temperature', in which we assume the velocity distribution to be both thermal and uniform throughout the cloud, will be sufficient for our purposes.

Using the calibrated EM CCD camera discussed above, we can determine the density of atoms as a function of position at any given time by taking an image of the atom cloud. The distance represented by each pixel on the EM CCD camera is determined by dividing the pixel size by the magnification of the lens system used to form the image. Figure 5.6 shows a sequence of images taken during the expansion of the MOT cloud, of the kind used to experimentally determine its temperature. This free expansion was achieved by 'deactivating' the MOT by using acousto-optic modulators to quickly switch off the laser beams. The magnetic field was left active as inductive effects prevented us from fully extinguishing the current in the MOT coils in a sufficiently short timescale, but in the absence of laser illumination the effect of a field of this magnitude is negligible. We must bear in mind that the camera only images in two dimensions, and that if, as in this case, the depth of field of the imaging system greatly exceeds the dimensions of the atom cloud, a line integral is taken along the third. We can therefore only determine a two-dimensional velocity distribution and, if we wish to determine a full, 3D velocity distribution, we must either make a separate measurement for the third dimension or have some grounds to assume that the atomic velocity distribution in this dimension does not differ significantly from that in the other two dimensions. This last criterion is met here by the assumption of a thermal velocity distribution.

Consider the effect of releasing the atom cloud from the MOT and then allowing unrestricted ballistic expansion of the cloud for a known time, before taking an image of the atom cloud in its expanded state. We define the density of the atoms per unit projected area at a position $\mathbf{r}$ and time t as $\rho_r(\mathbf{r}, t)$ and the probability density in two-dimensional velocity space for an atom at position $\mathbf{r}$ and time t to have velocity $\mathbf{v}$ as $\rho_v(\mathbf{v}, \mathbf{r}, t)$. We

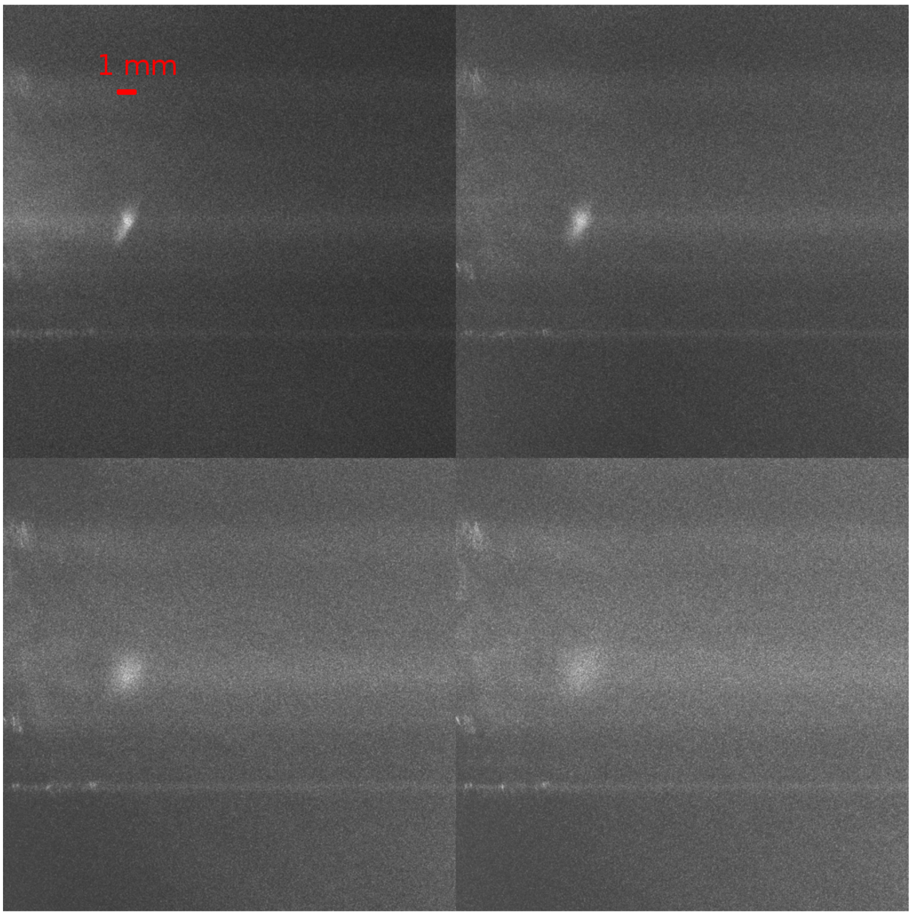

Figure 5.6: Image sequence showing the atom cloud formed by the Λ-MOT after 0, 2, 4 and 6 ms of unrestricted expansion after deactivation of the MOT. The exposure time used to capture the images was 0.2 ms. Here 'deactivation of the MOT' refers only to shutting off of the laser beams — the magnetic field was still applied, but its influence on the atoms was negligible in the absence of resonant illumination.

then find that, if the atoms are released at time $t = 0$ and allowed to expand until time $t = \tau$, we have:

$$\rho_r(\mathbf{r}, \tau) = \frac{1}{\tau^2} \int \mathrm{d}^2\mathbf{r}' \; \rho_r(\mathbf{r}', 0) \; \rho_v \left(\frac{\mathbf{r} - \mathbf{r}'}{\tau}, \mathbf{r}', 0 \right). \tag{5.4}$$

As $\rho_r(\mathbf{r}, \tau)$ and $\rho_r(\mathbf{r}, 0)$ can be measured directly, we can then set about using this result to calculate $\rho_v\left(\mathbf{v}, \mathbf{r}, 0\right)$. It is difficult to proceed from this point without making any approximations. Making the assumption that the velocity distribution is independent of position, and can therefore be written more simply as $\rho_v(\mathbf{v}, t)$, allows one to use the convolution theorem to recover $\rho_v\left(\mathbf{v}, 0\right)$.

In our case however, we instead measure the mean time taken for a trapped atom to expand out of a circular region around the trap centre. This allows us to ignore any atoms that have moved to very large distances from the trap centre, which is advantageous as the atom density at such distances is low and the signal to noise ratio from the relevant pixels on the EM CCD camera is consequently poor. Ignoring regions at very large distances from the trap centre also avoids systematic errors relating to incomplete saturation of the relevant transition in atoms that have moved far from the centre of the illuminating laser beams and are consequently not heavily illuminated.

Experimental determination of the mean time taken for an atom to leave the designated region, or 'mean escape time', was carried out by taking a sequence of images of the atom cloud at incrementally later times after its release and measuring, for each image, the total number of counts recorded on the subset of camera pixels that correspond to the designated circular region. The number of atoms lost from the region between each pair of consecutive images, as a proportion of the number in the initial cloud, can then be obtained. An example of such data, for a MOT with a temperature $\sim$100 μK, is shown in figure 5.7. By taking the product of the proportion of the atoms from the initial cloud escaping between any pair of images with the time since the deactivation of the MOT that corresponds to the centre of this time window, then summing over all such time windows, we obtain an approximation to the mean escape time. The accuracy of this approximation could be improved by linearly extrapolating the escape rate of the atoms between the available data points and then finding the mean escape time by taking the time-integral of the product of the time since the deactivation of the MOT with the escape rate at that time, or by using a smaller time interval between each pair of images. However, in our case there are other sources of systematic error that are more significant than those associated with our experimental determination of the mean escape time, and such a procedure is therefore unnecessary.

We now assume that the velocity distribution of our atoms is thermal and derive a theoretical value for the mean escape time, which we can compare with our experimental value to obtain the temperature of the atom cloud. Consider the time taken for an object to escape from a circular region of radius r. Figure 5.8 illustrates this situation. If the

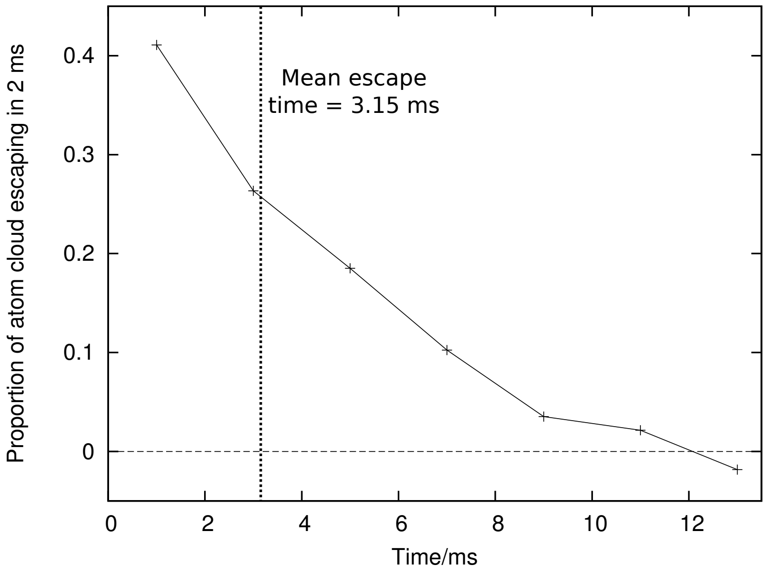

Figure 5.7: Proportion of atoms initially trapped in the MOT escaping from a circular region of radius 0.8 mm about the centre of the MOT during a 2 ms window, as a function of the time since the deactivation of the MOT that corresponds to the centre of the relevant time window.

object is not subject to any significant acceleration, and is released at a distance x from the centre of the region with a velocity of magnitude v at an angle of w to a line joining it to the centre, then the time τ that the object takes to cross the boundary is given by:

$$\tau = \frac{y}{v} = \frac{1}{v}\left(x\cos w + \sqrt{r^2 - x^2\sin^2 w}\right). \tag{5.5}$$

If the velocity distribution is isotropic, which a thermal velocity distribution will be, then there will be an equal probability density for the angle w to be anywhere between 0 and 2π. The mean escape time for an object with speed v, $\overline{\tau}_v$, is therefore given by (due to the symmetry of the situation we can reduce the integral to 0 to π):

$$\overline{\tau}_v = \frac{1}{\pi v}\int_0^{\pi} \mathrm{d}w \ \left(x\cos w + \sqrt{r^2 - x^2\sin^2 w}\right). \tag{5.6}$$

If x is assumed to be zero for all atoms, then this clearly simplifies to $\overline{\tau}_v = r/v$. However, if the atoms are initially distributed with uniform density over the entire region, the mean escape time $\overline{\tau}_{v,all}$ averaged over all atoms in the region with speed v is given by:

$$\overline{\tau}_{v,all} = \frac{1}{v\pi^2 r^2}\int_0^r \mathrm{d}x\ 2\pi x \int_0^{\pi} \mathrm{d}w\ \left(x\cos w + \sqrt{r^2 - x^2\sin^2 w}\right). \tag{5.7}$$

Numerical evaluation of this double integral yields the result that $\overline{\tau}_{v,all} = 0.855r/v$. We must now use this to calculate the mean escape time bounds for atoms with a thermal velocity distribution at temperature T. Since we are interested in the 2D velocity distribution for these atoms, the probability density $\rho(v)$ (per unit speed) for an atom to have speed v is given by:

$$\rho(v) = \frac{m}{kT}\ v\ \exp\left(-\frac{mv^2}{2kT}\right), \tag{5.8}$$

where m is the mass of the atoms. We then find that the mean escape time τ_{thermal} for the atoms in the uniform density case is given by:

$$\tau_{\mathrm{thermal}} = \int_0^{\infty} \mathrm{d}v\ \rho(v)\ \overline{\tau}_{v,all} = 0.855\int_0^{\infty} \mathrm{d}v\ \frac{r}{v}\ \rho(v). \tag{5.9}$$

In the point-like case, $\overline{\tau}_{v,all}$ is simply replaced by r/v. We then find that $\tau_{\mathrm{thermal}} = Ar\sqrt{\frac{\pi m}{2kT}}$, where A is equal to either 1 or 0.855 depending on which limit we are evaluating. Rearranging this yields the temperature as a function of τ_{thermal}:

$$T = \frac{\pi A^2 r^2 m}{2\tau_{\mathrm{thermal}}^2 k}. \tag{5.10}$$

The value derived for the temperature under the assumption of uniform density will therefore be equal to 0.731 times the value of the temperature derived if the atoms are assumed to originate from a single point. The size of the circular region from which we measure the atoms' mean escape time (0.8 mm in radius) was chosen such that very few of the atoms in the initial cloud were beyond the boundary. Although the atoms were spread out within the enclosed region, there was a bias towards increasing atom density at positions closer to the centre of the MOT. The actual atom distribution was therefore somewhere between the two extreme cases described above, and so the values derived

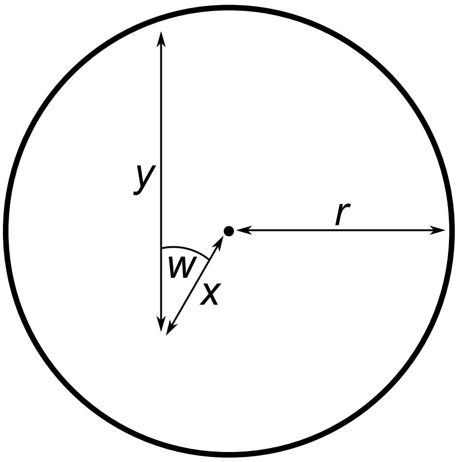

Figure 5.8: Diagram showing the distance y through which an object must move to leave a circular region of radius r when released at a point at distance x from the centre of the circle with a velocity that makes an angle w to the line joining it to the centre. It is assumed that any acceleration undergone by the particle during this time is negligible.

under these assumptions can be used to place upper and lower bounds on the actual temperature value. An appropriate approach to estimating both the temperature and the systematic error resulting from our measurement process is then to take the mean value of the two bounds as the temperature and set our error estimate equal to half the difference between this mean value and either of the bounding values (where we have assumed that the two bounds set on the true value correspond roughly to a 2σ deviation from the mean). We therefore find that the temperature we use will be taken to be 0.866 times that derived under the point-like assumption, and that the proportional systematic error will be taken to be $\pm(0.067/0.866) \simeq \pm 0.08$.

The effect of gravitational acceleration is taken into account by moving the origin of our coordinate system as for an object in free fall, i.e. rather than measuring $\mathbf{r}(t = \tau)$ relative to the same origin as we measured $\mathbf{r}(t = 0)$, we measure it relative to a point that has been displaced downwards from the original origin by a distance $\frac{1}{2} g \tau^2$.

As previously mentioned, the magnetic field used to produce the MOT remains active when the MOT is 'deactivated' by switching off the laser illumination, and its effect on the atoms during their expansion must therefore be considered. The magnetic field gradient across the atom cloud during its ballistic expansion is ~ 0.1 T m^{-1}, and a calculation of the maximum magnitude of the acceleration that any of the atoms could undergo in response to this yields a value of 6.5 ms^{-2}. As our expansion images were taken over the first 10 ms after the release of the atom cloud, the maximum resulting displacement of any atom by its interaction with the magnetic field is equal to 325 μm. This displacement is not insignificant relative to the size of the circular region from which we measure the atoms' mean escape time. However, we are nevertheless able to neglect the effect of the magnetic field on our results, as the continuous pumping of atoms in the MOT by beams with several different directions of propagation and polarisation states means that the atoms in the cloud should be fairly evenly distributed over the available m_f states. As a consequence, approximately equal numbers of atoms should be trapped and anti-trapped about the magnetic field zero. The cancellation between the effects of each on our temperature measurement can therefore be assumed to render the systematic measurement error resulting from the magnetic field gradient insignificant when compared to other sources of systematic error, and we are therefore able to neglect it.

5.2.4 Atom loading and loss rates for the Λ-MOT

The rate of change of the number of atoms N in a MOT is given by:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = L - \sum_{x=1}^{\infty} \alpha_x \int \mathrm{d}V \; n^x, \tag{5.11}$$

where L is the loading rate, n is the density of atoms at any given position within the MOT and α_x is the rate coefficient for x-body loss processes. In the case of the Λ-MOT, even with the coils in their modified orientation, the atom density is low compared to most standard magneto-optical traps. This means that losses are dominated by single atom loss processes, which correspond primarily to collisions with the background gas in the vacuum chamber resulting in the ejection of atoms from the MOT[4]. Other work has found that losses are dominated by background gas collisions even in MOTs that capture many more atoms, at higher densities, than our own [96], and our assumption that this is the dominant loss process for atoms in the Λ-MOT is therefore well justified.

---

[4] 'Single atom' here refers to the fact that only one trapped Rb atom is involved — such processes may well involve collisions with untrapped atoms.

We assume that evaporative losses from our MOT are minimal, as the characteristic timescale of thermalisation processes for atom clouds with densities similar to those found in our MOT is of the order of 1 s [97, 98], while the timescales associated with Doppler cooling of the trapped atoms are much shorter than this [59]. This allows (5.11) to be simplified to:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = L - \alpha_1 N. \tag{5.12}$$

Re-arranging and then integrating this equation before setting $N(0) = 0$ as a boundary condition yields:

$$N(t) = \frac{L}{\alpha_1}\left(1 - e^{-\alpha_1 t}\right). \tag{5.13}$$

By switching on the MOT and then recording the number of trapped atoms against time, it is possible to fit this equation to the experimental data with L and α_1 as free parameters in order to determine their values. We simplify this approach by measuring the atom number in the limiting regimes of very small t (in which losses from the trap have a negligible effect and $N \simeq Lt$) and very large t (in which $N \simeq L/\alpha_1$). In the Λ-MOT, the characteristic timescale associated with loading is expected to be of the order of a few seconds, and we therefore set our 'very small t' measurement at $t = 50$ ms and our 'very large t' measurement at $t = 20$ s.

5.2.5 Measuring laser detuning

In order to fully characterise the Λ-MOT, and in particular to determine how the temperature of the atom cloud varies with laser detuning, we must measure the detuning of our primary cooling/trapping laser relative to the transition it addresses. This laser is stabilised via the DAVLL technique [99], and the exact frequency to which the laser is locked can be modified by applying a constant scaling factor to the signal from one of the two photodiodes involved in the process — it is this frequency that we wish to determine.

The laser system is based on an external cavity diode laser [100], with a diffraction grating that returns some portion of the emitted light to the diode in a wavelength-selective way. The laser frequency can be modified by adjusting the angle of this grating via the use of a piezoelectric transducer. To stabilise the laser frequency, a servo loop adjusts the voltage applied to the piezo to maintain the value of the DAVLL signal at zero. The lock point of the laser (i.e. the frequency at which it is stabilised by the feedback system) can therefore be determined by establishing the wavelength at which the DAVLL signal is equal to zero.

We accomplish this by using saturated absorption spectroscopy (see [57] for a description) to observe a Doppler-free absorption spectrum from an Rb vapour cell. By scanning the laser frequency and recording the level of transmission of the probe beam as a function of time, it is possible to identify times at which the laser is on resonance with particular transitions. Measurement of the spacing of these transitions allows the scan rate of the laser to be determined, and if the DAVLL signal is measured simultaneously a knowledge of the laser's scan rate and the time at which the DAVLL signal crosses zero can then be used to determine the detuning of the lock point relative to its target transition. Figure 5.9 shows an example plot of the saturated absorption and DAVLL signals against time during a scan of the laser frequency, with the characteristic form of the saturated absorption signal, which allows straightforward identification of specific transitions, being clearly visible in the plot. The detuning in this case, which corresponds to the manually optimised MOT condition and is the laser detuning used in subsequent experiments in chapter 6, is approximately -24 MHz.

One source of error in this measurement was the uncertainty associated with visual identification of the position of resonance within each of the observed peaks of the saturated absorption spectrum. Another was the fact that our laser beam passed through an acousto-optic modulator before going to the MOT chamber, and we therefore had to subtract the frequency shift produced by this modulator from the detuning calculated from figure 5.9 in order to determine the laser detuning used in our experiment; the frequency shift produced by the AOM was measured by observing a beat signal between

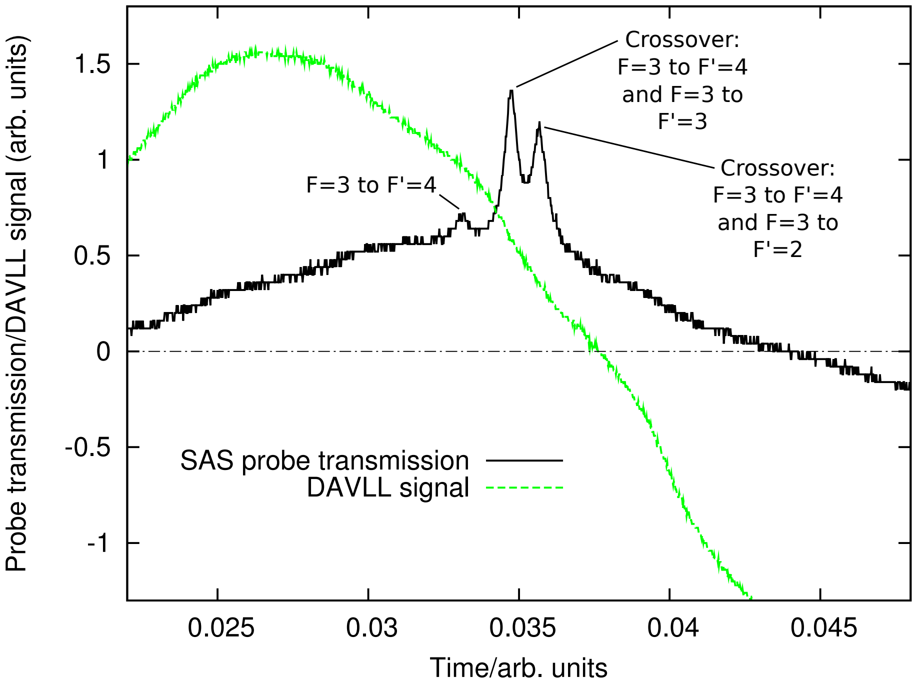

Figure 5.9: DAVLL signal and probe beam transmission (from a saturated absorption spectroscopy setup) during a scan of the laser frequency across the $5S_{1/2} \leftrightarrow 5P_{3/2}$ transition in ^{85}Rb. A frequency shift of ~128 MHz is applied to light from this laser via an acousto-optic modulator before it is sent to the vacuum chamber where our MOT is formed.

the shifted and un-shifted beams on a fast photodiode, and the uncertainty in this measurement also had to be taken into account.

5.3 Characterisation results for the Λ-MOT

The atom number, cloud temperature, loading rate and loss rate for the adapted Λ-MOT in its optimised state were found to be 270,000±30,000 atoms, 89±7.6 μK, 1.9×10^6 atoms per second and 7.1 s^{-1} respectively. This reflects an increase in the number of captured atoms by about a factor of 10 in response to rotation of the MOT coils. More details of these results and the procedures used to obtain them are given below. Multiple results were taken for the optimised case in order to reduce the random error in the calculated mean values. This also allowed the random contribution to the errors in these measurements to be derived statistically, which is fortunate since the causes of random error are too numerous and complex to model accurately. A plot of temperature against laser detuning is also shown in figure 5.10.

5.3.1 Number of trapped atoms in the Λ-MOT — results

The modified version of the Λ-MOT captured 270,000 atoms — about ten times more than had been captured with the coils in their previous orientation. Although this is still more than two orders of magnitude lower than the number typically captured in a standard magneto-optical trap with the same beam diameters [40], and even further from the number captured by a state of the art system such as that described in [70], this discrepancy is to be expected given that the trap geometry was designed so as to enable the trapped atoms to be brought close to a microstructured surface, not to maximise their numbers. The abnormally high rate at which atoms are lost from the Λ-MOT (see below) is also likely to be partially responsible for this result.

For the manually optimised condition, ten separate measurements of the atom number were made with a time separation between each measurement that greatly exceeded the characteristic timescale of atom number variations in the MOT. These measurements gave a mean value of 270,000 atoms and a standard deviation of 30,000 atoms, yielding a random error in the mean of ±10,000 atoms. Adding this in quadrature with the independent systematic error discussed earlier, we obtain a value of 270,000±30,000 atoms.

5.3.2 Temperature of trapped atoms in the Λ-MOT — results

Ten measurements of the temperature of the atom cloud in the optimised MOT, according to the technique previously described, gave a mean of 89 μK and a standard deviation of 8.5 μK, giving a random error in the mean of 2.8 μK. For the optimised MOT condition, we therefore obtain (after adding random and systematic errors in quadrature) a temperature of 89±7.6 μK. Though this temperature is a little higher than that typically quoted for standard magneto-optical traps — see, for example, the figures of 20–140 μK and 20–80 μK given in [68] and [69] — it remains below the Doppler limit for Rb atoms and is therefore consistent with the fact that we would expect our unusual MOT geometry to result in lower than normal levels of sub-Doppler cooling, as most sub-Doppler cooling mechanisms are dependent on the polarisation gradients resulting from the arrangement of counter-propagating beams found in a standard MOT [67].

Without modification of any other trap parameters, the laser detuning was set to a range of values between -8 and -45 MHz (with the detuning in the optimised condition being -24 MHz), and the cloud temperature was measured in each case, by the same method as described above. As only one temperature measurement was made in the non-optimised cases, the random error could not be statistically derived in each case, and so was assumed to have the same magnitude (as a fraction of the measured temperature) as that determined in the optimised case[5]. The results are plotted in figure 5.10, although it should be noted that to better clarify the scale of the relative errors between adjacent points the systematic contribution to the measurement errors has been omitted from the error bars on the plot, and should be taken account of as an additional, independent error contribution if the absolute value of any of the data points is to be considered.

The results are consistent with our expectation that there will be an optimum detuning for the minimisation of the temperature of the atom cloud. The position of this optimum can be seen to be somewhere between 13 and 28 MHz below the $5S_{1/2}$ F = 3 $\leftrightarrow$ $5P_{3/2}$ F = 4 transition, and this is once again a reasonable value to obtain, as the detuning is on a scale comparable to the linewidth of the atomic transition. The random errors in each measurement are substantial, and mean that little precise characterisation can be carried out from this data, but they are nevertheless not sufficiently large that the trends observed could realistically be explained by chance alone.

The optimum detuning range is a little further from resonance than the value of $-\Gamma/2$ expected for the cooling of atoms in low-intensity molasses [59]. However, the low intensity approximation is not appropriate in this case, and we must also allow for the influence of laser detuning on both sub-Doppler cooling mechanisms and multi-atom heating processes such as those investigated in [101]. Given these factors, the region to

---

[5] The random error component could not realistically be derived theoretically, as the factors contributing to the random error were too numerous and obscure to evaluate accurately, and thus a statistical derivation of the random error on each measurement was the only means available to us.

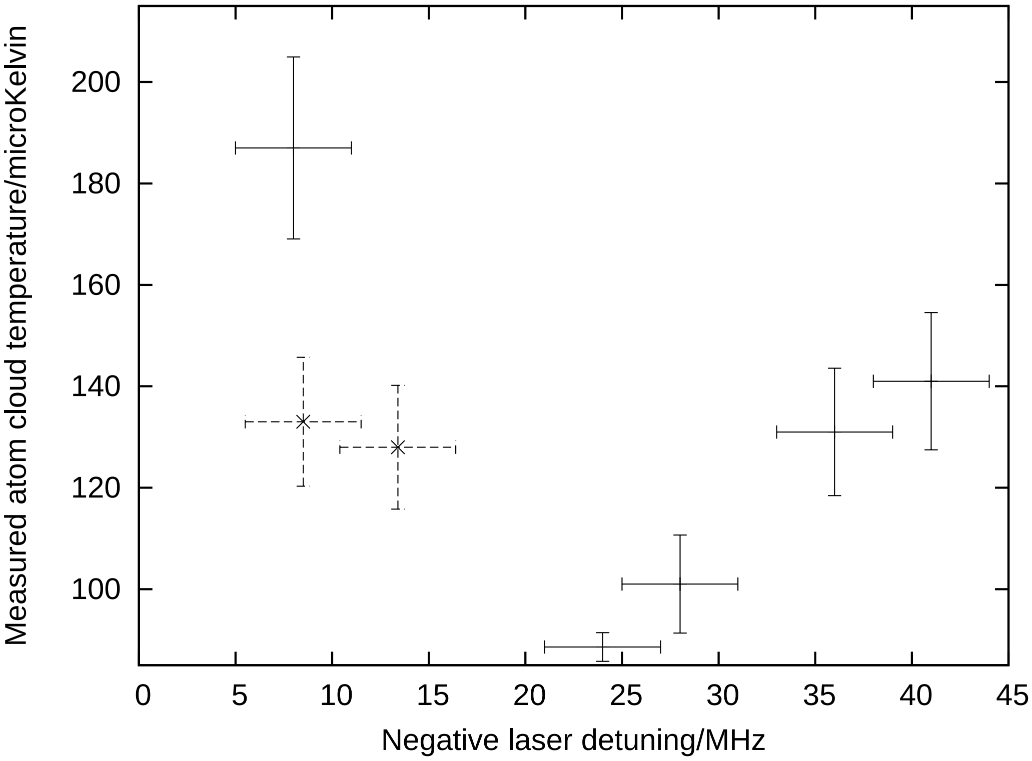

Figure 5.10: Atom cloud temperature against laser detuning in the modified Λ-MOT. The values on the horizontal scale indicate how far below the frequency of the $5S_{1/2}$ F = 3 $\leftrightarrow$ $5P_{3/2}$ F = 4 transition in ^{85}Rb the laser frequency is tuned. Note that only the random error in each temperature measurement is plotted here, and there is additionally a systematic error of ±8% in the overall scaling of the temperature values. The two point styles correspond to measurements made with two slightly different beam alignments, with the more numerous kind corresponding to the 'optimised' condition discussed earlier.

which our experimental results restrict the optimum laser detuning for minimising the temperature of the atom cloud appears to be consistent with basic physical reasoning.

5.3.3 Atom loading and loss rates in the Λ-MOT — results

The loading and loss rates for the modified Λ-MOT were found to be 1.9×10^6 atoms per second and 7.1 s^{-1} respectively. To obtain these results we proceeded in the same way as in sections 5.3.1 and 5.3.2, with the greatest source of systematic error in the measurement of loading rate being the same as that for atom number, and having the same fractional magnitude. There is no significant systematic source of error in the value of α_1, as we determine α_1 by dividing the loading rate by the steady state atom number, which are both ultimately derived from measurements of atom number and therefore share the same fractional systematic error. We therefore determine the magnitude of the error in α_1 by considering only the random errors in our measurements of loading rate and steady state atom number:

$$\Delta\alpha_1 = \sqrt{\left(\frac{\mathrm{d}\alpha_1}{\mathrm{d}L}\Delta L\right)^2 + \left(\frac{\mathrm{d}\alpha_1}{\mathrm{d}N_{ss}}\Delta N_{ss}\right)^2} = \sqrt{\left(\frac{\Delta L}{N_{ss}}\right)^2 + \left(\frac{\Delta N_{ss} L}{N_{ss}^2}\right)^2}, \tag{5.14}$$

where N_{ss} is the steady state atom number and the values of ΔL etc. include only the random component of the corresponding error. The values recorded for the optimised set of conditions were $L = 1.6\times10^6 \pm 1.9 \times 10^5$ atoms per second and $\alpha_1 = 6.0\pm0.36$ per second. The value of α_1 here is somewhat higher than expected, and we have to reconsider the approximation used in determining the loading rate that the fraction of atoms lost from the trap in 50 ms is negligible. The fraction of atoms lost in 50 ms with the value of α_1 determined from these results is $\sim$0.25, and the fact that we have not taken this into account will reduce both the measured loading and loss rates by a non-negligible amount. The values given above should therefore be considered as lower limits on the loading and loss rates, rather than their actual values.

In order to determine the actual values of α_1 and L, we set $t = t_0$ in (5.13), use the fact that we can determine the ratio $L/\alpha_1 = N_{ss}$ from our measurements of N_{ss} and re-arrange the equation to obtain:

$$\alpha_1 = -\frac{1}{t_0}\ln\left(1 - N(t_0)\frac{\alpha_1}{L}\right). \tag{5.15}$$

We can then determine L by using $L/\alpha_1 = N_{ss}$. By setting $t_0 = 50$ ms and using our measured values of the atom number when the MOT had been active for 50 ms, we obtain new values for L and α_1 of 1.9×10^6 atoms per second and 7.1 s^{-1} respectively, and we assume that the fractional error in these values is not significantly different from that calculated previously.

The loss rate for atoms in our MOT is not only higher than expected, but also significantly greater than those quoted by most others who have performed similar measurements — see for example [96]. This is not consistent with the fact that, when the Rb sources are inactive, the background pressure in our vacuum chamber (3×10^{-10} mbar) is somewhat lower than the background pressure quoted in most comparable studies. We therefore conclude that our high single-body loss rate must be accounted for not by collisions between trapped Rb atoms and other background gasses, but between trapped and untrapped Rb atoms.

The higher than expected value of α_1 determined here is of importance in determining the anticipated properties of microtrap arrays based on such a system, as the same loss mechanism will also be present for atoms held in the microtraps. It is desirable for atoms in the microscopic dipole traps formed in such an array to have as long a lifetime as possible, and if this lifetime is limited to less than 200 ms by collisions with the background gas alone then the utility of the system will be greatly reduced. It may prove to be the case that, in order to reduce the loss rate due to collisions with the background gas, a system such as that described in [102] would be desirable, in which the atom cloud is formed in a ‘trapping chamber’ with a high background Rb pressure before being captured in a dipole or magnetic trap and then transferred to a ‘science chamber’ with a lower background Rb pressure. Alternatively a scheme such as light-induced atomic desorption [103] could be used to provide a source of Rb atoms from which to load the Λ-MOT that could then be rapidly deactivated once the atoms have been transferred to the dipole trap array.

Chapter 6

Techniques for loading cold atoms into dipole trap arrays and imaging them once trapped

In this chapter we consider the transfer of atoms from the cloud formed by the Λ-MOT to the microtrap array. Such transfer processes are not only necessary for the loading of atoms into our microtrap arrays, but are also required to make the Λ-MOT a useful experimental system in its own right, as the purpose of the Λ-MOT is to bring large numbers of cold atoms close to microstructured surfaces. We first describe the construction and testing of the systems required for our experiment. This includes the presentation of experimental results showing the efficacy of an 'optical launch' technique that might be used for ballistic transfer of atoms from the cloud formed by the Λ-MOT to the vicinity of a structured surface. We also demonstrate an imaging technique for use on atoms close to microstructured surfaces, that is found to be capable of suppressing the background scatter received from the surface by a factor in excess of 10^6 while reducing the amount of light received from atomic flouresence by only a factor of 10. A brief test of the laser system we intend to use for dipole trapping is then carried out by constructing and loading a macroscopic dipole trap.

Having established the functionality of the experimental systems so far constructed, we consider the future of this research and discuss how one might best proceed with the work. To inform this discussion, we carry out a detailed calculation of the number of atoms we can expect to capture using a purely ballistic transfer mechanism, and conclude that this alone will not be sufficient. Modifications to our experimental system will therefore be required before a positive result can be obtained, and we propose a number of alternative approaches and enhancements to the existing system that represent some of the most promising future directions for this research.

6.1 Ballistic atom transfer from the Λ-MOT to a dipole trap array

Once a cloud of cold atoms has been obtained, these atoms must be transferred to the region where they are required with as little heating/expansion and as few losses as possible. The dipole traps can then be activated and capture any cold atoms currently within their trapping volume. In order to achieve this experimentally, we employ a technique that we call an ‘optical launch’ and find it to be capable of launching some portion of the initial atom cloud (about one third of the atoms) upwards towards the dipole trap array. These atoms should then decelerate under gravity such that they come approximately to rest in the plane of the dipole trap array.

In an optical launch, the return elements of the ‘W-beams’ (see figure 5.1) are blocked using the shutter shown in figure 5.1. The outward travelling elements, i.e. the beams that have not yet been retro-reflected, are then applied for a short time, during which period they impart a net momentum to the atoms in the upward direction. The mean magnitude of this momentum transfer can be calculated by multiplying the scattering force, given by equation (2.1) and §2.1.1.1, by the time duration of the launch pulse. In the limit in which the intensity of the launching light is very much higher than the saturation intensity for the relevant transition (appropriate in our case), it is a good approximation to assume that during the period of illumination the scattering rate of the atoms is equal to $\Gamma/2$. Under this assumption, the centre of mass velocity of a cloud of cold ^{85}Rb atoms launched using a unidirectional pulse with a duration t seconds will be approximately equal to $1.1{\times}10^5(t/\text{seconds})$ m s^{-1}. However, as seen in figure 5.1, our launch pulse is not unidirectional but rather made up of two beams both propagating at an angle θ to the same axis and at an angle of 2θ to each other. The mean velocity imparted on the atoms is therefore multiplied by a factor of $\cos\theta$, which in our case is around $1/\sqrt{2}$. For reasons of experimental convenience we also leave the horizontal beam of the Λ-MOT active during the launch, and as a consequence approximately one half of the spontaneous emission events undergone by the atoms can be attributed to absorption from this beam. This further reduces the expected launch velocity of the atom cloud by a factor of 2.

Using the above, it is possible to calculate the duration of a launch pulse that will result in the centre of mass of the atom cloud coming to rest in the plane of our dipole trap array. For a cloud that is initially 6 mm below the sample surface, a pulse length of $\sim$ 8.6 μs is required. A test launch was conducted using a pulse duration of 0.2 ms, for which the expected launch velocity was around 7.5 m s^{-1}. Figure 6.1 shows the results of this launch. A portion of the atoms from the cloud can be seen rising towards the sample with a constant speed of $\sim$8 m s^{-1}, with simultaneous ballistic expansion. Though this result does demonstrate that optical launching can produce the required ballistic atom transfer, and that the measured launch velocity is consistent with our theoretical predictions, there is one unexplained feature of our result: the fact that the majority of the atoms ($\sim$80%) are not launched.

The fact that there are two discrete groups of atoms rather than a continuous spread suggests that the explanation for this effect may relate in some way to the internal states of the atoms, for which there are number of distinct options, rather than the spatial distribution of the trapping light or the magnetic field structure. One potential cause of such an observation would be the accumulation of a large proportion of the atomic population in the $5S_{1/2}$ F = 2 state during the deactivation of the MOT, perhaps due to a slight difference in timing between the deactivation of the laser beams addressing the $5S_{1/2}$ F = 3 $\leftrightarrow$ $5P_{3/2}$ F = 4 ('cooling') and $5S_{1/2}$ F = 2 $\leftrightarrow$ $5P_{3/2}$ F = 3 ('repump') transitions, the second of which is employed to prevent this effect during normal MOT operation. The atoms in this state would not then be affected by a launch pulse employing only light from the cooling laser. As the repump laser is also active during our launch pulse, one might not expect this explanation to be viable. However, we hypothesised that the mechanical vibrations caused by the operation of the shutter used to extinguish the return elements of the 'W-beams' might be adversely affecting the wavelength-stability of the repump laser, thus preventing it from carrying out its usual role.

To test this hypothesis, we recorded the DAVLL signal from the repump laser's stabilisation electronics as a function of time after the activation of the shutter. We then related this to the detuning of the laser by measuring the gradient of the DAVLL signal with respect to changes in frequency. This was accomplished via the use of a plot similar to figure 5.9, in which the known frequency differences between identifiable atomic transitions allowed the rate of change of laser frequency, during a scan of the voltage supplied to the piezo, to be determined and related to the rate of change of the DAVLL signal. The results of this are plotted in figure 6.2. They show that, for several tens of ms after the activation of the shutter, the frequency of the repump laser deviates from resonance with its target transition by many MHz. Though a similar effect was observed for the cooling laser, its magnitude was found to be much smaller.

Given this result, we expected to observe a drastic increase in the number of atoms launched when the disturbance caused to the repump laser by the activation of the shutter was eliminated. When the shutter was re-mounted in such a way that it had no direct contact with the optical bench to which the repump laser was attached, no significant frequency variations were observed after activation of the shutter. Some increase in the number of atoms launched, to around one third of those found in the initial cloud, was also observed. However, were our original hypothesis as to the cause of this result correct, we would expect the number of unlaunched atoms to be negligible under such circumstances. We were therefore forced to conclude that our original suggestion was not sufficient to fully explain the observed results, and as yet have found no satisfactory explanation for our observations. A search of the literature has not revealed any similar results in the work of others.

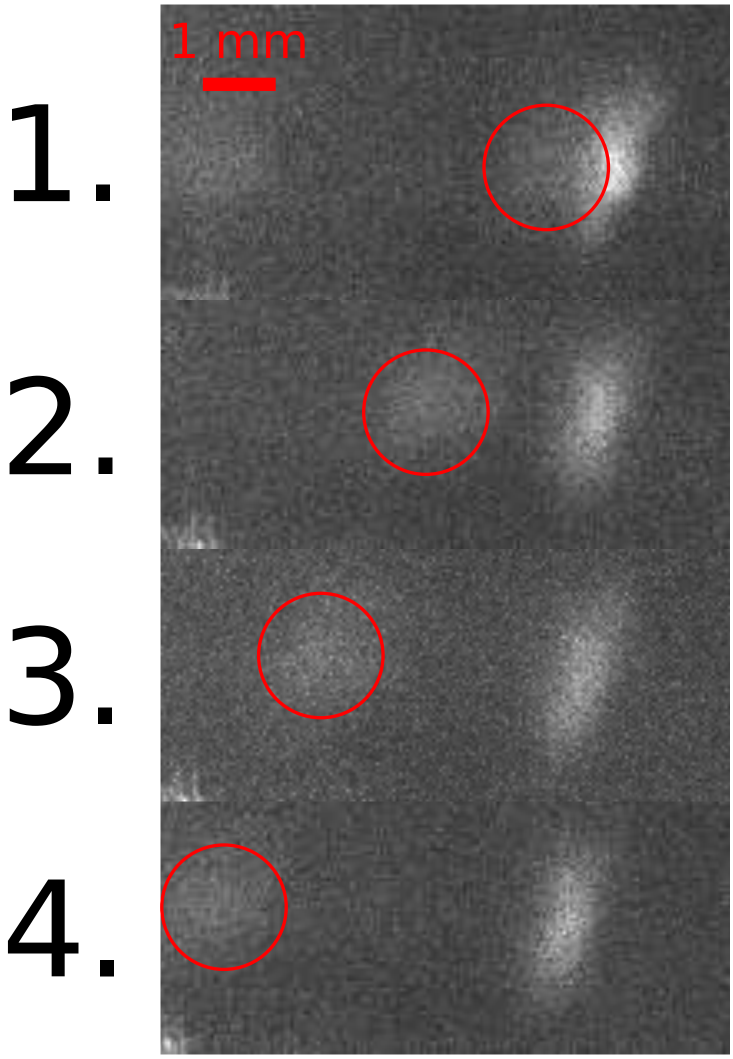

Figure 6.1: Motion of the atom cloud from the Λ-MOT after optical launching. The panels show a time sequence in which a fluorescence image (0.1 ms exposure time, $5P_{3/2} \leftrightarrow 5S_{1/2}$ transition heavily saturated by laser illumination) of the atom cloud is captured with the exposure starting 0,0.2,0.4 and 0.6 ms (top to bottom) after a 0.2 ms laser launch pulse is applied to the atom cloud. It can be seen that the majority of the atoms in the cloud (∼80%) are not launched but remain in their original position. The reason for this remains unknown.

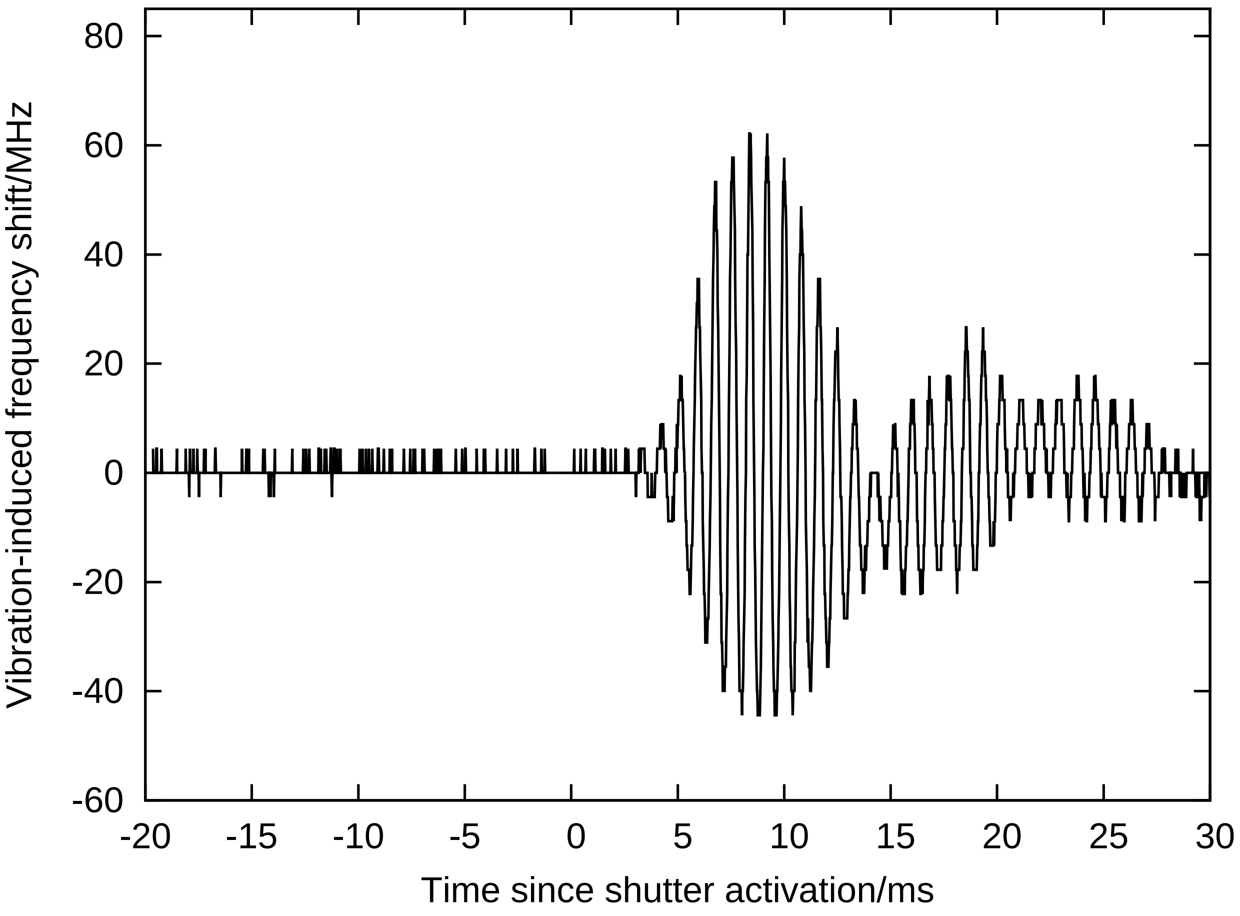

Figure 6.2: Plot of the deviation of the frequency of the repump laser from its set value as a function of time after the activation of the shutter.

The launching pulses also result in heating of the atomic sample via the inter-atom variation of the total number of photons scattered, the number scattered from each of the two beams, and the vector sum of the momenta of the photons spontaneously emitted from that atom. An upper limit to the heating produced by such a pulse can be determined by ignoring Doppler cooling effects and assuming that the kinetic energy gained by each atom is equal to the product of the photon recoil energy, i.e. the mean kinetic energy increase of an atom on spontaneous emission of a photon, which for a particle of mass m emitting a photon with a wavevector of magnitude k is equal to $\hbar^2 k^2/(2m)$, with the mean number of spontaneous photon emissions from each atom during the pulse (see for example the discussion of heating effects in optical molasses given in [59]). For an initial cloud temperature of 88.6 μK (as measured in §5.1), a pulse designed to bring the centre of mass of the atom cloud to rest close to the surface of the sample (when the cloud starts 6 mm below this surface) would not heat the atom cloud to more than ∼100 μK. With high beam intensities additional heating can result from repeated absorption and stimulated emission from different beams, which may cause heating beyond this 'upper limit'. However, if found to be necessary experimentally there is no reason why a longer pulse at lower laser intensity could not be used to reduce the importance of heating by repeated cycles of stimulated emission and absorption to negligible levels.

One potential problem with the use of this transfer mechanism is that the ballistic expansion of the cloud during the transit time is such that the atomic density at the sample surface will be small. For example, an atom cloud with a diameter of 1 mm and a temperature of 100 μK will expand by a factor of ∼6 in each dimension during the 32 ms required to move upwards 5 mm under only gravitational deceleration and come to rest at the top of its trajectory, resulting in a drop in density of about two orders of magnitude. Though we have not yet dealt with this problem experimentally, several means by which it could be addressed are proposed in §6.4.

6.2 Detecting and imaging atoms near structured surfaces

A microstructured surface is liable to scatter incident light across a wide range of angles on reflection. As a consequence, there is no direction from which such a surface can be viewed (if it is illuminated by a laser beam) without picking up significant scatter of laser light from the surface. This is a serious issue in any experiment that involves the use of cold atoms in proximity to a microstructured surface, as scattering from the surface will typically outweigh fluorescence from the atoms by several orders of magnitude.

A solution to this problem is to cause the atoms to emit light at a wavelength with which they are not being illuminated, and which can then easily be separated from the scatter from the surface using, for example, a coloured filter. In the case of ^{85}Rb, this can be done via the scheme described in [104], in which a 780 nm laser and a 776 nm laser are simultaneously employed to excite atoms from the $5S_{1/2}$ state to the $5D_{5/2}$ state via a two-photon transition. From this state, some portion of the atoms spontaneously decay via a pathway that results in the emission of a 420 nm photon. The fraction of the atoms decaying via this pathway can be determined by dividing the spontaneous decay rate for the relevant transition ($6P_{3/2}$ to $5S_{1/2}$) by the total decay rate from the $6P_{3/2}$ state, and is roughly equal to one third. Use of this effect for imaging purposes is described in [4].

Figure 6.3 shows the relevant part of the energy level structure for ^{85}Rb, with the data shown in the figure being sourced from [58, 105, 106, 107]. The scheme described above corresponds to exciting an atom from the ground state via transitions M and B[1], followed by spontaneous decay via transitions A and L. From this it is possible to see that the power emitted by an atom at 420 nm under this fluorescence imaging scheme, even if both the 780 nm and 776 nm lasers can supply enough power to heavily saturate the relevant transition, will be significantly less than that emitted in a standard imaging scheme in which the atoms are both illuminated and observed at 780 nm. There are two main reasons for this: firstly the low spontaneous decay rate between the $5D_{5/2}$ and $6P_{3/2}$ states restricts the rate at which atoms enter the relevant decay pathway, and secondly the poor branching ratio for the relevant transition when atoms decay from the $6P_{3/2}$ state causes only a small proportion of the atoms that do enter this pathway to emit a 420 nm photon. To quantify this fundamental difference in maximum efficacy, we shall assume that the lasers used to drive the relevant transitions perform perfectly and cause exactly equal division of the atomic population between the upper and lower states of the transitions they address. We shall also neglect the effects arising from accumulation of the atomic population in intermediate states, which will be small on account of the fact that all of the intermediate states shown in figure 6.3 spontaneously decay to lower energy states at a rate that is substantially greater than the rate at which they are populated by spontaneous decay from higher energy states. We can then find the ratio of the maximum optical power at 420 nm scattered by an atom under the multi-photon imaging scheme to that scattered by an atom illuminated only at 780 nm. Multiplying the ratio of the spontaneous decay rates for transitions A and M by the ratio of photon energies for transitions L and M and the branching ratio for

---

[1] In reality the tuning of our lasers meant that this step was more accurately viewed as a two-photon transition than two consecutive single-photon transitions.

decay of the $6P_{3/2}$ state to proceed via transition L, we find the overall power ratio to be 0.022:1. In practice (during the experimental work described in [4] and in subsequent experiments within the research group), it was found that the fluorescence rates obtained from this technique were even lower than would be expected in theory, perhaps due to a combination of insufficient optical power at 776 nm and a poor laser locking system that allowed significant deviations of the 776 nm laser from resonance.

To improve on this situation, we took advantage of a change in the availability of cheap laser diodes. At the time of the publication of [104] and of [4], laser diodes operating close to 420 nm were both rare and extremely costly. However, they have recently become much cheaper and more readily available. This means that illumination of our atoms with a laser operating at 420 nm is now a viable option.

Referring again to figure 6.3 and performing a similar calculation to the one previously described, we see that the total power emitted at either 780 or 795 nm from an atom under strong illumination at 420 nm should be approximately equal to 0.17 times the power scattered by an atom illuminated only at 780 nm, or just over seven times that emitted at 420 nm by an atom strongly illuminated at 780 and 776 nm. This significant increase in fundamental maximum efficacy is complemented by an increased ease of implementation (owing to the fact that only one wavelength of illumination is required and that a laser can be locked to this wavelength via spectroscopy on a transition involving the ground state of the atom) and an increased sensitivity of most commercially available cameras, photodiodes etc. at the wavelength(s) of observation.

This imaging technique was tested experimentally on the atom cloud formed by the Λ-MOT, and found to produce an integrated, background-subtracted count number on the EM CCD camera equal to 12% of that obtained for the same trapping conditions with direct illumination at 780 nm. As the filter we employed only allowed 780 nm light to pass without significant attenuation, the contribution of the 795 nm light was lost and a theoretical calculation of the expected result (via the method described above) suggests that we would expect the recorded count number to be equal to 12.5% of that obtained with direct illumination at 780 nm. This is very close to the experimental result, and the observed discrepancy could easily be explained by the fact that we

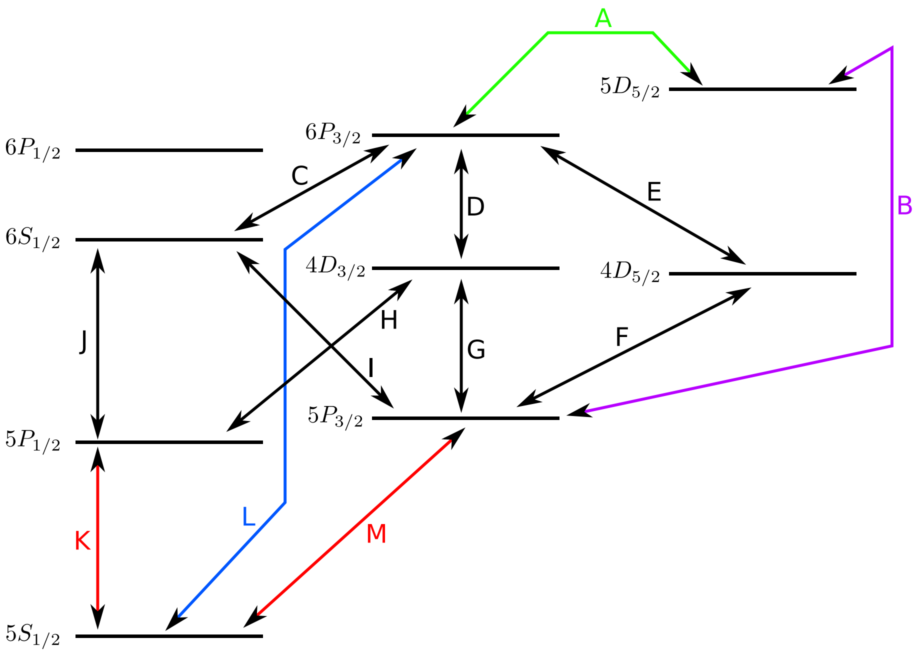

Transition	Rate/MHz	Wavelength/nm
A	1.45	5230
B	2.70	776
C	4.48	2720
D	0.19	2247
E	1.69	2247
F	11.9	1529
G	1.98	1529
H	10.7	1475
I	12.9	1368
J	6.62	1324
K	35.6	795
L	2.80	420
M	37.5	780

Figure 6.3: Full diagram showing all energy levels in ^{85}Rb that are on any possible decay pathway from the $5D_{5/2}$ state that involves only permitted electric dipole transitions. All permitted electric dipole transitions are shown, along with their wavelengths and spontaneous decay rates. To simplify the diagram the hyperfine structure is not shown.

ignored the accumulation of atomic population in intermediate states, which will tend to slightly reduce the efficacy of either of the multi-photon imaging schemes, in our theoretical calculation. The signal arising from direct scattering of the 420 nm laser, from both the atoms and the surface, was reduced to below the noise level, with nominal attenuation of transmitted scatter intensity (based on the manufacturer's specifications for the relevant coloured filters) being in excess of 10^6. Should further attenuation be required, an additional coloured filter could be used, giving a total attenuation of $\sim 10^{12}$ (these filters also attenuate the intensity of the desired imaging wavelength by a factor of ~ 2). A small background signal was detected due to leakage of 780 nm laser light through the (nominally inactive) acousto-optic modulator used to switch the laser beams used for the magneto-optical trap, but this could easily be eliminated via the use of a fast shutter.

One disadvantage of this approach, illustrated in figure 6.4, is that unless balanced, counter-propagating beam pairs are used it is fundamentally destructive, as it involves illumination with a strong, resonant laser beam. When used to image atoms contained in the microtrap array, the strong focusing of the reflected beam will mean that it is not possible to balance the scattering forces produced by counter-propagating beam elements, and the atoms will therefore be expelled from the microtraps during the imaging pulse unless the following two conditions are met: the maximum net scattering force exerted by the illuminating beam(s) must not exceed the maximum dipole force provided by the trapping potential, and the total heating caused by the imaging pulse must not exceed the difference between the trap depth and the initial energy of the trapped atom. The second criterion can be met by using a suitably short imaging pulse. For the first however, some calculation is required to determine whether or not we should expect it to be met when imaging atoms in our microscopic dipole traps. Figure 6.5 shows a theoretically derived plot of the component of the dipole force parallel to the cavity axis (as a function of its on-axis position) for an ^{85}Rb atom trapped in a microscopic dipole trap based on a microcavity synthesised using a 50 μm radius sphere with $\rho = 0.4$ and normally illuminated by 2 mW of optical power at 800 nm. The plot was produced using the same methods described in §4.1, and it can clearly be seen that the maximum force in either direction exceeds 4×10^{-21} N. We need only consider the on-axis force component

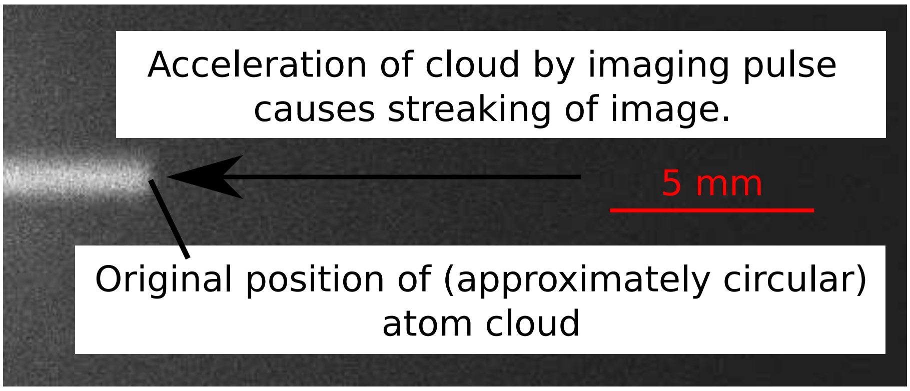

Figure 6.4: Fluorescence image of the atom cloud from the Λ-MOT, with imaging at 780 nm and illumination at 420 nm. The exposure time is 5 ms and the streaking is caused by the rapid acceleration of the atom cloud by the (unidirectional) imaging pulse.

as the symmetry of the system means that the scattering force from a normally incident imaging beam will be entirely axial. Comparing this to the maximum scattering force exerted by unidirectional illumination of the atom at 420 nm of $\sim 1.4\times 10^{-20}$ N, we see that the condition is not met for very strong illumination. We must also take account of the fact that the optical dipole force only has the form calculated in §4.1 when the atom is in the ground state, which will reduce the magnitude of the time-averaged dipole force on a strongly illuminated atom. However, it can be seen that a reduction in illumination intensity, such that the scattering rate is around one tenth of its maximum, will result in a reduction of the scattering force to below one half of the maximum dipole force on an atom in the ground state while simultaneously ensuring that the majority of the atomic population is in the ground state at any given time. The strength of the scattering force will therefore be below the level of the maximum time-averaged dipole force, and we shall make use of this fact in §6.2.2.

6.2.1 Laser stabilisation at 420 nm

The lasers used for magneto-optical trapping and Doppler cooling during the work described in this thesis are locked via the DAVLL stabilisation technique [99]. This method compares the relative levels of absorption of near-resonant laser beams of orthogonal circular polarisation states on passing through a heated vapour cell in a uniform magnetic

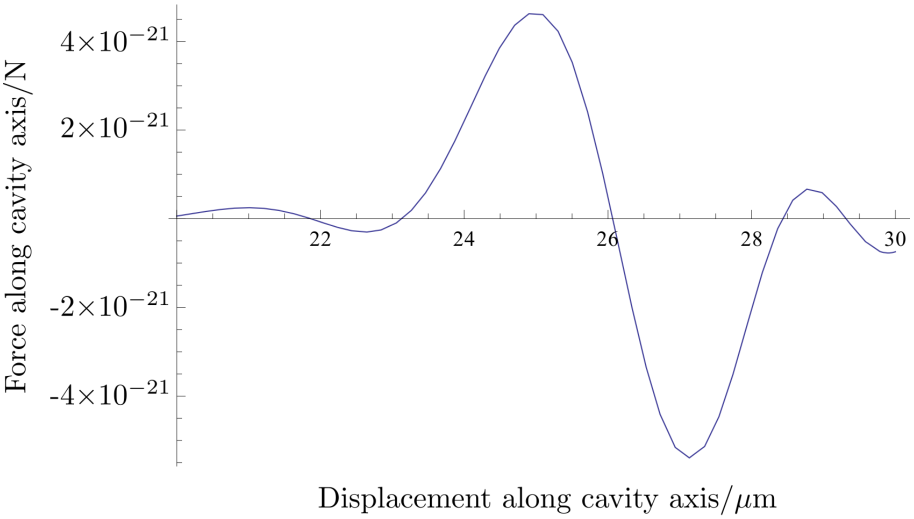

Figure 6.5: Axial component of the (theoretically derived) force on a ^{85}Rb atom, trapped in a microscopic dipole trap based on a microcavity synthesised using a 50 μm radius sphere with $\rho = 0.4$ and normally illuminated by 2 mW of optical power at 800 nm, as a function of its on-axis position. The origin is set at the point where the cavity axis intersects the metal surface.

field. The equal and opposite Zeeman shifts of the resonance frequency for the two polarisation states mean that, by subtracting the signal arising from one circular polarisation state from the signal generated by the other, a resultant signal can be generated that has finite gradient and takes a value of zero when the laser is exactly on resonance with the (unshifted) atomic transition. This signal is therefore appropriate for use in feedback stabilisation of the laser to a wavelength at or near that of the atomic transition.

Experimentally, this technique was found to be somewhat ineffective in the case of the 420 nm laser. Application of a uniform magnetic field to our vapour cell was found not to shift the central frequency of the resonance feature observed with circularly-polarised light, but merely to broaden it slightly [2]. This unexpected result renders the DAVLL technique inappropriate for laser stabilisation in this case, as it is no longer possible to generate a signal of a suitable form for use in feedback stabilisation of the laser to a position near the atomic resonance. An alternative stabilisation technique was therefore

---

[2] This is not a physically unreasonable result, as the Doppler-broadened resonance features used in DAVLL locking a laser contain within them sub-features corresponding to transitions involving several different hyperfine structure states. The net effect of applying a uniform magnetic field is therefore dependent on the Landé g factor of each of the relevant hyperfine states, and if these factors cause different hyperfine states to experience opposing Zeeman shifts when placed in a magnetic field then there may be very little shift in the central frequency of the overall resonance feature.

sought. Techniques based on acousto-optic or electro-optic modulation of the laser beam would have been the most desirable options — indeed a scheme similar to that which we will subsequently describe in §8.2.2.4 would have been ideally suited to this situation. However, as we had no such devices capable of operating at this wavelength, a locking scheme that did not require us to purchase and install additional apparatus was considered preferable. Direct laser current modulation was an option, but would necessarily have unwanted effects on the beam being sent to the experiment, as unlike techniques based on electro-optic or acousto-optic modulation it does not allow one to apply the modulation only to a sample taken from the main beam before the modulator.

We therefore opted for the arrangement is shown in figure 6.6. The 780 nm 'pump' beam is tuned 80 MHz below the frequency of the $5S_{1/2}$ F = 2 $\leftrightarrow$ $5P_{3/2}$ F = 3 transition. The 420 nm laser is tuned close to the $5S_{1/2}$ F = 3 $\leftrightarrow$ $6P_{3/2}$ transition. Atoms that are not heavily affected by the pump beam rapidly accumulate in the $5S_{1/2}$ F = 2 state, in which state they do not scatter light from the 420 nm laser beam. The strongest fluorescence signal is therefore not obtained when the 420 nm laser is exactly on resonance with transitions in stationary atoms, but rather at some frequency between that and the frequency at which resonance with both the pump beam and the 420 nm laser is achieved for the same atomic velocity. The result of this is that there is now a finite gradient in the fluorescence signal from the vapour cell with respect to small changes in the 420 nm laser's frequency about the value at which it is resonant with stationary atoms. This fluorescence signal is therefore appropriate for stabilisation of the laser to the desired wavelength. A similar method was employed in [108], in which it was also used as an alternative to a more conventional stabilisation technique that had been found to be ineffective.

Though somewhat inelegant, this technique was an effective practical solution that enabled the 420 nm laser to be stabilised to the wavelength required for fluorescence imaging. Figure 6.7 shows plots of vapour cell fluorescence against laser frequency for the arrangement shown in figure 6.6, both with and without the inclusion of the pump beam (in practice the data for both plots were recorded simultaneously by splitting the laser beam and sending it to two separate vapour cells). Though there is some noise on

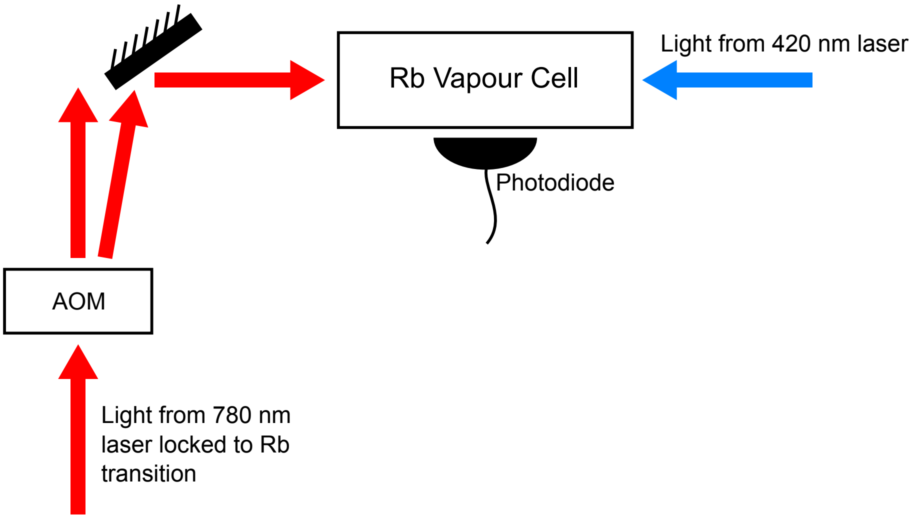

Figure 6.6: The experimental arrangement used for feedback stabilisation of the 420 nm laser to resonance with the $5S_{1/2} \leftrightarrow 6P_{3/2}$ transition in ^{85}Rb. The 780 nm laser beam shown is locked to the $5S_{1/2} \leftrightarrow 5P_{3/2}$ transition of ^{85}Rb, and the beam that enters the vapour cell is then displaced from this frequency by 80 MHz on passing through the acousto-optic modulator (AOM).

the signal, it can clearly be seen that the fluorescence signal takes a form that is both consistent with the above reasoning and suitable for feedback stabilisation of the laser to the atomic resonance.

6.2.2 Detecting single trapped atoms in microscopic dipole traps

As one of the possible uses for such a microtrap array is the creation of an array of individual trapped atoms, with each trap site limited to single occupancy via collisional blockade [44], it is important to consider whether or not our fluorescence imaging scheme will be capable of detecting single trapped atoms. As we are primarily interested in whether or not the experiment is viable, rather than its precise properties, we shall not pay too much attention to effects that may cause small inaccuracies in our estimate. A diagram of the situation is shown in figure 6.8, the features of which are explained in the subsequent text.

For this plausibility calculation, we shall assume the use of an array of microcavities similar to our prototype array (see Fig. 3.6) but with a slightly lower plating depth, such that the width of the opening of each cavity is equal to two fifths of the diameter of

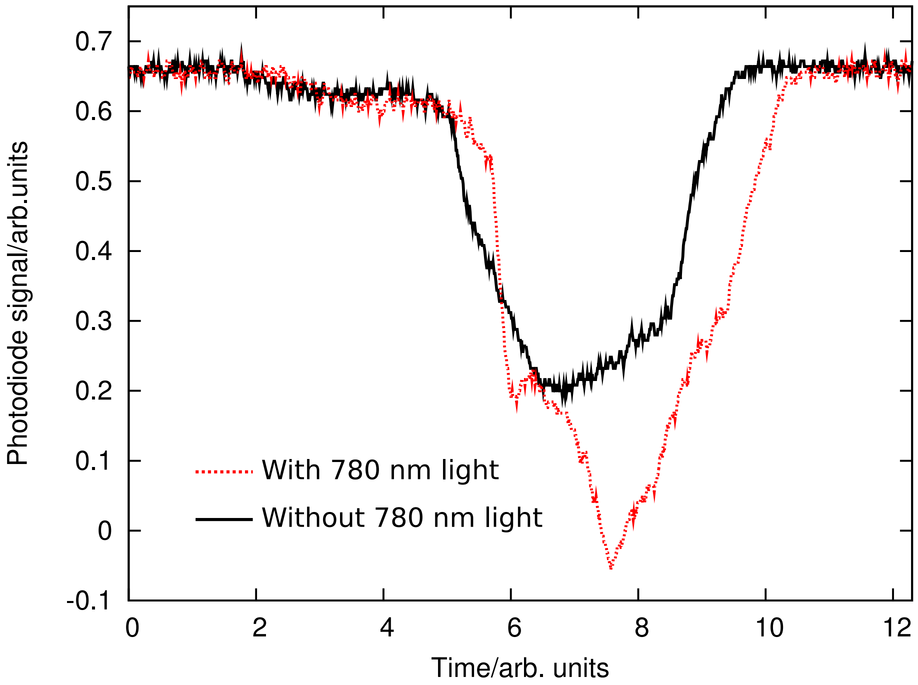

Figure 6.7: Fluorescence recorded on the photodiode shown in figure 6.6 as a function of time during a scan of the frequency of the 420 nm laser, both with and without the inclusion of the 780 nm light shown in the figure. Note that the photodiode output signal decreases with increasing illumination. It can clearly be seen that the inclusion of the 780 nm light results in a finite gradient of the fluorescence signal being recorded about the point at which maximum fluorescence was observed with illumination from the 420 nm laser alone.

the spheres used during the fabrication process. This allows some room for changes to the experimental system that may be made in the future, as a greater plating depth will increase the ease with which trapped atoms can be detected, and so the conditions used for our plausibility calculation will be somewhat less favourable than those actually produced by our prototype microcavity array. All other properties of the array — template sphere diameter of 100 μm, gold surface etc. — are assumed to be identical to those of our prototype array. If we suppose that we are using relatively shallow truncated cavities, in which spherical aberration effects are not especially large, then we can assume that, as the atoms are trapped at the foci of the cavities, all light emitted into the solid angle taken up by the cavity behind any trapped atom will be sufficiently collimated by the cavity that it will ultimately be gathered by the imaging system. This approximation is particularly good if we ensure that a large collecting lens is placed close to the viewport of the vacuum chamber, so that it is only a few cm from the microtrap

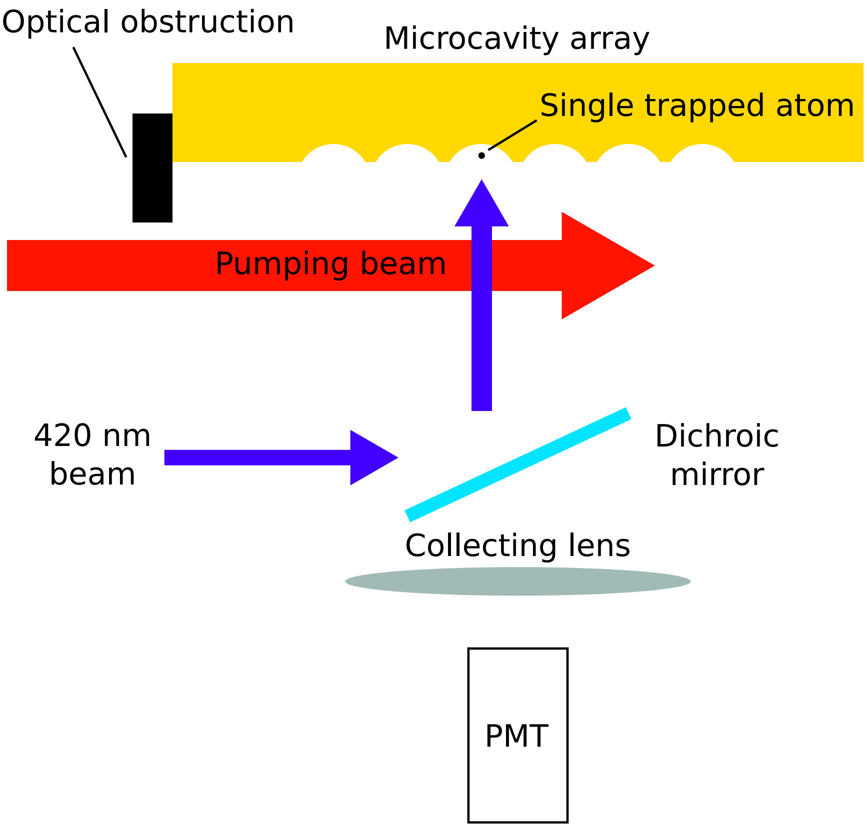

Figure 6.8: Diagram showing the proposed scheme for the detection of single trapped atoms. PMT: photomultiplier tube.

array. The fraction of the (4π total) solid angle that is subtended by a cavity with an opening width ρ equal to two fifths of the diameter of the sphere used to produce it (for an atom at the focus of the cavity) is equal to 0.035, meaning that the product of this with the transmission values of any intervening surfaces is the fraction of the total optical power emitted by a trapped atom that can be expected to reach the detection system. With four surfaces (two for the viewport, two for the collecting lens) each with intensity transmission coefficients of 0.96 (a reasonable figure for standard optics) and a 780 nm bandpass filter that attenuates the desired wavelength by a factor $\sim$2, the proportion of the total optical power reaching our detector will be about 0.015.

We can make a rough estimate of the number of photons emitted per atom before being expelled from the trap by dividing the trap depth by the photon recoil energy at 420 nm, as the atoms will only escape from the trap once their kinetic energy exceeds the trap depth. For this rough estimate we allow a factor of two in the number of photons emitted to account for the initial kinetic energy of the trapped atoms, effectively assuming that the mean initial kinetic energy of a trapped atom will be about half the trap depth. We make the assumption that the atoms remain trapped — i.e. we require that during the imaging pulse the scattering rate is sufficiently low, and the ground state population sufficiently high, that the maximum light intensity gradient in the microscopic dipole traps is large enough to generate a time-averaged dipole force that exceeds the scattering force exerted by the imaging beam. If this were not the case then the atoms would be almost instantly expelled from the trap. In order to ensure that this condition is met, we assume that the atoms are illuminated by a pulse with a sufficiently low intensity that the population of the $6P_{3/2}$ state is one tenth of that achieved under very high illumination intensities (see §6.2). This does not affect the mean number of scattering events undergone by an atom before its expulsion, but it does affect the time interval over which these events occur (it is increased by a factor of ten), and we must take this into account when considering the background count rate on our detector.

For some realistic experimental parameters — say 2 mW of optical power per trap site in the dipole beam at 800 nm — we find that around 5×10^4 absorption/decay cycles per atom can be expected before it is lost from the trap, roughly two thirds of which will result in the emission of a 780 nm photon. Using the experimental result from §6.2 that emission rates at 780 nm with 420 nm illumination are about 12% of those with 780 nm illumination, we can infer that (once we include the factor of ten from the preceding paragraph) the time taken for the emission of these 33,000 or so 780 nm photons will be about 0.15 s. The number of photons reaching our detector in this time will then be around 500.

If we suppose that as a detector we use a photomultiplier tube with a dark count rate of say 10 per second and a quantum efficiency of 10%, then we could expect a single atom to produce around 50 counts in the same time interval that would see an average of 1.5 dark counts. Provided that the optical shielding of the experiment from external light sources is sufficient, it is therefore realistic to expect that the fluorescence imaging scheme described above could detect single trapped atoms in a microcavity.

We must however also consider unwanted light from the experiment. Firstly there is the scattering of the laser beams from the surface. Some small portion of this will be transmitted through the bandpass filter and reach the detector. Estimating the magnitude of this effect is difficult as we do not know what fraction of the light incident on the array will be scattered into the detection optics. However, should direct laser scattering prove to be a significant noise source then an additional bandpass filter can be used to reduce the number of counts resulting from it by an additional factor of 10^6, at the expense of a loss of one half of the desired signal.

Secondly, there is fluorescence from the background gas of Rb atoms in the vacuum chamber. A typical pressure in our vacuum chamber during operation of the MOT is around 10^{-9} mbar, with a background when the Rb sources have been left inactive for a long time of $\sim 3\times10^{-10}$ mbar. The density of ^{85}Rb atoms (assuming that the Rb atoms behave as an ideal gas at 300 K, and multiplying the result for their density by the fractional abundance of ^{85}Rb in naturally occurring Rb) can then be determined to be about 1.4×10^5 atoms per cubic centimeter. It is reasonable to assume that light from regions more than ∼0.5 mm from the site to be imaged (in directions perpendicular to a line between the imaging site and the centre of the collecting lens) could be excluded by physical means. However, the number of atoms remaining in a column of these dimensions running from the exterior of the vacuum chamber to the region of the microtrap array will still be large.

In order to prevent the majority of these atoms from interacting with the 420 nm illumination, it is our suggestion that they could be optically pumped into the ground hyperfine state that is not addressed by the 420 nm laser shortly before atom detection/imaging takes place. Since we do not wish to optically pump the trapped atoms into this same state (or indeed perturb them in any way with the pumping laser), we must leave some small region close to the trap array un-illuminated. Once again, we use a rather generous estimate of 0.5 mm for the size of region we would have to leave unaffected by this beam in order to ensure minimal illumination of the trapped atoms. This gives us an approximate volume of one eighth of a cubic millimeter in which atoms of the background gas can interact with the imaging light and emit into the detection system. Multiplying this by the density of the atoms, we find that we can expect around 20 atoms to be present in the relevant region. These can be expected to produce a much lower total signal than a single atom in a microtrap. There are two factors that lead to this, both of which relate to the fact that the unwanted atoms are not at the focus of a microcavity. The first effect this has is that light emitted by these atoms is not collimated and sent to the imaging system by the microcavity, meaning that the solid angle factor for collection of photons emitted by these atoms is much lower than that for the trapped atoms. The second effect is that the illumination intensity will be much lower — indeed the simulations in §4.1 indicate that we can expect the light intensity at the foci of the cavities to exceed the initial illumination intensity by several orders of magnitude, particularly for larger cavities. As we have assumed already (in order to ensure that the atoms remain trapped) that the illumination intensity is not sufficient to heavily saturate the transition (even at the foci of the cavities), we can assume that the fluorescence rate from an atom is directly proportional to the intensity of the light illuminating it. We can therefore expect fluorescence rates from the background atoms to be lower than those from the trapped atoms by a factor equal to the ratio of the light intensities incident on each — about 150 for a 10 μm diameter cavity with $\rho = 0.4$ illuminated at 420 nm, for example. When these two factors are combined, the result is that the total signal from $\sim$20 background atoms would be expected to be much smaller than the signal from a single trapped atom.

Given this, we can proceed on the assumption that we will be able to detect single trapped atoms in our microtraps, even in the presence of unwanted laser scatter and fluorescence from the background gas. This is consistent with the fact that single trapped atoms have been observed and imaged in optical dipole traps in other systems [26, 44] in which background light scattering is less significant — a result that we could reasonably expect to be able to replicate when using an imaging system that heavily suppresses the effects of background light scattering in our system. In order to test this hypothesis experimentally, we must first construct and load the microscopic dipole traps, and the next section describes an experimental test of the laser system with which we intend to create these dipole traps.

6.3 Experimental test of the laser system to be used for dipole trapping

As a preliminary test of the laser system to be used for our dipole traps, we created a macroscopic dipole trap by using a one inch diameter lens, mounted outside our vacuum chamber, to bring the dipole beam to a focus (with a waist of ∼25 μm $1/e$ intensity radius) that overlapped directly with the atom cloud formed by the Λ-MOT. Unfortunately the power available was fairly limited, as the system had been designed to be used for the creation of microscopic dipole traps based on truncated spherical cavities, in which the focal power of the cavities allows high trap depths to be achieved with relatively modest laser powers. When used without the enhancement provided by these structures, the 200 mW of optical power available to us was not sufficient for the production of a long-lived dipole trap, as in order to achieve a sufficient trap depth the magnitude of the detuning had to be reduced to ∼0.15 THz, which resulted in fairly rapid heating of the trapped atoms due to unwanted scattering of the trapping light. This should theoretically have caused the lifetime of atoms in the trap to be limited by this heating to a timescale of around 80 ms.

The experimental test was conducted by switching off the Λ-MOT with the dipole beam active and then re-activating the MOT beams after some set time delay. These beams were left active for 10 ms to recapture any atoms still inside the trapping volume, following which an image of the remaining atoms was taken with the EM CCD camera. The number of atoms remaining was then determined as explained in §5.3.1, with the background signal resulting from carrying out the same procedure with the MOT initially inactive (so that it is only activated after the variable time delay) being subtracted from each measurement. The experiment was carried out with a range of time delays between 0 and 60 ms. The results are plotted in figure 6.9. In order to ensure that the results do not simply constitute a measurement of the number of atoms remaining in some region after a certain amount of ballistic expansion, we also made an identical set of measurements with a much larger detuning of the dipole laser (∼10 nm), for which the magnitude of the dipole potential was negligible compared to the thermal energies of the atoms. This set of data is also plotted in figure 6.9.

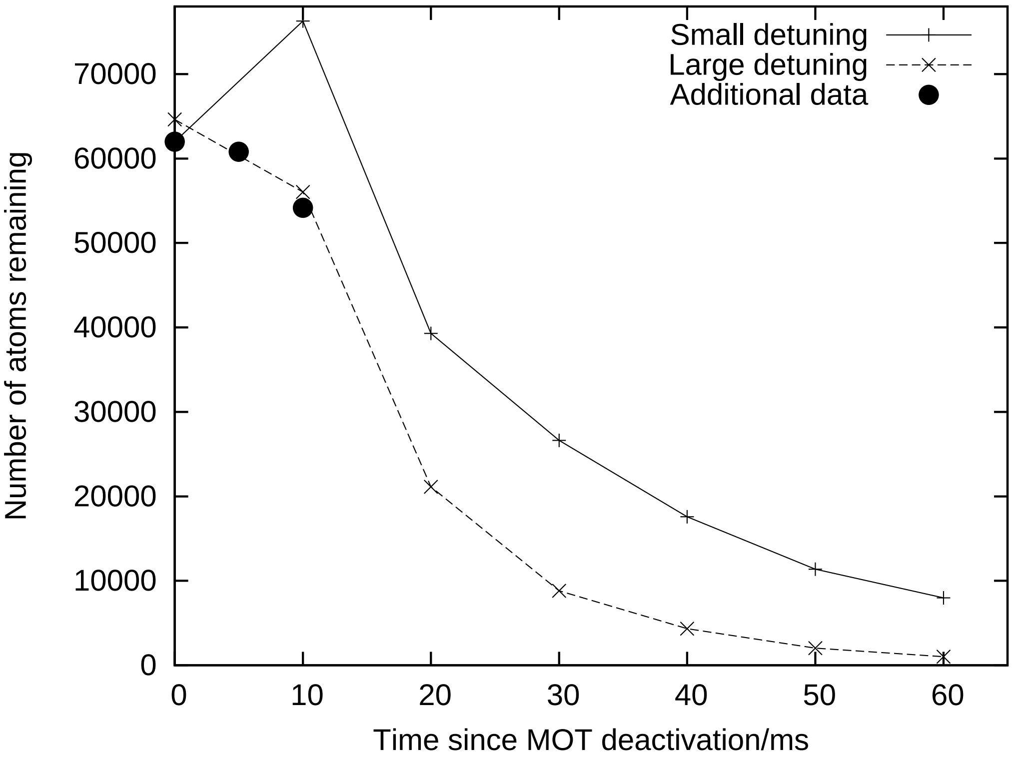

Figure 6.9: Number of remaining atoms against time after MOT deactivation in the presence of dipole beams with detunings of 0.3 nm (‘small’) and 10 nm (‘large’) to the red of the $5S_{1/2} \leftrightarrow 5P_{3/2}$ transition of ^{85}Rb. The beams have powers of ∼200 mW and $1/e$ intensity radii of around 25 μm at the position of the dipole trap. The points labeled as ‘Additional data’ were taken under the same conditions as those in the ‘Small detuning’ series, to ascertain whether or not the increase in trapped atom number seen during the first ten milliseconds of this plot was repeatable.

These results show that the number of remaining atoms does indeed decay more slowly when the magnitude of the dipole potential is significant compared to the temperature of the atom cloud than when it is not, indicating that some dipole trapping is taking place. In fact, the data points at larger time delays in figure 6.9 appear to suggest that there is some residual atomic population in the case of small dipole beam detuning that decays much more slowly than the rest of the atomic population — exactly the result one would expect if some small portion of the atoms in the initial cloud were captured in a dipole trap. The data are not sufficient for an accurate determination of the lifetime of the atoms in the trap, but are sufficient to show that the lifetime of atoms in the dipole trap is not significantly lower than the 80 ms determined theoretically.

One unexplained feature of the results was the initial increase in the number of trapped atoms seen after deactivation of the MOT for the case of small dipole beam detuning. In order to determine whether this was indicative of some genuine physical phenomenon, or whether it was merely a consequence of random variations in the number of trapped atoms in the initial atom cloud, we made repeat measurements of the number of atoms trapped 0, 5 and 10 ms after the deactivation of the MOT. The results of these measurements are shown as the 'Additional data' series in figure 6.9, and they do not show any evidence of the effect previously observed.

6.4 Future directions for work towards loading the dipole trap arrays

In order to determine the anticipated course of work on this experiment, we first question the efficacy of a purely ballistic atom transfer and then, in light of our result, we consider several possible modifications/enhancements to the experiment that may enable more efficient atom transfer.

6.4.1 Estimation of transferred atom number in ballistic transfer processes

The full formula for the number of atoms captured in any given trap site N_{cap} is based on a result similar to (5.4), in which all three spatial dimensions are included in the expansion and $-U(\mathbf{r})$ is the dipole potential resulting from that trap site as a function of position:

$$N_{cap} = \int \mathrm{d}^3\mathbf{r}' f(\mathbf{r}') \tag{6.1}$$

with

$$f(\mathbf{r}) = \frac{1}{\tau^3} \int \mathrm{d}^3\mathbf{r}' \; \rho_r(\mathbf{r}', 0) \; \rho_v \left(\frac{\mathbf{r} - \mathbf{r}'}{\tau}, \mathbf{r}', 0 \right) H \left(U(\mathbf{r}) - \frac{1}{2} m \left| \frac{\mathbf{r} - \mathbf{r}'}{\tau} - g\tau\hat{\mathbf{z}} \right|^2 \right). \tag{6.2}$$

We have taken $H(x)$ as the Heaviside step function and m as the mass of the atoms, and have once again kept the origin of our coordinate system in freefall, but as a consequence of our dipole traps being stationary have had to include a gravitational acceleration term (which we take to be in the negative z direction) in the final velocity of the atoms when determining whether or not they are captured by the trap. Here, we have effectively determined the density of atoms $\rho_r(\mathbf{r},\tau)$ that results from the initial distribution $\rho_r(\mathbf{r},0)$ and then calculated how many of those are trapped when the dipole trap is switched on at $t=\tau$.

To make this problem easier to solve, we shall first assume that the velocity distribution of the atoms is thermal and independent of the atoms' positions. We also neglect spatial regions at large distances from the trap centre, as the dipole potential (and consequently the number of atoms trapped) will be very small in these locations. Furthermore, we ignore the precise details of the initial atom distribution and assume that the atoms are initially evenly distributed with density ρ_r within a sphere of diameter 1 mm — a reasonable approximation given that the cloud is likely to expand to many times its original size during ballistic transfer; for a transfer time of the order of 30 ms as calculated in §6.1, the product of the characteristic thermal velocity of atoms at 100 μK with the transfer time is more than six times the initial cloud diameter. We also assume that the density and velocity distribution of atoms at time $t=\tau$ does not vary significantly over length scales $\sim$10 μm, and that for any given dipole trap site we can assume these to be equal to their values at the trap centre over all relevant spatial regions.

This makes the calculation straightforward, as we are now merely required to determine the spatial density and velocity of atoms at the centre of the dipole trap at time $t=\tau$ and then take an integral over the relevant region to determine the total number of atoms captured, which is found to be equal to

$$\frac{\rho_r}{\tau^3}\int \mathrm{d}^3\mathbf{r}\int \mathrm{d}^3\mathbf{r}' N_{initial} MB\left(\frac{\mathbf{t_c}-\mathbf{r}}{\tau}+\frac{1}{2}g\tau\hat{\mathbf{z}}-\mathbf{v_{com}}, T_{initial}\right) H\left(U(\mathbf{r}')-\frac{1}{2}m\left|\frac{\mathbf{t_c}-\mathbf{r}}{\tau}-\frac{1}{2}g\tau\hat{\mathbf{z}}\right|^2\right), \tag{6.3}$$

where $\mathbf{t_c}$ is the vector of the centre of the dipole trap from the centre of the initial atom cloud, $MB(\mathbf{v}, T)$ is the three-dimensional Maxwell–Boltzmann velocity distribution for particles at temperature T and $N_{initial}$ and $T_{initial}$ are the number and temperature of the atoms in the initial cloud from which the dipole traps are to be loaded (after the effects of any launching process). The origin of our coordinate system is now fixed rather than in freefall, and $\mathbf{v_{com}}$ represents the velocity of the centre of mass of the atom cloud immediately after launching, as our various launching techniques not only heat the cloud but also impart net momentum to it. Depending on the size of the dipole trap, some sensible limits can then be imposed on the integral over $\mathbf{r}'$ to avoid taking it over all space. The limits on the integral over $\mathbf{r}$ are the boundaries of a 1 mm diameter sphere that is co-centred with the initial atom cloud. From this we can obtain realistic estimates for the number of atoms we expect to capture via any given transfer technique, and we can use this information to optimise the technique — for example to select the correct length of pulse during a laser launch in order to optimally balance the higher atom densities that can be achieved with a faster transfer against the correspondingly increased velocity of the atoms on arrival at the dipole trap array.

A calculation of the efficacy of ballistic atom transfer, based on the approach described above (replacing the integrals with appropriate sums to allow the result to be computed in a timely manner), suggests that with 200 mW of power in an 800 nm laser beam illuminating an array of total area of 1 mm^2 of microcavities with $\rho = 0.6$ based on 100 μm diameter microspheres, around 7×10^{-6} atoms per site should be captured. This is not sufficient to make the experiment viable in its current state using a purely ballistic transfer process, and we have therefore considered some potential improvements to our current arrangements that might allow a greater number of atoms to be captured.

6.4.2 Potential improvements to the method of loading via ballistic transfer from a MOT

It should be possible to improve upon the result given above by increasing the atom density in the region of the dipole traps. This can be done by increasing the number of atoms in the original MOT, reducing their temperature or lowering the transfer time between the MOT and the dipole trap array. In practice several of these options may need to be combined before a worthwhile number of atoms can be loaded at each trap site.

Of these three options, the simplest is to increase the atom number. Our Λ-MOT captures at most about 3×10^5 atoms, while atoms numbers of 10^8 or more are not uncommon [86]. The most likely explanation for the comparatively low number of atoms in our trap is the size of the laser beams employed — our beams are initially around 3.4 mm in radius, but can effectively be considered to be somewhat smaller than this due to the unavoidable clipping that occurs on reflection from the limited space between the microcavity array and the edge of the mirror on which it was fabricated. Most MOTs in which very large atom numbers are obtained employ beams with diameters approaching 1 cm. Unfortunately, our experimental apparatus does not allow for any significant increase in beam diameter (due to the aforementioned clipping). However, a re-construction of the chamber and sample holder could allow much larger beams to be used. For an increase in beam diameter by a factor of two, for example, it is not unreasonable to expect a resulting increase in atom number by a factor of between one and two orders of magnitude, as theoretical and experimental work on this subject has shown that the number of atoms captured by a MOT typically scales as a power of the beam diameter that is between 3.6 and 6, depending primarily on the size of the trap [40].

Reducing the temperature of the starting atom cloud is somewhat less trivial. Our proposed approach here is informed by the fact that magneto-optical traps for ^{85}Rb with more standard geometries, that are otherwise very similar to the Λ-MOT (i.e. they have similar beam diameters and intensities, laser detunings etc.), often produce atoms clouds with substantially lower temperatures [68, 69]. We surmise that the geometry of the Λ-MOT is probably not as conducive to sub-Doppler cooling of the trapped atoms as more conventional forms. Given this, we propose that the creation of a MOT that operates on a different atomic transition, with a narrower linewidth, could substantially reduce the temperature of the atom cloud by reducing the Doppler limit. For example, a MOT based on the $5S_{1/2} \leftrightarrow 6P_{3/2}$ transition in ^{85}Rb could be expected to have substantially lower temperatures than our existing MOT, owing to a reduction of the Doppler cooling limit by more than a factor of ten. Similar schemes have been employed successfully with other atomic species [109].

In order to lower the transfer time, the deceleration of the atoms would have to exceed that due to gravity alone. The most practical method for doing this is via the application of an additional magnetic field. The laser launching process could then serve the dual purpose of imparting some initial velocity on the atoms and pumping them into the desired Zeeman sublevel (as only the population of one sublevel can be efficiently transferred via this technique). Application of a potential to the atoms via an external magnetic field could then increase deceleration rates to substantially above g, thus reducing the transfer time and hence lowering the amount of ballistic expansion undergone by the atoms during the transfer process.

Alternatively, the construction of microscopic magneto-optical traps of the kind previously discussed, such that the centres of these traps are coincident with the centres of the dipole traps in the array, would allow atoms to be cooled into the traps, meaning that atoms which moved into the trapping region while the trap was active could be collected and trapped. This should be somewhat more effective than merely trapping those atoms that happen to be inside the trapping volume at the moment of activation, and could also produce an increase in the number of atoms trapped per site.

6.4.3 Alternative approaches to dipole trap loading

Should further work be undertaken on this experiment, then instead of attempting to improve the existing loading method, it may be worthwhile to pursue an alternative. Informed by a brief consideration of any existing techniques that may be relevant, we discuss three possible alternatives.

6.4.3.1 Existing techniques

Before attempting to devise novel solutions to a problem, it is worth considering whether or not any existing methods may be applicable, and a search of the literature reveals that many experiments that involve either loading cold atoms into microscopic dipole traps or bringing them close to a surface have indeed been reported. However, remarkably few of these are applicable to our experiment. For example, many such experiments involve the loading of dipole traps that can simply be directly overlapped with a MOT [27, 110, 35], or bring atoms close to surfaces that are optically flat and allow transmission of a non-negligible proportion of the incident optical power, such that evanescent fields can be used to assist with the confinement of the atoms [111]. Other techniques that have been demonstrated are only capable of manipulating a few atoms at a time [112], and are therefore not appropriate for the loading of an entire microtrap array.

However, some relevant work has been reported, and is mentioned in the appropriate sections below. The most numerous reports of relevant procedures are those that involve the use of surface-based structures to assist with the confinement and manipulation of cold atoms via the production of strong magnetic field gradients near the surface, and this suggests that the approach discussed in §6.4.3.4 may be one of the most promising.

6.4.3.2 Dipole beam loading

Dipole beam loading involves the use of a dipole potential to enhance atom transfer from the Λ-MOT to the relevant region. There are two main ways in which this could be done: either the atoms could be loaded into a dipole trap which is then moved close to or into the plane of the microscopic dipole trap array (as in figure 6.10) or the atoms could be contained during ballistic transfer by a cylindrical dipole beam that runs along the direction of transfer (as in figure 6.11).

Transfer via a moving dipole trap does hold some promise. Atoms could be loaded directly from a magneto-optical trap into a mobile dipole trap that is then translated closer to the surface. The atoms could not be transported all the way into the plane of the dipole trap array via this method, as a sufficiently tight focus to confine the atoms in a position where they could enter the microscopic dipole traps but could not make contact with the (room temperature) surface about ten microns above their position would necessarily result in a Rayleigh range for the dipole beam that was on the order of a few tens of micrometers or less, and this approach could not therefore be used to

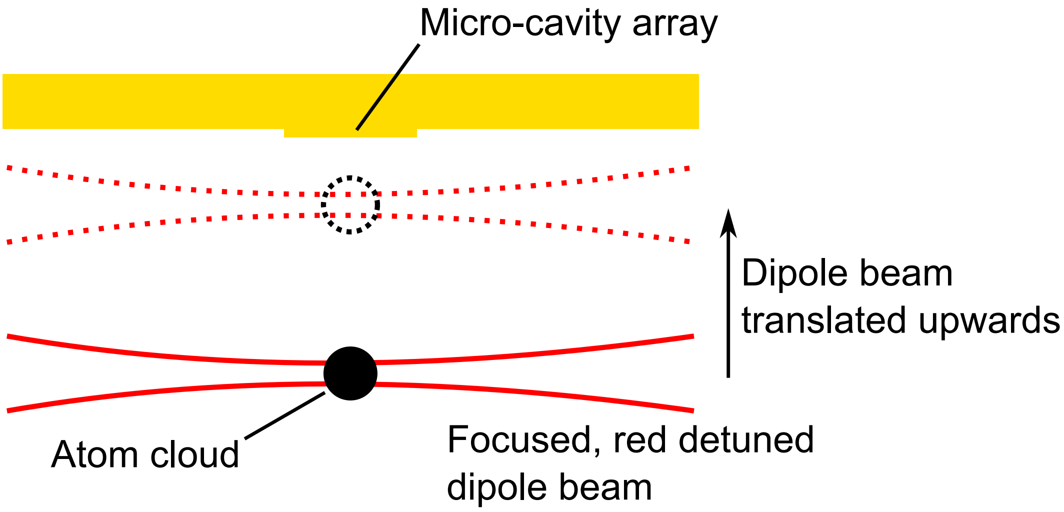

Figure 6.10: Method by which loading of atoms into the micro-trap array could be facilitated. Atoms from the Λ-MOT are captured in a dipole trap and then moved closer to the position of the array before ballistic transfer.

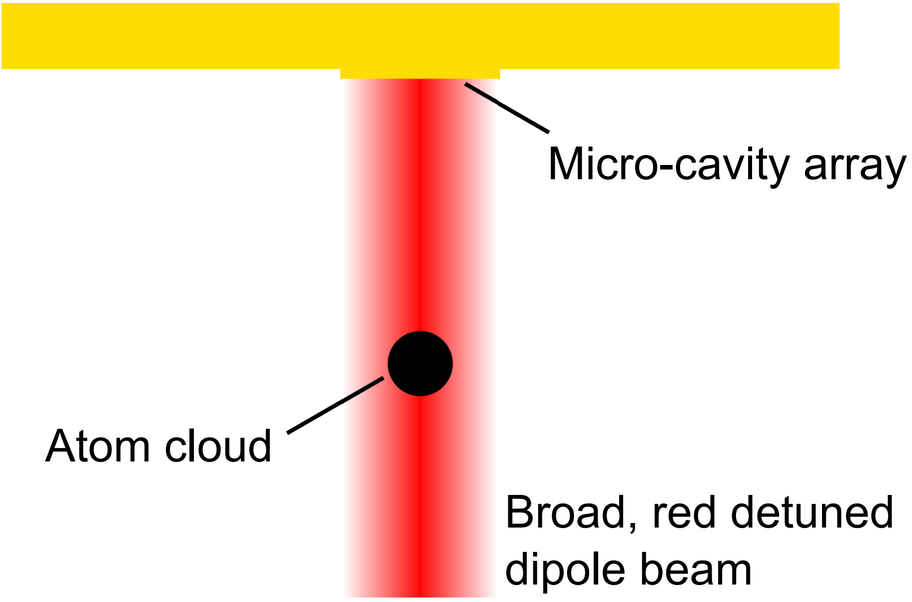

Figure 6.11: Method by which loading of atoms into the micro-trap array could be facilitated. Atoms from the Λ-MOT are contained by a cylindrical dipole beam during ballistic transfer, which prevents significant expansion of the atom cloud in the plane perpendicular to the dipole beam.

simultaneously load an entire array of microscopic dipole traps. However, transfer via a moving dipole trap could still be a useful option, as the atoms could easily be brought much closer to the sample surface than the Λ-MOT can before release — a distance of around half a millimeter is not unrealistic, which compares very favourably with the 6 mm displacement of the Λ-MOT from the surface. The atoms could even be accelerated to the required velocity for ballistic transfer over the remaining distance by the motion of the dipole trap, thus avoiding the necessity for any optical launch procedure.

Guiding the atoms with a cylindrical dipole beam is also theoretically sound, and is in fact very similar to the method demonstrated experimentally in [113], in which atoms that are horizontally confined by a cylindrical dipole beam are moved into the vicinity of a dielectric surface via the use of a magnetic potential. A similar procedure is also reported in [114], although in this case the atoms are merely guided through free space rather than brought close to a surface. The work described in [113] and [114] suggests that atoms can be guided over distances of many centimetres via this technique without significant losses, and a transport distance on the order of 5 mm should not present substantial difficulties.

One downside to these approaches is that the optical power required to make them viable exceeds that which can be provided by most commercially available laser devices operating in the relevant wavelength range. Consider using a 1 mm diameter beam (a significantly smaller beam would not encompass the majority of the atoms from the MOT cloud) tuned 0.3 nm from to the red of the $5S_{1/2} \leftrightarrow 5P_{1/2}$ transition in Rubidium 85. This detuning gives a lifetime against heating due to scattering of the trapping light of less than 100 ms, so it is not feasible to significantly increase the trap depth by reducing the detuning without causing substantial heating during the transfer process. In order to obtain a trap depth comparable to the temperature of the MOT cloud (about 100 μK), around 4 W of optical power would be required. The best way to obtain sufficient optical power at the required wavelength may be via a frequency-doubling scheme, such as that described in [115].

6.4.3.3 Magnetic transfer from a MOT

Similar to transfer via a dipole trap, the idea here is that atoms could be transported closer to the surface in a purely magnetic trap. Though there is no fundamental reason why this method should not be effective, and magnetic atom transport has been demonstrated experimentally [116, 117], it has no significant advantages when compared to transfer via a moving dipole trap and is likely to be more difficult to implement experimentally. It also comes with the additional disadvantage that one must either employ some additional optical pumping step to place the atoms from the Λ-MOT into trapped m_f states prior to magnetic trapping or suffer the loss of a substantial proportion of the atom cloud on activation of the magnetic trap. For these reasons, magnetic transfer is probably not the most effective way to proceed with this experiment. Indeed, the above statements are consistent with experimental evidence from the literature — for example [117] demonstrates transport of cold atoms over 33 cm with ~30% efficiency via the use of a moving magnetic trap, and the process involves an additional optical pumping step during the loading of the magnetic trap. Meanwhile, the much simpler system described in [114] achieves transport over a distance of 30 cm with 40% efficiency by guiding the atoms with a dipole potential.

Though direct magnetic transfer is unlikely to be useful, magnetic focusing of ballistically transferred atoms, perhaps employing techniques similar to those described in [118], may well allow an increase in the density of cold atoms that can be generated in the vicinity of the dipole trap array, thus enhancing loading rates. This option is therefore worth considering.

6.4.3.4 Transfer from a MOT via surface-based guide structures

Conveyance of cold atoms along magnetic guides embedded into a surface has been experimentally demonstrated [119]. Such a process could be used to transfer cold atoms from one region of a surface to another with minimal losses — in our case this would enable us to move cold atoms from a flat region of the mirror's surface into the region containing the dipole trap array. The atoms could then be released in this new position, where they would be much closer to the dipole trapping sites than they are in the Λ-MOT. Loading of the atoms into the waveguide could be undertaken from a mirror-MOT, as this can be formed very close to any surface that is a good planar reflector.

Indeed, the use of surface-based structures to produce magnetic micropotentials suitable for the trapping and manipulation of cold atoms is now extremely widespread and well documented — see for example [120, 121, 122]. As a result the techniques involved are well understood and documented. This represents a significant advantage of this approach over some of the alternatives proposed previously, which rely on methods that have not been studied so extensively. As a consequence, this is a very promising avenue for further research.

Chapter 7

Atom trap based on optical pumping: an ideal loading system for microscopic dipole trap arrays

In chapter 1 we described how many of the applications of cold atoms rely on an ability to produce a large and/or dense sample thereof, and how this has hitherto usually been achieved via the use of a magneto-optical trap as described in chapter 2. We also found that only atom traps based on the scattering force produced by illumination with near-resonant radiation tended to be well suited to this role. Furthermore, in the previous chapter we saw that the primary obstacle to the successful demonstration of our microscopic dipole trap arrays is the difficulty of loading a worthwhile number of atoms into each trap site.

In this chapter we propose and characterise a novel scattering-based trapping mechanism that could, in the long term, prove to be an extremely effective loading mechanism for our microscopic dipole trap arrays, and may also be more appropriate than magneto-optical trapping for some other applications [3]. The new trapping scheme allows the spatial dependence of the trapping force to be tailored with a high degree of flexibility, making it suitable for the loading of atoms into unusually shaped conservative trapping potentials, of which an array of microscopic trap sites is but one example. It is also capable of producing traps with much greater spring constants than are typical of MOTs, thus potentially allowing higher densities of cold atoms to be produced and so permitting more efficient loading of dipole traps with a small capture volume, such as the microscopic dipole traps formed in our truncated cavity arrays.

7.1 Generalised trapping scheme

The most general way to describe the trapping scheme is to consider an atom with two or more long-lived or 'metastable' electronic states. The rate at which the atom scatters light from a laser beam tuned to a transition involving one of these states will be dependent on the probability of the atom being in that state at any given time. Further laser beams can therefore be used to alter the scattering force resulting from this initial beam by optically pumping the atoms into or out of the relevant metastable state. This allows one to circumvent the optical Earnshaw theorem [123], which is responsible for the extremely limited selection of spontaneous-force based atom traps proposed or demonstrated to date [52], and so to create a purely optical atom trap. As this trap is based on optical pumping between metastable states, we shall refer to it as a 'metastable optical pumping' (MOP) trap.

One instance of such a trap has been studied previously in [53], in which optical pumping between Zeeman sublevels was used to produce an atom trap based on the scattering force. However, the work described in [53] was largely restricted to the specific pumping and trapping scheme used in their experimental implementation. We found that, within the very general scheme set out above, there were other possible implementations that we considered to have more favourable properties.

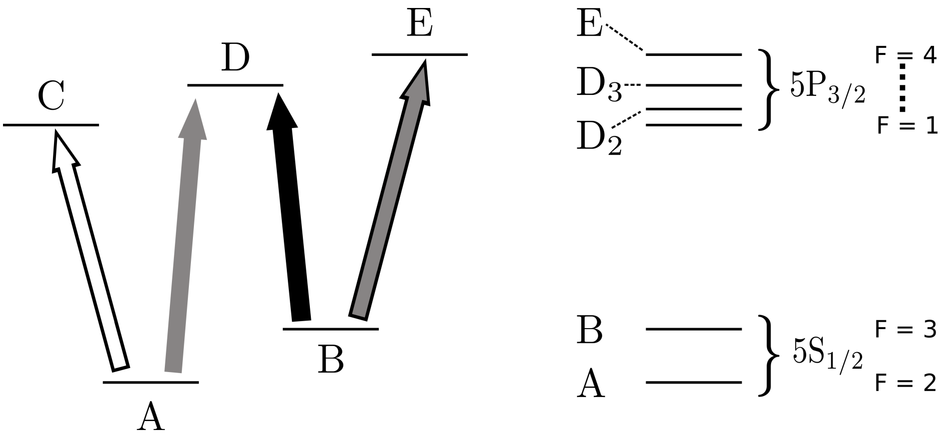

Figure 7.1: Pumping scheme for a MOP trap. The leftmost part shows a generalized pumping scheme and dipole-allowed transitions between the metastable levels A, B and unstable levels C, D and E. The corresponding levels of ^{85}Rb used in our experimental implementation are shown to the right. The $5P_{3/2}$ F=2 and F=3 states of ^{85}Rb are collectively represented by the single level D in the generalised scheme.

7.2 Specific implementation

We examined the specific case illustrated in figures 7.1 and 7.2. Here, we considered a one-dimensional atom trap based on optical pumping between two metastable states that we label A and B. We assume also the existence of unstable states C, D and E, that undergo rapid spontaneous decay into the two metastable states. The permitted electric dipole transitions between the different states are shown in figure 7.1. This situation is accurately mirrored in alkali metal atoms, which represent the most commonly trapped atomic species, and the correspondence of levels A-E to the actual energy levels of ^{85}Rb used in our experimental work is also shown in figure 7.1.

The trap operates as follows: ‘trapping’ beams address transitions A-C and B-E, and are arranged such that the net scattering force produced by all beams addressing the A-C transition opposes that caused by the beams addressing the B-E transition. ‘Pumping’ beams propagate across the trapping beams as shown in figure 7.2, and provide spatially-dependent population transfer between states A and B. Retro-reflection of the pumping beams ensures that the net scattering force exerted by them is negligible. The balance of the atomic population between states A and B then determines the relative magnitudes of the forces exerted by the different sets of trapping beams, and so the overall force

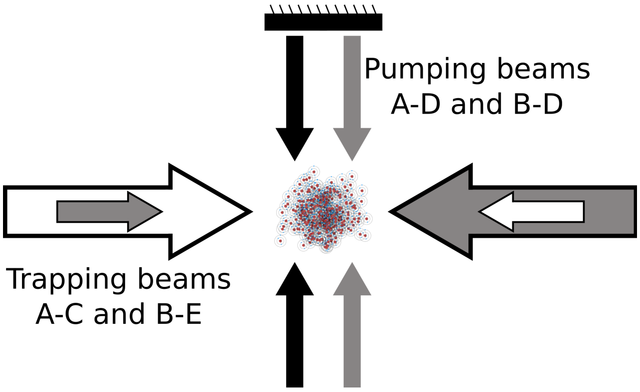

Figure 7.2: Beam geometry for a 1D MOP trap. The transition addressed by each beam is that shown in figure 7.1 by the same arrow colour and border. Larger arrows for the horizontal ‘trapping’ beams indicate higher intensities. The intensity distributions of the vertical ‘pumping’ beams determine the force profile of the trap, while the balance between the outgoing and reflected elements of these beams ensures that they exert no net force on the atoms.

exerted on an atom by the trapping beams is determined by the local intensities of the pumping beams. Adjusting the intensity profiles of the pumping beams, for example by using a mask or a spatial light modulator (SLM), therefore allows the scattering force along the axis of the trapping beams to be set as a user-selected function of the atom's position along this axis, which we shall take to be the z axis. By setting the z component of the force to have a negative gradient with respect to changes in z about a position at which it is equal to zero, it is then possible to create a one-dimensional atom trap about this position.

To reduce the difficulty of realising this experimentally, we made a few simplifications to the above scheme. Firstly, we combined the roles of the A-C trapping beam and A-D pumping beam, such that they were both replaced by a single A-D beam incident on the trapping region from the left. We then produced the imbalanced, counter-propagating elements of the B-E trapping beam shown in figure 7.2 by using a single beam incident from the right that was then attenuated and retro-reflected. The B-D pumping beam was implemented as shown in figure 7.2, and all laser beams were tuned about two linewidths to the red of the relevant atomic transitions to allow them to Doppler cool the trapped atoms.

7.3 Experimental results

Using this scheme, we produced a one-dimensional atom trap for ^{85}Rb atoms, with magneto-optical trapping providing confinement in the orthogonal directions. In this way, about 10^4 atoms were captured at a temperature of 93 ± 8 μK — these values were obtained via the same measurement procedures described already for the Λ-MOT in chapter 5. When the MOP trap was replaced by magneto-optical trapping in the third dimension, $\sim 3 \times 10^4$ atoms were captured[1], and the temperature of the atom cloud in this MOT was measured as 110 ± 40 μK [4]. This result is similar to that obtained in [53], which also found that the number of captured atoms was substantially lower than with magneto-optical trapping but the temperature of the atom cloud was comparable. However, in this case we believe the magnitude of the reduction in the number of captured atoms to be significantly increased as a result of our highly non-ideal experimental implementation, in which we were heavily constrained by the time and equipment available to us. In general we would expect the number of atoms captured by a trap of this form to exceed the number captured by a trap of the kind described in [53], as while both traps lack the Zeeman-assisted slowing effects present in a MOT, the trap described in [53] also had a greatly reduced spring constant. We shall subsequently see that a trap of the kind we describe above should be capable of producing spring constants that match and exceed those typical of magneto-optical traps.

An image of the atom cloud produced in this experiment is shown in Fig. 7.3. Meanwhile, Fig. 7.4 provides a demonstration of our proposed method for the creation of arrays of multiple trap sites: by placing a thin wire across the centre of the B-D beam we were able to create, in effect, two B-D beams separated by a small region of low B-D intensity. A trap site was therefore produced on the edge of each of these two beams, with a reduced atom density between the two sites. This density variation is illustrated more clearly by figures 7.5 and 7.6, which show the two-dimensional atom density as a function of position along the axis of the MOP trap. The density is measured per unit area in the plane of the images in figures 7.3 and 7.4, with a line integral being taken perpendicular to this plane. It should be noted that the nominal axis of the MOP trap

---

[1] At the time when this experiment was performed, the coils for the Λ-MOT had not yet been rotated into their new orientation.

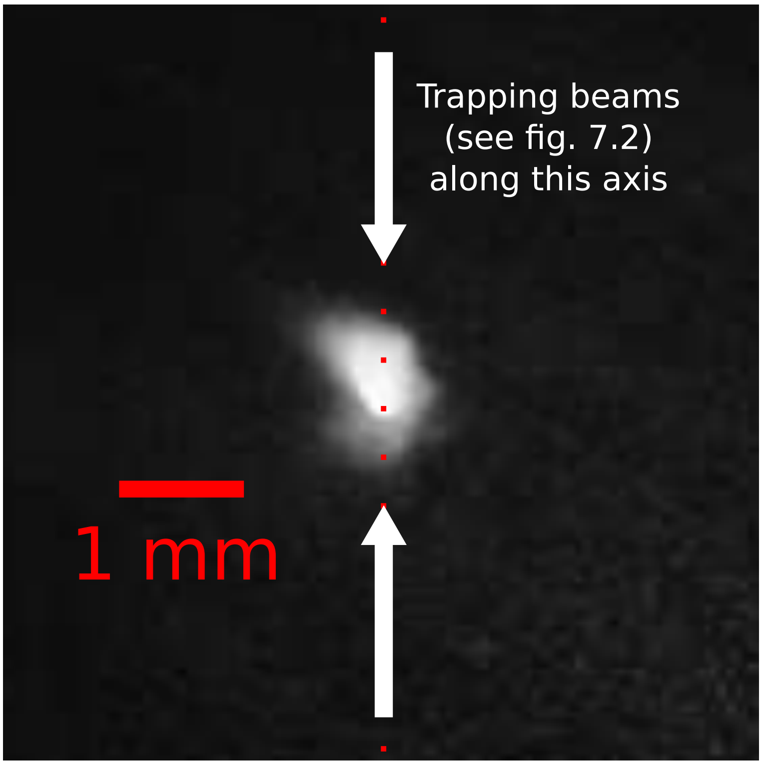

Figure 7.3: Atom cloud formed by a one-dimensional MOP trap combined with a MOT in the other two dimensions. The dotted red line indicates the axis of the MOP trap.

(which corresponds to the horizontal direction in Fig. 7.2) is along the vertical direction in figures 7.5 and 7.6, as indicated by the dotted red lines. The shape of the cloud seen in these figures appears to have been rotated slightly away from this axis — a small effect which we attribute to imperfect beam alignment.

7.4 Theoretical modeling

In our theoretical model, we ignore the atomic coherences and neglect the coupling of the laser beams to non-target transitions, as the influence of both of these factors has been shown to be negligible in similar situations [60]. We are also able to determine the behaviour of the trapped atoms accurately by considering steady-state populations and time-averaged forces, as the timescales associated with optical pumping are much shorter than those associated with motion of the atoms within the trap. This leaves us with a simple, rate-equation based model in which we consider a five-level system

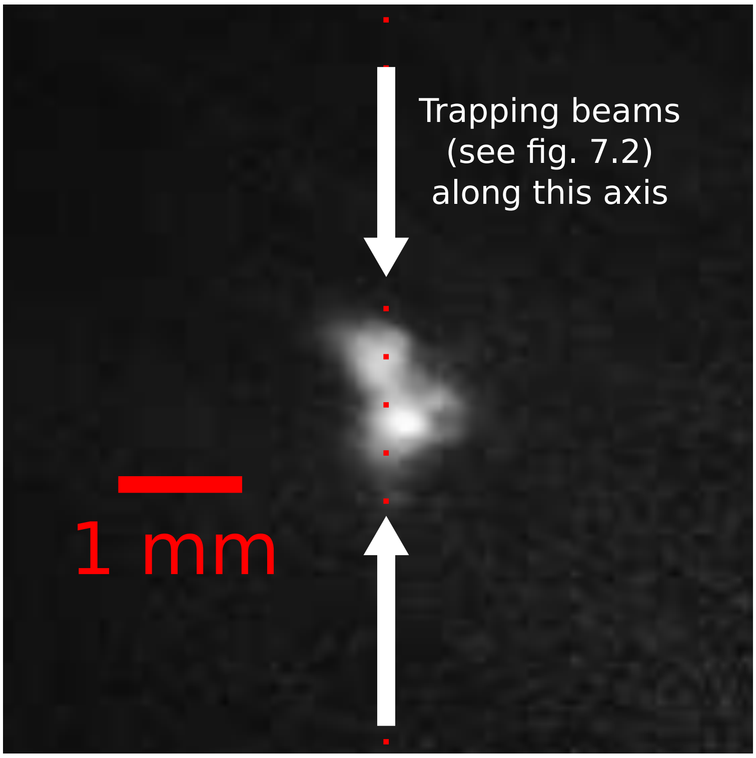

Figure 7.4: Sculpted atom cloud resulting from obstruction of the B-D beam with a thin wire. While the result is not spectacular, this extremely crude method nevertheless demonstrates the validity of the proposed approach for creating multiple trap sites. The dotted red line indicates the axis of the MOP trap.

consisting of the states A, B, D_2, D_3 and E as shown in the rightmost part of figure 7.1. We define a set of rate coefficients, τ_{ij} and Γ_{ij}, such that the stimulated and spontaneous transition rates between, for example, states E and B are given by $\tau_{\mathrm{EB}} I_{\mathrm{EB}}$ and Γ_{EB} respectively, where I_{EB} is the intensity of the laser light tuned to the relevant transition. The spontaneous decay rates for the relevant transitions are already known accurately — see for example [58]. To determine the rate coefficients for stimulated transitions, we equate the steady-state results for the upper state population produced by our rate equation model to those produced by solving the full optical Bloch equations for a two level system. For a transition with spontaneous decay rate Γ, illumination of detuning δ and intensity I, with Rabi frequency Ω, we obtain

$$\frac{\Omega^2/4}{\delta^2 + \Omega^2/2 + \Gamma^2/4} = \frac{\tau I}{2\tau I + \Gamma}. \tag{7.1}$$

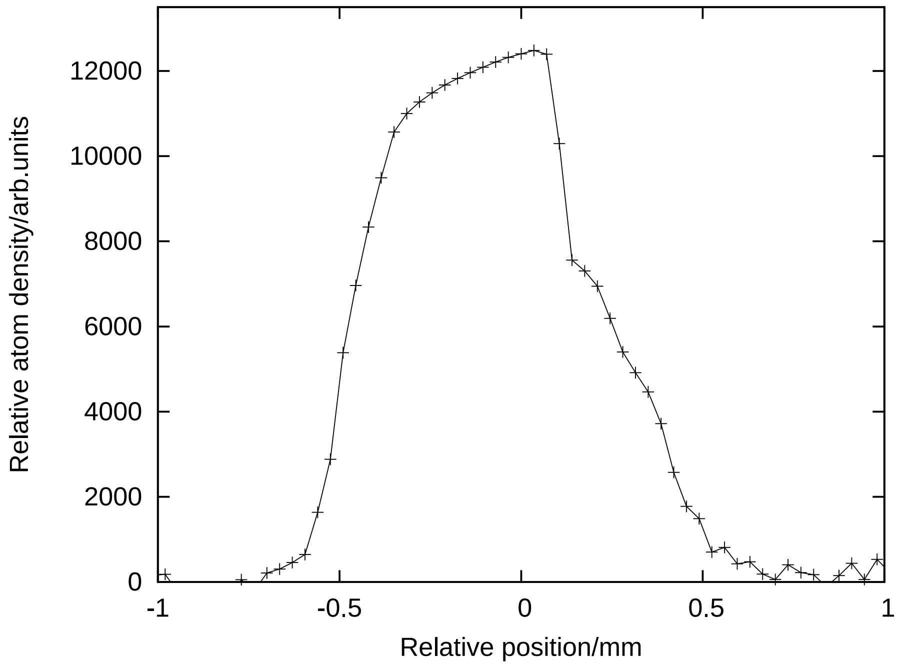

Figure 7.5: Relative (two-dimensional) atom density as a function of position along the axis of the MOP trap, for the atom cloud shown in figure 7.3. The dotted red line in figure 7.3 indicates the axis of the MOP trap.

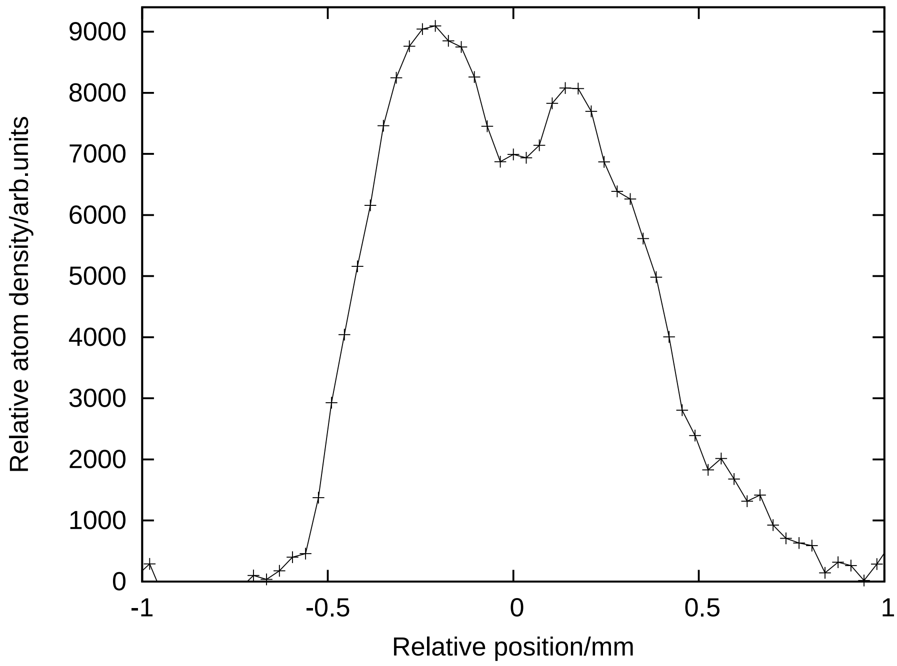

Figure 7.6: Relative (two dimensional) atom density as a function of position along the axis of the MOP trap, for the atom cloud shown in figure 7.4. The dotted red line in figure 7.4 indicates the axis of the MOP trap.

Therefore, labeling the dipole matrix element $\langle \mathrm{E}|x|\mathrm{B}\rangle$ between two levels as X_{EB}, we find that

$$\tau_{\mathrm{EB}} = \frac{\Omega_{\mathrm{EB}}^2 \Gamma_{\mathrm{EB}}}{4 I_{\mathrm{EB}}(\delta_{\mathrm{EB}}^2 + \Gamma_{\mathrm{EB}}^2/2)} = \frac{e^2 \left|X_{\mathrm{EB}}\right|^2 \Gamma_{\mathrm{EB}}}{2\hbar^2 c \epsilon_0 (\delta_{\mathrm{EB}}^2 + \Gamma_{\mathrm{EB}}^2/2)} \tag{7.2}$$

and similar results for the other transitions. Hence the rate equations governing the system are:

$$\frac{\mathrm{dE}}{\mathrm{dt}} = (\mathrm{B} - \mathrm{E})\tau_{\mathrm{EB}} I_{\mathrm{EB}} - \mathrm{E}\Gamma_{\mathrm{EB}}, \tag{7.3}$$

$$\frac{\mathrm{dD_3}}{\mathrm{dt}} = (\mathrm{A} - \mathrm{D_3})\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}} - \mathrm{D_3}\Gamma_{\mathrm{D_3B}} - \mathrm{D_3}\Gamma_{\mathrm{D_3A}}, \tag{7.4}$$

$$\frac{\mathrm{dD_2}}{\mathrm{dt}} = (\mathrm{B} - \mathrm{D_2})\tau_{\mathrm{D_2B}} I_{\mathrm{D_2B}} - \mathrm{D_2}\Gamma_{\mathrm{D_2B}} - \mathrm{D_2}\Gamma_{\mathrm{D_2A}}, \tag{7.5}$$

$$\frac{\mathrm{dB}}{\mathrm{dt}} = (\mathrm{E} - \mathrm{B})\tau_{\mathrm{EB}} I_{\mathrm{EB}} + (\mathrm{D_2} - \mathrm{B})\tau_{\mathrm{D_2B}} I_{\mathrm{D_2B}} + \mathrm{E}\Gamma_{\mathrm{EB}} + \mathrm{D_3}\Gamma_{\mathrm{D_3B}} + \mathrm{D_2}\Gamma_{\mathrm{D_2B}}, \tag{7.6}$$

and

$$\frac{\mathrm{dA}}{\mathrm{dt}} = (\mathrm{D_3} - \mathrm{A})\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}} + \mathrm{D_3}\Gamma_{\mathrm{D_3A}} + \mathrm{D_2}\Gamma_{\mathrm{D_2A}}, \tag{7.7}$$

where A-E represent the populations of the states A-E shown in figure 7.1. In order to simplify the results we are to obtain, the following notation is introduced:

$$K_1 = \frac{\tau_{\mathrm{EB}} I_{\mathrm{EB}}}{\tau_{\mathrm{EB}} I_{\mathrm{EB}} + \Gamma_{\mathrm{EB}}}, \tag{7.8}$$

$$K_2 = \frac{\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}}}{\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}} + \Gamma_{\mathrm{D_3B}} + \Gamma_{\mathrm{D_3A}}}, \tag{7.9}$$

$$K_3 = \frac{\tau_{\mathrm{D_2B}} I_{\mathrm{D_2B}}}{\tau_{\mathrm{D_2B}} I_{\mathrm{D_2B}} + \Gamma_{\mathrm{D_2B}} + \Gamma_{\mathrm{D_2A}}}, \tag{7.10}$$

and

$$K_4 = \frac{-K_3 \Gamma_{\mathrm{D_2A}}}{\Gamma_{\mathrm{D_3A}} K_2 + (K_2 - 1)\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}}}. \tag{7.11}$$

We then set all the time derivatives in (7.3)–(7.7) to zero to obtain the steady-state populations of the five levels, which are given by

$$\mathrm{B}_s = (1 + K_1 + K_2 K_4 + K_3 + K_4)^{-1}, \tag{7.12}$$

$$\mathrm{E}_s = K_1 \mathrm{B}_s, \tag{7.13}$$

$$\mathrm{D}_{3s} = K_2 K_4 \mathrm{B}_s, \tag{7.14}$$

$$\mathrm{D}_{2s} = K_3 \mathrm{B}_s, \tag{7.15}$$

and

$$\mathrm{A}_s = K_4 \mathrm{B}_s, \tag{7.16}$$

where the subscript s indicates the steady-state population of a given level. We can then obtain the time-averaged force on an atom using

$$\overline{\mathbf{F}} = (\mathrm{B}_s - \mathrm{E}_s)\tau_{\mathrm{EB}} I_{\mathrm{EB}} \mathbf{p}_{\mathrm{EB}} + (\mathrm{A}_s - \mathrm{D}_{3s})\tau_{\mathrm{D_3A}} I_{\mathrm{D_3A}} \mathbf{p}_{\mathrm{D_3A}} + (\mathrm{B}_s - \mathrm{D}_{2s})\tau_{\mathrm{D_2B}} I_{\mathrm{D_2B}} \mathbf{p}_{\mathrm{D_2B}}, \tag{7.17}$$

where $\mathbf{p}_{\mathrm{EB}}$ etc. are the (vector) mean photon momenta for light in the corresponding laser beams. A full expression for the time-averaged force as a function of position, $\overline{\mathbf{F}}(x, y, z)$, can now be derived for any given set of experimental parameters. Setting $\overline{\mathbf{F}}(x, y, z) = \mathbf{0}$ then allows the coordinates of the trap centre to be identified. The spring constant along a given axis can also be calculated by computing $\left.\frac{\mathrm{d}\overline{F}_i}{\mathrm{d}x_i}\right|_{\text{trap centre}}$. Performing these calculations for a realistic set of experimental parameters allows the restoring force and its gradient to be derived and compared with values typical of a MOT. We consider a one-dimensional trap using the same geometry and pumping scheme as our experimental prototype, assuming 10 mW of optical power in each beam and 3 mm beam diameters. With a 30% attenuation of the B-E beam on reflection, all beams red-detuned from their target transitions by $\sim 2\Gamma$ and the z-axis defined to be along the axis of the trapping beams, the relevant light intensity distributions and photon momenta become

$$I_{\mathrm{EB}}(x, y, z) = 1.7 I_0 \exp\left[-2(x^2 + y^2)/r_0^2\right], \tag{7.18}$$

$$I_{\mathrm{D_3A}}(x, y, z) = I_0 \exp\left[-2(x^2 + y^2)/r_0^2\right], \tag{7.19}$$

$$I_{\mathrm{D_2B}}(x, y, z) = 2 I_0 \exp\left[-2(x^2 + z^2)/r_0^2\right], \tag{7.20}$$

$$\mathbf{p}_{\mathrm{EB}} = \frac{-0.18h}{\lambda}\hat{\mathbf{z}}, \qquad \mathbf{p}_{\mathrm{D_3A}} = \frac{h}{\lambda}\hat{\mathbf{z}} \qquad \& \qquad \mathbf{p}_{\mathrm{D_2B}} = \mathbf{0}, \tag{7.21}$$

where $I_0 = 2.8 \times 10^3$ W m^{-2} and $r_0 = 1.5$ mm. We also assume the trapping light to be unpolarised. Using these parameters, we find the trap to have a spring constant around 4×10^{-18} N m^{-1}. This exceeds the spring constant of a typical MOT (with a magnetic field gradient of $\sim$0.1 T m^{-1}) by around a factor of 10, and is approximately two orders of magnitude greater than the spring constant measured in [53] during the only previous experimental demonstration of an optical-pumping based atom trap of which the author

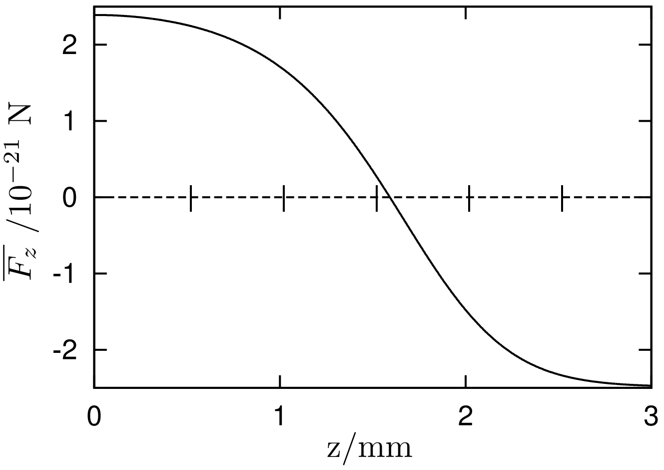

Figure 7.7: Time-averaged z component of the radiation pressure force on an atom as a function of its position on the z axis for a one dimensional MOP trap based on the scheme used in our experimental prototype and the parameters given in (7.17)–(7.21). The origin is taken to be on the axis of the B-D beam.

is aware. Such large spring constants make the MOP trap an excellent candidate system for the production of very high densities of trapped atoms. We also find that the trap centre is about 1.6 mm from the centre of the B-D beam, and the force profile of the trap is plotted in Fig. 7.7.

7.5 Extension to three dimensions and potential applications

So far, we have only demonstrated a one-dimensional prototype MOP trap. However, there are two obvious ways in which the scheme could be extended to produce full, three-dimensional trapping. Firstly, MOP trapping can be achieved along one axis for each independent pair of metastable states involved in the pumping scheme, and the inclusion of a third metastable state would therefore allow a 3D atom trap to be produced. In the absence of an appropriate third metastable level, an alternative is to use time-multiplexing, such that independent MOP traps operate along orthogonal axes for mutually exclusive time periods. The rapid switching of the laser beams required for this could be achieved, for example, by using a Pockels cell and polarising beam splitter. Fig. 7.8 illustrates how time-multiplexing could be used to produce a 2D MOP

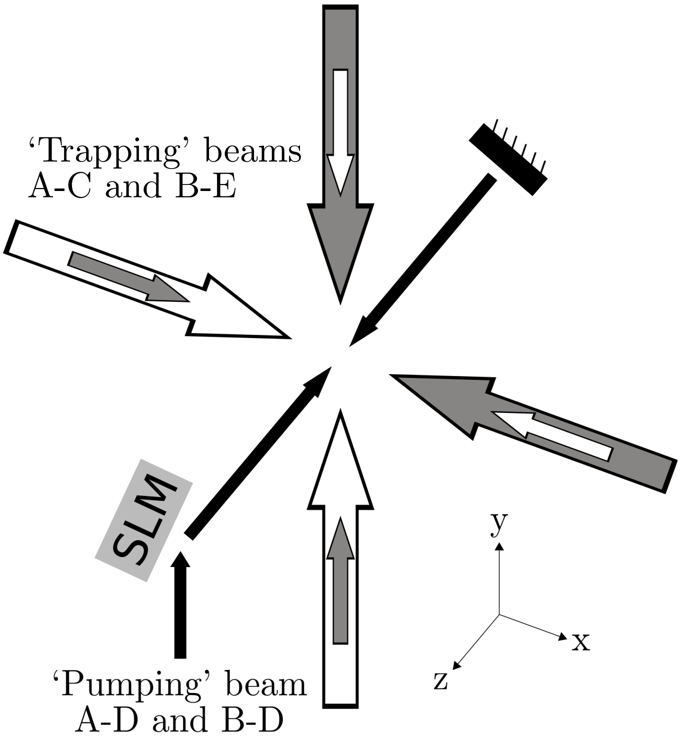

Figure 7.8: Proposed beam geometry for 2D trapping via time-multiplexing. The pairs of opposing 'trapping' beams would only be active along either the x or y axis at any given time. The spatial light modulator (SLM) would be used to switch the intensity profile of the 'pumping' beam, with simultaneous sculpting of the A-D and B-D beams being possible through, for example, the use of distinct spatial regions of the SLM.

trap, and it is then trivial to see how this could be extended to a third dimension by a repeated application of the same principle.

While, from our perspective, the loading of our microscopic dipole trap arrays may be the most important application of this trapping scheme, it is worth briefly considering its potential utility in other fields of research. Firstly, there are many other scenarios besides our own research in which it would be useful to efficiently load an array of conservative atom traps — see, for example, [27, 120]. Secondly, the trapping force in a MOP trap can be modified rapidly without encountering the problems caused by inductive effects when using rapidly varying magnetic fields. Furthermore, the atoms in the trap can be spin-polarised to some degree while remaining trapped, as the trapping mechanism is not dependent on population transfer between Zeeman sublevels — an advantage already obtained with some alternative trapping schemes [54], but not with any that have been widely adopted by experimentalists.

A disadvantage of MOP traps when compared to MOTs is that their experimental complexity will be greater than that of a MOT if they are implemented in more than one dimension. Zeeman assisted slowing and Sisyphus cooling are also not inherently present in a MOP trap, meaning that, when used on its own, a MOP trap is likely to trap fewer atoms at a higher temperature than a typical MOT. The optical pumping based atom trap reported on in [53] was found to capture between 5 and 20 times fewer atoms than a MOT with the same beam parameters, although this may in part result from the low spring constant of the trap, which was found to be about ten times smaller than that for the equivalent MOT. However, as the MOP trapping mechanism places no restrictions on the polarisation state of the trapping light, there is no reason why the beam polarizations could not be chosen so as to promote sub-Doppler cooling to reduce the temperature of the trapped atoms — under these conditions, we could expect performance similar to that obtained from the trap described in [53], which reaches a temperature of 40 μK in the absence of a magnetic field. This highlights one further benefit: because the MOP trap does not require the trapping light to be in any specific polarisation state, MOP traps can also function normally in the presence of unwanted magnetic fields, which could be an important advantage in applications where atoms must be brought close to magnetised structures or surfaces.

On account of the many useful properties of MOP traps discussed above, we believe that there may well be several applications beyond our own research for which they offer significant advantages over the MOT, or can be used in combination with a MOT to further enhance its trapping effects. For example, the option to partially spin-polarise the atoms while keeping them trapped might be of use in experiments involving magnetic traps, such as [116, 124], as it could allow a more efficient transfer of atoms into the trapped m_f states of the magnetic trap. Alternatively, combining a MOP trap with a MOT to enhance the atom density near the trap centre might be beneficial in experiments employing Bose-Einstein condensates, such as [125, 126], since a higher initial atom density will make it easier to use evaporative cooling techniques to achieve the phase-space densities required to produce such condensates. Indeed, as the combination of a MOP trap with a MOT would allow the spring constant of the trap to be substantially increased without causing a corresponding drop in the loading rate of the trap, it is reasonable to expect such a system to be capable of producing atom densities significantly higher than those obtained in 'compressed' MOTs (see chapter 2) — given that [72] reports that trapped atom density was increased by a factor of 10 when using the 'compressed' MOT technique, we could realistically expect an even greater increase in trapped atom density if an appropriate MOP trap were combined with a standard MOT. We therefore expect that MOP traps will prove to be of use in a number of atomic physics experiments.

Chapter 8

Light generation for atomic physics experiments

A key part of the motivation for our work on microscopic dipole trap arrays is their potential for use in the experimental realisation of quantum information processing. Atoms held in such traps provide individually addressable quantum systems with long lifetimes against decoherence. However, performing even single qubit operations on such atoms requires coherent optical manipulation of the atoms' electronic states. The use of single photon transitions for such manipulations is rendered impractical by the large spontaneous decay rates typical for such transitions, which would result in rapid decoherence of any atomic qubit under manipulation. It is therefore preferable to use multi-photon Raman transitions between states that are both long-lived against spontaneous decay. Driving such transitions optically typically requires the use of two closely-spaced frequency components [56], and the frequency difference between these components must often be set with a greater degree of accuracy than can be achieved with two independently stabilised laser sources. For example, the linewidths for Raman transitions between the hyperfine ground states of ^{85}Rb are on the order of several kHz, while the frequency of a typical DAVLL stabilised diode laser might vary by as much as 1 MHz. The ability to generate two optical frequency components with a small but well defined frequency difference is therefore key to the utility of our atom trap arrays in the event of a successful demonstration.

There are several ways to achieve an accurately determined frequency separation between different spectral components with similar wavelengths — most notably acousto-optic and electro-optic modulation of the output of a single laser source. Direct current modulation of a semiconductor laser is also an option, as is phase-locking [127] of a "slave" laser to a master laser source. All of these schemes have been successfully applied in experimental situations — see for example [128, 129, 130].

Electro-optic and laser current modulation are probably the most widely used of these methods, and also have the potential to allow one highly-stable 'master' laser to generate a large number of regularly-spaced optical frequency components [131, 132] that can in turn be used in the stabilisation of other laser sources via phase-locking. However, both of these modulation techniques place all resulting frequency components into the same spatial mode. It is therefore useful to have the option of separating the various frequency components from one another. The most commonly required operation is to remove the unaltered frequency component (known as the 'carrier') from a phase-modulated beam. The ability to perform this operation efficiently and with a minimum of expense and experimental complexity therefore has the potential to be useful in almost all atomic physics experiments, including those described in the preceding chapters.

In this chapter we propose and experimentally demonstrate two novel techniques for the removal of the carrier wave from a phase-modulated laser beam [1, 2]. One of these schemes separates out the carrier wave by means of its polarisation state and is applicable only to phase-modulated beams produced via electro-optic modulation, while the other employs an actively-stabilised fibre-optic Mach–Zehnder interferometer in order to split beam components according to their wavelengths. The core experimental and theoretical work on the first of these techniques was carried out by the author, while for the second of these techniques most of the experimental work relating to the 'frequency-shift spectroscopy' stabilisation technique discussed in §8.2.2.3 was carried out by Jonathan Woods, but all other work described herein (including the proposal and theoretical analysis of the 'frequency-shift spectroscopy' technique) was done by the author.

In addition to their importance for quantum information applications involving dipole trap arrays, the techniques we describe herein are likely to be applicable within other experiments within the field, as many atomic physics experiments and laser locking schemes rely on the use of two or more closely spaced optical frequency components [133, 134]. Even the standard magneto-optical trap described in chapter 2 requires the use of two closely spaced wavelength components in order to prevent accumulation of the atomic population in an un-addressed hyperfine state of the $5S_{1/2}$ level. The wider applicability of these techniques should therefore not be overlooked.

8.1 Polarisation based carrier elimination

In this section we explain the theory behind the removal of the carrier wave from an electro-optically phase-modulated laser beam based on its polarisation state, following which an experimental demonstration of the technique is provided.

Most electro-optic modulators work by applying a time-varying electric field to a crystal, thereby modifying its refractive index (and hence the phase accumulated by light passing through it) in a time-dependent way. Since the electric field is usually only applied along a single axis, such modulators typically only apply phase-modulation to a single linear polarization component of the light passing through the crystal[1]. Consider such an EOM with driving frequency Ω that produces phase modulation with depth[2] m in a single linear polarization component of incident light with frequency ω. For now we assume that the incident beam is linearly polarized with real amplitudes $A\cos\vartheta$ and $A\sin\vartheta$ in the modulated and unmodulated directions respectively, although we shall subsequently show that no extra calculation is required to generalize our treatment to any input polarization. The phase-modulated output field $\mathbf{E}$ may then be written as

$$\mathbf{E} = A \begin{pmatrix} e^{i(\omega t + m\cos\Omega t)}\cos\vartheta \\ e^{i\omega t}\sin\vartheta \end{pmatrix}. \tag{8.1}$$

---

[1] In our case, for example, New Focus Model 4431 employing MgO:LiNbO3

[2] The maximum modulation depth achievable with most commercial EOMs is of the order of 1 radian.

The Jacobi-Anger identity can then be used to decompose the modulated polarisation component into its constituent frequencies, yielding

$$\mathbf{E} = A \begin{pmatrix} J_0(m)\cos\vartheta \\ \sin\vartheta \end{pmatrix} e^{i\omega t} + A \begin{pmatrix} \cos\vartheta \\ 0 \end{pmatrix} e^{i\omega t} \sum_{n\neq 0} i^n J_n(m) e^{in\Omega t}, \tag{8.2}$$

where J_n is the nth order Bessel function of the first kind. The carrier wave, which corresponds to the first term in (8.2), can be removed using a linear polarizer. The effect of this polarizer, which must be aligned so that none of the carrier wave power is transmitted, is then to leave the transmitted beam composed entirely of the sidebands produced by phase modulation, which are represented by the second term in (8.2). Following this polarizer, the electric field of the remaining portion of the sidebands is given by

$$\mathbf{E}_{sidebands} = \begin{pmatrix} -\sin\phi \\ \cos\phi \end{pmatrix} A\cos\vartheta\sin\phi \sum_{n\neq 0} i^n J_n(m) e^{i(\omega+n\Omega)t}, \tag{8.3}$$

where $\phi = \arctan\{\tan\vartheta/J_0(m)\}$ is the angle between the axes of the modulated polarization component of the laser beam and the linear polarizer. It is then possible to calculate the proportion of the incident power emerging in each of the first-order ($n = \pm 1$) sidebands, which we label P_{S1}:

$$P_{S1} = \cos^2\vartheta \sin^2\phi \; J_1^2(m) = \frac{\cos^2\vartheta\tan^2\phi}{1+\tan^2\phi} J_1^2(m) = \frac{\cos^2\vartheta J_1^2(m)}{\{1+[J_0(m)\cot\vartheta]^2\}}. \tag{8.4}$$

For a modulation depth m this is maximized when $\cos^2\vartheta = 1/(1+|J_0(m)|)$. The sideband power emerging from the system, as a proportion of the sideband power that could be obtained if carrier removal were not required, is then equal to $1/(1+|J_0(m)|)^2$. This is plotted in Fig. 8.1.

One potential pitfall of this approach is that most commercially available EOMs exhibit a birefringence that has a strong temperature dependence. This birefringence introduces an additional phase difference between the polarisation components of the light that lie along the modulated and unmodulated axes, thus changing the resulting polarisation state of the carrier wave after phase modulation and so preventing it from being cleanly

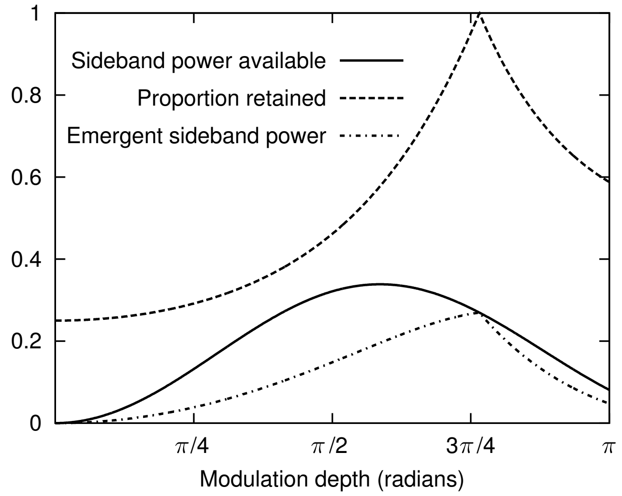

Figure 8.1: Power-efficiency of the carrier removal scheme as a function of the EOM's modulation depth. The three lines plotted correspond to the maximum proportion of the input power that could be placed into the 1st order sidebands without carrier removal, the proportion of this power that can be retained using polarisation-based carrier elimination and hence the maximum proportion of the input power that can be placed into carrier-free 1st order sidebands.

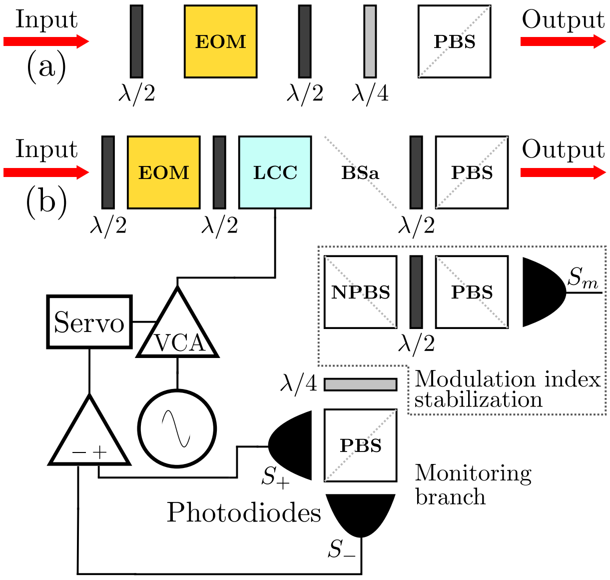

Figure 8.2: Experimental systems used for polarisation-based carrier removal. Part (a) shows a suitable passive setup that can provide short-term carrier suppression, while part (b) illustrates an actively stabilised system employing feedback via a liquid crystal cell, as described below. BSa: non-polarizing beam sampler. EOM: electro-optic modulator. LCC: liquid crystal cell. (N)PBS: (non-)polarizing beam splitter. VCA: voltage controlled amplifier. In practice the electronic feedback was implemented digitally [1].

removed using a linear polarizer[3] In order to solve this problem, further polarisation optics must be introduced to compensate for this phase difference, and a suitable passive arrangement is shown in Fig. 8.2(a). However, the temperature dependence of the birefringence means that, unless very effective temperature stabilisation is applied to the EOM, its birefringence will tend to drift over time and continuous adjustment of the system will be required to deal with these birefringence variations.

To overcome this problem, we use an E7 liquid crystal cell[4] as a voltage-controlled wave plate, similar to that discussed in [136], to actively compensate for the birefringence of the EOM[5]. Fig. 8.2(b) shows the experimental setup used for this. We found that the birefringence of the liquid crystal cell could be varied by changing the amplitude of an AC voltage (we used a 1 kHz sine wave) applied across it, and that this variation was approximately linear for peak-to-peak voltages between 1.2 and 2.6 V.

In order to control the voltage applied to the liquid crystal cell it was necessary to generate a signal that depended on the phase difference between the horizontal and vertical polarisation components of the carrier wave, based upon which the peak-to-peak voltage applied to the liquid crystal cell could be set. We opted for a feedback based system in which an error signal was generated that took a value of zero and had finite gradient (with respect to changes in the phase difference) about the desired lock point, this being the point at which the phase difference was an integer multiple of 2π and the carrier wave was restored to the correct polarisation state for removal by the linear polariser. This was done by using a quarter-wave plate and polarizing beam splitter to separate the two circular polarization components present in the transmitted beam, which are then directed onto the pair of photodiodes shown in Fig. 8.2(b), in an approach similar to the spectroscopy arrangement used by Hänsch and Couillaud [138] for cavity stabilisation. If the axes of the quarter-wave plate are set at an angle of $\pi/4$

---

[3] We now see that any phase relationship between the horizontally and vertically polarised components of the input light can be subsumed into this phase difference, and hence our calculations are valid for any input polarization.

[4] The full chemical composition of this widely used liquid crystal mixture is given in [135].

[5] Our cell was produced via the method detailed in [137] and was kindly provided to us by Mark Herrington. However, they are also available commercially — see, for example, Thorlabs Product LCR-1-NIR.

to the axes of the polarizing beam splitter, the two photodiode signals are

$$S_{\pm} = KA^2[1/2 \pm \cos\vartheta \sin\vartheta J_0(m) \sin\delta], \tag{8.5}$$

where K and δ represent the photodiode sensitivity and the net birefringent phase difference incurred during passage through both the EOM and the liquid crystal cell. Subtracting one from the other then gives the error signal as

$$S_E = S_+ - S_- = 2KA^2 \cos\vartheta \sin\vartheta J_0(m) \sin\delta, \tag{8.6}$$

which has a finite gradient and takes a value of zero at $\delta = 0$. A servo controller then provides feedback based on this error signal by controlling the amplitude of the AC signal applied to the liquid crystal cell[6].

8.1.1 Benefits of a polarisation-based approach

Established methods for carrier removal generally operate by separating light according to its wavelength. One downside of this is that sophisticated equipment is needed to provide the wavelength-scale path stabilisation that is essential for interferometric techniques [139, 140], as well as to provide the system with sufficient mechanical isolation. Secondly, wavelength-dependent techniques are not appropriate for use in situations in which the modulation frequency undergoes rapid variation — indeed many schemes are only able to operate at one pre-defined modulation frequency [139, 141]. Separation of the carrier wave based on its polarisation state offers a way to avoid these difficulties.

8.1.2 Experimental Results

The approach produced instantaneous carrier suppression in excess of 30 dB, while the active stabilisation technique shown in Fig. 8.2(b) was able to maintain time-averaged carrier extinction of 28.8 dB over a period of 2.5 hours. This is comparable to or better than the results obtained with most existing carrier removal techniques — see

---

[6] This was implemented using an Atmel ATMEGA328P-PU microprocessor mounted on an Arduino Uno board and Analog Devices DAC8562 12 bit digital-to-analog converter.

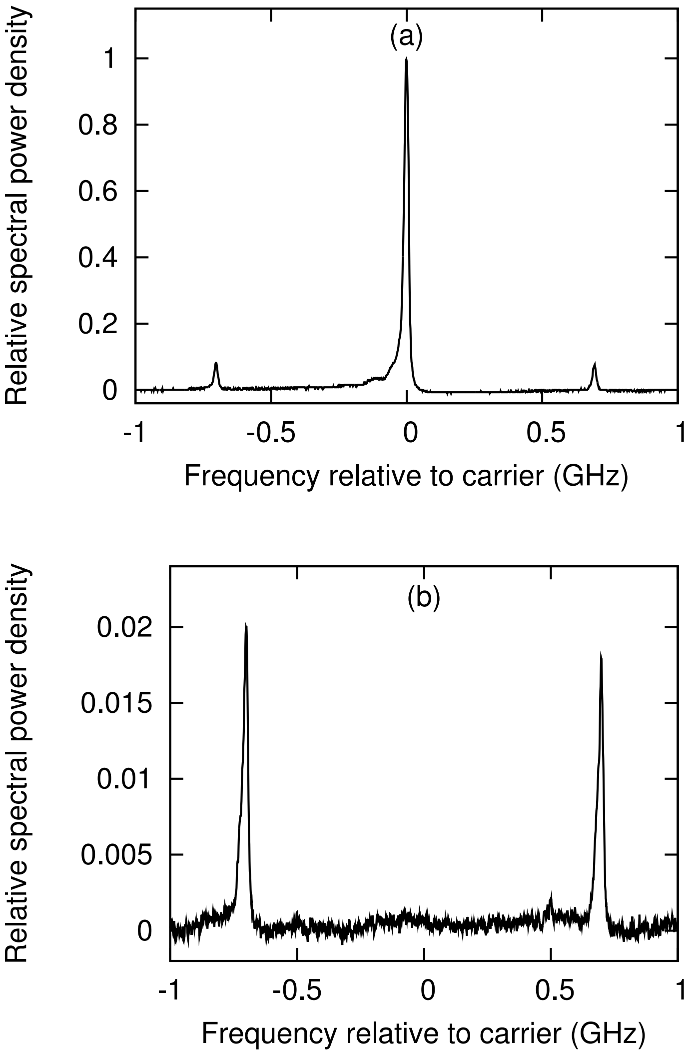

Figure 8.3: Spectra of light (a) leaving the EOM and (b) emerging from the carrier removal system. The modulation frequency is 2.7 GHz and our optical spectrum analyzer has a free spectral range of 2 GHz, hence the apparent appearance of the first order sidebands at a relative frequency of ± 700 MHz.

the 18 dB carrier extinction produced in [141], for example. Without active stabilisation carrier transmission was observed to increase over time, exceeding 2% after five minutes of unstabilised operation. Instantaneous spectra of the light before and after carrier removal are shown in Fig. 8.3. Attenuation of the sidebands was found to be 6.5 dB, which is consistent with (8.4).

Experimental characterisation of the error signal, S_E, was also carried out, and the results are shown in Fig. 8.4. These results are consistent with expectations based on the underlying theory, and show S_E crossing zero at the minimum of carrier transmission

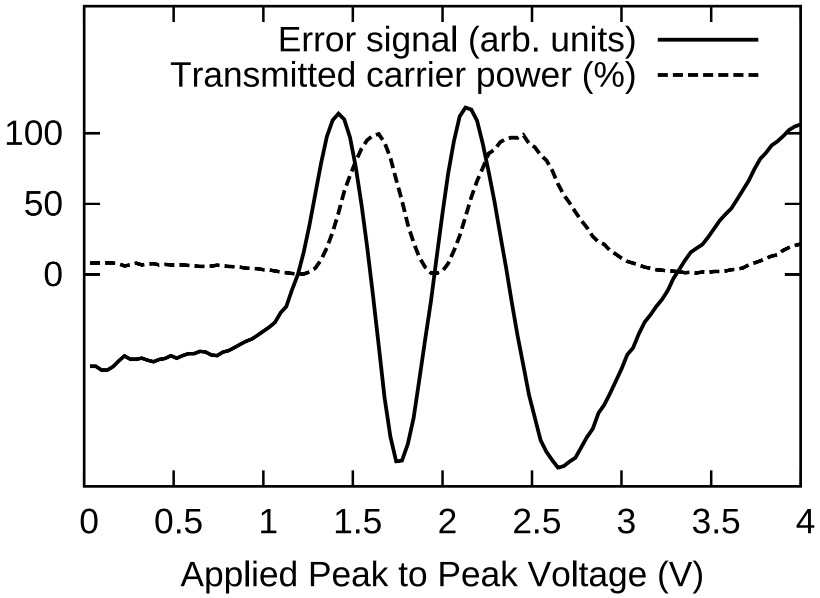

Figure 8.4: Error signal and carrier transmission (as a percentage of its maximum value) during a sweep of the peak to peak voltage applied to the LCC.

as well as varying approximately sinusoidally with changes in the peak to peak voltage applied to the liquid crystal cell in the region 1.2 V $< V_{pp} <$ 2.6 V. We attribute the slight asymmetry between the positive and negative going parts of the error signal to interference-induced transmission variations in the liquid crystal cell, which was not anti-reflection coated.

8.1.3 Compensation for variations in modulation depth

It was previously observed that the time-averaged carrier extinction provided by the device over prolonged periods was not as effective as the instantaneous carrier removal could be, even when active stabilisation was used. One of the primary reasons for this is that changes in the modulation depth of the EOM, while they do not change the lock point of the active birefringence compensation, do cause the polarisation axis of the emerging carrier wave to vary over time. As we did not implement any system capable of compensating for such variations, they therefore allowed a slight increase in carrier transmission.

Although this was a relatively minor concern for our particular application, there may well be scenarios in which changes in modulation depth will be produced intentionally as part of an experiment, or will arise as an unintended consequence of altering the modulation frequency. We therefore propose that the following method could be used to monitor the modulation depth, thereby providing a control signal that could be used to adjust the electrical power applied to the EOM in order to maintain a constant modulation depth, or to control additional components that could be installed in the system to allow it to cope with changes in modulation depth. We propose that the beam on the monitoring branch be split as shown in Fig. 8.2(b). Provided that δ remained locked to zero by the birefringence compensation and the axes of the polarising beam splitter were set at an angle of $\pi/4$ to the modulated polarisation component the signal, S_m, produced by the additional photodiode could be normalised to give:

$$\hat{S}_m = \frac{S_m}{(S_+ + S_-)} = \frac{KA^2[1/2 + \cos\vartheta \sin\vartheta J_0(m) \cos\delta]}{KA^2} = \frac{1}{2} + \cos\vartheta \sin\vartheta J_0(m), \quad (8.7)$$

from which the value of $J_0(m)$ could then be extracted.

In scenarios in which modulation depth variations are an essential part of an experiment, and cannot simply be avoided by controlling the signal applied to the EOM, the system must be modified to cope with modulation depth variations. If these variations take place slowly then the polarising beam splitter could be physically rotated, for example in a motorised rotation mount, so that the transmission axis remains orthogonal to the polarisation of the emerging carrier wave. However, for rapid changes in modulation depth such approaches are unlikely to produce a sufficiently fast response. Most polarizing components capable of a very fast response operate via an electrically induced birefringence, and therefore allow one to modify the phase term added to some particular polarisation component, rather than the polarisation axis of linearly polarised light. In order to allow the polarisation axis of the emerging carrier to be maintained during rapid variations in modulation depth, we have therefore devised the following scheme.

After passing through the liquid crystal cell and being sampled onto the relevant feedback optics, the light passes through two variable phase retarders (such as Pockels cells) that introduce phase differences of ϕ and ψ between orthogonal linear polarisation components: one has its slow axis aligned to the desired polarisation axis of the emerging carrier and the other has its slow axis at an angle of $\pi/4$ to this. Let us now suppose that the polarisation state of the carrier changes slightly due to a variation in modulation depth, such that it now has (real) field components a and b along the axes perpendicular and parallel to the desired polarisation state of the emerging carrier respectively. If this passes through the phase retarders in the order in which they are given above the output field from this system $\mathbf{E}_{out}$ is given by:

$$\mathbf{E}_{out} = \begin{pmatrix} 1+e^{i\phi} & 1-e^{i\phi} \\ 1-e^{i\phi} & 1+e^{i\phi} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & e^{i\psi} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}. \tag{8.8}$$

The component of $\mathbf{E}_{out}$ perpendicular to the desired polarisation axis of the emerging carrier is then equal to:

$$a\left(1+e^{i\phi}\right) + be^{i\psi}\left(1-e^{i\phi}\right). \tag{8.9}$$

From an examination of this formula it is clear that the modulus of this component can be set to zero by a suitable selection of ϕ and ψ, thus returning the emerging carrier wave to the required linear polarisation state for elimination: a correct selection of ϕ can make the moduli of the two terms equal and their arguments can then be set to be exactly out of phase by choosing ψ. Setting the squares of the moduli of the two terms in (8.9) to be equal and then re-arranging provides the required value of ϕ:

$$\cos\phi = \frac{b^2/a^2 - 1}{1 + b^2/a^2}. \tag{8.10}$$

It is then possible to determine the arguments of the two terms:

$$\tan\left(\mathrm{Arg(first\ term)}\right) = \frac{\sin\phi}{1+\cos\phi} \tag{8.11}$$

and

$$\tan\left(\mathrm{Arg(second\ term)}\right) = \frac{-\sin\phi}{1-\cos\phi}. \tag{8.12}$$

The required value of ψ is then equal to $\pi + \mathrm{Arg(first\ term)} - \mathrm{Arg(second\ term)}$.

The experimental complexity of this scheme means that it is not likely to be very widely used — especially given that one of the primary advantages of polarisation based carrier elimination over alternative techniques is the simplicity of the required system. However, there may be some specialised applications in which very rapid changes in both modulation depth and modulation frequency must be accommodated while maintaining high levels of carrier wave extinction. In these cases, the rapid changes in modulation frequency may rule out many alternative methods for carrier removal, leaving a system such as this as the only viable option. It is therefore possible that this scheme may be of relevance for certain specialised applications.

8.2 Fibre-optic Mach-Zehnder interferometry

A Mach-Zehnder interferometer (MZI) is a device in which a wave is split into two components that then traverse separate paths before subsequent recombination, as seen in figure 8.5. MZIs are not only useful for separating or combining coherent light beams of nearly equal frequency [142], but also have other important applications in optics [143], quantum information experiments [144] and metrology [145].

In this section we describe the use of a fibre-optic MZI for the removal of the carrier wave from a phase-modulated laser beam. There are several benefits associated with the use of a fibre-optic device instead of free space optics: the quality of the output beam and the ease of overlapping the two interfering beams accurately are both improved by the use of single mode fibres, the fact that the optical paths are contained within a solid structure reduces the sensitivity of the device to certain sources of unwanted phase change (such as air currents in and around the optical paths), and the device is easily integrated with other fibre-optic systems.

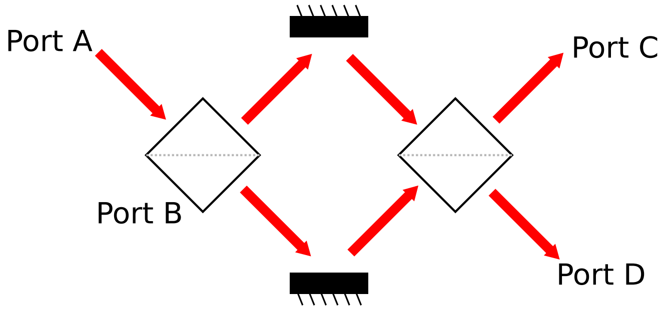

Figure 8.5: A Mach-Zehnder interferometer. For simplicity, only a single input beam and the resulting output is shown, though light could in principle enter the interferometer via any of the four ports.

Below, we first discuss the use of MZIs in splitting or combining nearly equal frequencies and then consider the specific case of our fibre-optic device and its intended application. We subsequently investigate possible methods for the stabilisation of our fibre-optic MZI to a state in which it has an inter-arm optical path difference appropriate for the separation of the carrier wave from the first order sidebands of a 780 nm laser beam that has been phase-modulated at 2.7 GHz.

8.2.1 Use in splitting/combining nearly equal frequency components

Consider the effect of sending a coherent, monochromatic and linearly polarised light beam into the interferometer in figure 8.5, such that it enters the device at position A as shown. Energy conservation requires a $\pi/2$ phase shift during transmission through the beam splitters. For an input field E_0, we therefore find that (up to an overall phase factor for each) the output fields at C and D are equal to

$$\frac{1}{2}E_0(1+e^{i(\pi+\frac{\omega}{c}(\eta_{low}(\omega)l_{low}-\eta_{upp}(\omega)l_{upp}))}) \tag{8.13}$$

and

$$\frac{1}{2}E_0(1+e^{i(\frac{\omega}{c}(\eta_{low}(\omega)l_{low}-\eta_{upp}(\omega)l_{upp}))}) \tag{8.14}$$

respectively, where l_{upp} and l_{low} are the lengths of the upper and lower paths in the interferometer and $\eta(\omega)$ is the refractive index of the medium through which light propagates in these arms. The power emerging from each output is therefore proportional to the square of the modulus of the relevant term above. We can greatly simplify the above expression by labeling the difference in phase accumulated by light passing through opposite arms as $\delta(\omega)$:

$$\delta(\omega) = \frac{\omega}{c}(\eta_{low}(\omega)l_{low} - \eta_{upp}(\omega)l_{upp}) \tag{8.15}$$

We then see that the proportion of the power entering at position A that emerges at position D (assuming no losses), which we shall label T_{AD}, is equal to $\frac{1}{4}\left|1 + e^{i\delta(\omega)}\right|^2$. The symmetry of the device means that this must be the power transmission coefficient for any diagonal traverse of the interferometer, while energy conservation requires that the power transmission coefficient for non-diagonal traverses be equal to $1 - T_{AD}$.

Now consider the effect of the interferometer on a beam containing two frequency components, ω_0 and ω_1, where $\Delta_\omega = \omega_1 - \omega_0$. It is generally a good approximation that

$$\delta(\omega_1) \simeq \delta(\omega_0) + \Delta_\omega \left.\frac{d\delta(\omega)}{d\omega}\right|_{\omega_0} \quad \forall \quad |\Delta_\omega| \ \ll \ \omega_0, \tag{8.16}$$

and from (8.15) we find that

$$\left.\frac{d\delta(\omega)}{d\omega}\right|_{\omega_0} = \frac{\delta(\omega_0)}{\omega_0} + \frac{\omega_0}{c}\left(\frac{d\eta_{low}(\omega)}{d\omega}l_{low} - \frac{d\eta_{upp}(\omega)}{d\omega}l_{upp}\right). \tag{8.17}$$

If we define ω_s such that $\delta(\omega_0+\omega_s) = \delta(\omega_0)+\pi$, we find from our expression for T_{AD} that when $|\Delta_\omega| = \omega_s$ frequency components ω_0 and ω_1 can be fully split or combined by the interferometer. Provided that $\delta(\omega_0)$ is an integer multiple of π, if a beam containing both components is placed into the same input port they will emerge from opposite output ports, while if the two frequency components are placed into opposite input ports they will emerge in a single beam from the same output. The value of ω_s can be calculated from (8.17):

$$\omega_s = \pi \left[\frac{\delta(\omega_0)}{\omega_0} + \frac{\omega_0}{c} \left(\frac{d\eta_{low}(\omega)}{d\omega} l_{low} - \frac{d\eta_{upp}(\omega)}{d\omega} l_{upp} \right) \right]^{-1}. \tag{8.18}$$

For a known value of ω_0 and arms containing transmission media with known $\eta(\omega)$, it is therefore possible to produce an interferometer with the desired value of ω_s by setting l_{low} and l_{upp} appropriately.

In most real devices changes in $\eta(\omega)$ over the relevant frequency region are negligible, and both the upper and lower paths can be approximated as having an equal, constant refractive index η. It is then the case that ω_s can be set to a good approximation by using:

$$\Delta_l = |l_{low} - l_{upp}| = \frac{\pi c}{\eta \omega_s}. \tag{8.19}$$

In our case we use a fibre-optic interferometer, in which light traverses the same paths as shown in figure 8.5 but is directed through an optical fibre rather than being guided via reflection from a mirror.

We wish to separate components with a frequency spacing of 2.7 GHz, and the relevant refractive index of the optical fibres is 1.53, as a consequence of which we find that a length difference of 36 mm is required between the fibres forming the two arms. Our interferometer was constructed according to this specification by OzOptics, and figure 8.6 shows the results of using it to separate the carrier wave from the first order sidebands of a 780 nm laser beam after phase-modulation at 2.7 GHz. No observable difference was found between the conditions required for maximum transmission of the carrier to a given output port and minimum transmission of the first order sidebands to the same, confirming that the interferometer is acting as a beam splitter/combiner for components with a frequency separation of 2.7 GHz.

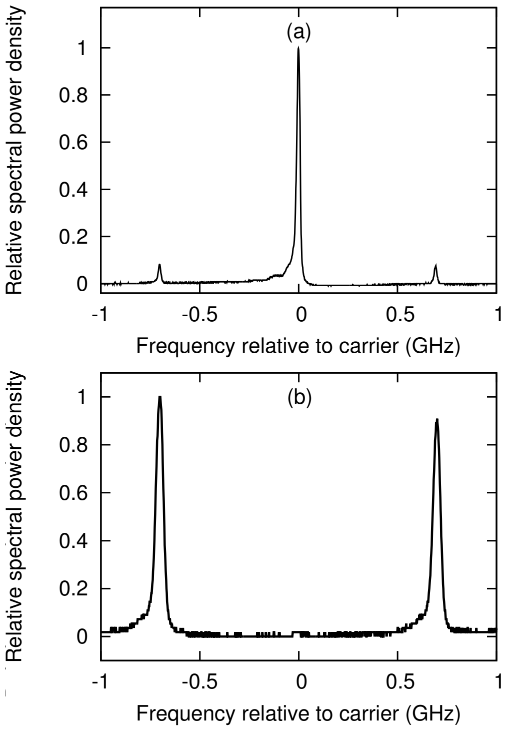

Figure 8.6: Spectra of phase-modulated light before (a) and after (b) being passed through the interferometer while it was at a position of minimum carrier transmission to the relevant output. The modulation frequency is 2.7 GHz and our spectrum analyser is based on a cavity with a free spectral range of 2 GHz, hence the apparent appearance of the first order sidebands at ±700 MHz relative to the carrier. The scale is not consistent between the panels as each is individually normalised — the interferometer reduces carrier intensity by over 30 dB while attenuating the sidebands by ∼4 dB (this attenuation is due to poor input beam quality preventing efficient coupling of light into the fibre, and is not a fundamental limitation of the device).

8.2.2 Stabilisation methods

Continuous use of the MZI as a beam splitter or combiner requires stabilisation of the value of $\delta(\omega_0)$ to an integer multiple of π, while many other applications also need $\delta(\omega_0)$ to be maintained at some particular value[7]. $\delta(\omega_0)$ can be modified by changing the optical path length of one of the arms, for example by stretching the fibre using a piezoelectric transducer, or in our case modifying its temperature (and therefore its refractive index) via a thermo-electric cooler (TEC).

---

[7] In reality it is almost always the case that $\delta(\omega_0) \gg \pi$ and changes to $\delta(\omega_0)$ on the scale of a few π therefore have a negligibly small effect on ω_s.

However, in order to use such a device for feedback stabilisation of $\delta(\omega_0)$ an 'error signal' must be derived that provides information about both the magnitude and direction of the feedback required. To do so the error signal must have a finite gradient, with respect to changes in $\delta(\omega_0)$, at the desired value of $\delta(\omega_0)$. It is also advantageous if the error signal takes a value of zero here, as the value of $\delta(\omega_0)$ to which the feedback stabilises the interferometer, known as the 'lock point,' is not then affected by changes in total signal amplitude. Due to the form of T_{AD}, when the desired value of $\delta(\omega_0)$ is equal to 0 or π (henceforth we shall neglect integer multiples of 2π), this can't be achieved by simply monitoring the level of transmission from a particular input to a given output, as these values of $\delta(\omega_0)$ will correspond to either a maximum or minimum of transmission for light at $\omega_0 \pm n\omega_s$ (for integer $n \ll \omega_0/\omega_s$).

Existing devices intended for use as beam splitters/combiners typically use dithering [146] for stabilisation. However, dithering involves varying either the state of the interferometer or the frequency of the input light in a way that may have unwanted effects on the ongoing output light. Devices used for other applications also frequently employ stabilisation schemes with undesirable properties, such as side of fringe locking [147], in which the lock point is sensitive to variations in total input power. Below, we examine three locking schemes that aim to avoid these drawbacks.

We found that the easiest of these schemes to implement in practice was based on comparing the relative levels of transmission of two frequency-shifted beams generated via the use of an acousto-optic modulator, and this was therefore the scheme that we ultimately adopted. Locking schemes based on polarisation effects and phase-modulation spectroscopy were also investigated. Although our work on these techniques was ultimately abandoned in favour of the use of the aforementioned alternative, it did bring to light some of the experimental challenges associated with these approaches to the stabilisation of fibre-optic MZIs, and a foreknowledge of these issues may be helpful when considering future work with such devices.

8.2.2.1 Polarisation based approach

One option is to lock the interferometer via a polarisation-based scheme similar to that used by Hänsch and Couillaud for cavity stabilisation [138] and closely related to the stabilisation method used for a free-space MZI in [139]. To facilitate this, one of the arms of the interferometer contains a built-in linear polariser, aligned to transmit light polarised along the slow axis of the polarisation-maintaining fibres (which we take to be the x axis). Light polarised along the fast axis of this fibre is completely blocked, and as a result this polarisation component is transmitted only via the other fibre arm (having lost half its power) and then divided equally at the second beam splitter. Consider the output at port C (E_c) — if the light entering port A with amplitude E_0 and frequency ω_0 is linearly polarised at an angle θ to the x axis then the x component will obey (8.13) while the y component is attenuated in amplitude by a factor of 2 and accumulates a phase difference of ψ compared to the x component as a result of the fibre birefringence, so that

$$\vec{E}_c = \frac{1}{2}\begin{pmatrix} E_0(1+e^{i\delta(\omega_0)})\cos\theta \\ E_0 e^{i\psi}\sin\theta \end{pmatrix}. \tag{8.20}$$

The optics shown in figure 8.7 can then be used to derive an error signal from this output. Light from port C is passed through a quarter waveplate set at an angle of $\pi/4$ to the fibre axes followed by a polarising beam splitter aligned to the fibre axes, with the two components separated at the beam splitter being directed onto the two photodiodes shown. The overall Jones matrices for the optical components on the paths to the two photodiodes are therefore

$$\frac{1}{2}\begin{pmatrix} 1+i & 1-i \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad \frac{1}{2}\begin{pmatrix} 0 & 0 \\ 1-i & 1+i \end{pmatrix}. \tag{8.21}$$

Applying these to E_c and calculating the transmitted optical powers, we find that the difference between the signals recorded at the two photodiodes is proportional to

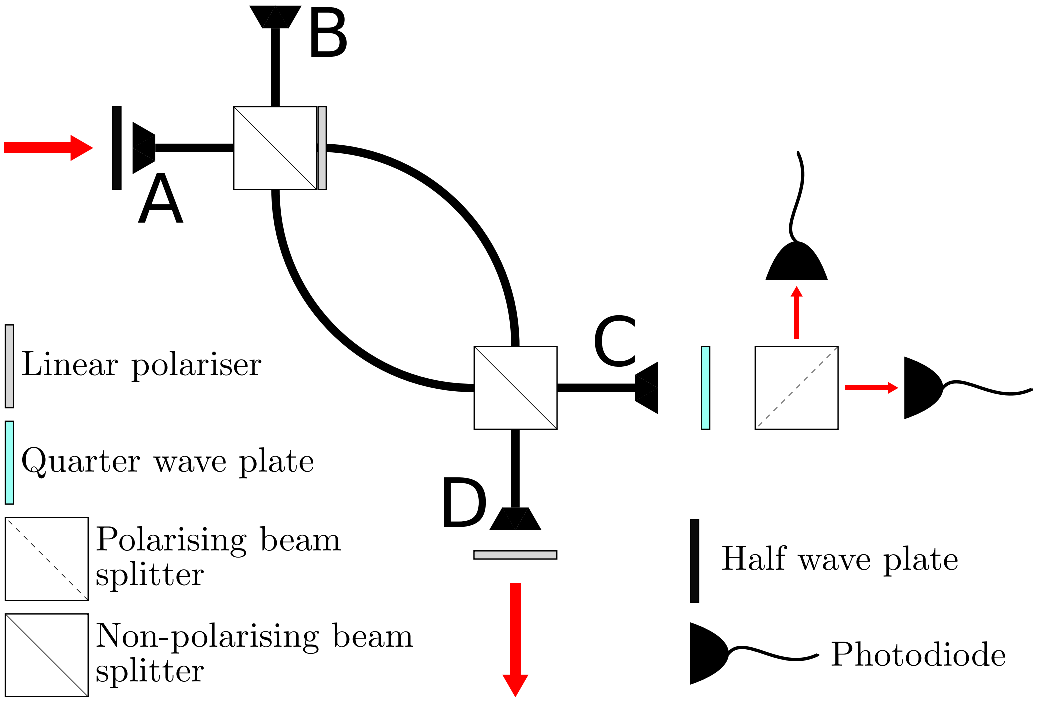

Figure 8.7: Optics for polarisation based stabilisation of a fibre-optic Mach-Zehnder interferometer. The half wave plate on the input is used to regulate the angle θ (see main text). The two linear polarisers and the polarising beam splitter are all aligned to transmit light polarised along the slow axis of the fibres, while the quarter wave plate on the upper output port is set at an angle of $\pi/4$ to the fibre axes.

$$\sin\theta\cos\theta\left\{\cos\psi\sin\delta(\omega_0)-\sin\psi\left[1+\cos\delta(\omega_0)\right]\right\}. \tag{8.22}$$

Provided that $\cos\psi$ and $\sin\theta\cos\theta$ are non-zero, this yields an error signal with a finite gradient when $\delta(\omega_0)$ is equal to zero or π (and a value of zero when $\delta(\omega_0)=\pi$). As such it is a suitable signal with which to stabilise the interferometer in the state required for separation of the carrier from the first order sidebands of a phase-modulated spectrum. However, there are two further points that require consideration.

Firstly, we arrived at (8.22) by assuming that only a single frequency component (ω_0) entered the interferometer at port A, when in fact the light entering this port is a phase-modulated spectrum. However, provided that all components of the spectrum are separated from the carrier by integer multiples of ω_s (as will be the case in an interferometer designed for use as a beam splitter or combiner), we can derive the form of the signal resulting from each spectral component by substituting $\delta(\omega_0) \pm n\pi$ into (8.22) in place of $\delta(\omega_0)$. This shows that these additional spectral components all produce signals with a gradient of the same, finite magnitude about the two desired values of $\delta(\omega_0)$ (the sign is reversed in the case of those separated from the carrier by odd multiples of ω_s). The properties of the phase-modulated spectrum may therefore affect the form of the overall error signal, which we obtain by summing the signals resulting from each spectral component, but provided that power division between the odd and even sideband components of the spectrum is not exactly equal the signal will always be appropriate for locking.

Secondly, the polarisation-maintaining fibres used in our interferometer exhibited a birefringence that depended very strongly on temperature. As data were not readily available regarding their thermal coefficient of birefringence, an initial estimate had been made by taking the product of the birefringence with the thermo-optic coefficient. This produced a value that would not have resulted in significant changes in ψ in response to variations of a few K in a temperature of $\sim$300 K, as might be expected to occur in a non-temperature-stabilised device operating in a typical laboratory environment. However, experimentally we found that significant changes in ψ did take place over timescales $\sim$10 minutes, and as such the long-term stability of this locking technique was seriously impaired. We hypothesise that the higher than expected thermal coefficient of birefringence arises from the mode of operation of the polarisation maintaining fibres: rods of a different material are built into the fibres and exert mechanical stress across one axis of the fibre, thus inducing birefringence. If the thermal coefficient of expansion for these rods differs significantly from that for the surrounding material, then small changes in temperature could result in large changes in the stress applied to the fibre by the rods and thus cause significant birefringence variations.

The likelihood of a system such as our fibre-optic MZI being of use to other researchers in the field, given that alternative devices capable of performing the same function already exist, is strongly dependent on the ease with which the necessary stabilisation systems can be constructed and used. Although there are several viable methods by which the aforementioned difficulty could be overcome, they would increase the complexity of the device and thereby reduce the significance of our result. We therefore decided to pursue an alternative approach in which this complication did not arise.

8.2.2.2 Phase-modulation spectroscopy

Phase-modulation spectroscopy is a well understood technique that has been applied to a range of stabilisation and measurement schemes [148]. Applying phase-modulation with depth m and frequency Ω to a laser beam (using a Pockels cell, for example) adds a time dependent phase, $\phi = m\cos\Omega t$, to the light field. The frequency of the laser beam is given by $\omega = \frac{\mathrm{d}}{\mathrm{d}t}(\text{phase})$, and this time-dependent phase term therefore produces a time-dependent frequency shift, $\Delta\omega = \omega - \omega_{\text{unmodulated}} = -m\Omega\sin\Omega t$. Combining this with (8.16) and the definition of ω_s reveals that, if such modulation is applied to a beam with frequency ω_0, we then obtain

$$\delta(\omega) = \delta(\omega_0 + \Delta\omega) \simeq \delta(\omega_0) - \frac{\pi m\Omega\sin\Omega t}{\omega_s}. \tag{8.23}$$

Combining this with our existing expression for T_{AD} and making the approximation that $2\pi m\Omega \ll \omega_s$ (appropriate in most cases) yields:

$$\begin{aligned} T_{AD} &\simeq 2(1+\cos\delta(\omega_0)) - \frac{\pi m\Omega\sin\Omega t}{\omega_s}\frac{d}{d\vartheta}\bigg|_{\vartheta=\delta(\omega_0)} 2(1+\cos\vartheta) \\ &= 2(1+\cos\delta(\omega_0)) - \frac{2\pi m\Omega}{\omega_s}\sin(\Omega t)\sin\delta(\omega_0). \end{aligned} \tag{8.24}$$

Differentiating the input signal to the Pockels cell gives a term proportional to $\sin\Omega t$. Taking the product of this with T_{AD} (obtained by sampling the beam onto a photodiode both before and after the interferometer and then determining the ratio of the measured optical powers, or by measuring only the power after the interferometer if the input power is known and stays very nearly constant over time) and time-averaging therefore yields a signal given by

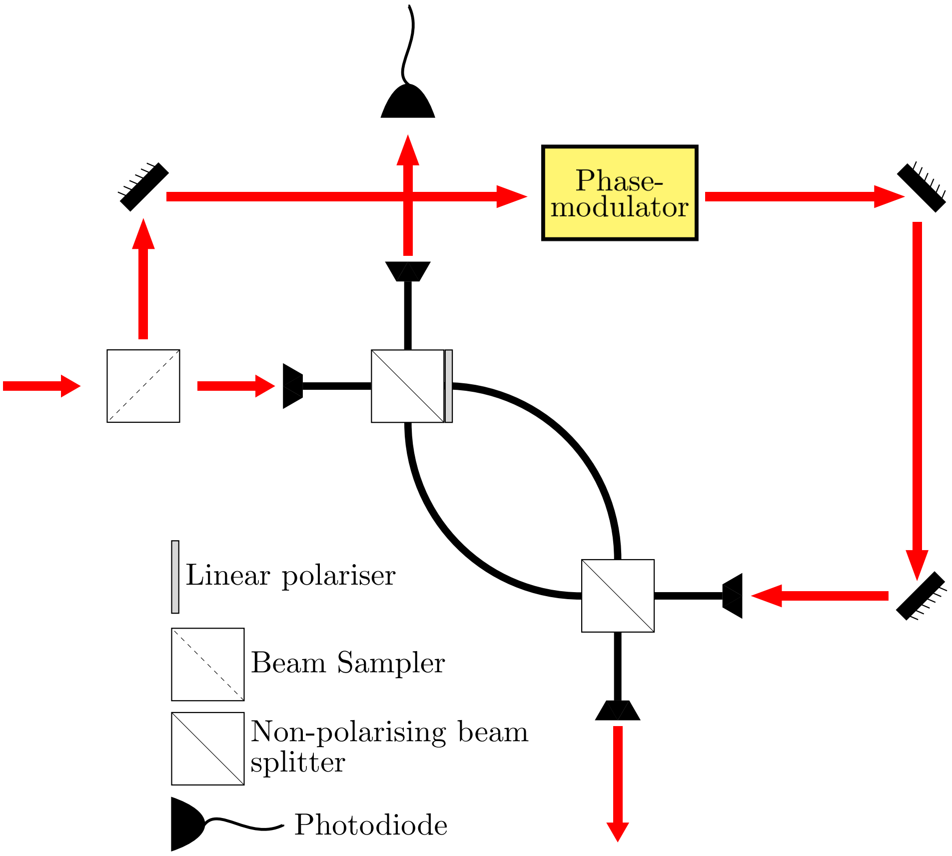

Figure 8.8: Experimental setup of the Mach-Zehnder interferometer and surrounding optics for stabilisation based on phase-modulation spectroscopy.

$$\mathrm{Sig} \simeq \frac{-2\pi m\Omega}{\omega_s} \overline{\sin^2(\Omega t)} \sin\delta(\omega_0), \tag{8.25}$$

which is of an appropriate form for locking of $\delta(\omega_0)$ to zero or π.

In order to avoid interfering with the ongoing output light, the beam used for this locking scheme is generated by sampling the main beam before it enters the interferometer. This sampled light is then phase-modulated and fed into the interferometer in the reverse direction via port C, as shown in figure 8.8, relying on the fact that $T_{CB} = T_{AD}$ to ensure that the signal produced is equal to that given in (8.25).

If the interferometer is being used to separate the components of an electro-optically modulated laser beam, the light used for locking can be extracted from the main beam prior to this modulation. If it is being used to separate the frequency components of the output of a current-modulated laser, the same argument made in §8.2.2.1 (regarding the use of light that is phase-modulated prior to entry into the locking system) can be applied here — indeed in this case the form of the error signal is totally unchanged by phase modulation with $\Omega = \omega_s$, with only its magnitude being modified by changes in the modulation depth. The lock point is therefore not affected by such changes.

However, we found that devices capable of producing phase-modulation with a sufficient depth-frequency product usually operate via an electrically-induced birefringence variation. As a result, they often also introduce unwanted polarisation modulation, and as this occurs at the modulation frequency it is a significant impediment to the generation of a useful error signal. Though such polarisation variations can in principle be eliminated by correctly adjusting the polarisation state of the light entering the device, temperature-induced variations of the modulator's birefringence mean that long-term elimination of the unwanted polarisation modulation requires either accurate temperature stabilisation of the modulator or active feedback to the polarisation state of the incoming light. Once again, this was considered to be impractical and so an alternative approach was sought.

8.2.2.3 Frequency-shift spectroscopy

Frequency-shift spectroscopy achieves the same end as phase-modulation spectroscopy without encountering the difficulty discussed in the last paragraph of §8.2.2.2. The scheme was proposed and theoretically analysed by the author, while most of the experimental construction work and data collection was carried out by Jonathan Woods (under guidance from both the author of the this thesis and the remaining co-authors of [2]). The mechanism is illustrated schematically in figure 8.9 — light is extracted from the main beam and passed through an acousto-optic modulator, which produces spatially distinct beams at ω_0 and $\omega_0 \pm \omega_{AOM}$ (in our case $\omega_{AOM} \simeq 2\pi{\times}80$ MHz). The beam at ω_0 is dumped, while the other two beams are passed through a beam chopper, aligned such that when one beam is blocked the other is passed and vice versa. The two beam paths are then combined at a beam splitter and fed into the interferometer via port C. A photodiode beyond port B monitors the output from this port while another photodiode receives a beam sampled from one of the two frequency-shifted beam paths

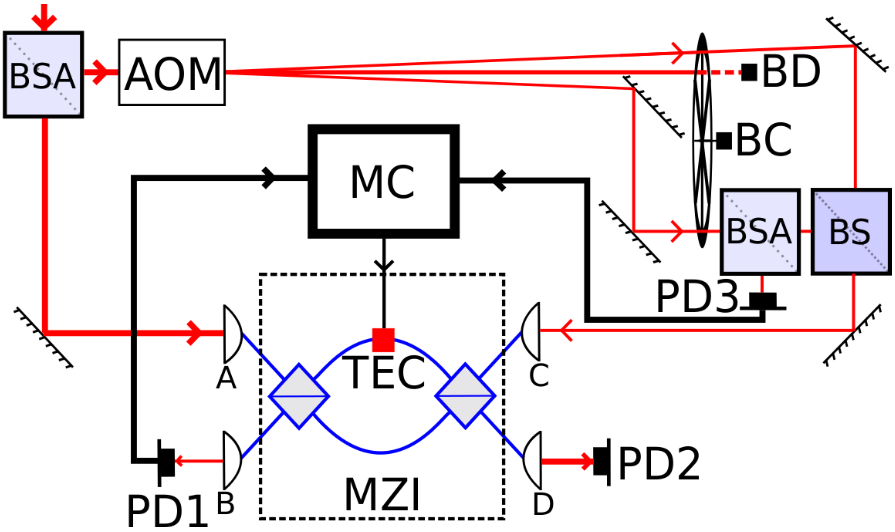

Figure 8.9: Experimental setup of the Mach-Zehnder interferometer and surrounding optics for stabilisation based on frequency-shift spectroscopy. BS: beam splitter, BSA: beam sampler, AOM: acousto-optical modulator, BD: beam dump, BC: beam chopper, TEC: thermoelectric cooler, MC: microcontroller (including a transistor used to regulate current through the TEC). This figure was produced by Jonathan Woods and is used here with his permission.

before they are recombined but after the beam chopper, thus serving as a triggering reference to determine which beam is being transmitted by the beam chopper at any given time. Using this in conjunction with the photodiode beyond port B, it is possible to determine the amount of power from each of the two frequency-shifted beams that emerges from port B. If it is ensured that equal amounts of power are coupled into the interferometer from each of these beams, then taking the difference of the powers recorded beyond port B will yield an error signal with the following form:

$$\text{Error Signal} \quad \propto \quad T_{CB}(\omega_0 + \omega_{AOM}) - T_{CB}(\omega_0 - \omega_{AOM}), \tag{8.26}$$

which, combined with our expression for T_{AD} and the definition of ω_s, gives:

$$\text{Error Signal} \quad \propto \quad \left|1 + e^{i(\delta(\omega_0)+\pi\omega_{\mathrm{AOM}}/\omega_s)}\right|^2 - \left|1 + e^{i(\delta(\omega_0)-\pi\omega_{\mathrm{AOM}}/\omega_s)}\right|^2. \tag{8.27}$$

This, as with phase-modulation spectroscopy, is suitable for locking to the required positions and does not change form in response to alterations in the modulation depth.

There are distinct similarities between this technique and phase-modulation spectroscopy. While phase-modulation spectroscopy employs a probe beam with a sinusoidally time-varying frequency, frequency-shift spectroscopy uses a beam with an effective frequency variation given by a square wave. Furthermore, in the limit in which the maximum frequency deviation produced by the modulator is very much less than ω_s (appropriate in most cases, including ours) both techniques yield an error signal proportional to $\sin\delta(\omega_0)$, with one being a constant multiple of the other.

Application of this technique produced a stable lock to a state of minimum T_{AD}. While this lock was applied, the time-averaged value of T_{AD} over a period of two hours was measured at 6.3×10^{-4} (corresponding to extinction of 32.0 dB). This was limited by the response rate of the feedback — a fact that was confirmed by observing that measurably non-zero values of T_{AD} were almost always accompanied by a corresponding deviation of the error signal from zero. Further evidence that this is the case comes from figure 8.10, which shows the Fourier transforms of $T_{AD}(t)$ for periods during which the interferometer was either actively stabilised (a) or allowed to drift freely (b). This figure clearly shows that the low frequency components of $T_{AD}(t)$ were suppressed most heavily by active locking, although it is not appropriate to attempt to derive a bandwidth value from this figure as by locking the interferometer to a minimum of $T_{AD}(t)$ we were also minimising the gradient of T_{AD} with respect to changes in the optical path lengths of the arms, and this meant that active stabilisation was capable of suppressing the effect of small path length variations even at frequencies well above those at which the feedback response could have any significant effect.

The current through the TEC was regulated by controlling the gate voltage of a field effect transistor (FET) connected in series with it. A digital micro-controller system based around the Arduino Uno board was employed to control this voltage, thus allowing us to use 12-bit pulse-width modulation to regulate the time-averaged current through the FET and consequently reducing the amount of power dissipated in the transistor.

A rapid scan of the interferometer over several fringes was performed by making a sudden change to the current through the TEC, and both the error signal and T_{AD} during this scan are plotted against time in figure 8.11. No systematic deviations were

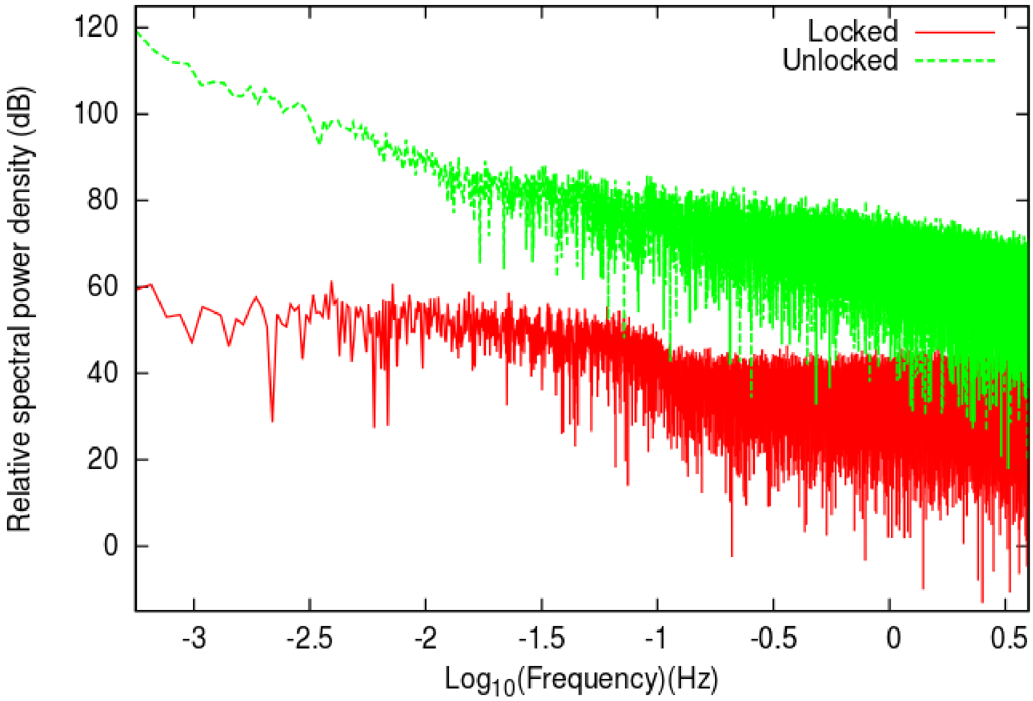

Figure 8.10: Fourier transforms of $T_{AD}(t)$ for the interferometer while actively stabilised to a position of minimum T_{AD} ('locked') and while allowed to drift freely ('unlocked'). This figure was produced by Jonathan Woods and is used here with his permission.

found between the measured form of the error signal and that predicted by (8.27), with most of the random deviations being the result of electrical noise in the photodiode amplifier circuits.

Based on these results, we concluded that frequency-shift spectroscopy is an effective approach to the stabilisation of a fibre-optic Mach-Zehnder interferometer at a position of maximum or minimum transmission of the dominant frequency component. It therefore allows fibre-optic Mach-Zehnder interferometers stabilised via this scheme to be used for continuous removal of the carrier wave (and potentially other frequency components) from a phase-modulated laser beam. As such, it may prove to be a useful component in any future experimental applications of our microscopic dipole trap arrays.

8.2.2.4 Other applications of frequency-shift spectroscopy

As discussed in the previous section, spectroscopy based on the use of the positively and negatively frequency-shifted components produced by an AOM can achieve results very

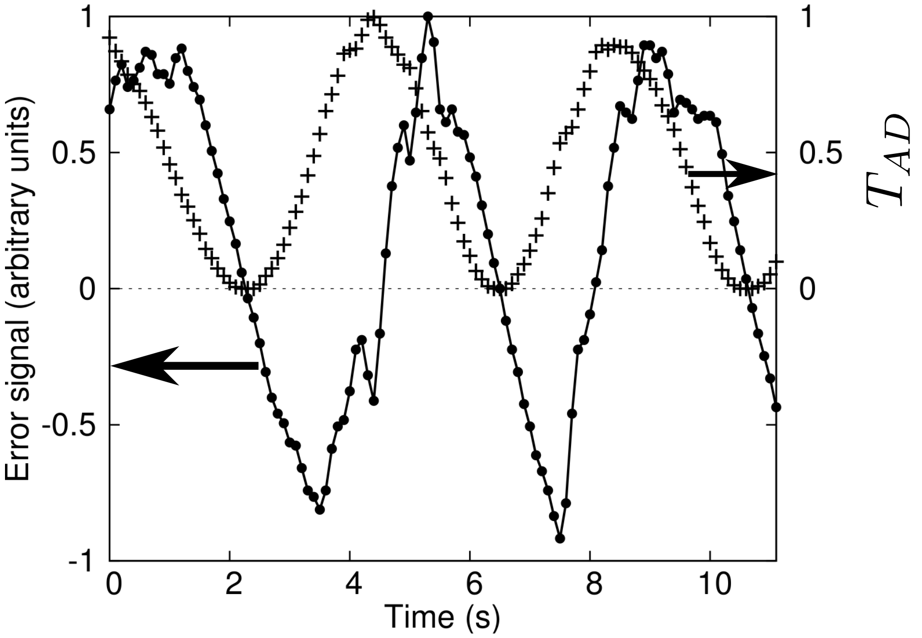

Figure 8.11: Plot of T_{AD} (proportion of total optical output power that emerges from port D when port A is the input) and error signal against time during a rapid scan of the interferometer over several fringes. The noise on the error signal is primarily electrical in origin. This figure was produced by Jonathan Woods and is used here with his permission.

similar to those obtained with phase-modulation spectroscopy. However, it is often more straightforward to implement in practice, as it requires only low frequency demodulation electronics and does not necessitate the use of high frequency electro-optic phase-modulators. These are frequently impractical owing to their temperature-dependent birefringence properties and their tendency to introduce unwanted polarisation and/or amplitude modulation. Phase-modulation spectroscopy is widely used for both cavity and laser stabilisation [149], and it is possible that for many such schemes frequency-shift spectroscopy may provide a more practical (if less elegant) alternative.

In the particular implementation discussed above frequency-shift spectroscopy comes with the drawback of a lower maximum feedback bandwidth, as this would be limited by the rate at which the beam chopper could switch between transmitting the different beams: usually something of the order of 1 kHz. However, there are many applications, such as stabilising a laser to an atomic or cavity resonance, for which the switching used above could be omitted altogether, thus avoiding this drawback and further simplifying the system. Examples of how this could be done are shown in figure 8.12. For stabilisation to a Doppler broadened[8] atomic transition the two frequency-shifted beams can simply take distinct spatial paths through an atomic vapour cell, with the transmission of both beams being continuously measured on two separate photodiodes. For stabilisation to a cavity resonance the two beams could be placed into orthogonal polarisation states before being combined at a polarising beam splitter, passed through the cavity and then directed onto separate photodiodes by a second polarising beam splitter. In either of these cases the frequency shift produced by the AOM must be comparable to the width of the feature to which the laser is to be stabilised, but as commercial AOMs producing a wide range of frequency-shifts are available selection of a suitable device for the desired application will usually be possible. In situations such as these the ease of experimental implementation and robust stabilisation offered by this approach make it an attractive alternative to true phase-modulation spectroscopy.

8.3 An acousto-optic frequency-comb

As mentioned at the start of this chapter, electro-optic and laser current modulation can be employed to produce a set of regularly-spaced optical frequency components, known as a ‘frequency-comb’, from a single laser source [131, 132]. It occurred to the author that the same effect should be achievable via the use of an acousto-optic modulator in a ring cavity containing a gain medium, and a search of the literature revealed that a similar experiment (without the gain medium) had in fact been carried out, with a promising result [150]. A very cursory investigation of this possibility was therefore carried out, with a view to determining whether or not future work on this system might prove worthwhile.

The system illustrated in figure 8.13 was constructed and the output monitored using an optical spectrum analyser. A frequency-comb with 80 MHz tooth spacing (set by the drive frequency of the AOM) and a half-power width of about 400 MHz was produced, where we have defined the ‘half-power width’ as the frequency-shift that results in a

---

[8] Locking to a Doppler-free atomic spectrum is also theoretically possible using this technique, but most commercially available AOMs operate at frequencies that make them more appropriate for use with a Doppler broadened feature.

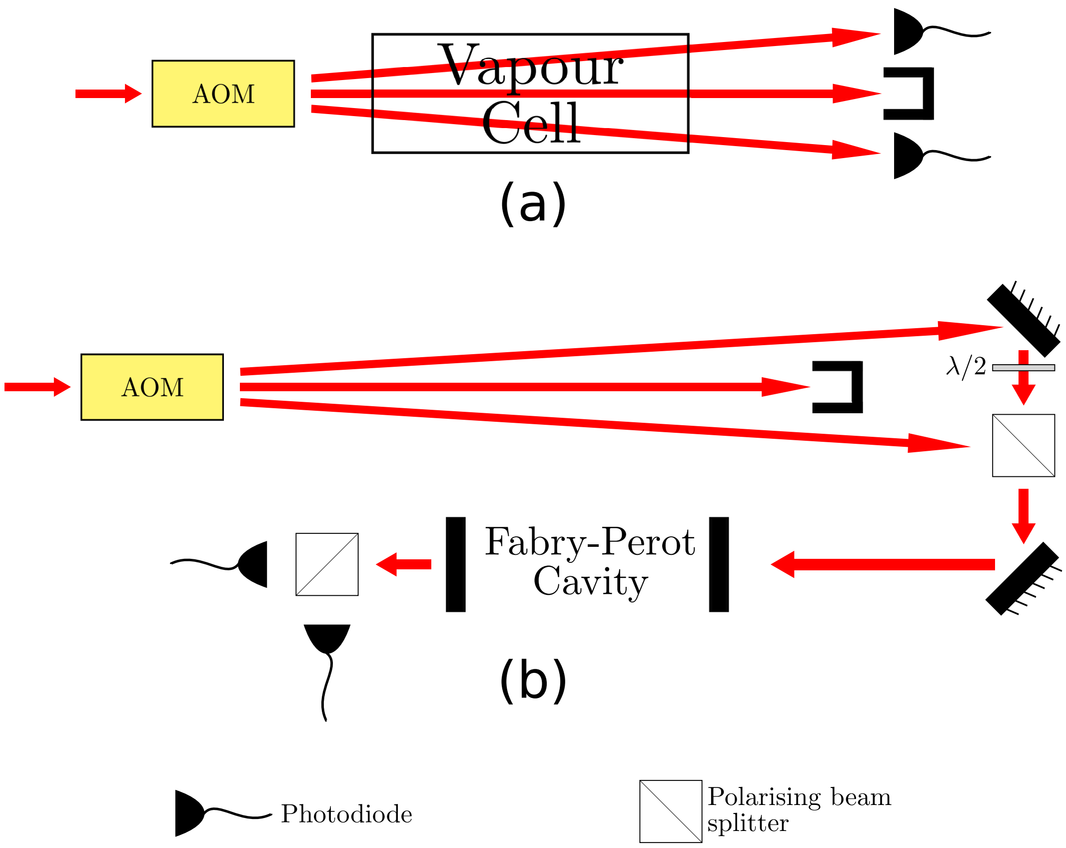

Figure 8.12: Suggested experimental arrangements for laser stabilisation via frequency-shift spectroscopy. The input light comes directly from the laser to be stabilised and an appropriate signal for stabilisation to an atomic (a) or cavity (b) resonance is generated by taking the difference of the two photodiode outputs. AOM: acousto-optic modulator.

halving of the optical power in each tooth. This performance is already sufficient to make the device useful, as the frequency-width of the comb is comparable to the energy differences resulting from the hyperfine structure of alkali metal atoms.

However, we also believe that the performance could be improved considerably by using a higher frequency AOM or reducing the re-amplification of amplified spontaneous emission from the gain medium inside the ring cavity, as this currently limits the half-power width of the frequency-comb by making it impossible to further increase the amplification produced by the gain medium without damaging it via the application of excessive optical power. This could be done by employing a longer optical path or some form of spatial filtering within the ring cavity, or placing an etalon with a free spectral range equal to the drive frequency of the AOM into the ring cavity.

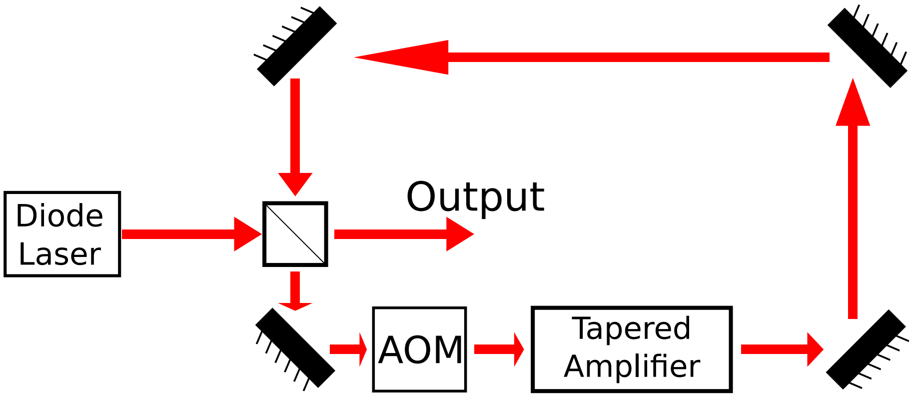

Figure 8.13: Experimental arrangement for the creation of an acousto-optic frequency comb.

Figure 8.14 shows a plot of the data recorded on a oscilloscope connected to the output of an optical spectrum analyser, into which the output beam from the frequency comb was coupled. The vertical axis is a linear scale showing the photodiode output signal voltage (proportional to the light intensity in the cavity of the spectrum analyser, and therefore a measure of the relative spectral power density of the input beam at the frequency of the cavity resonance) and the horizontal axis is time, but via the scanning of a piezoelectric transducer that controls the length of the cavity it corresponds linearly with the resonance frequency of the cavity. The plot is therefore effectively one of spectral power density against relative optical frequency. The tooth spacing of the comb is 80 MHz and the free spectral range of the spectrum analyser is around 1.5 GHz, with the output of the comb being seen here to run over into another free spectral range, hence the appearance of the smaller teeth to the left of the plot.

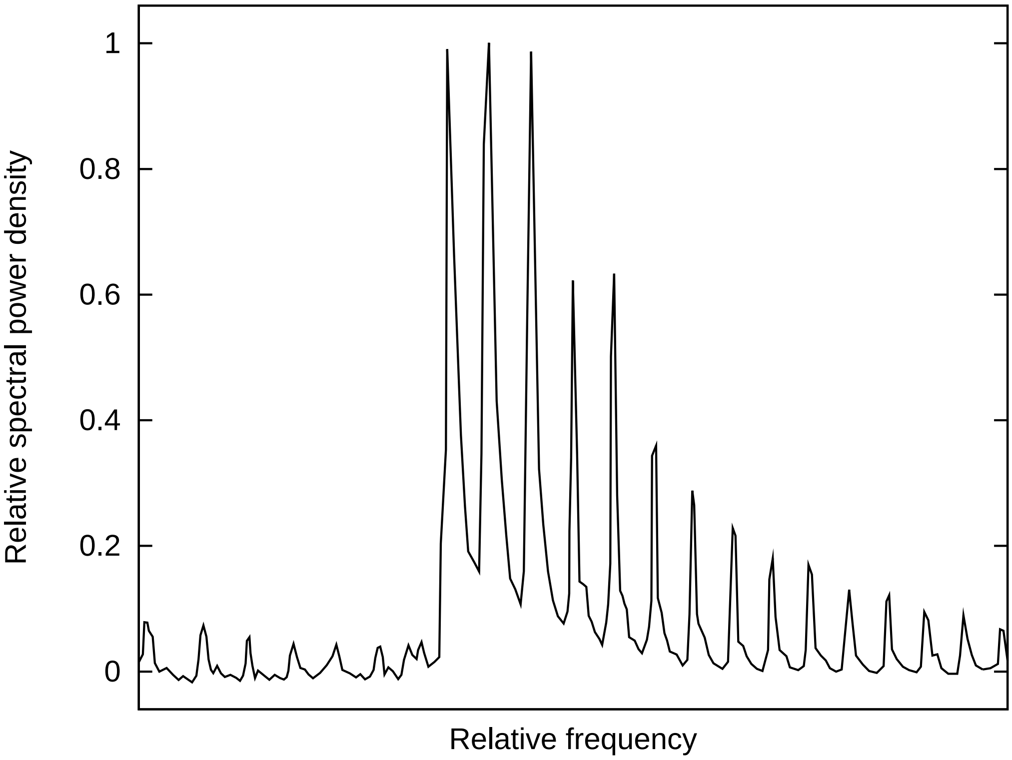

Figure 8.14: Example output spectrum of the acousto-optic frequency comb. The tooth spacing is 80 MHz.

Chapter 9

Conclusions

The work described herein has investigated a range of techniques that may improve our ability to trap and manipulate cold atoms. Some of these methods have already shown demonstrable utility in atomic physics experiments [151]. Below, we give a brief discussion of the outcomes associated with each technique and the prospects for further research in the area, followed by a summary of the place of this work within the wider field and a discussion of the likely direction of future research on these topics. We also examine how these techniques might be combined with one another and/or other existing techniques to further enhance their utility.

9.1 Outcomes

9.1.1 Microscopic dipole trap arrays

Arrays of microscopic truncated spherical cavities have been fabricated, one such array has been placed in a vacuum chamber and a magneto-optical trap for ^{85}Rb atoms has been formed ~5 mm from the microcavity array. Following this, we have constructed and demonstrated the efficacy of several of the experimental systems required to use this structure for the creation of an array of dipole traps. A macroscopic dipole trap has been formed to confirm the functionality of the laser system we intend to use to provide the requisite light for the formation of dipole traps. Additionally, an 'optical launch' of up to one third of the atom cloud formed by the Λ-MOT has been performed, and it has been shown that this technique could be used for ballistic transfer of cold atoms into the region of the microstructured surface. A multi-photon imaging system for atoms close to microstructured surfaces has been proposed and experimentally demonstrated, with both experimental and theoretical results showing that it can produce around 12% of the optical output power per atom that could be achieved with direct, single-photon imaging while suppressing background scatter from the surface by a factor of $\sim 10^6$, making it many times more effective than related schemes that have previously been employed for this purpose [4]. A feasibility calculation was performed to confirm that, using this imaging scheme, it should be possible to detect individual trapped atoms within the dipole traps formed by an array of truncated spherical microcavities.

In addition to this progress towards an experimental demonstration of an array of microscopic dipole traps, a theoretical investigation of the challenges associated with the loading of dipole trap arrays based on microstructured surfaces has resulted in the proposal of several viable means by which these difficulties could be overcome. Cold atoms could be guided along a surface from a flat region to a microstructured one, using techniques such as that described in [119]. Alternatively, the creation of microscopic magneto-optical traps as described in chapters 3 and 4 would allow continuous cooling of transferred atoms into the dipole traps of the array, thus increasing the number of atoms captured during ballistic transfer. The same effect could be achieved by starting with a larger or colder atom cloud, for example by increasing the size of the beams used to form the Λ-MOT or by creating a MOT based on a transition with a narrower linewidth, such as the open $5S_{1/2} \leftrightarrow 6P_{3/2}$ transition in ^{85}Rb. If a significantly greater total optical power could be obtained at a wavelength appropriate for dipole trapping of Rb, perhaps via a frequency-doubling scheme such as that described in [115], then horizontal confinement of cold atoms during ballistic transfer via the use of a large diameter dipole beam would be a realistic option. The same is true of the suggestion made in §6.4.3.2 that a macroscopic dipole trap could be used to bring the atom cloud closer to the surface before initiating ballistic transfer. In addition, there are further options available for increasing the density and reducing the temperature of the atoms in the initial cloud used for ballistic loading that we have not yet considered, such as the use of a dark MOT [152] or evaporative cooling of the atoms within a dipole or magnetic trap.

9.1.2 Microscopic magneto-optical traps

A method for the fabrication of microstructures that might be suitable for the creation of arrays of magneto-optical traps has been proposed, and the first step of this process has been demonstrated experimentally. The likely properties of such microscopic MOTs have been investigated theoretically, and we have derived a scaling law that shows that the loading rate of small magneto-optical traps is approximately proportional to the sixth power of the linear trap size. This is consistent with experimental results reported elsewhere, as well as other theoretical models based on different reasoning [40, 38].

9.1.3 The Λ-MOT

The efficacy of the Λ-MOT as a source of cold atoms that could be brought close to a microstructured surface had already been demonstrated in [4] before the work described in this thesis was begun. However, there were two key issues with the original implementation of the Λ-MOT that reduced its utility as an experimental tool, both of which stemmed from the orientation of the coils used to generate the magnetic field. Firstly there was the low number of atoms trapped by the Λ-MOT, and secondly there was a perceived restriction on how close the atom cloud could be formed to a microstructured surface. Modification of the coil orientation in the Λ-MOT has therefore not only increased the atom capture number by an order of magnitude, to around 270,000 atoms, but also brought to light the fact that it may in principle be capable of forming an atom cloud much closer to a structured surface than was previously believed to be the case. Given these improvements, the Λ-MOT may prove to be a useful system in the context of bringing atoms close to microstructured surfaces.

9.1.4 MOP traps

The principle behind MOP traps has been shown to be theoretically sound and its efficacy has been confirmed by experiment. The results of our theoretical examination of the mechanism show that the spring constants achievable with MOP traps should greatly exceed those typical of magneto-optical traps under most experimental conditions. When combined with the ease with which the spatial dependence of the trapping force in a MOP trap can be sculpted, this makes the MOP trap an ideal loading system for microscopic dipole trap arrays. A consideration of the properties of the trapping mechanism also reveals that it may posses several further advantages over magneto-optical trapping, which are summarised in §9.2.

9.1.5 Light generation for atomic physics experiments

Polarisation based carrier elimination has been shown to be an effective and practical solution to a common experimental problem, giving long-term carrier suppression in excess of 28 dB. As such, it has already found use in other experiments within the research group [151]. It is expected to prove particularly important in experiments that require continuous carrier removal during rapid changes in modulation frequency, especially if the enhancement suggested in §8.1.3 is employed to compensate for any unwanted variations in modulation depth that may come about as a consequence of changing the modulation frequency.

Our work on fiber-optic Mach-Zehnder interferometry has shown that frequency-shift spectroscopy using an acousto-optic modulator is a simple and robust means of providing the feedback stabilisation necessary to maintain the path difference between the two arms of a Mach-Zehnder interferometer (MZI), and may in fact have applications beyond interferometry, for example in the stabilisation of diode lasers to the wavelengths of atomic or cavity resonances. Fiber-optic MZIs stabilised via frequency-shift spectroscopy have been shown to be an effective means of splitting or combining laser beams of nearly equal frequency, with experimental work on the use of such a device to remove the carrier wave from a phase-modulated laser beam yielding long-term carrier suppression of over 30 dB. In addition, we have obtained useful information regarding the nature of the experimental challenges associated with stabilising fiber-optic MZIs via polarisation based techniques or methods dependent on phase-modulation spectroscopy, and this may help to inform those contemplating future work with such devices.

Finally, we have carried out a very brief demonstration of the fact that a large number of regularly-spaced wavelength components can be obtained from a single laser source via the use of an acousto-optic modulator in a ring cavity containing a gain medium, and have seen that the wavelength range covered by the resulting optical spectrum is sufficient that such a device might be a useful tool in some atomic physics experiments and consequently merits further investigation.

9.2 Summary and Outlook

We have developed a set of novel techniques for the trapping and manipulation of ultracold atoms. Several of these techniques have been investigated thoroughly enough to merit dissemination of the results, and are likely to find applications in a range of atomic physics experiments in the near future. Significant progress has also been made with other methods, and means by which research on these might be continued and ultimately brought to a successful conclusion have been proposed.

Microscopic dipole trap arrays based on microstructured surfaces remain a suitable candidate system for the experimental implementation of quantum information processing. Although single-site addressability has been shown to be possible in other systems since this research began [26, 27], dipole trap arrays based on microstructured surfaces retain several unique advantages over alternative systems; they make very efficient use of the available optical power, produce dipole traps of a sufficiently small size that the use of collisional blockade effects might allow single-occupation of each trap site to be ensured [44], may allow modification of the spontaneous decay rates of the atoms contained therein via an appropriate choice of cavity depth [46] and permit stronger coupling of the trapped atoms to incident light than can usually be achieved without placing the atoms into a conventional optical cavity [45]. We have constructed and tested most of the experimental systems required to create, load and detect the contents of these microtrap arrays, and have also proposed several mechanisms by which the remaining obstacles to a successful demonstration of such a microtrap array might be overcome.

There may also be some benefits to pursuing work on the fabrication of microcavities containing a central hole, besides those directly associated with the formation and loading of microscopic atom traps. Since the first arrays of truncated spherical cavities were fabricated via the process reported in [43], there has been interest in the plasmonic properties of such cavity arrays [153], and the creation of arrays in which each cavity has a central hole (of a size that can be chosen by controlling the depth to which copper is electroplated) would provide an interesting related system for experimental investigation. In fact, the similarity of such an array to some of the structures already investigated in the context of extraordinary optical transmission [154] suggests that arrays of this kind are indeed likely to have noteworthy optical and plasmonic properties. As the principles of our fabrication process have been shown to be sound by our experimental results, it should not be difficult to further refine the process so that it produces higher quality structures of the same form, particularly if the work were taken up by a researcher with a stronger background in electrochemistry than atomic physics.

The Λ-MOT, already shown to be a useful system for bringing cold atoms close to structured surfaces, has been significantly improved. The number of atoms captured by the Λ-MOT has been increased by an order of magnitude, and it has been shown that it should be possible to form the Λ-MOT much closer to a microstructured surface than was previously believed to be the case. An experimental demonstration of the latter of these two improvements would be a useful direction for future research on the Λ-MOT to take. Given the widespread interest in the interactions between atoms and structured surfaces, the Λ-MOT may well prove to be of interest as a candidate system for the generation of dense clouds of ultracold atoms in close proximity to such surfaces.

The principle of MOP trapping has been shown to be valid and a scheme has been proposed by which it could be employed to produce three-dimensional, spontaneous-force based traps in which the spatial variation of the trapping force can be dynamically tailored with a high degree of flexibility. Furthermore, it has been shown to be likely that MOP traps could be used to trap partially spin-polarised atom clouds, could be formed in the presence of unwanted magnetic fields without significant disruption and could be combined with magneto-optical trapping to produce very high densities of cold atoms. For these reasons, combined with the aforementioned fact that MOP traps would typically have stronger spring constants than MOTs, MOP traps are likely to find many applications within the field. The extreme importance of spontaneous-force based atom traps within atomic physics means that, even if there are only a few niche areas of research for which MOP traps prove to be a more appropriate tool than the MOT, this may well still be the most significant result presented in this thesis.

Finally, our work on light generation techniques has produced two novel methods for the removal of the carrier wave from a phase-modulated laser beam, both of which demonstrate performance that is comparable with or better than existing techniques — compare the levels of carrier extinction we obtain from either of these schemes (about 30 dB) with the 18 dB suppression quoted in [141] or the 20 dB extinction achieved in [139], for example. The polarisation-based approach to carrier removal is of particular interest because, in contrast to all existing techniques of which the author is aware, it does not separate the carrier wave from the sidebands on the basis of their relative frequencies, and as such it can continue to operate normally even in the presence of rapid changes of the modulation frequency.

We have also shown that all of these techniques might also benefit from being used in combination with each other or existing techniques. For example, MOP traps should allow spontaneous-force based atom traps with highly flexible shapes to be produced, and as such could be used to efficiently and selectively load specific sites within a dipole trap array based on an optical lattice or microlens array. Once atoms have been loaded into these dipole traps, the light required to coherently manipulate them via Raman transitions could be generated by employing one of the carrier removal schemes described in chapter 8 in combination with an electro-optic modulator. The tight focusing of the trapping light that occurs in dipole traps based on truncated spherical cavities might allow trap depths comparable to those that can be achieved in alternative systems to be obtained with a much larger laser detuning, thus reducing the rate at which the atoms scatter light from the trapping beam. Combining this with the use of one of the methods from chapter 8 to generate highly stable light sources at all of the necessary wavelengths, the atoms held in these traps could then be trapped and manipulated for long periods of time with minimal decoherence. In conclusion, the techniques we have demonstrated and worked towards in this thesis are not only of interest individually, but when combined with each other and with existing results they represent a useful extension of the overall toolkit of experimental methods available to atomic physicists.

Appendix A

List of abbreviations

AOM	acousto-optic modulator
CCD	charge-coupled device
DAVLL	dichroic atomic vapour laser lock
EM CCD	electron-multiplying charge coupled device
EOM	electro-optic modulator
MOP	metastable optical pumping
MOT	magneto-optical trap
PMMA	poly(methyl methacrylate)

References

[1] Nathan Cooper, James Bateman, Alexander Dunning, and Tim Freegarde. Actively stabilized wavelength-insensitive carrier elimination from an electro-optically modulated laser beam. *J. Opt. Soc. Am. B*, 29(4):646–649, April 2012.

[2] N. Cooper, J. Woods, J. E. Bateman, A. Dunning, and T. Freegarde. Stabilized fiber-optic Mach-Zehnder interferometer for carrier-frequency rejection. *Applied Optics*, **52**(5713), 2013.

[3] Nathan Cooper and Tim Freegarde. Trapping of 85Rb atoms by optical pumping between metastable hyperfine states. *J. Phys. B*, **46**(215003), 2013.

[4] Hamid Ohadi, Matthew Himsworth, Andre Xuereb, and Tim Freegarde. Magneto-optical trapping and background-free imaging for atoms near nanostructured surfaces. *Optics Express*, 17(25):23003–23009, December 2009.

[5] Masao Takamoto, Feng-Lei Hong, Ryoichi Higashi, and Hidetoshi Katori. An optical lattice clock. *Nature*, 435(7040):321–324, May 2005.

[6] Susannah M. Dickerson, Jason M. Hogan, Alex Sugarbaker, David M. S. Johnson, and Mark A. Kasevich. Multi-axis inertial sensing with long-time point source atom interferometry. *Phys. Rev. Lett.*, **111**(083001), 2013.

[7] Zhou Lin, Xiong Zong-Yuan, Yang Wei, Tang Biao, Peng Wen-Cui, Wang Yi-Bo, Xu Peng, Wang Jin, and Zhan Ming-Sheng. Measurement of Local Gravity via a Cold Atom Interferometer. *Chinese Physics Letters*, 28(1):013701+, January 2011.

[8] J. P. Davis and F. A. Narducci. A proposal for a gradient magnetometer atom interferometer. *Journal of Modern Optics*, 55(19-20):3173–3185, November 2008.

[9] Pierre Cladé, Estefania de Mirandes, Malo Cadoret, Sa"ida G. Khélifa, Catherine Schwob, François Nez, Lucile Julien, and François Biraben. Determination of the Fine Structure Constant Based on Bloch Oscillations of Ultracold Atoms in a Vertical Optical Lattice. *Physical Review Letters*, 96:033001+, January 2006.

[10] Andreas Wicht, Joel M. Hensley, Edina Sarajlic, and Steven Chu. A Preliminary Measurement of the Fine Structure Constant Based on Atom Interferometry. *Physica Scripta*, pages 82+, January 2002.

[11] J. B. Fixler, G. T. Foster, J. M. McGuirk, and M. A. Kasevich. Atom Interferometer Measurement of the Newtonian Constant of Gravity. *Science*, 315(5808):74–77, January 2007.

[12] C. Raman, M. Köhl, R. Onofrio, D. S. Durfee, C. E. Kuklewicz, Z. Hadzibabic, and W. Ketterle. Evidence for a Critical Velocity in a Bose-Einstein Condensed Gas. *Physical Review Letters*, 83(13):2502–2505, September 1999.

[13] Naaman Tammuz, Robert P. Smith, Robert L. D. Campbell, Scott Beattie, Stuart Moulder, Jean Dalibard, and Zoran Hadzibabic. Can a Bose Gas Be Saturated? *Physical Review Letters*, 106:230401+, June 2011.

[14] Iulia Buluta and Franco Nori. Quantum Simulators. *Science*, 326(5949):108–111, October 2009.

[15] Thaddeus D. Ladd, Fedor Jelezko, Raymond Laflamme, Yasunobu Nakamura, Christopher Monroe, and Jeremy L. O'Brien. Quantum Computing. *Nature*, 464(7285):45–53, September 2010.

[16] G. Wilpers, T. Binnewies, C. Degenhardt, U. Sterr, J. Helmcke, and F. Riehle. Optical Clock with Ultracold Neutral Atoms. *Physical Review Letters*, 89:230801+, November 2002.

[17] G. A. Kazakov, A. N. Litvinov, V. I. Romanenko, L. P. Yatsenko, A. V. Romanenko, M. Schreitl, G. Winkler, and T. Schumm. Performance of a 229Thorium solid-state nuclear clock. *New Journal of Physics*, 14(8):083019+, August 2012.

[18] S. Merlet, Q. Bodart, N. Malossi, A. Landragin, Pereira Dos Santos, O. Gitlein, and L. Timmen. Comparison between two mobile absolute gravimeters: optical versus atomic interferometers. *Metrologia*, 47(4):L9+, August 2010.

[19] J. E. Debs, P. A. Altin, T. H. Barter, D. Döring, G. R. Dennis, G. McDonald, R. P. Anderson, J. D. Close, and N. P. Robins. Cold-atom gravimetry with a Bose-Einstein condensate. *Physical Review A*, 84:033610+, September 2011.

[20] Matthew L. Terraciano, Spencer E. Olson, Mark Bashkansky, Zachary Dutton, and Fredrik K. Fatemi. Magnetically controlled velocity selection in a cold-atom sample using stimulated Raman transitions. *Physical Review A*, 76:053421+, November 2007.

[21] Yannick Bidel, Olivier Carraz, Renée Charrière, Malo Cadoret, Nassim Zahzam, and Alexandre Bresson. Compact cold atom gravimeter for field applications. *Applied Physics Letters*, 102(14):144107+, 2013.

[22] David P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschr. Phys.*, 48(9-11):771–783, September 2000.

[23] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys. Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade. *Physical Review Letters*, 104(1):010502+, January 2010.

[24] Farrokh Vatan and Colin Williams. Optimal quantum circuits for general two-qubit gates. *Physical Review A*, 69:032315+, March 2004.

[25] Immanuel Bloch. Ultracold quantum gases in optical lattices. *Nature Physics*, 1(1):23–30, October 2005.

[26] Waseem S. Bakr, Jonathon I. Gillen, Amy Peng, Simon Folling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature*, 462(7269):74–77, November 2009.

[27] Jens Kruse, Christian Gierl, Malte Schlosser, and Gerhard Birkl. Reconfigurable, site-selective manipulation of atomic quantum systems in two-dimensional arrays of dipole traps. *Phys. Rev. A*, **81**(060308(R)), 2010.

[28] N. Poli, F. Y. Wang, M. G. Tarallo, A. Alberti, M. Prevedelli, and G. M. Tino. Precision Measurement of Gravity with Cold Atoms in an Optical Lattice and Comparison with a Classical Gravimeter. *Phys. Rev. Lett.*, 106:038501+, January 2011.

[29] G. M. Tino, L. Cacciapuoti, K. Bongs, Ch Bordé, P. Bouyer, H. Dittus, W. Ertmer, A. Görlitz, M. Inguscio, A. Landragin, P. Lemonde, C. Lammerzahl, A. Peters, E. Rasel, J. Reichel, C. Salomon, S. Schiller, W. Schleich, K. Sengstock, U. Sterr, and M. Wilkens. Atom interferometers and optical atomic clocks: New quantum sensors for fundamental physics experiments in space. *Nuclear Physics B - Proceedings Supplements*, 166:159–165, April 2007.

[30] Holger Muller, Achim Peters, and Steven Chu. A precision measurement of the gravitational redshift by the interference of matter waves. *Nature*, 463(7283):926–929, February 2010.

[31] Supurna Sinha and Joseph Samuel. Atom interferometry and the gravitational redshift. *Classical and Quantum Gravity*, 28(14):145018+, July 2011.

[32] Peter Wolf, Luc Blanchet, Christian J. Borde, Serge Reynaud, Christophe Salomon, and Claude Cohen-Tannoudji. Atom gravimeters and gravitational redshift. *Nature*, 467(7311):E1, September 2010.

[33] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. Optical dipole traps for neutral atoms. *Adv. At. Mol. Opt. Phys.*, 42:95–170, February 2000.

[34] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23):2631–2634, December 1987.

[35] Silvia Bergamini, Benoît Darquié, Matthew Jones, Lionel Jacubowiez, Antoine Browaeys, and Philippe Grangier. Holographic generation of microtrap arrays for

single atoms by use of a programmable phase modulator. *J. Opt. Soc. Am. B*, 21(11):1889–1894, November 2004.

[36] David G. Grier and Yael Roichman. Holographic optical trapping. *Appl. Opt.*, 45(5):880–887, February 2006.

[37] R Dumke, M Volk, T Muther, F Buchkremer, G Birkl, and W Ertmer. Micro-optical realization of arrays of selectively addressable dipole traps: a scalable configuration for quantum computation with atomic qubits. *PRL*, 89:097903, 2002.

[38] S. Pollock, J. P. Cotter, A. Laliotis, F. Ramirez-Martinez, and E. A. Hinds. Characteristics of integrated magneto-optical traps for atom chips. *New Journal of Physics*, 13(4):043029+, April 2011.

[39] S. Pollock, J. P. Cotter, A. Laliotis, and E. A. Hinds. Integrated magneto-optical traps on a chip using silicon pyramid structures. *Opt. Express*, 17(16):14109–14114, August 2009.

[40] Gregory W. Hoth, Elizabeth A. Donley, and John Kitching. Atom number in magneto-optic traps with millimeter scale laser beams. *Opt. Lett.*, 38(5):661–663, March 2013.

[41] Y Jiang-Ping, G Wei-Jian, and H Jian-Jun. Arrays of microscopic magnetic traps for cold atoms and their applications in atom optics. *Chinese Phys.*, 11:472, 2002.

[42] J Yin, W Gao, J Hu, and Y Wang. Magnetic surface microtraps for realizing an array of alkali atomic BoseEinstein condensates or Bose clusters. *Optics Communications*, 206:99–113, 2002.

[43] Philip N. Bartlett, Peter R. Birkin, and Mohamed A. Ghanem. Electrochemical deposition of macroporous platinum, palladium and cobalt films using polystyrene latex sphere templates. *Chem. Commun.*, 0(17):1671–1672, 2000.

[44] N. Schlosser, G. Reymond, and P. Grangier. Collisional Blockade in Microscopic Optical Dipole Traps. *Physical Review Letters*, 89(2):023005+, June 2002.

[45] Meng K. Tey, Zilong Chen, Syed A. Aljunid, Brenda Chng, Florian Huber, Gleb Maslennikov, and Christian Kurtsiefer. Strong interaction between light and a

single trapped atom without the need for a cavity. *Nat Phys*, 4(12):924–927, December 2008.

[46] G. Hétet, L. Slodička, A. Glätzle, M. Hennrich, and R. Blatt. QED with a spherical mirror. *Physical Review A*, 82(6):063812+, December 2010.

[47] Jingzhi Wu and Xiaosong Gan. Three dimensional nanoparticle trapping enhanced by surface plasmon resonance. *Opt. Express*, 18(26):27619–27626, December 2010.

[48] Andy Rundquist, Anatoly Efimov, and David H. Reitze. Pulse shaping with the Gerchberg-Saxton algorithm. *J. Opt. Soc. Am. B*, 19(10):2468–2478, October 2002.

[49] Gavin Sinclair, Jonathan Leach, Pamela Jordan, Graham Gibson, Eric Yao, Zsolt Laczik, Miles Padgett, and Johannes Courtial. Interactive application in holographic optical tweezers of a multi-plane Gerchberg-Saxton algorithm for three-dimensional light shaping. *Opt. Express*, 12(8):1665–1670, April 2004.

[50] Helmar Bender, Christian Stehle, Claus Zimmermann, Sebastian Slama, Johannes Fiedler, Stefan Scheel, Stefan Y. Buhmann, and Valery N. Marachevsky. Probing atom-surface interactions by diffraction of Bose-Einstein condensates. *Phys. Rev. X*, **4**(011029), 2014.

[51] J. D. Carter, O. Cherry, and J. D. D. Martin. Electric-field sensing near the surface microstructure of an atom chip using cold Rydberg atoms. *Physical Review A*, 86:053401+, November 2012.

[52] D. E. Pritchard, E. L. Raab, V. Bagnato, C. E. Wieman, and R. N. Watts. Light Traps Using Spontaneous Forces. *Physical Review Letters*, 57(3):310–313, July 1986.

[53] P. Bouyer, P. Lemonde, M. Ben Dahan, A. Michaud, C. Salomon, and J. Dalibard. An Atom Trap Relying on Optical Pumping. *EPL (Europhysics Letters)*, 27(8):569+, July 2007.

[54] T. Walker, P. Feng, D. Hoffmann, and R. S. Williamson. Spin-polarized spontaneous-force atom trap. *Physical Review Letters*, 69:2168–2171, October 1992.

[55] E. A. Hinds and I. G. Hughes. Magnetic atom optics: mirrors, guides, traps, and chips for atoms. *Journal of Physics D: Applied Physics*, 32(18):R119–R146, September 1999.

[56] Nir Davidson, Heun J. Lee, Mark Kasevich, and Steven Chu. Raman cooling of atoms in two and three dimensions. *Physical Review Letters*, 72(20):3158–3161, May 1994.

[57] C. J. Foot. *Atomic Physics.* Oxford University Press, 2005.

[58] Steck D 2012 http://steck.us/alkalidata (revision 2.1.5).

[59] H. J. Metcalf and P. van der Straaten. *Laser Cooling and Trapping.* Springer, 1999.

[60] M. Himsworth and T. Freegarde. Rubidium pump-probe spectroscopy: Comparison between *ab initio* theory and experiment. *Physical Review A*, 81(2):023423+, February 2010.

[61] P. F. Barker. Doppler Cooling a Microsphere. *Physical Review Letters*, 105(7):073002+, August 2010.

[62] E. S. Shuman, J. F. Barry, and D. DeMille. Laser cooling of a diatomic molecule. *Nature*, 467(7317):820–823, October 2010.

[63] T. H. Bergeman, Patrick McNicholl, Jan Kycia, Harold Metcalf, and N. L. Balazs. Quantized motion of atoms in a quadrupole magnetostatic trap. *J. Opt. Soc. Am. B*, 6(11):2249–2256, November 1989.

[64] Wolfgang Petrich, Michael H. Anderson, Jason R. Ensher, and Eric A. Cornell. Stable, Tightly Confining Magnetic Trap for Evaporative Cooling of Neutral Atoms. *Physical Review Letters*, 74(17):3352–3355, April 1995.

[65] R. Dubessy, K. Merloti, L. Longchambon, P. E. Pottie, T. Liennard, A. Perrin, V. Lorent, and H. Perrin. Rubidium-87 Bose-Einstein condensate in an optically plugged quadrupole trap. *Physical Review A*, 85:013643+, January 2012.

[66] B. Ohayon and G. Ron. New approaches in designing a Zeeman Slower. *Journal of Instrumentation*, 8(02):P02016+, February 2013.

[67] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6(11):2023–2045, November 1989.

[68] C. D. Wallace, T. P. Dinneen, K. Y. N. Tan, A. Kumarakrishnan, P. L. Gould, and J. Javanainen. Measurements of temperature and spring constant in a magneto-optical trap. *J. Opt. Soc. Am. B*, 11(5):703–711, May 1994.

[69] P. Kohns, P. Buch, W. Süptitz, C. Csambal, and W. Ertmer. On-Line Measurement of Sub-Doppler Temperatures in a Rb Magneto-optical Trap-by-Trap Centre Oscillations. *EPL (Europhysics Letters)*, 22(7):517+, June 1993.

[70] A. Ridinger, S. Chaudhuri, T. Salez, U. Eismann, D. R. Fernandes, K. Magalhães, D. Wilkowski, C. Salomon, and F. Chevy. Large atom number dual-species magneto-optical trap for fermionic 6Li and 40K atoms. *EPJ D*, 65(1-2):223–242, 2011.

[71] C. Townsend, N. Edwards, K. Zetie, C. Cooper, J. Rink, and C. Foot. High-density trapping of cesium atoms in a dark magneto-optical trap. *Physical Review A*, 53(3):1702–1714, March 1996.

[72] Wolfgang Petrich, Michael H. Anderson, Jason R. Ensher, and Eric A. Cornell. Behavior of atoms in a compressed magneto-optical trap. *J. Opt. Soc. Am. B*, 11(8):1332–1335, August 1994.

[73] Nir Davidson, Heun Jin Lee, Charles Adams, Mark Kasevich, and Steven Chu. Long Atomic Coherence Times in an Optical Dipole Trap. *Physical Review Letters*, 74(8):1311–1314, February 1995.

[74] A. Laliotis, M. Trupke, J. P. Cotter, G. Lewis, M. Kraft, and E. A. Hinds. ICP polishing of silicon for high-quality optical resonators on a chip. *Journal of Micromechanics and Microengineering*, 22(12):125011+, December 2012.

[75] See for example www.thermoscientific.com/ecomm/servlet/productscatalog.

[76] Hai-Chau Chang and Lih-Chung Wang. A Simple Proof of Thue's Theorem on Circle Packing. Repository publication only at http://arxiv.org/abs/1009.4322, September 2010.

[77] S. Coyle, G. V. Prakash, J. J. Baumberg, M. Abdelsalem, and P. N. Bartlett. Spherical micromirrors from templated self-assembly: Polarization rotation on the micron scale. *Applied Physics Letters*, 83(4):767–769, July 2003.

[78] S. Coyle, M. Netti, J. Baumberg, M. Ghanem, P. Birkin, P. Bartlett, and D. Whittaker. Confined Plasmons in Metallic Nanocavities. *Physical Review Letters*, 87(17):176801+, October 2001.

[79] Leo K. E. B. Serrona, A. Sugimura, R. Fujisaki, T. Okuda, N. Adachi, H. Ohsato, I. Sakamoto, A. Nakanishi, and M. Motokawa. Magnetic and structural properties of NdFeB thin film prepared by step annealing. *Materials Science and Engineering: B*, 97(1):59–63, January 2003.

[80] J. Goldwin and E. A. Hinds. Tight focusing of plane waves from micro-fabricated spherical mirrors. *Opt. Express*, 16(22):17808–17816, October 2008.

[81] J. D. Jackson. *Classical Electrodynamics*. Wiley, 1998.

[82] Thanassis Speliotis, Dimitris Niarchos, Vassil Skumryev, Yong Zhang, and George Hadjipanayis. Effect of post deposition annealing on the hysteresis loops of sputtered NdFeB films. *Journal of Magnetism and Magnetic Materials*, 272-276:E877–E879, May 2004.

[83] See for example www.optisourcellc.com/waveplates_ut.html.

[84] Magnus Haw, Nathan Evetts, Will Gunton, Janelle Van Dongen, James L. Booth, and Kirk W. Madison. Magneto-optical trap loading rate dependence on trap depth and vapor density. *J. Opt. Soc. Am. B*, 29(3):475–483, March 2012.

[85] K. Lindquist, M. Stephens, and C. Wieman. Experimental and theoretical study of the vapor-cell Zeeman optical trap. *Physical Review A*, 46(7):4082–4090, October 1992.

[86] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard. Phase-space density in the magneto-optical trap. *Physical Review A*, 52(2):1423–1440, August 1995.

[87] S. R. Muniz, K. Magalh F. M. Aes, Ph, M. A. Perez, L. G. Marcassa, and V. S. Bagnato. Measurements of capture velocity in a magneto-optical trap for a broad range of light intensities. *Physical Review A*, 65:015402+, December 2001.

[88] J. Reichel, W. Hänsel, and T. Hänsch. Atomic Micromanipulation with Magnetic Surface Traps. *Physical Review Letters*, 83(17):3398–3401, October 1999.

[89] M. A. Clifford, G. P. T. Lancaster, R. H. Mitchell, F. Akerboom, and K. Dholakia. Realization of a mirror magneto-optical trap. *Journal of Modern Optics*, 48(6):1123–1128, May 2001.

[90] R Folman, P Kruger, D Cassettari, B Hessmo, T Maler, and J Schmiedmayer. Controlling Cold Atoms using Nanofabricated Surfaces: Atom Chips . *PRL*, 84:4749, 2000.

[91] M Hyodo, K Nakayama, R Ohmukai, K Kurihara, and M Watanabe. Mirror magneto-optical trap using circularly polarized light-emitting optical fibers. *Applied Optics*, 15:3629–3633, 2006.

[92] M Hyodo, K Nakayama, M Watanabe, and R Ohmukai. Mirror magneto-optical trap exploiting hexapole-compensated magnetic field. *PRA*, 76:013419, 2007.

[93] J. D. Weinstein and K. G. Libbrecht. Microscopic magnetic traps for neutral atoms. *Physical Review A*, 52:4004–4009, November 1995.

[94] Richard J. Cook and Richard K. Hill. An electromagnetic mirror for neutral atoms. *Optics Communications*, 43(4):258–260, October 1982.

[95] A. S. Arnold and P. J. Manson. Atomic density and temperature distributions in magneto-optical traps. *J. Opt. Soc. Am. B*, 17(4):497–506, April 2000.

[96] Umakant D. Rapol, Ajay Wasan, and Vasant Natarajan. Loading of a rubidium magneto-optic trap from a getter source. *Physical Review A*, 64(2):023402+, June 2001.

[97] N. R. Newbury, C. J. Myatt, and C. E. Wieman. s-wave elastic collisions between cold ground-state 87rb atoms. *Physical Review A*, 51:R2680–R2683, April 1995.

[98] Dipankar Nath, R. Kollengode Easwaran, G. Rajalakshmi, and C. S. Unnikrishnan. Accelerated Thermalisation of 39K atoms in a Magnetic Trap with Superimposed Optical Potential. *J. Phys. B*, **46**(155303), 2013.

[99] Kristan L. Corwin, Zheng-Tian Lu, Carter F. Hand, Ryan J. Epstein, and Carl E. Wieman. Frequency-Stabilized Diode Laser with the Zeeman Shift in an Atomic Vapor. *Appl. Opt.*, 37(15):3295–3298, May 1998.

[100] K. C. Harvey and C. J. Myatt. External-cavity diode laser using a grazing-incidence diffraction grating. *Opt. Lett.*, 16(12):910–912, June 1991.

[101] C. J. Cooper, G. Hillenbrand, J. Rink, C. G. Townsend, K. Zetie, and C. J. Foot. The Temperature of Atoms in a Magneto-optical Trap. *EPL (Europhysics Letters)*, 28(6):397+, November 1994.

[102] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Görlitz, S. Gupta, D. E. Pritchard, and W. Ketterle. Transport of Bose-Einstein Condensates with Optical Tweezers. *Physical Review Letters*, 88:020401+, December 2001.

[103] M. Meucci, E. Mariotti, P. Bicchi, C. Marinelli, and L. Moi. Light-Induced Atom Desorption. *EPL (Europhysics Letters)*, 25(9):639+, March 1994.

[104] Alexander M. Akulshin, Russell J. McLean, Andrei I. Sidorov, and Peter Hannaford. Coherent and collimated blue light generated by four-wave mixing in Rb vapour. *Opt. Express*, 17(25):22861–22870, December 2009.

[105] O. S. Heavens. Radiative Transition Probabilities of the Lower Excited States of the Alkali Metals. *J. Opt. Soc. Am.*, 51(10):1058–1061, October 1961.

[106] J. E. Sansonetti. Wavelengths, Transition Probabilities, and Energy Levels for the Spectra of Rubidium (Rb I through Rb XXXVII). *Journal of Physical and Chemical Reference Data*, 35(1):301–421, 2006.

[107] National institute for standards and technology's atomic line database, www.nist.gov.

[108] Alexander Akulshin, Christopher Perrella, Gar-Wing Truong, Russell McLean, and Andre Luiten. Frequency evaluation of collimated blue light generated by wave mixing in Rb vapour. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(24):245503+, December 2012.

[109] Low-temperature high-density magneto-optical trapping of potassium using the open transition at 405 nm. *Physical Review A*, 84:063420+, December 2011.

[110] Nicolas Schlosser, Georges Reymond, Igor Protsenko, and Philippe Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027, June 2001.

[111] Jacek Fiutowski, Dobrosława Bartoszek-Bober, Tomasz Dohnalik, and Tomasz Kawalec. Evanescent wave mirror for cold atomsA quasi-resonant case. *Optics Communications*, 297:59–64, June 2013.

[112] D. E. Chang, J. D. Thompson, H. Park, V. Vuletić, A. S. Zibrov, P. Zoller, and M. D. Lukin. Trapping and Manipulation of Isolated Atoms Using Nanoscale Plasmonic Structures. *Physical Review Letters*, 103(12):123004+, September 2009.

[113] A. L. Marchant, S. Händel, T. P. Wiles, S. A. Hopkins, and S. L. Cornish. Guided transport of ultracold gases of rubidium up to a room-temperature dielectric surface. *New Journal of Physics*, 13(12):125003+, December 2011.

[114] Laurence Pruvost, Dan Marescaux, Olivier Houde, and Hong T. Duong. Guiding and cooling of cold atoms in a dipole guide. *Optics Communications*, 166(1-6):199–209, August 1999.

[115] 11 W narrow linewidth laser source at 780nm for laser cooling and manipulation of Rubidium. *Optics Express*, 20(8915), 2012.

[116] S. Handel, A. L. Marchant, T. P. Wiles, S. A. Hopkins, and S. L. Cornish. Magnetic transport apparatus for the production of ultracold atomic gases in the vicinity of a dielectric surface. *Review of Scientific Instruments*, 83(1):013105–013105–10, January 2012.

[117] Markus Greiner, Immanuel Bloch, Theodor W. Hänsch, and Tilman Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Physical Review A*, 63(3):031401+, February 2001.

[118] Matthew J. Pritchard, Aidan S. Arnold, David A. Smith, and Ifan G. Hughes. Single-impulse magnetic focusing of launched cold atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 37(22):4435+, November 2004.

[119] Saijun Wu, Wilbert Rooijakkers, Pierre Striehl, and Mara Prentiss. Bidirectional propagation of cold atoms in a “stadium”-shaped magnetic guide. *Physical Review A*, 70:013409+, July 2004.

[120] József Fortágh and Claus Zimmermann. Magnetic microtraps for ultracold atoms. *Reviews of Modern Physics*, 79(1):235–289, February 2007.

[121] H. Ott, J. Fortagh, G. Schlotterbeck, A. Grossmann, and C. Zimmermann. Bose-Einstein Condensation in a Surface Microtrap. *Physical Review Letters*, 87(23):230401+, November 2001.

[122] M. P. A. Jones, C. J. Vale, D. Sahagun, B. V. Hall, C. C. Eberlein, B. E. Sauer, K. Furusawa, D. Richardson, and E. A. Hinds. Cold atoms probe the magnetic field near a wire. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 37(2):L15+, January 2004.

[123] A. Ashkin and J. P. Gordon. Stability of radiation-pressure particle traps: an optical Earnshaw theorem. *Opt. Lett.*, 8(10):511–513, October 1983.

[124] *Inductively Coupled Magnetic Ring Trap for Atom Interferometry*, OSA Technical Digest (online). Optical Society of America, October 2012.

[125] Jin Y. Zhang, Si C. Ji, Zhu Chen, Long Zhang, Zhi D. Du, Bo Yan, Ge S. Pan, Bo Zhao, You J. Deng, Hui Zhai, Shuai Chen, and Jian W. Pan. Collective Dipole Oscillations of a Spin-Orbit Coupled Bose-Einstein Condensate. *Physical Review Letters*, 109:115301+, September 2012.

[126] J. C. Jaskula, G. B. Partridge, M. Bonneau, R. Lopes, J. Ruaudel, D. Boiron, and C. I. Westbrook. Acoustic Analog to the Dynamical Casimir Effect in a Bose-Einstein Condensate. *Physical Review Letters*, 109:220401+, November 2012.

[127] Marino A.M. and Stroud C.R. Jr. Phase-locked laser system for use in atomic coherence experiments. *Review of Scientific Instruments*, 79(1):013104+, 2008.

[128] Mark Kasevich, David S. Weiss, Erling Riis, Kathryn Moler, Steven Kasapi, and Steven Chu. Atomic velocity selection using stimulated Raman transitions. *Physical Review Letters*, 66(18):2297–2300, May 1991.

[129] Michael J. Snadden, Roger B. M. Clarke, and Erling Riis. Injection-locking technique for heterodyne optical phase locking of a diode laser. *Opt. Lett.*, 22(12):892–894, June 1997.

[130] J. Ringot, Y. Lecoq, J. C. Garreau, and P. Szriftgiser. Generation of phase-coherent laser beams for Raman spectroscopy and cooling by direct current modulation of a diode laser. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 7(3):285–288, October 1999.

[131] F. Quinlan, S. Ozharar, S. Gee, and P. J. Delfyett. Harmonically mode-locked semiconductor-based lasers as high repetition rate ultralow noise pulse train and optical frequency comb sources. *Journal of Optics A: Pure and Applied Optics*, 11(10):103001+, October 2009.

[132] S. Ozharar, I. Ozdur, F. Quinlan, and P. J. Delfyett. Optical frequency comb generation by direct modulation of CW light. *SPIE Proceedings*, 6975, March 2008.

[133] P. Palm, D. Hanke, W. Urban, and M. Mürtz. Ultrahigh-resolution spectrometer for the 5 micrometer wavelength region. *Opt. Lett.*, 26(9):641–643, May 2001.

[134] J. I. Thorpe, K. Numata, and J. Livas. Laser frequency stabilization and control through offset sidebandlocking to optical cavities. *Opt. Express*, 16(20):15980–15990, September 2008.

[135] Ana R. E. Brás, Teresa Casimiro, Jorge Caldeira, and Ana Aguiar-Ricardo. Solubility of the Nematic Liquid Crystal E7 in Supercritical Carbon Dioxide. *J. Chem. Eng. Data*, 50(6):1857–1860, October 2005.

[136] Shin-Tson Wu, Uzi Efron, and LaVerne D. Hess. Birefringence measurements of liquid crystals. *Appl. Opt.*, 23(21):3911–3915, November 1984.

[137] M. Kaczmarek, A. Dyadyusha, S. Slussarenko, and I. C. Khoo. The role of surface charge field in two-beam coupling in liquid crystal cells with photoconducting polymer layers. *Journal of Applied Physics*, 96(5):2616–2623, 2004.

[138] T. W. Hansch and B. Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. *Optics Communications*, 35(3):441–444, December 1980.

[139] J. E. Bateman, R. L. D. Murray, M. Himsworth, H. Ohadi, A. Xuereb, and T. Freegarde. Hänsch-couillaud locking of mach-zehnder interferometer for carrier removal from a phase-modulated optical spectrum. *JOSA B*, **27**(8), 2010.

[140] C. E. Liekhus-Schmaltz, R. Mantifel, M. Torabifard, I. B. Burgess, and J. D. D. Martin. Injection-locked diode laser current modulation for Pound-Drever-Hall frequency stabilization using transfer cavities. *J. Opt. Soc. Am. B*, 29(6):1394–1398, June 2012.

[141] Richard P. Abel, Ulrich Krohn, Paul Siddons, Ifan G. Hughes, and Charles S. Adams. Faraday dichroic beam splitter for Raman light using an isotopically pure alkali-metal-vapor cell. *Opt. Lett.*, 34(20):3071–3073, October 2009.

[142] D. Haubrich, M. Dornseifer, and R. Wynands. Lossless beam combiners for nearly equal laser frequencies. *Review of Scientific Instruments*, 71(2):338+, 2000.

[143] L. B. Soldano, A. I. de Vreede, M. K. Smit, B. H. Verbeek, E. G. Metaal, and F. H. Green. Mach-Zehnder interferometer polarization splitter in InGaAsP/InP. *IEEE Photonics Technology Letters*, 6(3):402–405, March 1994.

[144] Artur K. Ekert, John G. Rarity, Paul R. Tapster, and G. Massimo Palma. Practical quantum cryptography based on two-photon interferometry. *Physical Review Letters*, 69:1293–1295, August 1992.

[145] Ping Lu, Liqiu Men, Kevin Sooley, and Qiying Chen. Tapered fiber Mach–Zehnder interferometer for simultaneous measurement of refractive index and temperature. *Applied Physics Letters*, 94(13):131110+, April 2009.

[146] I. Dotsenko, W. Alt, S. Kuhr, D. Schrader, M. Müller, Y. Miroshnychenko, V. Gomer, A. Rauschenbeutel, and D. Meschede. Application of electro-optically generated light fields for Raman spectroscopy of trapped cesium atoms. *Applied Physics B: Lasers and Optics*, 78(6):711–717, April 2004.

[147] G. B. Xavier and J. P. von der Weid. Stable single-photon interference in a 1 km fiber-optic Mach-Zehnder interferometer with continuous phase adjustment. *Opt. Lett.*, 36(10):1764–1766, May 2011.

[148] Axel Schenzle, Ralph G. DeVoe, and Richard G. Brewer. Phase-modulation laser spectroscopy. *Physical Review A*, 25:2606–2621, May 1982.

[149] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(97), 1983.

[150] F. V. Kowalski, J. A. Squier, and J. T. Pinckney. Pulse generation with an acousto-optic frequency shifter in a passive cavity. *Applied Physics Letters*, 50(12):711–713, 1987.

[151] Alexander Dunning, Rachel Gregory, James Bateman, Mathew Himsworth, Nathan Cooper, and Tim Freegarde. Broadband zeeman-degenerate velocity-sensitive raman pulses with composite rotations in an extended atomic cloud (in preparation).

[152] Wolfgang Ketterle, Kendall B. Davis, Michael A. Joffe, Alex Martin, and David E. Pritchard. High densities of cold atoms in a *dark* spontaneous-force optical trap. *Physical Review Letters*, 70(15):2253–2256, April 1993.

[153] T. A. Kelf, Y. Sugawara, R. M. Cole, J. J. Baumberg, M. E. Abdelsalam, S. Cintra, S. Mahajan, A. E. Russell, and P. N. Bartlett. Localized and delocalized plasmons in metallic nanovoids. *Physical Review B*, 74(24):245415+, December 2006.

[154] T. W. Ebbesen, H. J. Lezec, H. F. Ghaemi, T. Thio, and P. A. Wolff. Extraordinary optical transmission through sub-wavelength hole arrays. *Nature*, 391(6668):667–669, February 1998.