Integration of Magneto Optical Traps in Atom Chips

Samuel Pollock

Centre for Cold Matter
Department of Physics
Imperial College London

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy of the University of London
and the Diploma of Membership of Imperial College February 2010

Summary

Integration of Magneto Optical Traps in Atom Chips

Samuel Pollock

This thesis describes the manufacture and demonstration of microfabricated hollow pyramidal mirrors for the purposes of integrated laser cooling and trapping of atoms on an atom chip. A single incident circularly polarised laser beam is reflected inside the pyramid hollow creating all the required beams of the correct polarisation to create a magneto optical trap (MOT). Many pyramids can be manufactured on the same device to produce many cold atom sources as the etching process is intrinsically scalable, with the number of simultaneous traps possible only limited by the size of the input beam.

The pyramids created in silicon have an apex angle of 70° , and the principles of a MOT created from the resultant irregular beam geometry were studied in a macroscopic glass model of a 70° pyramid. The scaling of atom number with trap dimensions at small scales was experimentally found to be well described by a power law with an exponent of 6. This value was confirmed with simple theoretical considerations and numerical simulations of the motion of incoming atoms into the pyramids.

An imaging scheme capable of resolving the atomic fluorescence from a strong background of scattered light from the mirror surfaces was developed which offers the prospect of imaging small MOTs containing on the order of 100 atoms inside the pyramids.

The processes required to create the pyramids in silicon by anisotropic etching in potassium hydroxide are refined to produce structures on a mm-scale. The resultant surfaces are smoothed using a process of isotropic plasma etching and finally coated with a metallic layer of suitable reflectivity to create mirrors. These devices were tested and the cooling and capture of rubidium atoms from a background vapour was demonstrated to result in a MOT of several 1000 atoms at a temperature of $\sim 170\mu$K in a 4.2mm wide (3.0mm deep) silicon pyramid.

Acknowledgments

Firstly I would like to thank my supervisor Ed Hinds, not only for giving me this fantastic opportunity to explore science but also for being an inspirational scientist and leader. Ed's remarkable insight and intuition has been a great motivator and I thank him for the support and encouragement he has shown in enabling me to reach the point of being able to sit down and write this thesis.

Primarily my thanks has to be expressed to Athanasios Laliotis and Joe Cotter. We all expended a great volume of blood, sweat and tea in our small corner of Bay 2, and it was an absolute pleasure and a joy to work with you both. I cannot thank you both enough for the guidance and advice offered in the composition of this thesis. I would also like to express my gratitude to my predecessor Fernando Ramirez-Martinez for all his help at the start of my PhD.

I regret there isn't the space to thank all the students and postdocs working in the Centre for Cold Matter individually. Not only has every single person who has worked there offered help and assistance when asked, they have done it willingly and often gone the extra mile to really tackle a problem. I credit the good humour and friendships formed at CCM for making my time here incredibly enjoyable, even in the darkest hours of research frustration I could always rely on someone to lift my spirits with some good banter over a cuppa.

Special mentions must go to the efforts of Jon Dyne and Steve Maine in the CCM workshop, who produced works of art from hastily drawn schematics that I had scrawled on the back of an envelope. Jon especially deserves a mention for the invaluable help he has been with pumps and vacuum systems. The final person I would like to thank from Imperial is Sanja Maricic, for without her tireless efforts and assistance most of this work would not have been possible.

Outside of the lab my thanks go to my family and my muse and inspiration Jenny for her encouragement in the difficult times and sharing the celebration in the times of prosperity.

Contents

List of Figures

List of Tables

Chapter 1

Atom Chips

This chapter details the motivation behind the research presented in this thesis. Initially an overview of the field of ‘atom chips’ will be provided. I have chosen to break this down into a discussion of ‘generations’ of experiments, where each generation builds upon the advances and discoveries made in the previous one. It is not the intention to imply that experiments labelled as an early generation are inferior to later generations, but instead that later generations use a lot of now well-established techniques developed by early generations.

1.1 The Zeroth Generation - Pioneers

The focus of this thesis is research and development of new techniques for the advancement of atom chip technology. The basic concept of the atom chip is a miniaturised device to perform experiments on cold trapped atoms, essentially being an atomic-physics lab on a chip.

Atom chips arose from the field of atom-optics, where the motion of neutral atoms is controlled using optical, magnetic and electric fields. The field of atom-optics was characterised by the celebrated experiments of Stern & Gerlach, Dunoyer and Estermann in the early part of the 20th century. It underwent a revolution in the late 1970s with proposals from Hänsch and Schawlow [1], and Wineland and Dehmelt [2] that laser light could be used to cool atoms and reduce their kinetic energy. This was experimentally realised for neutral atoms in 1982 by Philips and Metcalf who were able to demonstrate reduction of the atomic velocity by 40% [3]. The same researchers went on to apply laser cooling to the practical realisation of magnetic trapping of neutral atoms, first demonstrated in 1985 [4]. The interaction of alkali atoms (such as Li, K, Cs etc.) with a magnetic field is particularly strong due to the presence of an unpaired electron which results in an atomic magnetic moment of the order of the Bohr magneton. The interaction energy between an atom and a magnetic field ($\approx \mu_B |B|$) is still significantly weaker than an atom’s typical thermal energy, but the new laser cooling techniques of Hänsch and Schawlow allowed for the reduction of this thermal energy to the below the magnitude of the trapping potential. The ability to manipulate atoms in magnetic traps in this way eventually led to the first realisation of Bose-Einstein condensation (BEC) in a dilute gas of alkali atoms[5].

Magnetic traps on a chip using microfabricated conductors were originally proposed as one of many novel magnetic trap concepts in mid-1998 [6]. The small dimensions of the microscopic conductors create steep gradients and large curvature leading to tightly confining potentials for atoms a few microns away from the wires with much lower power consumption [7]. This is a great advantage for the preparation of BECs, produced using radio-frequency evaporation of atoms in magnetic traps, as the rate at which evaporation can proceed is faster in tighter traps. Magnetically confined atoms in traps formed with microfabricated wires were experimentally demonstrated in 1999 by Reichel et al. [8], and in 2000 by Folman et al.[9]. At roughly the same time, teams at JILA and Harvard [10] [11] had been able to demonstrate guiding of atoms on a chip.

The term 'atom-chip' was quickly adopted to describe integrated atom-optical surface devices, created using state of the art microelectronic, micro electromechanical (MEMS) and photonic microfabrication techniques. The hope was to eventually confine, control and manipulate ultra cold atoms entirely using these techniques, allowing for a new generation of sensors and quantum information devices. A definition of what could compromise an ideal-atom chip was given by Folman et al [9].

> *"A final integrated atom chip should have a reliable source of cold atoms, for example, a BEC, with an efficient loading mechanism, single mode guides for coherent transportation of atoms, nanoscale traps, movable potentials allowing controlled collisions for the creation of entanglement between atoms, extremely high resolution light fields for the manipulation of individual atoms, and internal state sensitive detection to read out the result of the processes that have occurred (e.g., the quantum computation)."*

Developments in fabrication of microwires using lithographic techniques over the next few years enabled demonstration of new magnetic tools for the manipulation of atoms, such as splitters, switches, linear atom colliders and novel magnetic transporters.

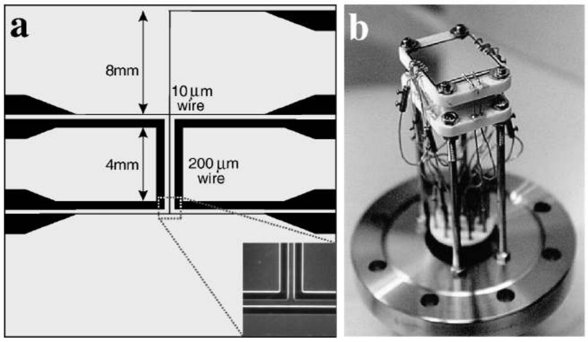

Figure 1.1: One of the original atom chips as demonstrated by Folman et al. in 2000[9]

1.2 Generation I - Towards BEC

The next logical step marking the boundary of a new generation of atom-chips, was to take a sample of cold atoms and cool them to quantum degeneracy using the early atom chips. The possibility of on-chip production of BECs had been strongly debated as it was thought that the local environment of the atom-chip would be too hostile for condensation. However, in 2001 two groups, one at the Max-Planck Institute in Munich and one at Tübingen University, within 4 days of each other successfully produced the first atom-chip Bose-Einstein Condensate [12][13].

Early experiments discovered two major undesirable features arising from magnetic traps created by microwires. Firstly, thermal fluctuation of the charges in the wire generates magnetic field noise causing decoherence through magnetic dipole spin-flip transitions[14]. Secondly, current in microwires does not flow perfectly along the length of the wire but meanders, resulting in uneven magnetic traps which cause a BEC to break up in a process known as 'fragmentation'[15]. The fragmentation of BECs in wire traps has been subsequently reduced by up to two orders of magnitude by reduction in the geometric roughness of the wires using improvements in the microfabrication process[16].

1.2.1 Permanent Magnet atom chips

An alternative to producing magnetic trapping fields on an atom chip using microfabricated conductors is to use microscopic structures of permanent magnetisation [17]. This offers several substantial advantages in that there is no power dissipation, no fluctuations in magnetic field from temporal or spatial current variation, and near field noise from thermal electron movement is avoided [18][1]. Audio tape, floppy discs, videotape, magnetic and magneto-optical films have been used to create microtraps which have all successfully held atoms at trap frequencies up to 1MHz (for comparison, one of the original BEC on a chip experiments created wire traps with transverse frequencies of 6.2kHz [12]).

A chip made from a piece of commercial videotape has been demonstrated by Sinclair et al.[18]. Magnetic traps are formed on the tape by writing a sinusoidal pattern of magnetization along the length of the videotape. A bias field is added to create atomic waveguides, but due to the low susceptibility and high coercivity of the videotape this does not result in erasure of the patterns. BECs have been created in these waveguides and the magnetic traps have been manipulated to form a conveyor belt for cold atoms, transporting them several mm across the chip surface.

It has been demonstrated [19] [20] [15] that patterns of magnetization for atom traps can be created in magneto optical films. Two notable examples of magneto optical materials used are magnetically hard ferrite-garnet materials deposited on a dielectric substrate and Co/Pt thin films which can create traps on the order of 1μm in size. Traps are created by uniformly

---

[1] Permanent magnet atom chips do however still feature fragmentation which arises from imperfections in the magnetic medium

magnetising the film, then locally heating it with a laser in the presence of a weak field applied in the opposite direction to the initial magnetization. Using this technique it is possible to create any desired magnetization pattern in the film, traps can also be reconfigured in-situ, as the only things required are an external magnetic field and a focused laser beam. The magneto-optical films used by Shevchenko et al. [20] are transparent giving superb advantages in optical access, and the ability to create a standard MOT by passing one MOT beam through the chip itself.

1.2.2 Multi-Layer Magnetic traps on Chips

A recent development of note is the formation of a multi-layer atom chip, namely the combination of a 'carrier-chip' with an 'atom-optics' chip. By separating the heavy-duty current carrying structures responsible for trapping and guiding and the micron-sized conductor structures responsible for precision atom-optics manipulation, the atoms can be effectively shielded from the harmful effects of Johnson noise and fragmentation (the atom-optics layer with which the atoms interact only requires a tiny amount of metal). Another benefit with the multilayer structure is that the chip does not need bulky and inaccurate external bias fields, as these can all be generated on the chip[21].

1.3 Generation II - 'Atom-Chip Devices'

As the techniques for production of on-chip BECs began to become established, experiments could begin to exploit the unique features of atom-chips to create miniature integrated atom-optical devices.

1.3.1 Atom Interferometers

The properties of atoms at ultracold temperatures have been used to demonstrate matter-wave interference. Sensitive measurements of quantum phase using an interferometer means highly sensitive devices can be created from atom-chips. For example condensates in interferometers have been proposed for a measurement of the Casimir-Polder force[22], and to make a very sensitive measurement of $\delta g/g$[23]. Atom interferometry offers substantial advantages over current techniques, in the case of Sagnac gyroscopes, an atomic implementation is theoretically capable of improving the best conventional SNR by up to 10^{11}[24].

There are several techniques for realising the atomic implementation of the traditional interferometer component, the beamsplitter, by means of dynamically splitting atoms trapped in a single potential well into a double well (and back again). Atoms can be split using the optical dipole force, or by using the effect of condensate fragmentation arising from inhomogeneities in permanent magnetic trap [23]. However, these processes are generally not coherent, preserving little phase repeatability between experiments [25].

The current most popular technique for atom chips is the use of Radio Frequency (RF) dressed potentials to split a magnetic trap [26]. A BEC of 1.1 μm has been split to separations from 3.4 μm to 80μm using a 50 μm wire. The potential barrier between the two wells can be controlled with extremely high precision, creating either entirely separate, isolated wells or allow access to the tunneling regime[25].

1.3.2 Non-interferometric sensing with BECs

Non-interferometric techniques have also been used to measure fields and forces using atom-chip devices. For example a technique has been devised to sense the nature of electric and magnetic fields close to a current carrying conductor. A one dimensional BEC close to the surface of an atom chip has been used to measure variations in $\vec{E}$ or $\vec{B}$ from small perturbations to the trapping potential causing fragmentation. This allowed for precision measurements to be made close to an accuracy of 10^{-9}T in magnetic field, or the electric field induced by a charge of $10e$ at a distance of 10μm[27]. The close proximity of the atom chip surface has also been used to make a sensitive measurement of the Casimir-Polder force at distances as far away as 5μm[28].

The use of atom chips for sensing offers many advantages over other more tradition methods. For example current magnetometers either offer high sensitivity at low spatial resolution (e.g. with SQUIDs) or high spatial precision with poor sensitivity (Magnetic force microscopy), but not both simultaneously. Sensing with BECs offers a good degree of spatial resolution and sensitivity which conveniently straddles the region between these two current magnetometry techniques[27].

1.3.3 Magnetic Lattices

Periodic optical lattices are extensively used for manipulating ultracold atoms and for performing fundamental physics experiments such as Mott insulator to superfluid transitions [29]. Periodic lattices also have potential application in quantum information science since they may provide registers of single-atom qubits ,and can act as ‘quantum-simulators’ of condensed matter systems [30]. Arrays of magnetic traps, or ‘magnetic lattices’ offer several advantages over their optical counterparts. As the size of the traps is not fixed, the spacing between sites can be made as large as required and is not limited to the wavelength of the laser forming the lattice. This also implies much easier addressing of single sites with lasers. Each individual site typically can also hold larger numbers of atoms ($> 10^5$ atoms per site) than optical lattices. Magnetic lattices in 1D, 2D and 3D in two separate experiments have been demonstrated by Grabowski et al.[31]. The traps can be achieved in an arrangement of overlapping layers of perpendicular wires by varying the current directions in each wire. A figure detailing these configurations is displayed in [31]. Permanent magnetic lattices have been demonstrated by Singh et al. [19] and Whitlock et al. [32]. Whitlock created traps in a 300 nm-thick FePt film, patterned using optical lithography to create an array of Ioffe-Pritchard traps with a density of 1250 traps per square mm. Up to 500 atom clouds have been trapped in each site, and have been cooled to degeneracy. By manipulating the external fields used to complete the traps, the clouds can be transported across the chip surface in an ‘atomic shift register’. This will potentially allow the transfer between storage, interaction and readout areas on a fully integrated device.

1.4 Generation IIb- Let there be light

While not a technological shift warranting definition of a new generation, a significant milestone however has been the integration of optical components onto atom chips[33]. Optical components are essential in atomic physics, primarily being the tools by which atoms are detected, via scattered fluorescence and absorption, or by using light to prepare and probe the internal states of atoms. A second use is the optical dipole force as an alternative or complement to magnetic fields for trapping and manipulation of ultracold atoms, where fixed optical microstructures offer significant advantages in precision and stronger confinement.

1.4.1 Detecting Atoms

Fluorescence and absorption imaging using off chip optical techniques is restricted by the efficiency of detection. Imaging components, such as collection optics, can be integrated onto the chips so by combining these with macroscopic detectors, which offer close to unity quantum efficiency, the optical sensing of atoms can be improved drastically, offering the prospect of detecting single atoms with certainty. If we want to observe small numbers of atoms, they can either be seen through destructive processes (absorption or spontaneous emission) or by non-destructive measurements of the atoms effect on the driving optical field via phase changes[34].

Fluorescence detection can be achieved with optical fibres by mounting two fibres in a perpendicular configuration as illustrated in figure 1.2. A single mode tapered fibre delivers light for pumping while spontaneous emission is detected by a multimode fibre at 90° to the other fibre. Atoms can be guided into the gap between the fibres and probed using this emitter-detector setup [35]. The fibres are affixed to the chip using lithography to create mounting structures. By using extra-thick layers of photoresist a kind of scaffolding can be erected to give the structural support needed to affix the fibre. This arrangement has been shown to have an efficiency of 66% with a high signal-to-noise ratio[36].

Another technique is to use a single fibre with a miniature aspheric lens on the front. Atoms are held in the focus of the fibre lens using a dipole trapping beam passed along the fibre, and simultaneously fluorescence is collected by the same aspheric lens and coupled into the fibre [37].

1.4.2 Cavities

Cavities can enhance the interaction between atoms and a light field. Photons have many chances to interact with an atom as they take multiple trips between the cavity mirrors, increasing the SNR dramatically even in cavities with moderate finesse. The improvement of detection with a cavity offers enhancement of fluorescence and absorption by a factor of $2F/\pi$

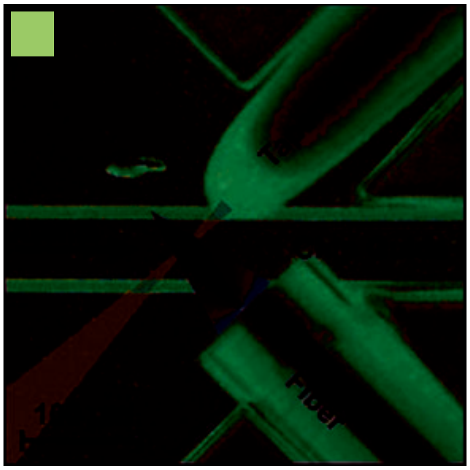

Figure 1.2: Fibre detector as detailed in [36]

where $F \approx \frac{2\pi}{1-\rho}$ is the finesse of a symmetric cavity of high reflectivity ρ [38]. A major advantage of integrating these cavities onto an atom chip, is the improved scaling. As we reduce the dimensions of a cavity, the mode waist of the cavity will concomitantly decrease leading to an increase in the atom-light interaction strength[35].

At Imperial there have been successfully demonstrated micro-fabricated cavities (illustrated in figure 1.3) which have the potential to be incorporated in scalable arrays. Each plano-concave cavity consists of an etched spherical mirror and a coated single-mode optical fibre. Clouds of cold atoms have been dropped through the cavities and the reflected light monitored revealed that in cavities of moderate finesse ($F = 280$) it was possible to detect single atoms [39].

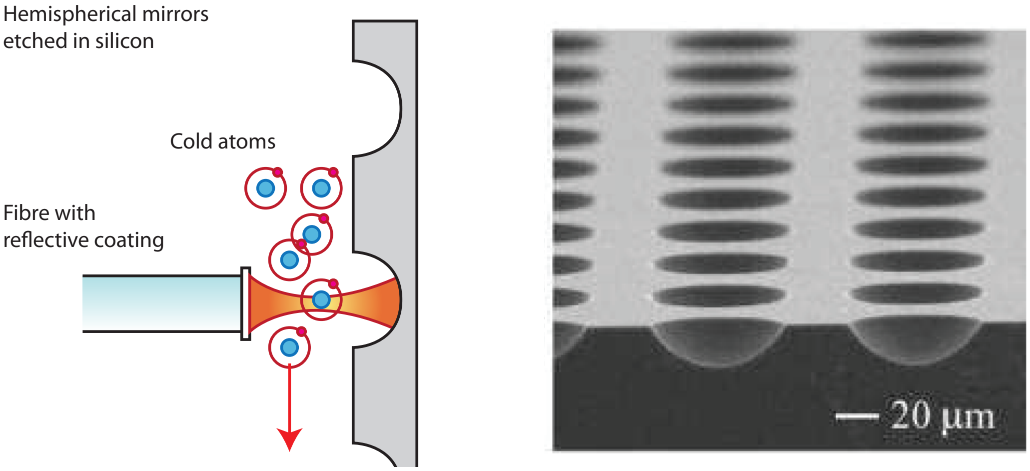

Figure 1.3: Left: Schematic of microfabricated hemispherical mirrors forming cavities as described in [39] Right: SEM image of arrays of the mirrors etched in silicon

Another example of an on-chip tunable optical cavity has been demonstrated by Steinmetz et al. [40]. The cavity is a Fabry-Perot type (two parallel reflecting surfaces) and formed by a transfer technique, which involves lifting a high-quality mirror from a convex dielectric substrate directly onto a fibre tip. These cavities typically have a finesse of around 1000. This can be operated either in a fibre/fibre configuration, or by placing a single fibre in front of a plane mirror (the atom chip surface).

Finally strong coupling of a BEC to an on-chip cavity has been demonstrated by Reichel et al. [41]. A Fabry Perot cavity is formed between two pairs of fibres glued to the chip, and a series of microwires on the chip allow deterministic positioning of a BEC within the cavity to within a micron-precision.

1.4.3 Microlenses

Two dimensional arrays of microlenses can be used to trap atoms in a two dimensional array. In the implementation described by Birkl et.al [42], a 50 × 50 array of microlenses with a diameter and lateral separation of 125μm are manufactured lithographically in quartz. If this array is illuminated with a red-detuned laser, an array of diffraction-limited foci are formed in the focal plane, trapping atoms in the dipole potential. The trap position and spacing can be changed by re-imaging to move the focal plane, or demagnification which reduces the spacing between traps. Illuminating an array with two laser beams separated by a relative angle of propagation produces two interleaved arrays of dipole traps. These array sites can be overlapped by reducing the angle between the beams thus demonstrating the principle behind a controlled cold collision based two-qubit gate.

1.4.4 Photonic Structures

One new development of note is a combined photonic/atom-chip structure which features numerous parallel silica optical waveguides, with a 16 micron wide trench halfway along the waveguide lengths, leaving room for a cold atomic cloud. There the atoms can be addressed by light delivered via the waveguides for a wide range of processes including resonant absorption measurements, fluorescence collection, off-resonant phase-shift detection, optical-dipole trapping or 2-photon coherent manipulations [43].

1.5 Generation III - Integration

To make atom chips truly versatile with real commercial applicability, and accessible for the non-physicist, a number of significant obstacles need to be overcome. A fully functional atom chip will require a full laser setup with associated spectroscopy for stabilisation and ultra-high vacuum systems to name but the primary requirements.

Although it is not a cold atom experiment, Schwindt et al. [44] have demonstrated a ‘chip-scale magnetometer’. This can sense magnetic fields at a sensitivity of 5 $\text{pT}/\sqrt{\text{Hz}}$ for a bandwidth from 1 to 100 Hz. Although still a factor of 1000 above the best SQUID magnetometers, this experiment contains all the components, namely lasers, alkali vapour cell, vapour cell heater, waveplates and RF coils in a volume not exceeding 25mm^3 (see figure 1.4). Obviously this is a staggering achievement, and is probably the best example of what will comprise a third generation atom chip. It is hoped that similar scale reductions of peripheral experimental components can be achieved with cold-atom chip experiments.

Unlike vapour cell experiments, cold atom experiments have to be carried out in Ultra-High Vacuum (UHV) which requires a typically bulky vacuum chamber and other associated pumping equipment. An atom chip in a simplified vacuum system has been built by Du et al. [45] offering the potential to create a chip-based BEC experiment in a volume occupying less than 0.5l. The atom chip itself is used to seal the vacuum cell, simplifying the vacuum and feed-through construction. UHV is provided by only a nonevaporable getter and a compact, 8 Ls^{-1} ion pump. A similar design has been adopted for an early commercial atom chip product, manufactured by ColdQuanta [46].

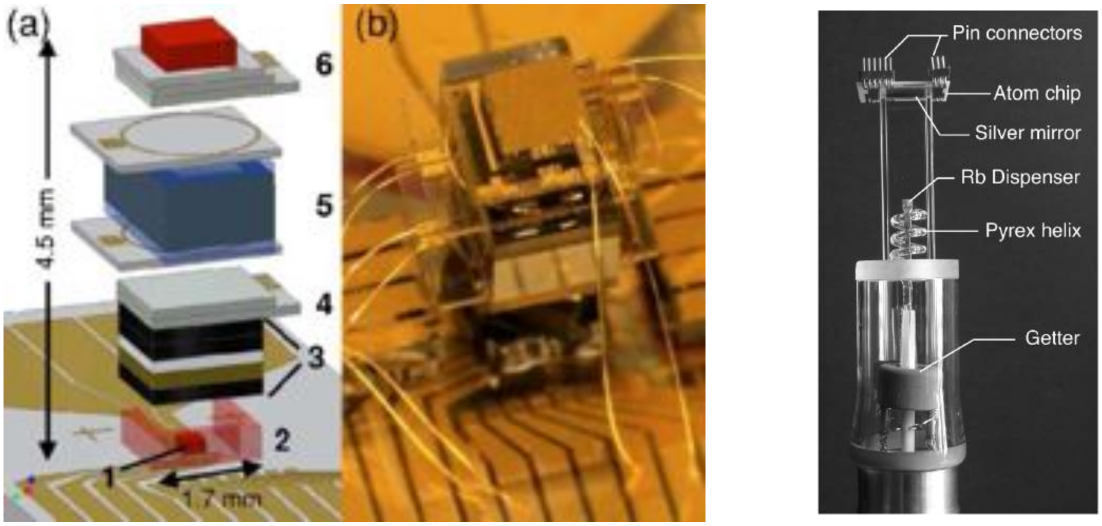

Figure 1.4: Left: A chip scale magnetometer featuring 1) VCSEL laser 2) polyimide spacer 3)optics package 4) ITO heater 5) Rb vapour cell with RF coils 6) ITO heater and photodiode in a volume not exceeding 25mm^3 [44] Right: An atom chip in a portable vacuum cell [45]

1.5.1 Cold atom sources for atom chips

With the advances in photonics, and engineering ingenuity in reducing the size of vacuum equipment, it still remains, quoting Folman et al. again, to make ‘a reliable source of cold atoms’. Currently the favoured method of preparing a cold atomic sample is to cool and trap atoms in a reflection MOT [8] (also known as a mirror MOT). In a reflection MOT, the 6-beam MOT geometry is rotated by 45°. One pair of beams propagate parallel to the chip surface, and two beams are incident at 45° to the surface. If the surface is reflective (in most cases the surface is partially coated in gold) the two 45° beams are reflected with the correct polarisation to create a MOT. This allows for atoms to be trapped a few millimetres from the chip surface, and potentially transferred into the magnetic traps on the chip. (see section 1.5.1 for a full description). This arrangement still requires the balancing of intensity, direction and polarisation of four beams, and a quadrupole field which is either generated by external coils, or by an external U-wire trap.

When the MOT is more than about 2mm from the chip surface, the performance of a reflection MOT is comparable to that of a standard MOT. However, at this distance the chip wires cannot generate sufficient field to create a magnetic trap. Closer to the chip however, the atom number trapped in the MOT decreases as a function of height. Therefore the compromise is made to load the MOT far away from the chip, and then quickly (< 100ms) translated to the surface chip trap circumventing most losses. The process can be simplified by creating an intermediate quadrupole field using chip wires[47].

The atoms are sometimes further cooled to a few tens of μK using optical molasses, before being captured in a weak magnetic trap to form a large atom cloud, typically 1 mm in size. At this point the atoms still have to be handed over to the microscopic magnetic traps on the chip, a process that involves further compression of the cloud and very accurate positioning of the atoms. This sequence of loading and transfer is complicated and could be largely eliminated if the MOT were integrated into the chip. A first step was made by Grabowski et al. [31] where an array of square current-carrying wires was used to create a 2D array of magnetic quadrupole fields and a 2×2 array of reflection magneto-optical traps. Similarly using a transparent magneto-optical film, arrays of quadrupole fields were created, forming multiple MOTs from a standard 6-beam configuration [20]. It is not clear however how a transparent film can be incorporated into a more complex chip of higher functionality.

One proposed alternative is the concept of the micropyramid MOT[48], in which the silicon wafer of the chip is processed to make pyramidal hollows wherever atoms are required. A single incident light beam is then multiply reflected in each pyramid to form the set of appropriately-polarized light beams needed for magneto-optical trapping in a small scale analogue of the pyramid MOT as demonstrated by Lee et al. [49]. This has the virtue that atoms are cooled and collected directly on the chip at the locations of interest. The fabrication process is intrinsically scalable, meaning that multiple MOTs can be formed in an array of structures and illuminated with the same incident laser beam. Successful demonstration of trapping in these devices was first reported in [50] and is the main focus of this thesis.

1.6 Beyond Generation III

The development of basic techniques to prepare, manipulate and probe cold atomic samples on a chip, and the understanding of the science involved in fields such as atom - surface interactions could lead to a deeper probing of the process of decoherence. As well as providing research with strong experimental tools, it is not unreasonable to see in the not too distant future implementation of existing atomic technologies to produce highly accurate portable atomic clocks, gravity gradiometers, or precise acceleration sensors for applications in navigation, seismic studies or motion sensing. Another promising application of atom chips is in precision sensing of magnetic fields, which has applications in non-invasive medical imaging or geomagnetic surveying.

On a more esoteric level, the field of quantum encryption is rapidly gaining in popularity, with commercial products already being targeted at governments and financial corporations. Quantum computing, while still many years away from commercial devices, is a promising field for application of cold neutral atoms as they offer a weaker coupling to the environment than charged ions or solid state qubits. Photons are also very promising for quantum information, however a problem arises when the photon qubits require storing. By reversibly transferring information from photons to the atoms [51], in principle a cold atomic sample can act as long term qubit memory. Along with coupling to quantum information stored in light, there are hybrid quantum system experiment being performed studying the coupling between ultracold atoms and nano-resonators [52], superconducting qubits [53] and cavity QED [39] all on atom chips. Even further down the line the prospect of a ‘molecule-chip’[54] is beginning to show promise.

Crucially as of mid-2009, we are yet to see fully integrated atom-chip devices, one major obstacle being the lack of a reliable system for cold atom preparation. The reflection MOT is the most widely used technique, but the atom cloud is still unhelpfully distant from the chip, and has to be transferred onto the chip before the experiment can begin.

In conclusion it is proposed that the creation of a system for trapping the atoms *in situ* on the chip would be a large step forward and hence pursuing research into the micropyramid magneto-optical trap hopefully will prove prosperous[48]. In terms of applicability, the provision of cold atoms is a universal problem faced by every experiment and hence should the micro-pyramid technique be feasible it is possible to envision in the future many silicon based atom-chips could use variants of this technology.

1.7 Thesis Outline

This thesis presents the development of integrated magneto optical traps for atom chips from silicon micropyramids. I present studies behind the scaling of the magneto optical trap from the typical macroscopic MOTs down to the size useful for atom chips. The fabrication of these silicon devices is presented and a successful integrated MOT device is demonstrated.

This thesis continues in Chapter 2 with an overview of the theory involved in laser cooling and trapping of neutral atoms, specifically the alkali-metal group. The concept of the magneto-optical trap (MOT) is introduced and explored, with some discussion of the practical implementations of this technique including the Pyramid MOT.

The experimental realisation of a rubidium MOT is discussed in Chapter 3. Here we discuss the optics and stabilization techniques used for producing the required laser light, and the vacuum system used to produce the UHV environment necessary. Finally there is some discussion of the detection equipment, specifically charge-coupled device (CCD) imaging is presented along with an overview of the computer system used to control and obtain measurements from the experiment.

In chapter 4, the principle behind the 70° pyramid is introduced. The properties of a 70° MOT in a macroscopic model are studied and analysed. From this implications are drawn for the production of a successful silicon pyramid MOT. The unusual beam geometry is discussed, and the requirements this imposes on the mirror surfaces are presented in both theory and experiment. The initial prototype micropyramid chip is introduced, and this motivates refinements in detecting small numbers of atoms inside a pyramid. As the prototype did not perform as expected, an investigation was carried out into the scaling of the number of atoms trapped in the pyramid as a function of pyramid size. Constraints are placed on the size and properties of a successful silicon micropyramid based on the derived scaling law and detection limitations.

Chapter 5 details the processes used to make the micropyramids in silicon. The basics of anisotropic etching are presented, along with an outline of the development process followed in order to make an optimal pyramid chip. The properties of the surfaces following initial etching are evaluated and a number of microfabrication techniques for smoothing these faces are proposed and evaluated. The resultant 'optimal' devices are analysed with surface metrology techniques and the performance as mirrors for a pyramid MOT evaluated and discussed.

A successful process was developed and a pyramid MOT was demonstrated inside a silicon wafer. Chapter 6 contains the properties of this MOT, namely the dependence on laser and magnetic field parameters as well as the effect of the confinement within a small volume. Direct experimental measurement of the temperature, damping and spring constants of the trap are also presented. The chapter concludes with further experiments in different scales of devices, with the scaling compared to the model hypothesised in Chapter 4.

Based on these results, an introduction to the theory behind the scaling law is presented in Chapter 7. A full semi-classical computational simulation of the capture process, containing many of the intrinsic features of the 70° geometry, is detailed and results compared to the observed behaviour in the silicon pyramids. From this a simple analytical description of the scaling law can be derived to predict capture rates in the micropyramids.

Chapter 8 summarises the findings of the thesis, and tries to indicate areas for possible further improvement of the devices. It also proposes a number of suggested directions for further research along with challenges that need to be overcome in order to reach these goals.

Chapter 2

Cooling atoms with light

2.1 Force on a two-level atom

This thesis is primarily concerned with the force exerted by light on atomic systems, particularly the use of coherent laser light to trap and cool atoms. In order to understand how this can happen, this chapter begins with an overview of the process behind the laser cooling process.

Firstly, what is meant by the 'force exerted'? School level physics tells you a force is a 'push or a pull' but it is hard to imagine how light can either push or pull an object such as an atom. We define the force from Newton as

$$\vec{F} = \dot{\vec{p}}, \tag{2.1}$$

where $\vec{p}$ is the momentum. When an atom absorbs a photon, it experiences a recoil due to the momentum of the photon $\hbar k$, and this recoil occurs in the direction of propagation of the light. The subsequent emission of the photon also leads to a identical recoil but with a random symmetric distribution, which means the contribution from many emissions averages out to zero. As a result the atom experiences net change in momentum from many absorption processes, which can be defined in terms of a force as the product of the photon momentum and the rate of photon scattering.

2.1.1 The two-level atom

In order to understand the rate at which an atom absorbs and emits photons, we approximate the atom as a 2-level system interacting with a classical monochromatic electromagnetic field (the laser). The behaviour of the 2-level atom is given by the Optical Bloch Equations, which are related to the Bloch equations for nuclear magnetic resonance which earned Felix Bloch his Nobel prize (shared with Purcell) in 1952. They define the time evolution of the two-level system, and we can use them to derive the steady state population of the excited state as

$$\rho_{ee} = \frac{1}{2}\frac{s_0}{1 + s_0 + (2\delta/\gamma)^2}, \tag{2.2}$$

where $\delta = \omega_l - \omega_{eg}$ is the detuning of the laser frequency ω_l from the atomic resonant frequency ω_{eg} and γ is the rate of spontaneous emission (the reciprocal of the excited state lifetime). The parameter $s_0 = I/I_s$ is the on-resonance saturation parameter, and is the ratio of the intensity of the incident light I to the saturation intensity I_s, itself defined as

$$I_s = \frac{\pi}{3}\frac{hc\gamma}{\lambda^3}. \tag{2.3}$$

where λ is the wavelength of the resonant transition.

To produce a full and concise explanation of how equation (2.2) is derived would only be covering ground that has been well covered before, and hence as a result I recommend a number of textbooks [55] [56] [57] for an interested reader wanting the full derivation.

In the steady state, the rate of decay from the excited state, $\gamma\rho_{ee}$ is balanced with the rate of excitation. Using equation 2.2 we can henceforth define the total rate of absorption followed by spontaneous emission, the *scattering rate*, to be

$$\gamma_{Scatter} = \gamma\rho_{ee} = \frac{\gamma}{2}\frac{s_0}{1 + s_0 + (2\delta/\gamma)^2}. \tag{2.4}$$

It is obvious to see that for large intensities ($I/I_s >> 1$) ρ_{ee} approaches 1/2 and $\gamma_{Scatter} = \gamma/2$, meaning the population is equally divided between the ground and excited state and that total scattering rate saturates at half the decay rate. The width of this Lorentzian lineshape is given by

$$\Delta\gamma_{FWHM} = \gamma\,(1 + s_0)^{1/2}. \tag{2.5}$$

Thus, as the intensity increases, the linewidth of the scattering resonance is increased, in a process known as power broadening. Another way of considering it is to see that at $\delta = 0$, the transition saturates with half the atoms in the excited state. In wings, the transition is not yet saturated and the rate will continue to rise with increasing intensity, hence broadening the curve.

2.1.2 The scattering force

As we now have the rate of photon absorption, we can derive the total force as

$$F_{Scatter} = \hbar k\gamma_{Scatter} = \frac{\hbar k\gamma}{2}\frac{s_0}{1 + s_0 + (2\delta/\gamma)^2}. \tag{2.6}$$

As expected from the scattering rate, this force saturates at $\frac{\hbar k\gamma}{2}$, but for something as light as an atom, and considering γ to be on the order of 10^7, this is a substantial force capable of producing acceleration or decelerations as large as $10^4 g$.

Here I have only described one way in which light may exert a force on an atom. Another fundamental light-induced force is possible known as the **dipole force**, and can be thought of as a conservative force arising from absorption-stimulated emission cycles. This force is frequently used to manipulate cold atoms and produce atom-optical devices. It is not particularly relevant to this thesis though, and too complicated to do justice to in a brief description so I will not go much further into detail and recommend further reading for interested parties [58].

2.2 Use of the scattering force to slow atoms

Considering initially the simplest case of motion in one dimension, we superimpose a laser beam travelling from the left with one travelling from the right. If both beams are red-detuned from resonance, an atom will be Doppler shifted to the blue into resonance with a beam if it is travelling towards it, and red Doppler shifted further away from resonance with the beam propagating in the direction of motion. This leads to more photons scattered from the laser beam opposing the motion and a net force decelerating the atom almost to rest. When an atom is sufficiently slowed such that the Doppler shift is negligible, it will scatter photons equally from both beams, leading to no net force and hence no acceleration.

To consider the process analytically we assume the atom does not undergo stimulated emission, and the scattering rate is low enough that the two beams can be treated independently. The force from each laser beam can be added, where each beam sees a different detuning due to Doppler shift $\vec{k_{\pm}}.\vec{v} = \pm |k|\,|v|$,

$$F_{1D} = \frac{\hbar k\gamma}{2}\frac{s_0}{1+s_0+4(\delta_0-|k|\,|v|)^2/\gamma^2} - \frac{\hbar k\gamma}{2}\frac{s_0}{1+s_0+4(\delta_0+|k|\,|v|)^2/\gamma^2}, \tag{2.7}$$

where δ_0 is the detuning of the laser fields from the atomic transition. With some algebraic manipulation it can be shown that

$$F_{1D} = -\frac{8hk^2v\gamma^3 s_0\delta_0}{\left(4k^2v^2+\gamma^2+\gamma^2 s_0\right)^2 + 8\left(-4k^2v^2+\gamma^2+\gamma^2 s_0\right)\delta_0^2 + 16\delta_0^4} \approx \alpha v, \tag{2.8}$$

where a Taylor expansion about $v = 0$ yields a series in odd powers of v where we neglect terms of v^3 and above. This approximation shows a force, linearly proportional to velocity. The proportionality constant α is given by

$$\alpha = \frac{8\left(\hbar k^2 s_0\delta_0\right)}{\gamma\left(1+s_0+\frac{4\delta_0^2}{\gamma^2}\right)^2}. \tag{2.9}$$

If the laser detuning is negative (red detuned) this coefficient is negative leading to a force which will oppose the velocity hence being a viscous-damping type force, hence the term ‘Optical-Molasses’ being coined to explain the behaviour of atoms in this situation. As can be seen from figure 2.1, an increase in detuning reduces the gradient near to the origin but extends the range of the force (with the width approximately equal to $\frac{kv}{\gamma} = \delta$).

2.2.1 Doppler Limit

Consideration of equation (2.8) suggests that the atomic motion can be damped to $v = 0$, but this does not consider heating effects arising from the scattering force. The momentum gained from spontaneous emission averages out to zero, but fluctuates because of the random nature of the emission.

The mean number of photons scattered per second is given by $R_{scatter}$ but the uncertainty in this factor is $\sqrt{R_{scatter}}$. The momentum kicks from this fluctuation could be equally likely to be from either direction leading to a random walk of momentum with step size $\hbar k$. This

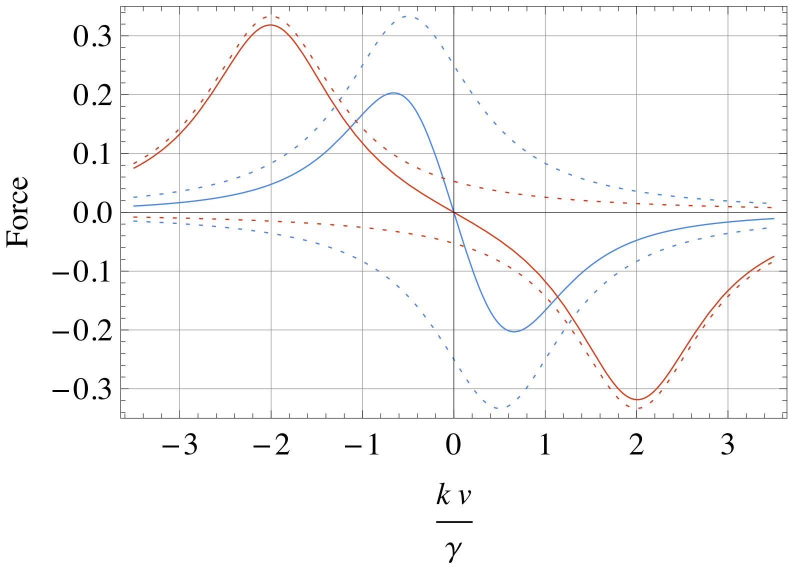

Figure 2.1: The damping force with varying detunings $\delta = -0.5\gamma$ (blue line), -2.0γ (red line). Dotted lines show the contribution from individual beams

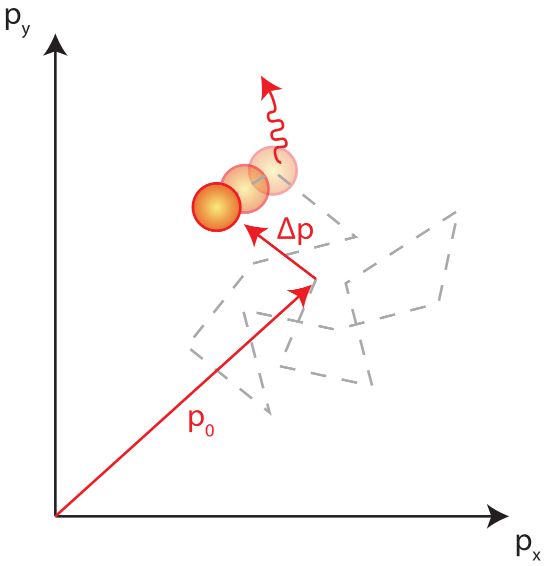

Figure 2.2: Doppler limit: $\vec{p_0}$ is the momentum of the atom and $\Delta\vec{p}$ is the momentum gained from the random walk process

random walk leads to a momentum spread, and after a given time interval the atom has gained momentum $\Delta p = \hbar k\sqrt{R_{scatter}\Delta t}$ where R is the scattering rate. The average energy is given by

$$\langle\frac{1}{2m}(\vec{p_0} + \Delta\vec{p})^2\rangle = \frac{1}{2m}(\vec{p_0}^2 + \Delta\vec{p}^2) = \frac{1}{2m}(\vec{p_0}^2 + \hbar^2k^2R_{scatter}\Delta t) \tag{2.10}$$

and this allows us to derive a heating rate

$$\dot{E} = \frac{\hbar^2k^2}{2m}R_{scatter} \tag{2.11}$$

The cooling rate can be described by the rate at which work is done by the laser beams $\dot{E}_{cool} = F_{cool}.v = -\alpha v^2$ and at equilibrium the rates of heating and cooling are balanced leading to

$$\dot{E}_{heat} - \dot{E}_{cool} = \frac{\hbar^2k^2}{2m}R_{scatter} - \alpha v_{limit}^2 = 0 \tag{2.12}$$

which, assuming optimal $\alpha = \frac{\hbar k^2}{2}$ and $R_{scatter} = \frac{\gamma}{2}$, gives the limiting velocity

$$v_{limit} = \sqrt{\frac{\hbar\gamma}{2m}}, \tag{2.13}$$

which is known as the Doppler limit. An alternate way to consider the Doppler limit is as a temperature, and taking the mean thermal energy as $1/2 k_BT$ per degree of freedom we obtain the Doppler temperature to be

$$T_{Doppler} = \frac{\hbar\gamma}{2k_B} \tag{2.14}$$

2.2.2 Sub-Doppler Cooling

> *We have measured the temperature of a gas of sodium atoms released from 'optical molasses' to be as low as* $43 \pm 20\mu K$*. Surprisingly, this strongly violates the generally accepted theory of Doppler cooling which predicts a limit of* $240\mu K$*.* Phys. Rev. Lett. 61, 169 - 172 (1988) Observation of Atoms Laser Cooled below the Doppler Limit. Lett et. al

The above observation, made by Litt et al. in 1988 suggests that there is more to optical molasses than the simple explanation provided above. It is now understood that atoms are not simple two level systems, but instead possess multiple ground and excited state sublevels. A combination of optical pumping and the light shift of these sublevels combines to create extra cooling mechanisms responsible for the sub-Doppler temperatures observed in optical molasses.

Referring to figure 2.3 which shows for the simplest case of two identical plane waves with orthogonal linear polarisations, the polarisation varies in space from linear at circular and back again over half a wavelength. The polarisation gradients lead to spatially dependent optical pumping rates, but also assuming a $J = 1/2$ ground state, the different interaction of the two ground states with different polarisations means the lowest-lying state is also spatially varying.

Optical pumping acts to pump the system into the lowest ground state (equivalent to the atomic dipole lining up with the field). Following optical pumping, as the atom is in motion, the chances are it will have moved into a different region of light. In this different region this ground state now has a higher energy and it will then be optically pumped into the new lowest energy state with the excess energy carried away in the spontaneously emitted photon. This conversion of internal polarisation energy into kinetic energy is reminiscent of the Greek myth where Sysiphus was condemned by the gods to eternally roll a rock up a mountain, whereupon reaching the summit the stone would roll back to the bottom. This explains why this process of Sub-Doppler cooling can also be known as *'Sisyphus Cooling'*.

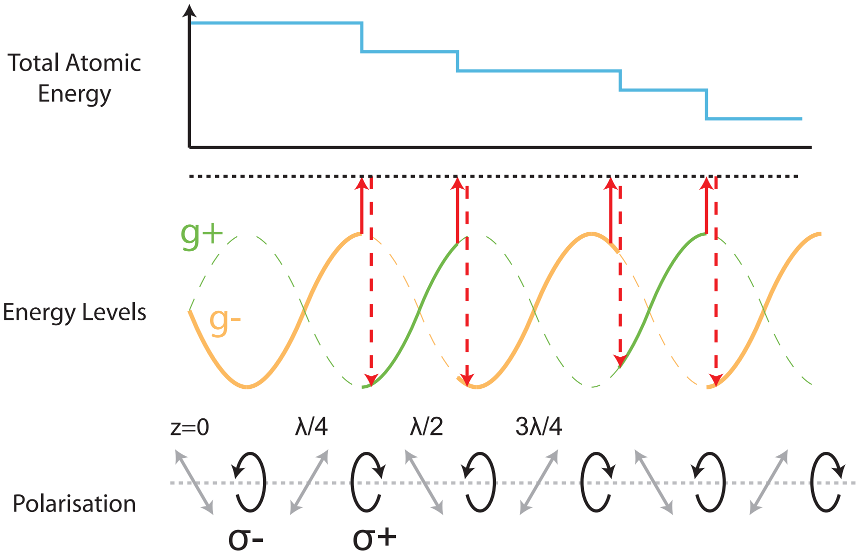

Figure 2.3: Sub-Doppler Cooling, adapted from [57]

At large detunings the rate of scattering is low and as a result the fluctuations due to Poissonian statistics giving rise to heating are reduced and the sub-Doppler effects can cool effectively. The lowest possible momentum achievable is equal to the recoil of a single photon emission, which corresponds to the 'recoil-temperature' and the theoretical limit to sub-Doppler cooling

$$T_{recoil} = \frac{\hbar^2 k^2}{k_B m} \tag{2.15}$$

2.2.3 Optical Molasses in three dimensions

In reality atoms will have motion in all 3 dimensions, and so three pairs of laser beams are required, one for each dimension. With more than one beam, especially in the case of 3D molasses, the effect of saturation cannot really be ignored. With six beams, treating each beam independently would only be valid if the intensity of each beam $I/I_s \leq 1/6$. Accounting for this saturation comprehensively would require solving a large number of coupled Bloch equations, and so an approximation is sought to derive a force for a large number of beams all near saturation. We have the following requirements

- The atom can never scatter at a rate higher than $\gamma/2$

- The approximation should reduce to the 1D form discussed earlier.

To this end, we adapt the form of the force given in equation (2.7) and introduce a correction factor in the denominator where the factor $s_0 = I/I_s$ is replaced with the total intensity $\sum_j s_j = \sum_j I_j/I_s$ giving

$$\vec{F} = \sum_i \frac{\hbar\vec{k_i}\gamma}{2} \frac{s_i}{1 + \sum_j s_j + 4\left(\frac{\delta_0 - \vec{k_i}.\vec{v}}{\gamma}\right)^2} \tag{2.16}$$

This approximation has been made by several previous researchers [59] [60] [61] and it will prove essential when approximating the trapping and cooling properties of our pyramid MOTs, where over 9 beams are interacting with the atoms at most points. It does not reduce to the 1D form we discussed previously, as it retains a factor of $2s$ in the denominator, but this is consistent with the discussion of Adams et al. [57].

We can calculate the effect on the damping coefficient, α to be as follows

$$\alpha' = \frac{8\left(\hbar k^2 n s \delta_0\right)}{\gamma\left(1 + Ns + \frac{4\delta_0^2}{\gamma^2}\right)^2}, \tag{2.17}$$

where N is the total number of beams interacting with the atom, and $n = \sum_i \vec{k_i}.\vec{v}$ is the number of beams with components in the direction of the atomic motion. For large detunings the effect of saturation reduces, and the equation reduces to the form of equation (2.9).

2.3 Magneto-optical trapping

2.3.1 Zeeman Effect

In the presence of a magnetic field, the magnetic hyperfine sublevels, labelled by m_F, are no longer degenerate. The energy shift of a level at low fields is directly proportional to the applied field strength

$$\Delta E = g_F \mu_B m_F B \tag{2.18}$$

where B is the applied field, μ_B is the Bohr magneton and g_F is the Landé g-factor including the nuclear spin. This Zeeman shift is exploited in a magneto optical trap to produce a spatially dependent force that is missing from optical molasses. As always, the dynamics are best understood by confining ourselves to a one dimensional argument. A magnetic field is arranged so it linearly increases from left to right, with a zero crossing at the centre. Consider two counterpropagating laser beams (See figure 2.4), both equally red detuned from resonance with the $F = 0 \rightarrow F = 1$ transition. It is worth noting at this point that although we assume a $F_g = 0 \rightarrow F_e = 1$ transition, this technique will work for any $F_g \rightarrow F_e = F_g + 1$ system as atoms are optical pumped to the outermost m_f substates, forming a closed system between the $|F_g, m_g = +F_g\rangle \rightarrow |F_e = F_g + 1, m_e = +F_g + 1\rangle$ substates. Both beams have the same helicity, but as they are coming from opposing directions the atom will see opposite angular momentum carried from each beam. On the left hand side of the trap, the magnetic field and

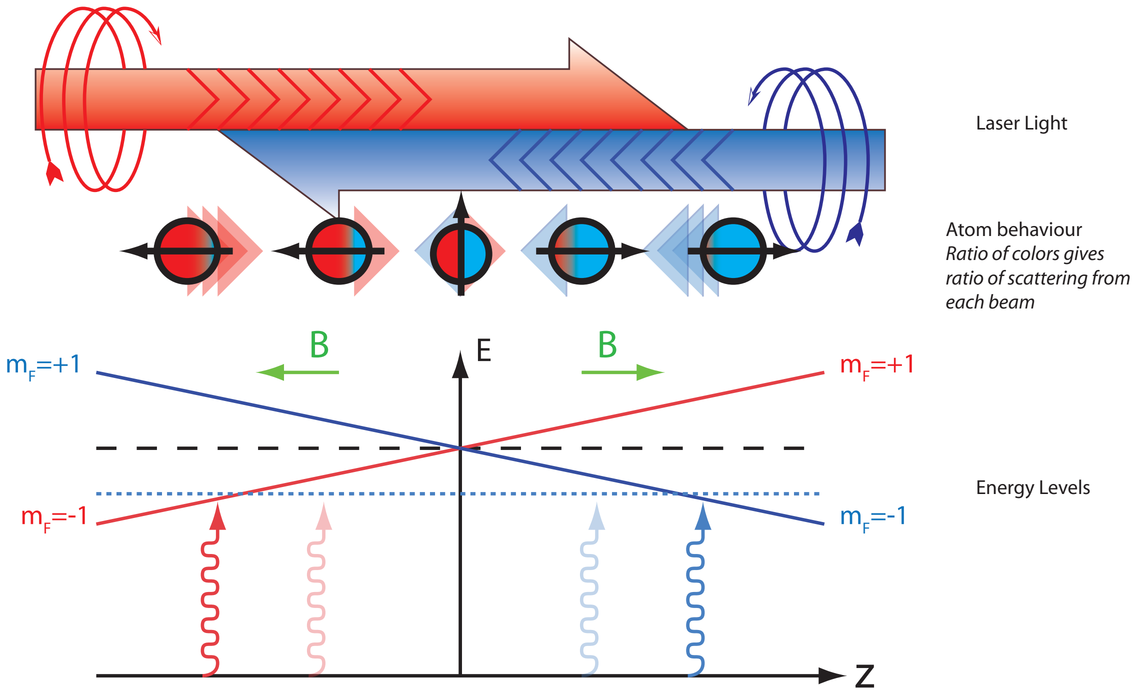

Figure 2.4: Behaviour in a 1D MOT. Both laser beams are identical in frequency, intensity and polarisation and only differ in propagation direction

hence the atomic polarisation vector points in the −ve direction, causing the projection of the left and right propagating beams to appear σ_+ and σ_- respectively. On the right hand side of the trap, the magnetic field and polarisation vector flip and the opposite polarisations are perceived.

Considering an atom initially on the left of the trap moving away from the centre, the energy of the $|F_g, m_g\rangle \rightarrow |F_g + 1, m_g - 1\rangle$ transition will *decrease* (as $\Delta E \propto -B$ for σ_- light) and it will come into resonance with the red-detuned laser beam opposing its motion, scattering more light from this beam than the other beam (which appears σ_+ polarised and is shifted away from resonance). Henceforth the atom traveling left will be driven backwards towards the centre of the trap.

If the atom crosses the centre of the trap, the atom now sees σ_- coming from the right and σ_+ from the left. The Zeeman shift will again bring the atom into resonance with the $\Delta m_F = -1$ transition, scattering more photons from the beam propagating left and driving it back towards the centre of the trap. At the centre of the trap the atom will scatter equally from both beams leaving it in equilibrium. As the beams are red-detuned from resonance, they also will Doppler-cool the atoms leading to simultaneous trapping and cooling which makes the magneto-optical trap such a powerful experimental tool.

For a simple analytical description of the motion inside the trap we once again treat the force on an atom as the sum of the forces arising from the left and right propagating beams.

$$F_{1D} = \frac{\hbar k \gamma}{2} \frac{s_0}{1 + s_0 + 4(\delta_L/\gamma)^2} - \frac{\hbar k \gamma}{2} \frac{s_0}{1 + s_0 + 4(\delta_R/\gamma)^2} \tag{2.19}$$

As each beam has different wavevectors and drives different $\sigma_\pm$ transitions, the detuning arising from the Doppler and Zeeman effects will be opposite in each beam $\delta_L = -k.v + \mu' B/\hbar$ and $\delta_R = k.v - \mu' B/\hbar$ where

$$\mu' = (g_e m_e - g_g m_g)\mu_B \tag{2.20}$$

is the effective magnetic moment for the transition between the ground and excited hyperfine states m_g and m_e. Once again the force from each beam is summed, and by only considering small velocities $k\,|v| < \gamma$ and small displacements $\mu' \frac{\partial B}{\partial z} z/\hbar < \gamma$ a Taylor expansion about $v = 0$ and $z = 0$ can be performed to show that

$$F = -\alpha v - \kappa z \tag{2.21}$$

The form of the above equation is analogous to a damped harmonic oscillator, α is the same damping coefficient encountered in optical molasses (see equation (2.9)) and in accordance with this definition we refer to κ as the 'spring constant' of the MOT where

$$\kappa = \frac{\mu' \frac{\partial B}{\partial z}}{\hbar k}\alpha. \tag{2.22}$$

Typical parameters indicate overdamped motion of atoms in the MOT, with oscillation frequencies $\omega = \sqrt{\frac{\kappa}{m}}$ of a few kHz and damping times $\tau_{damp} = \frac{m}{\alpha}$ of tens of μs.

2.3.2 The 3D MOT

To extend the issue to 3 dimensions, as with 3D optical molasses we require cooling forces in all three dimensions for which the simplest configuration is again three counter propagating pairs of laser beams. Similarly we require confinement in all three dimensions and hence seek a suitable magnetic field which has a zero at the centre of the trap and varies linearly in all directions. Such a field can be created by a pair of circular coils (with N turns) running equal currents in opposite directions arranged with their centres along the z-axis spaced by half their radius (the Anti-Helmholtz configuration). Working with the additive properties of magnetic fields, and starting with the well known equation for the on-axis field of a current loop we can calculate the total on-axis field from both coils to be

$$B_z = B_z^+ + B_z^- = \frac{N\mu_0}{4\pi}\left[\frac{2\pi R^2 I}{((z-R/2)^2+R^2)^{3/2}} - \frac{2\pi R^2 I}{((z+R/2)^2+R^2)^{3/2}}\right] \approx \frac{48}{25\sqrt{5}}\frac{N\mu_0 I}{R^2} z \tag{2.23}$$

for small displacements about the centre of the coil arrangement. In order to work out the field in the radial $x-y$ plane, we make use of symmetry to equate the form of B_x and B_y and invoke Maxwells first equation, the conservation of magnetic flux

$$\nabla.B = \frac{\partial B_x}{\partial x} + \frac{\partial B_y}{\partial y} + \frac{\partial B_z}{\partial z} = 0 \tag{2.24}$$

to show that

$$\frac{\partial B_x}{\partial x} = \frac{\partial B_y}{\partial y} = -\frac{1}{2}\frac{\partial B_z}{\partial z}. \tag{2.25}$$

Thus the total field can be described at small displacements by

$$\{B_x, B_y, B_z\} = \frac{48}{25\sqrt{5}} \frac{N\mu_0 I}{R^2} \left\{-\frac{x}{2}, -\frac{y}{2}, z\right\} \tag{2.26}$$

As the strength of the magnetic field in the x and the y direction decreases with x and y, the beams propagating along these axes require opposite helicity to that of the beams in the z direction (which increases with z).

We include the effects of saturation in the same fashion as in the consideration of optical molasses in 3 dimensions and hence take the spring constant along direction x_i to be $\kappa \to \kappa'_{x_i} = \frac{\mu' \frac{\partial B}{\partial x_i}}{\hbar k}\alpha'$ where α' is as defined in equation (2.17).

For simple models of the MOT, equation (2.22) for each of the 3 dimensions independently (with $\kappa'_{x,y} = {}^1\!/\!{}_2\kappa'_z$) can be enough to explain most typical MOT behaviour but for more rigorous treatments of the MOT we must consider the angle between the light propagation direction and the local magnetic field direction. As the angle of the light vector to the local magnetic field vector θ varies the relative probability of driving a $m_F = \pm 1, 0$ transitions with $\sigma_\pm$ light are

$$P^{\sigma\pm}_{\Delta m=+1} = \frac{1}{4}(1 \pm \cos\theta)^2 \qquad P^{\sigma\pm}_{\Delta m=-1} = \frac{1}{4}(1 \mp \cos\theta)^2 \qquad P^{\sigma\pm}_{\Delta m=0} = \frac{1}{2}\sin^2\theta \tag{2.27}$$

We can expand the equation detailed in equation (2.16), covering the force experienced by an atom in three dimensional optical molasses to take account of the varying transition probabilities.

$$\vec{F} = \sum_i \sum_{\Delta m=-1,0,+1} P_{i,\Delta_m} \frac{\hbar \vec{k_i} \gamma}{2} \frac{s_i}{1 + \sum_j s_j + 4\left(\frac{\delta_{i,\Delta m}}{\gamma}\right)^2} \tag{2.28}$$

Here P_{i,Δ_m} is as defined above, with $\cos\theta_i = \hat{B}.\hat{k_i}$ and the detuning δ_{i,Δ_m} given as

$$\delta_{i,\Delta_m} = \delta_0 - \vec{k_i}.\vec{v} - \Delta_m \mu' B/\hbar \tag{2.29}$$

This means that there exists a finite probability of the atom scattering light from the 'wrong' beam and hence reducing the effective spring constant in some regions of the trap. There is no easy way to calculate the spatial dependence of the trapping force analytically, but it can be calculated numerically as shown in figure 2.5.

The above model is of course only an approximation. It does not correctly treat the details of all the scattering forces produced by the multiple beams. In addition, it ignores the coherent redistribution of photons between beams that gives rise to Sisyphus effects at high intensity. Nevertheless it is sufficient for the purpose of understanding many of the magneto optical traps that will be encountered in the course of this thesis.

2.4 Cooling and trapping Rubidium

> *"Therefore we propose for this alkali metal, in respect to those two remarkable dark red lines, the name Rubidium with the symbol Rb from Rubidus which was used by the ancients to designate the deepest red."* G. Kirchhof & R. Bunsen, Chemische Analyse durch Spectralbeobachtungen. Annalen der Physik 113 (1861), pp. 337-381

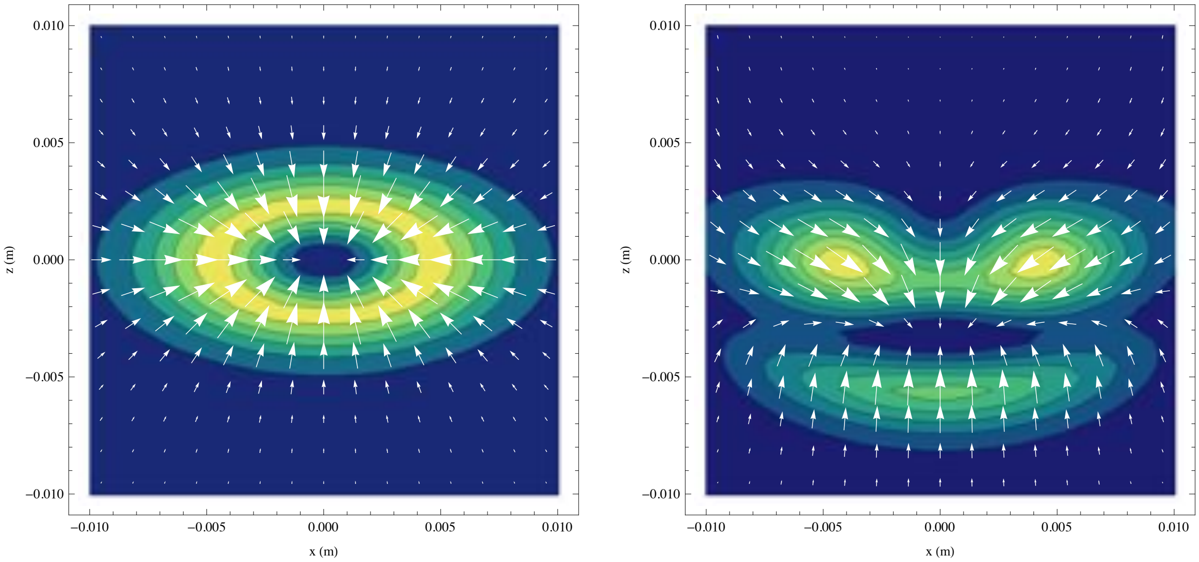

Figure 2.5: Force on an atom in a MOT, taking account of non-circular polarisation effects Left: stationary atom. Right: atom moving upwards at 8ms^{-1}

The discovery of of new spectral lines in the mineral lepidolite, announced the discovery of rubidium by Bunsen and Kirchof in the mid 19th century. It was also the first element to be condensed to quantum degeneracy [5] with the weakly repulsive interactions in ^{87}Rb aiding the condensation process. In this experiment we use both the prominent naturally occurring isotopes, ^{85}Rb and ^{87}Rb, which have relative abundances of 3 to 1 respectively. The valence electron resides in the $5s$ subshell and the spectra of the two atoms differs primarily due to the different nuclear spin, $I_{85} = 5/2$ and $I_{87} = 3/2$. The energy levels are shown in figure 2.6.

In ^{85}Rb we use the D2 transition $5^2S_{1/2} \rightarrow 5^2P_{3/2}$ for laser cooling, specifically the $F_g = 3 \rightarrow F_e = 4$ cycling transition. This transition has a lifetime of $\tau = 26.63$ns giving a maximum scattering rate of $\gamma/2 = 2\pi \times 2.99$MHz. The saturation intensity is 1.67mW/cm^2 for the $\langle F = 3, m_F = \pm 3| \rightarrow \langle F = 4, m_F = \pm 4|$ cycling transition with $\sigma_\pm$ light, but taking into account all the different m_F states and polarisations a more appropriate value may be 3.9mW/cm^2 thus requiring approximately 6mW in a beam of 1cm $1/e^2$ waist to give an average intensity of $I/I_s = 1$. This means that decelerations of up to $11500g$ are possible with moderate laser powers. The Doppler temperature of rubidium is 143μK and the recoil limit temperature is only 1.03μK.

As the $F_e = 3$ and $F_e = 4$ states of the upper $5^2P_{3/2}$ level are separated by only 120MHz, there exists a small possibility for the $F_g = 3 \rightarrow F_e = 3$ transition to be excited. The atom can then undergo spontaneous emission to the $F_g = 2$ ground state which is not resonant with the light and will neither be cooled nor trapped. As a result a second ‘repump’ laser is required, resonant with the $F_g = 2 \rightarrow F_e = 3$ transition to bring the atom back into the cycling transition.

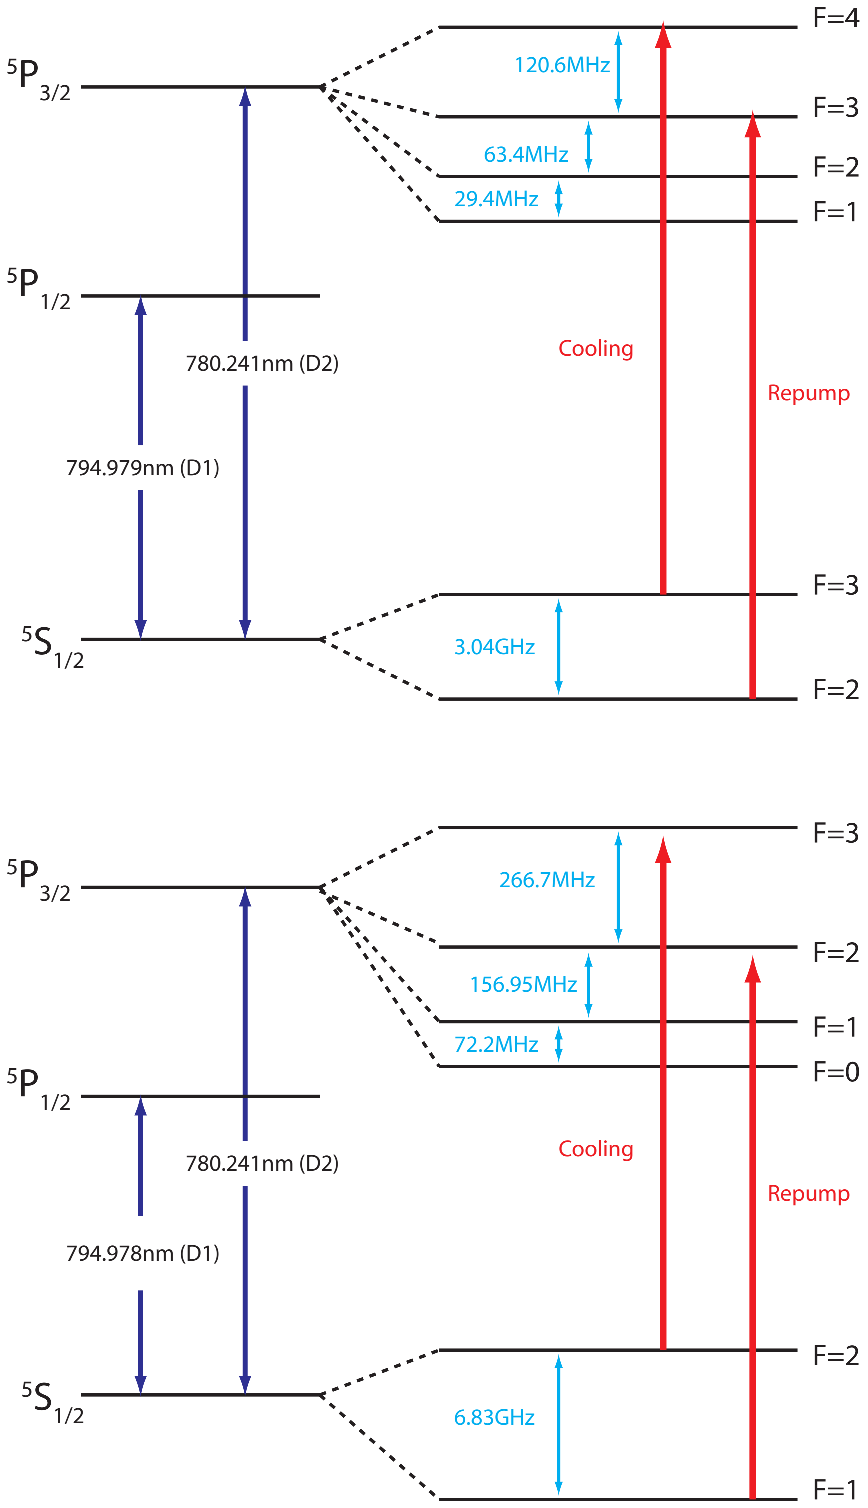

Figure 2.6: Energy levels of Top:^{85}Rb and Bottom:^{87}Rb. Hyperfine structure of the D1 line not shown for clarity

2.5 Typical MOT behaviour

2.5.1 Loading

Atoms can be loaded into a MOT in two ways, they can either be captured directly from a room temperature thermal vapour, or from a precooled atomic beam [1]. Either way there is a finite upper limit to the velocities of atoms that can be successfully cooled and trapped within the MOT region.

The capture process can be approximated by considering a thermal distribution of atoms given by the Maxwell-Boltzmann distribution. For a given surface element dS, the rate of atoms entering the surface with $v < v_c$(the capture velocity of the MOT) can be approximated by

$$n_0 dS \int_0^{v_c} v f(v) dv = n_0 dS \int_0^{v_c} \frac{4}{\sqrt{\pi}} \left(\frac{m}{2\pi k_B T} \right)^{3/2} v^3 e^{\left[-\frac{mv^2}{2k_BT}\right]} dv \approx \frac{v_c^4}{\pi \overline{v}^3} n_0 dS. \tag{2.30}$$

where $\overline{v} = \sqrt{\frac{8k_BT}{\pi m}}$. If we take A as the trap surface area and $n_0 = n_{Rb}$ as the density of rubidium atoms we obtain a total trapping rate

$$R = \frac{n_{Rb} A}{\pi} \left(\frac{v_c^4}{\overline{v}^3} \right) \tag{2.31}$$

Some authors include a constant factor in the rate equation to compensate for geometrical factors, but here it will be left within the definition of the trap area A.

The performance of a trap is dominated by the factor of v_c^4, and the determination of what this ‘capture velocity’ is for a given MOT is subject to a number of factors, primarily the available distance for atoms to be brought to rest (either the diameter of the trapping laser beams, or in the case of the pyramid MOTs the size of the pyramidal hollows used to form the trap) and the force that the atoms experience. Being able to determine how the capture rate scales with trap size is of great importance to us, as the pyramid MOTs we are attempting to construct are much smaller than typical MOTs used in the past and it is not clear how these devices will perform at small scales.

The form of the force as a function of velocity was shown in figure 2.1 and shows that the force is initially linear, saturates at approximately $v = \pm\delta/k$ and rolls off down to a negligible force at high velocity. The size of the trap determines which of these regions will dominate the capture process, and hence how the number of atoms captured will scale as a function of trap size.

As the capture process in a MOT is not just a function of velocity, we consider the trajectories of atoms under the influence of a force given by the form of equation (2.28). We can numerically solve this equation for typical laser parameters to derive the maximum initial velocity that can be stopped in a given distance, i.e. the maximum capture velocity for a given trap size. We choose to solve the equation for the case of rubidium atoms, with laser

---

[1] A cold atomic beam can either be produced with a Zeeman slower [3] or from a Low-Velocity Intense Source (LVIS) [62] which is a preliminary MOT with a hole in it to let an intense cold atom beam out

parameters $\delta_0 = -\Gamma$, ${}^{dB}/{}_{dz} = 15\text{G/cm}$ and $I/I_s = 1$. We also include the effects of a further 4 beams contributing to the saturation of the transition, but they do not contribute to the force in order to make the situation approximate to a typical 6-beam MOT. Results are shown in figure 2.7 and show on the left, the capture velocity as a function of trap size $v_c(L)$, and on the right predicted atom numbers which we derive from the fact that $N \propto L^2 v_c(L)$ (we use a normalisation constant to bring the numbers in line with expected values). There are three regimes of behaviour apparent from the figures.

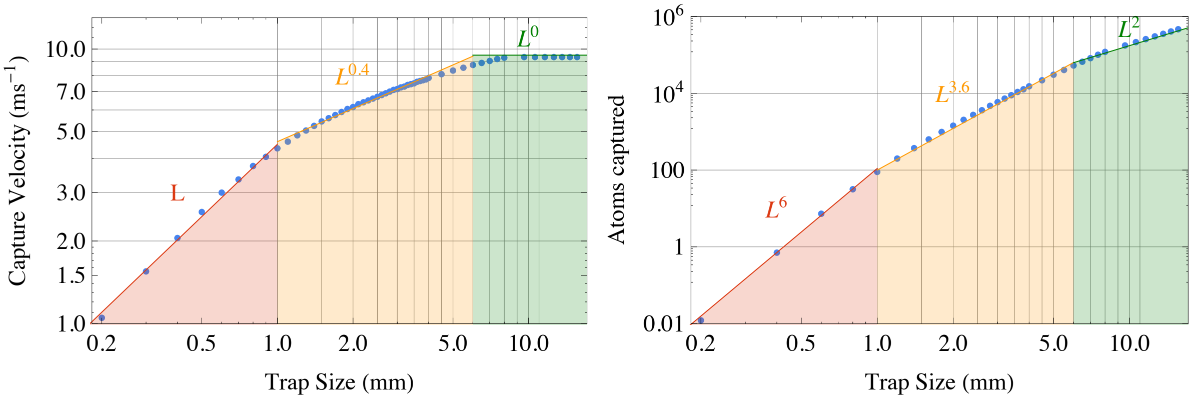

Figure 2.7: Left: Capture velocity as a function of trap size. Right: Atom number as a function of trap size. Text annotations indicate power law drawn for each regime (solid lines)

Small traps (Red Region) At small distances there is a linear increase in capture velocity with trap size. This corresponds well to the behaviour of the force as a function of velocity, which appears linear for small velocities. This can be understood in terms of considering a basic force

$$m\ddot{z}(t) = -\alpha\dot{z}(t), \tag{2.32}$$

which can be analytically solved for the initial conditions $z(0) = 0, \dot{z}(0) = v_0$ to derive the position of the atom as

$$z(t) = \frac{mv_0}{\alpha}\left(1 - e^{-\frac{\alpha}{m}t}\right). \tag{2.33}$$

Thus from this the atom is brought to rest in a distance $\frac{m}{\alpha}v_0$, which is proportional to the initial velocity as required. Turning this argument around, this means the capture velocity increases linearly with the scale of the trap in this regime, which implies from equation 2.31 the capture rate is proportional to L^6.

Medium traps (Orange Region) The proportionality of the force to velocity only increase up to a point where the force saturates and begins to roll off at approximately $v = \delta/k$. In this region it was suggested by Steane et al.[63] that atoms starting in this region will on average will feel half of the maximum scattering force $a \approx \hbar k\gamma/4m$ overall as they are brought to rest giving the capture velocity for atoms stopped in a physical trap size L as

$$v_c = \sqrt{L\hbar k\gamma/2m}. \tag{2.34}$$

This would imply that atom numbers would scale as L^4. In our simulation results we find that the relation between capture velocity and size is closer to $v_c^{0.4}$, and this gives rise to a scaling of atom number $L^{3.6}$ (Orange line), which is identical to the experimental observations made by Lindquist et al [64].

Large traps (Green Region) Finally after a point there is no further increase in capture velocity with size. This occurs when the magnitude of the velocity is such that the slowing force above this is negligible, which can be argued to occur at $v > 2\delta/k$. In this region atom number gain only arises from an increased surface area for capture. It is not certain whether this exponent is observed experimentally as before this scale can be reached the MOT dynamics begin to be dominated by effects arising from multiple scattering [65].

Table 2.1 summarises the main points of each regime that we expect to see

Regime	Velocity Range	$v_c(L)$	$N(L)$
Small	$0 < v_c < \frac{\delta}{k}$	$\propto L$	$\propto L^6$
Medium	$\frac{\delta}{k} < v_c < \frac{2\delta}{k}$	$\propto L^{0.4}$	$\propto L^{3.6}$
Large	$v_c > \frac{2\delta}{k}$	Constant	$\propto L^2$

Table 2.1: MOT Scaling Regimes

This is obviously a very limited approximation, as we only consider atomic trajectories which are initially directed towards the trap centre and along the axis of the laser beams. In reality this makes up a very small fraction of all possible incoming atomic trajectories for a MOT capturing from a thermal background vapour. It is expected that in 3D the requirement to damp motion in all 3 dimensions before the atom leaves the trapping region will reduce the mean capture velocity for a given scale of MOT indicating that the boundaries between the regions will move.

2.5.2 Choosing parameters for cooling

The damping coefficient α can be maximised for $\delta = \frac{\gamma}{2}$ and $I/I_s = 1$. However the range of velocities over which the force is effective is not necessarily optimised by the same parameters. For larger detunings a greater range of velocities will be captured increasing the loading rate in the MOT. The reduced scattering rate at higher detunings means that the optimal detuning will be a compromise between the range of the force and the potential for deceleration. In order to understand this a little better we use a simple 1D simulation (similar to the one used in the previous section) and calculate the maximum velocity that can be slowed to rest in a given distance for a given laser intensity and detuning.

For the simulation we consider a rubidium atom acting under a force given by equation (2.16) for the case of six orthogonal beams. We set the initial velocity of the atom and look at the distance taken to bring the atom to rest. We define the ‘capture velocity’ in this case as the velocity where the atom comes to rest in 5mm. Figure 2.8 shows the results of the simulation, and shows that the optimal detuning increases as intensity is increased. The maximal capture velocity occurs at $\delta \approx 1.5\gamma$ and $I \approx 2I_{sat}$ corresponding to $\delta = 9$Mhz and I =8mW cm^{-2} in rubidium.

These values, as we will see later in the thesis, correspond well to what we observe in a real MOT but this simulation does not take account of the final temperature of the trapped atoms, which from consideration of the Doppler cooling limit alone is lower for lower intensities and higher detunings.

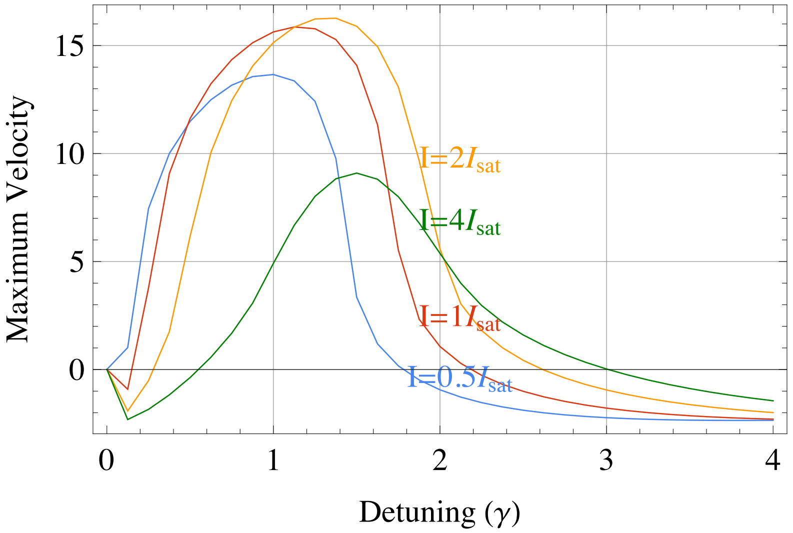

Figure 2.8: Maximum velocity of ^{85}Rb atoms that can be stopped in 5mm determined by numerical simulation as a function of detuning and intensity.

2.5.3 Choosing a magnetic field gradient

Although the laser detuning and intensity affect the spring constant κ, it is in the same way as α (as $\kappa \propto \alpha$ for the Doppler theory of the MOT). One independent parameter unique to the trapping force is the magnetic field gradient however and this controls two factors. Firstly the force is spatially limited by the range where the Zeeman shift is less than the linewidth $\omega_B = \frac{\mu_B}{\hbar}\frac{dB}{dz}z < \gamma$ and this defines the trapping volume. As seen in figure 2.9, moderate gradients (>10G/cm) are desirable so the trapping volume is entirely inside the spatial extent of the beam, but high gradients will start to constrain the volume so that fewer atoms are trapped. This effect at extremely high gradients (>1kG/cm) is exploited in some cases to produce 'single-atom' MOTs [66].

The second effect is the physical size of the trapped atom cloud, as the gradient is typically inversely proportional to the square of the atom cloud radius (see section 2.5.6), and higher densities of trapped atoms are frequently desirable.

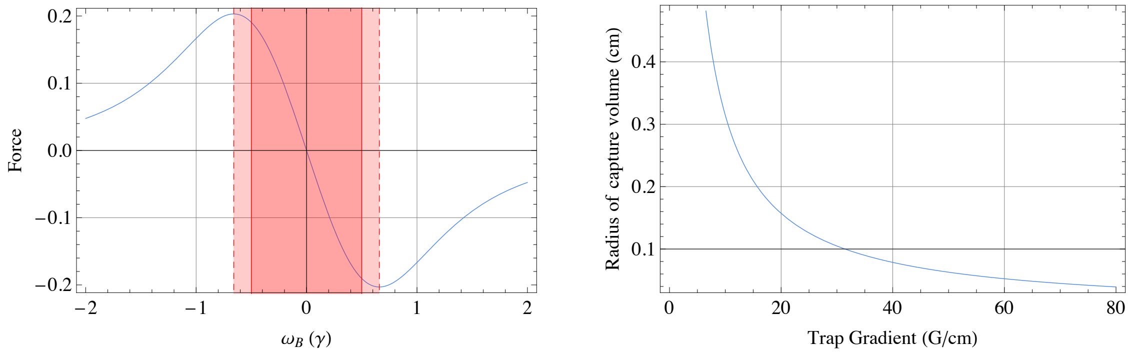

Figure 2.9: Left: Force on a stationary atom as a function of zeeman detuning for $s = 2$, $\delta = 0.5\gamma$. Solid line is $\omega_B = \pm 0.5\gamma$ while the dashed line indicates where $dF/dz = 0$. Right: Radius of trapping volume as a function of gradient for ^{85}Rb and $s = 2$, $\delta = 0.5\gamma$

2.5.4 MOT Number Dynamics

We can consider the capture and loss dynamics to obtain a rate equation for the number of atoms in the MOT as

$$\frac{dN}{dt} = R - \Gamma N - \beta N(N-1) \tag{2.35}$$

where Γ is the one-body loss rate of atoms from the trap (discussed below), and β is the density dependent two-body loss rate. For relevant values of N we can neglect the inelastic collision losses between trapped atoms ($\beta N << \Gamma$) and solve the equation:

$$N(t) = N_0(1 - e^{-\Gamma t}) \tag{2.36}$$

where $N_0 = R/\Gamma$ is the steady state atom number. From this it can be seen that from an initially empty trap, the MOT will fill to $1/e$ of its steady state population in a characteristic time $1/\Gamma$.

2.5.5 Loss

There are three primary loss mechanisms which all MOT experiments encounter, collisions with background gas atoms of a density n_{bg}, collisions with hot untrapped rubidium atoms of density n_{Rb} and collisions between cold atoms which as previously stated will not be significant unless the number of atoms is very large[67]. The loss rate (per trapped atom) from the first two mechanisms can be written as

$$\Gamma_{bg} = n_{bg}\sigma_{bg}\overline{v}_{bg} \qquad \Gamma_{Rb} = n_{Rb}\sigma_{Rb}\overline{v}_{Rb}. \tag{2.37}$$

The cross sections for collisions with untrapped atoms σ_{Rb} can either be calculated theoretically from the Van der Walls interaction or determined experimentally and a typical value is on the order of 10^{-13}cm^2 [63]. An analytical value for the cross section for collision with background atoms σ_{bg} requires knowledge of the composition of the background gas in the vacuum, and so this value is often simpler to estimate experimentally.

Other processes which can lead to losses from the trap include unbalanced pushing beams, holes in the trapping potential, or losses to nearby surfaces. For simplicity these loss rates unrelated to collisions with the background gas will be combined into a single one-body loss rate Γ_{other}. The one-body loss processes can be summed to give the total loss rate per trapped atom as

$$\Gamma = \Gamma_{Rb} + \Gamma_{bg} + \Gamma_{other} \tag{2.38}$$

2.5.6 MOT properties

For a stable MOT it is reasonable to assume a Gaussian distribution of position and momentum in equilibrium at a temperature T. Thus in terms of the MOT spring constant κ the mean potential energy is

$$\frac{1}{2}\kappa r^2 = \frac{1}{2}k_B T, \tag{2.39}$$

where r is the rms extent of the cloud. Since κ is direction-dependent (due to the magnetic field) the MOT is an ellipsoid, with equatorial radii $r_z = r_{x,y}/\sqrt{2}$. In the absence of interatomic interaction, the radius is independent of N and therefore the density of the MOT is proportional to N. Experimental work has shown that the temperature achievable is comparable to achieved in molasses with identical laser parameters but no field gradient [68]. This regime is known as the 'Temperature Limited' regime, and is only valid when atoms can be assumed to behave independently. This is an appropriate approximation for the atom clouds featured in this thesis which are typically optically thin. A discussion of the properties of larger MOTs (where reabsorption of scattered photons becomes significant) is covered by Townsend et al.[65].

2.5.7 MOTs in hollow mirror systems

90° **Pyramidal mirror MOT**

A pyramidal mirror with an apex angle of 90° was first used to form a MOT in 1996 [49]. The strength of this technique is that a single laser beam incident on the aperture of the pyramid automatically generates three pairs of beams with the correct polarization configuration as shown in figure 2.10. A circularly polarised beam incident from the top produces two pairs of beams counterpropagating in the horizontal direction having opposing circular polarization. The second reflection returns the light along the vertical direction, with opposite polarisation to the incident light. At any point in the pyramid region therefore there will exist three pairs of counterpropagating beams with the correct polarization configuration for a MOT.

The number of atoms that can be trapped is ultimately limited by the volume of the pyramid. A group at JILA have trapped up to 2×10^{10} ^{87}Rb atoms in a pyramid with a side length of 10cm and 200mW of trapping light[69].

An even simpler form of the pyramid is to remove two of the mirrors, thus creating a V-shaped groove. This requires one laser beam vertically to produce beams in the vertical and the axis across the groove. Two more laser beams along the length of the length of the groove are required for MOT operation. This trap can be constructed from a pair of gold coated prisms and is ideally suited for experiments in 1D confinement and guiding[70].

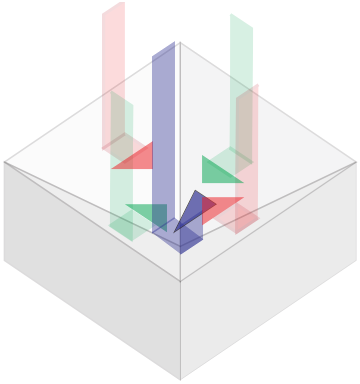

Figure 2.10: Diagram illustrating principle behind pyramid MOT

Conical mirror traps

In the same paper that proposed the pyramidal mirror trap[49], it was also proposed that an axicon mirror (a hollow cone) could produce a MOT. This works on exactly the same principle as the pyramid, but instead of 2 pairs of beams created perpendicular to the vertical, there are infinitely many pairs of concentric beams. The first experimental realization of an axicon-type trap was by Kim et al. at Seoul National University in 1996 [71]. They created a hollow cone from a polished glass substrate, which was then coated with aluminium to create a mirror. Some studies of axicon-magneto optical traps were performed during the course of this thesis and a brief overview of the work is provided in appendix C.

Other pyramid geometries

An alternative configuration to the free space 6-beam MOT is to arrange four beams of equal intensity tetrahedrally, with pairs of beams at angles of $\arccos(-1/3) \approx 109.5^\circ$ to each other[72]. This has the advantage of leaving a large solid angle free of laser beams for access to the trapped atoms. Recently this configuration of beams has been formed from a set of three mirrors arranged in a tetrahedron (or a glass block which forms the beams from refraction) and this has been demonstrated to trap atoms by Vangeleyn et al. [61]. This is unique as this is the first demonstration of a supra-plane pyramid MOT, i.e. the atoms are trapped outside the pyramid volume.

The main part of this thesis concerns micropyramids formed from silicon, which have an apex angle of $\arccos(1/3) \approx 70.5^\circ$ as a result of the silicon crystal structure. This geometry gives rise to a range of reflections from a single incoming beam, and under appropriate conditions this can effect stable trapping and cooling. The use of this geometry for a pyramid MOT is discussed in detail in chapter 4.

Chapter 3

Experimental Apparatus

In order to laser cool and trap neutral atoms, two main components are required. Firstly we require appropriate laser systems with narrow linewidths, with good control over of the precise frequency of light. Secondly a vacuum system which is capable of reaching pressures where collisions with background gas are adequately reduced is essential. In this chapter I describe the current apparatus used to produce a rubidium MOT at Imperial College, with details of the laser and UHV systems developed and refined during this thesis. Finally I introduce the CCD imaging system used, and describe the techniques used to control the experiment and acquire data with computer control.

3.1 Lasers

Three lasers are used: a reference laser to provide an absolute spectroscopic reference, a main trapping laser and a repumping laser. All these are tuned to the wavelength of the D2 transition in rubidium (780.2nm).

3.1.1 Reference Laser

This laser is a homebuilt External-Cavity Diode laser (ECDL) Littrow configuration based upon the Sussex design [73]. It features a laser diode held in a kinematic mount in a collimation tube with a grating mounted at an angle such that first-order diffracted light is coupled back into the diode. The mount allows rotation about the two grating axes for vertical alignment of the external cavity and for selection of the wavelength fed back into the diode. The zeroth-order reflected from the grating is directed out of the exit aperture. Typically 10mW of power is output for use in the experiment.

Power for the diode is supplied primarily by a 12V standard battery, continuously charged by a mains supply. The battery ensures uninterrupted operation with the added benefit of removing a significant proportion of mains noise on the incoming supply. A homebuilt circuit allows AC current to be added for high-frequency locking.

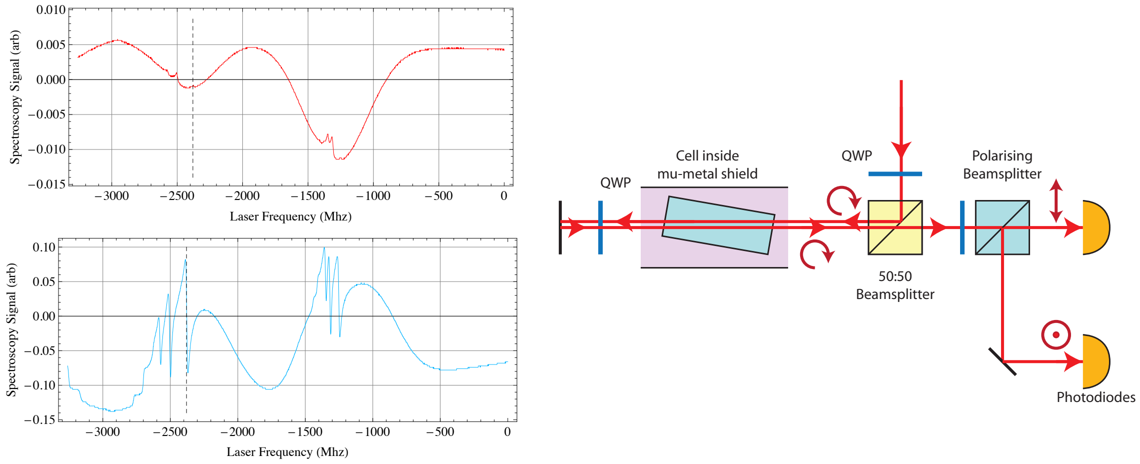

Figure 3.1: Left: Saturated absorption (top) and polarisation spectroscopy (bottom) error signal for the ^{87}Rb $|F = 2\rangle$ and ^{85}Rb $|F = 3\rangle$ transitions. Right: Schematic of polarisation spectroscopy. Vertical dotted line indicates typical locking point.

Locking via Polarisation Spectroscopy

The laser is stabilised on the desired transition by means of an error signal derived from a polarisation spectrometer, shown on the right of figure 3.1,which provides a error signal for feedback control of the grating angle (via a piezoelectric actuator). Polarisation spectroscopy is based on the principle of introducing birefringence in an atomic sample using a strong pump beam and interrogating it with a weak counter-propagating probe beam[74].

This produces a series of dispersive lines, illustrated in the bottom left of figure 3.1. For comparison, the upper graph shows the corresponding saturated absorption spectroscopy signal for the same transition. This lineshape is ideal for locking the laser to the centre of the transition. We lock the reference laser to the ^{87}Rb $|F = 2\rangle \rightarrow |F' = 3\rangle$ line which is approximately 10MHz wide.

Polarisation spectroscopy is incredibly sensitive to stray magnetic fields. We observed that the shim coils located on the main chamber were creating a field on the order of 10mG at the position of the cell. This small field was sufficient to change the position of the lock point by up to 0.5MHz and as a result the cell was placed within a mu-metal cylinder to minimise the field variation, and the setup is orientated perpendicular to the axis of the experiment.

3.1.2 Main Trapping Laser

The primary laser for the experiment (henceforth after referred to as the MOT laser) is a Sacher Lasertechnik LION laser, in Littman/Metcalf configuration. This provides up to 500mW of power with good scan range (typically scans across all 4 D2 transitions without mode hops).

Locking via frequency offset

We use the ‘side of filter’ technique to tune the primary laser to the ^{85}Rb MOT transition, which is approximately 1.1GHz away from the reference laser, locked to the ^{87}Rb MOT transition.

This scheme uses the principle of 'offset-locking', whereby a beat note, equal to the frequency difference between two lasers, is detected by a photodiode and stabilised electronically. Since the desired frequency difference is quite large (i.e. 1.1GHz), it is convenient to mix the beat note down to a lower frequency which can be electronically handled more easily. This lock can be precisely tunable over a large range, determined by the tuning range of the local oscillator.

There are a number of techniques for generating the error signal for locking from this mixed-down signal, for example use of frequency-to-voltage converters [75] or by using an electronic delay line (several metres of coaxial cable) to produce interferometer fringes[76]. The technique we have employed uses the amplitude response of an RF high-pass filter, taking advantage of the fast roll off of the filter to produce strong frequency discrimination[77].

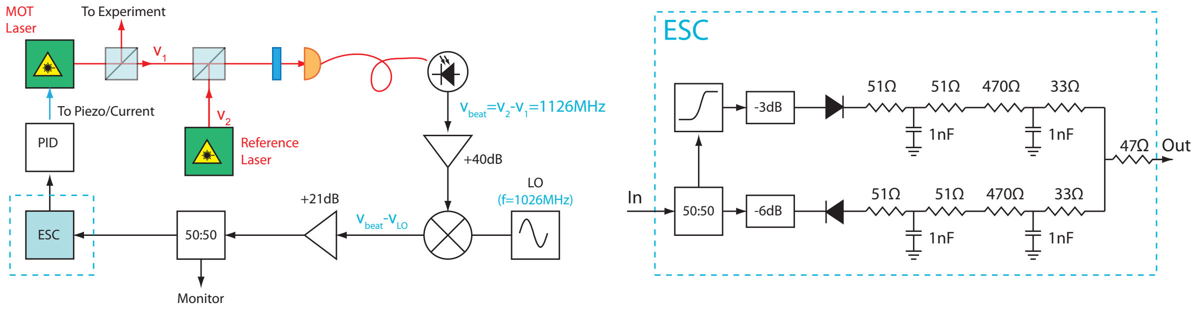

Figure 3.2: Schematic of the offset lock

The offset is the frequency difference between the ^{87}Rb $|F = 2\rangle \rightarrow |F' = 3\rangle$ and ^{85}Rb $|F = 3\rangle \rightarrow |F' = 4\rangle$ transitions, corresponding to 1126GHz. We simultaneously couple the reference (locked to the ^{87}Rb $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition) and MOT laser into a single mode fibre to ensure the two laser modes are completely overlapped, and the optics setup for this is illustrated in the left of figure 3.2. We use a Thorlabs DET02AFC fast fibrecoupled Si-detector (bandwidth=1200MHz) to pick up the beatnote. We typically obtain up to -30dBm (before the amplifier stage) at 1126GHz beat frequency with a 50Ω termination with input laser powers not exceeding 1mW.

The beatnote, ν_{beat}, is initially amplified by 40dB (MiniCircuits ZX60-2534M+) before being mixed (Minicircuits ZX05-153) with a local oscillator of frequency ν_{LO} provided by a voltage controlled oscillator (Minicircuits ZX95-1303-5t , 960MHz-1300MHz at $\approx$ 3dBm) to produce an intermediate frequency $\Delta\nu = |\nu_{LO} - \nu_{beat}|$ on the order of 100MHz. This has a portion split off for monitoring (Minicircuits ZFSC-2-1), and the other part is amplified by +21dB(ZX60-3018G-S+) before being fed into the error-signal circuit (ESC) which is detailed on the right of figure 3.2. In the ESC, the signal is split into two parts of equal power. One part is passed through a high-pass filter and then attenuated by a further 3dB, and the other arm is attenuated by 6dB solely for the purpose of balancing the two signal amplitudes. Each arm is rectified (one positive, and one negative) and the resultant two DC signals combined to produce an error signal with the zero crossing at the 3dB frequency of the filter. The zero crossing of this error signal is hence power independent [77]. The typical form of the error signal as a function of difference frequency $\Delta\nu = |\nu_{LO} - \nu_{beat}|$ and input power is shown in

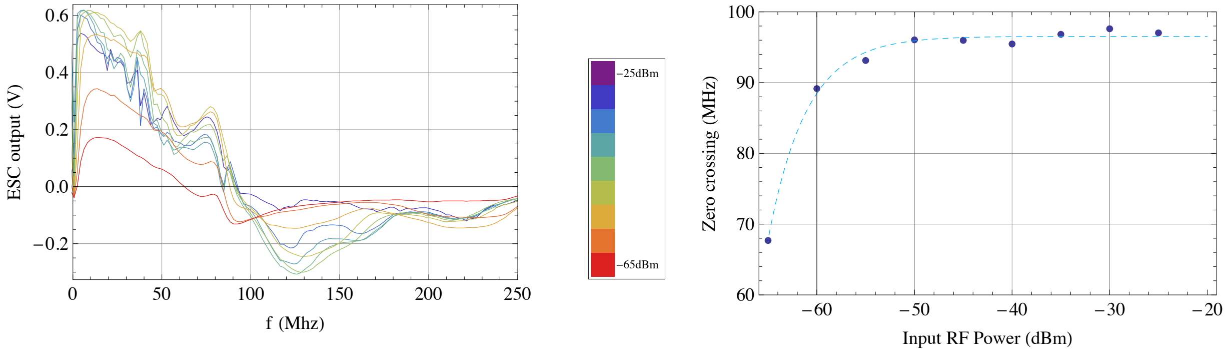

Figure 3.3: Left: ESC signal as a function of input frequency for a number of input powers. Right: Position of the zero crossing as a function of input power

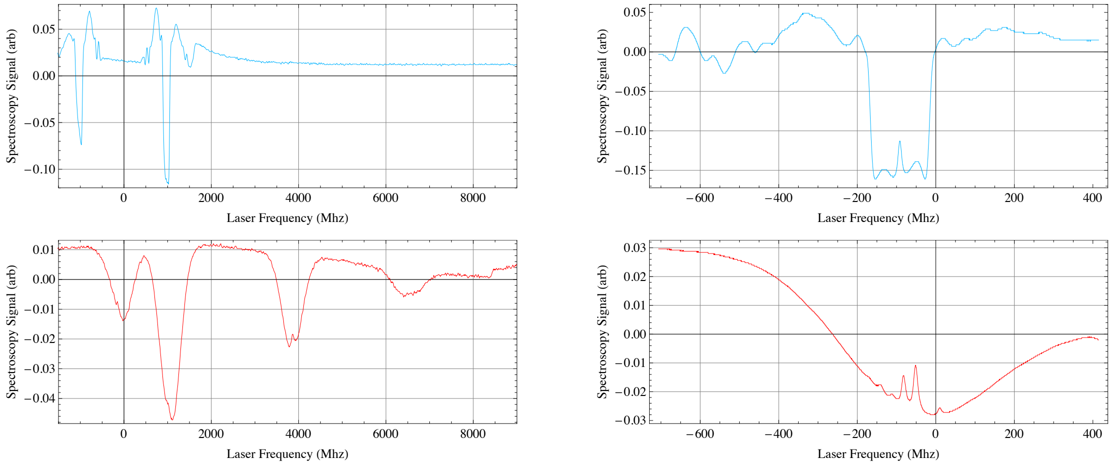

Figure 3.4: Saturated absorption (top) and offset lock (bottom) error signal (reference locked on ^{87}Rb $|F=2\rangle \rightarrow |F'=3\rangle$) Left: All D2 transitions, $f=0$MHz defined as lock point of reference laser Right: close up on ^{85}Rb $|F=3\rangle \rightarrow |F'=4\rangle$, $f=0$MHz occurs at 10MHz red detuned from centre of transition

figure 3.3.

It is found that the stability of the lock is excellent, requiring little or no adjustment in day to day operation. The primary unknown is the absolute position of the lock (set by the tuning voltage of the VCO). This can vary due to slight DC offsets in the error signal generated by RF interference (which can result in locking to a slightly different part of the feature) or variation in the lock point of the polarisation spectroscopy of the reference laser. This typically means the tuning voltage requires calibration once a day and this can be done by picking off the signal from the saturated absorption setup and scanning across accessible ^{85}Rb hyperfine transitions by ramping the VCO tuning voltage. By measuring the absolute positions of the $|F=3\rangle \rightarrow |F'=4\rangle$ and $|F'=3\rangle/|F'=4\rangle$ crossover peaks, the detuning of the MOT can be set with sub-MHz precision. The error signal observed when locking the laser is shown in figure 3.4.

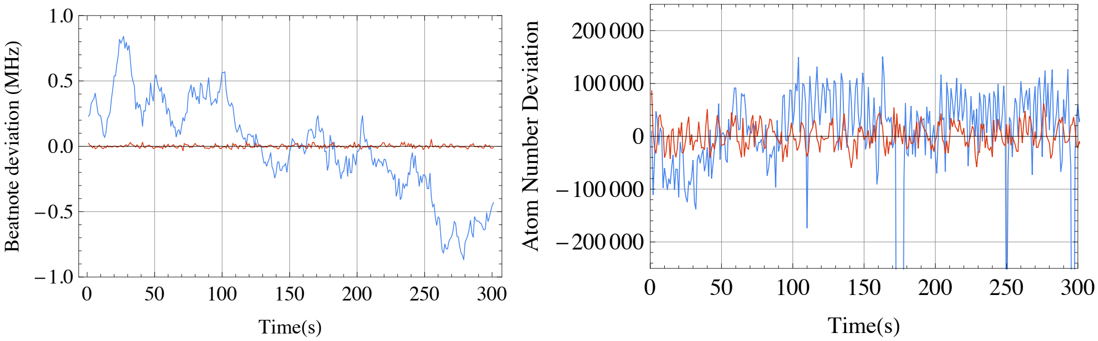

Figure 3.5: Performance of offset lock (red) vs. DAVLL (blue) for beatnote stability (left) and atom number (right). Mean atom number was approximately 6×10^5 and the two tests were carried out separately.

Effectiveness of offset locking

The offset lock was built to replace the previous locking scheme, as we wanted to improve the performance of the experiment. The MOT laser was previously stabilised using Dichroic Atomic Vapour Laser Lock (DAVLL, see repump laser section below). We found this unsuitable for the primary MOT laser, as it was subject to drifts caused by temperature, magnetic field or intensity variations which would cause the MOT detuning to have to be continually adjusted to maintain optimal MOT performance[1].

As a result we wanted to test if the new lock offered any improvement on the old setup. Two criteria were used to judge the effectiveness of the locks, firstly the beatnote for each lock (when locked 10MHz detuned from the ^{85}Rb $|F = 3\rangle \rightarrow |F' = 4\rangle$ transition) was compared. The beatnote for the DAVLL locked laser was taken by heterodyning it with the reference laser and the beatnote for the offset lock was obtained by monitoring the difference frequency with the VCO. Secondly, the number of atoms trapped in the MOT as a function of time was recorded when the trapping laser was locked with the two schemes. Results are shown in figure 3.5. It is apparent that the offset lock is far superior in terms of beatnote stability, leading to a much improved MOT stability. Of course the MOT is also affected by factors such as intensity fluctuations and repump stability which account for the remaining fluctuation of the atom number.

3.1.3 Repumping Laser

This is a commercial Sacher LYNX laser, in the ECDL Littrow configuration. It originally had a AR coated diode for extremely narrow linewidth, but these are costly to replace and an unnecessary luxury for a repumping laser. As a result when the original diode failed, we modified the mount to accept standard diodes (Sanyo DL-7140) which produce up to 30mW at 780nm. This mount is relatively free from drift and the cavity does not require realignment with time.

---

[1] Although DAVLL is subject to this issue of drift, we find the magnitude of the drifts insignificant for the performance of the repump laser.

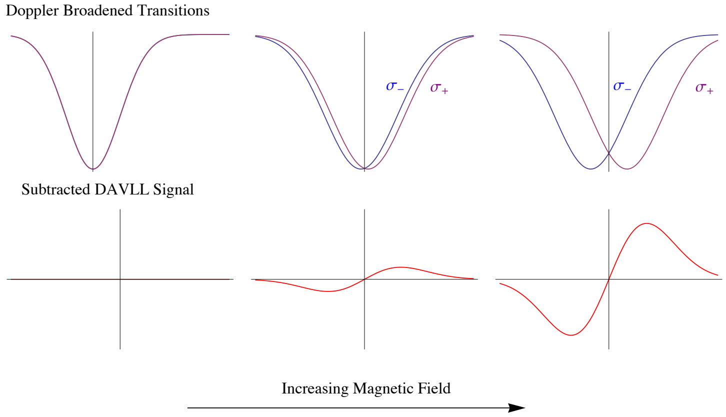

Figure 3.6: Figure showing how as the field increases, the splitting between the two opposing hyperfine transitions leads to the DAVLL lineshape

Locking laser via DAVLL

We use the DAVLL scheme [78] to lock the laser on the ^{85}Rb $F2 \rightarrow F' = 3$ repump transition. DAVLL works by inducing dichroism in an atomic vapour with a magnetic field. The resonant frequencies for the $\Delta m_F = \pm 1$ transitions are (to first order) shifted linearly with applied field and in *opposite* directions.

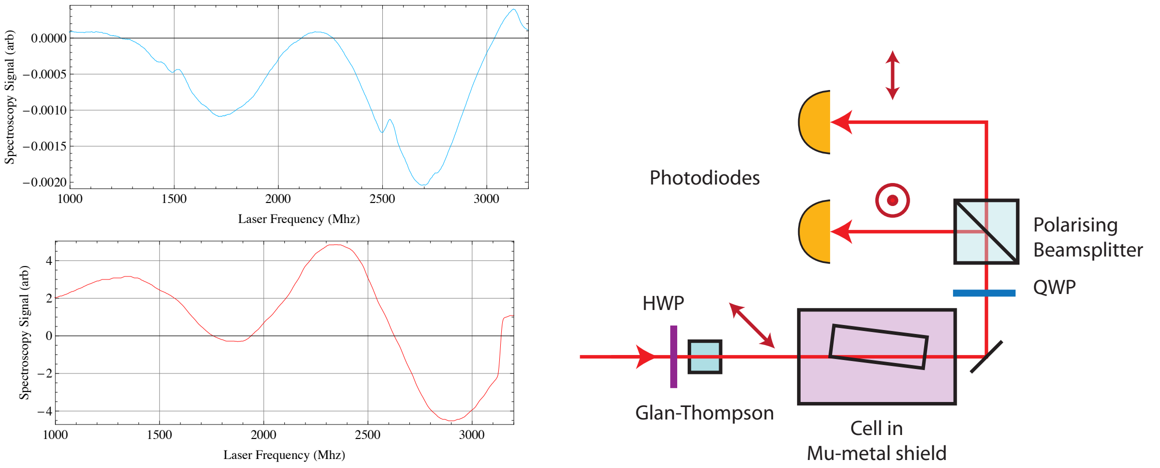

Figure 3.7: Left: Saturated absorption (top) and DAVLL (bottom) error signal for ^{87} and ^{85}Rb MOT transitions. Right: Schematic of DAVLL arrangement

The optics setup used is illustrated in the right of figure 3.7. A weak beam passes through a Glan-Thompson filter to create linearly polarised light. Linearly polarised light can be thought of as an equal superposition of left and right handed circularly polarised light. This light passes through a reference vapour cell mounted inside a pair of strong permanent magnets (producing a field of several gauss), and the two components are separated by passing through a quarter wave plate producing vertical and horizontal light from left and right handed circularly polarised light respectively. These two orthogonal polarisations are separated by a polarising beamsplitter onto two photodiodes and the signals from each photodiode are subtracted in a difference amplifier to produce the locking signal shown in the left of figure 3.7.

In order to obtain the strongest signals, the Zeeman splitting is made roughly equal to the Doppler width, in this case approximately 500MHz. This gives a similarly broad error signal. Using the saturated absorption spectroscopy and the Fabry Perot, the lock point on the DAVLL slope is chosen to centre the laser on the desired transition. In typical operation the lock does not require much tuning and will stay within a few MHz of the transition over the course of a day. The width of the transition means the laser can be quickly detuned several hundreds of MHz from the transition by addition of a voltage offset. This allows us to rapidly turn repumping on and off while keeping the laser locked.

3.1.4 Optics design philosophy

The layout of the optics on the table was recently redesigned to make the setup more structured and easy to follow as illustrated in figure 3.8. Each laser is mounted along the bottom edge of the table, and is raised up on legs so the beam exits exactly 100mm above the table.

Each laser initially passes through an optical isolator, before passing through a wedge prism which picks off two weak reflections (at approx 2% each). One of these reflections is passed to a small compact saturated absorption spectroscopy arrangement placed directly next to the laser for useful lock diagnostics, and another is sent to the far side of the table to be combined with the other lasers and directed into a Fabry Perot cavity (described in the next section) for diagnostic purposes.

The light for use in the experiment is formed from a mix of the repump and the main MOT laser. These lasers are combined on a polarising beam cube and are coupled into a polarisation-maintaining fibre. The collimated beam from the output coupler of the fibre has a $1/e^2$ radius of 3.6mm and passes through a quarter wave plate to become circularly polarised just before it reaches the main chamber viewport. If a larger beam is required, a pair of lenses are mounted on the post to form a 4x telescope. This allows easy access to any set of pyramids in the chamber without having to realign mirror arrangements.

Fabry Perot Diagnostics

We use a scanning Fabry Perot cavity (Toptica FPI100, controller SC100) primarily to ensure the lasers are all in single-mode operation. This cavity has a free spectral range of 1GHz and a finesse of approximately 200, giving a resolution of 5MHz. A useful application of the Fabry Perot arises from the fortuitous spacing of the transitions in rubidium, where the frequency differences are such that the free-spectral range of the cavity (1GHz) means when all lasers are locked, they appear within 100MHz of each other in the scan. This allows for fine tuning of the locks relative to one another to ensure optimal trapping performance. The downside of this arises when the cavity is imperfectly aligned. The free spectral range of the cavity is given by

$$\Delta\lambda = \frac{\lambda_0^2}{2nl\cos(\theta)} \tag{3.1}$$

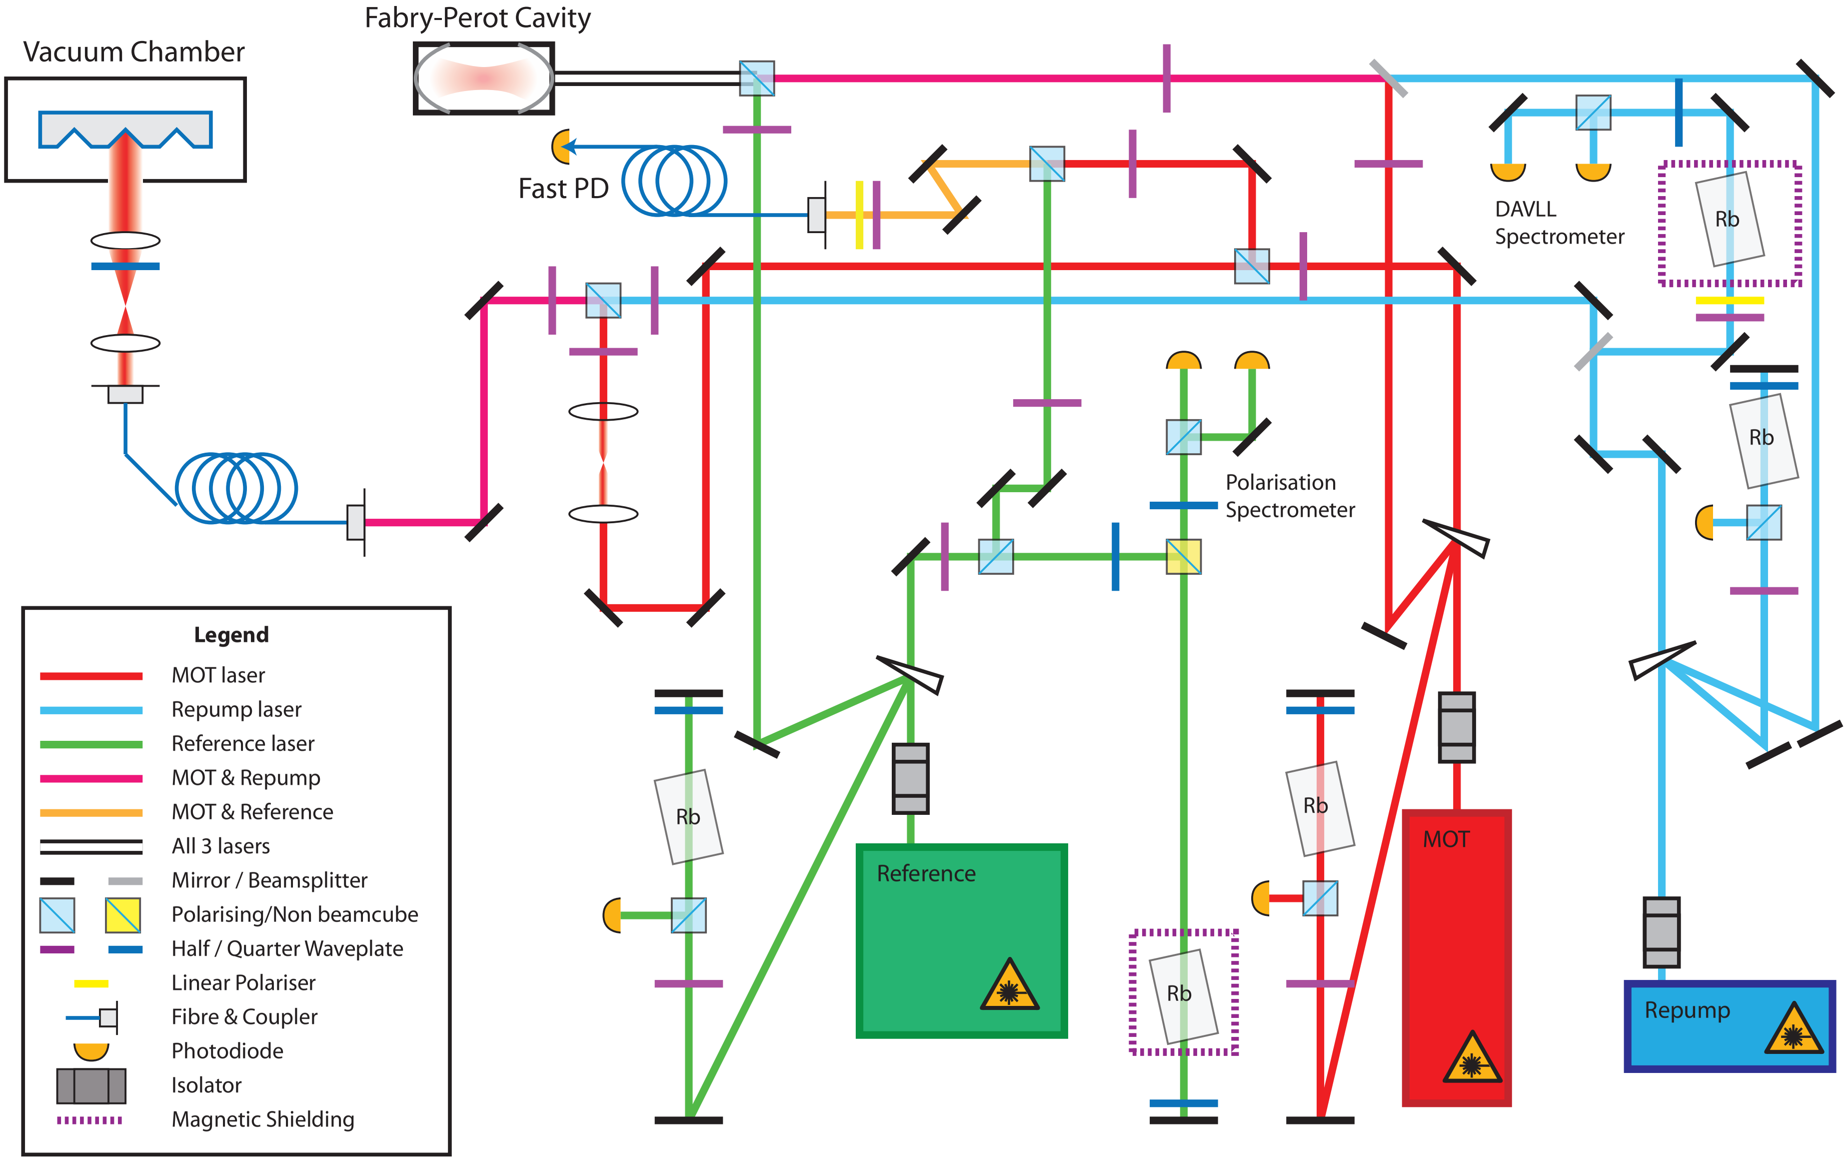

Figure 3.8: Optics plan of the table

where λ_0 is the central wavelength, l is the etalon spacing, n is the refractive index in the cavity and θ is the angle of incidence. There is a 1MHz change in the free spectral range if the angle varies by only 40mrad. Care must be taken to align the cavity to ensure the back reflection from the cavity is collinear with the incoming beam. The lasers are protected from the back reflection by isolators.

3.2 Vacuum System

3.2.1 Vacuum Chamber

The vacuum chamber is built around a Kimball Physics extended octagon (MCF800-EO200080.16-A) unit. This has eight 2^3/4" CF flanged ports around the circumference, 16 1^1/3" CF flanged ports above and below the equator of the chamber, and two 8" CF flanged ports on the top and bottom. By placing an 8" anti-reflection coated viewport on the chamber, we have exception optical access to the chamber. The entire atom chip apparatus is assembled onto the other 8" flange as shown in figure 3.9 which makes changing the experiment exceptionally easy. A four-way cross is attached to one of the 2^3/4" flanges, which connects the ion pump (20 l/s Leybold ion pump) and an all metal angle valve to the main chamber. All remaining 2^3/4" flanges have viewports attached for imaging access parallel to the chip surface. Of the 16 1^1/3" CF flanged ports, half have viewports attached and there is one with a ten-pin feedthrough connected (for rubidium dispensers).

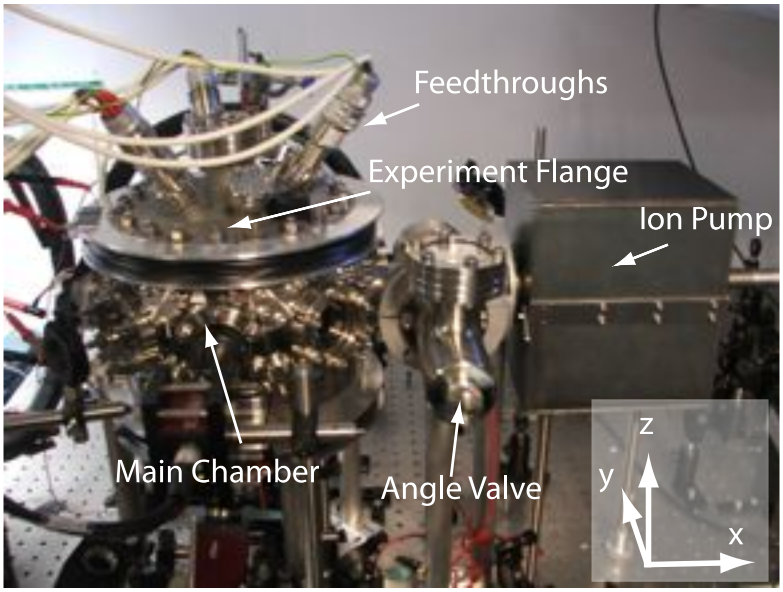

Figure 3.9: Main features of the vacuum system

3.2.2 Main experiment flange

The experiment is all mounted onto one multiport 8" flange displayed in figure 3.10(a). This flange has six 1^1/3" ports spaced around a central 2^3/4" port. The experiments, displayed in

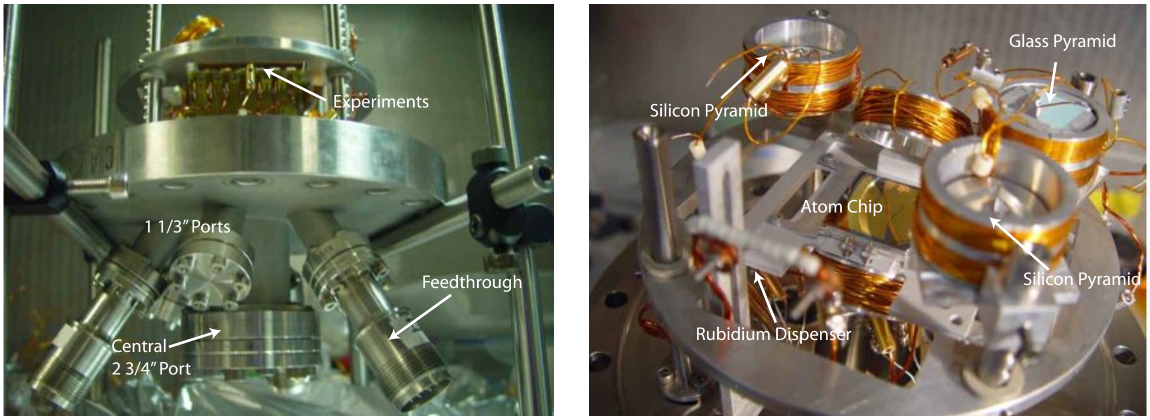

(a) (b)

Figure 3.10: (a) A view of the feedthroughs on the top of the flange. (b) The main experiment flange showing arrangements of pyramids mounted inside coils on the posts and a prototype atom chip in the centre of the flange.

3.10(b), are either mounted onto one of four stainless steel posts or onto a central aluminium ring which mounts onto all four posts. Devices are affixed to the posts by means of threaded holes, which allow for location into one of the dimple holes at 5mm spacing along the length of the post. Three 10-pin feedthroughs are affixed to three of the $1\frac{1}{3}$” ports and provide all the required current for the chips, coils and rubidium dispensers. This arrangement allows for easy changing of components in the chamber, as the entire flange can be removed, worked on in isolation and then returned to the chamber.

3.2.3 Obtaining UHV

We use turbomolecular pumps backed by oil-free roughing pumps to reduce the pressure from atmosphere to a point where the ion pump can be started. In order to achieve a base pressure below 10^{-9}mbar, parts need to be scrupulously cleaned, and a process of baking is generally required. Full details of the procedures undertaken are detailed in appendix A.

3.2.4 Rubidium Dispensers

The source of rubidium for the experiment is a commercial alkali metal dispenser (SAES RB/NF/3.4/12/FT10) which contain an inert source of rubidium, released by direct resistive heating.

These have to be outgassed initially (in the same fashion as other current carrying components), and then if it is the first time they are used they have to be activated (typically half an amp above normal operating current). By passing a laser beam on resonance with a rubidium transition through the chamber and monitoring fluorescence with a camera, one can easily tell when the dispensers have activated. The yield of the devices is exponential with current above the threshold current, and they generally require approximately one minute to reach normal dispensing rates when started from cold. They can either be operated in continuous operation, or pulsed at higher currents to provide a short burst of rubidium in situations where low background pressures need to be maintained. It is possible to reuse dispensers, they should not degrade if exposed to atmosphere following activation but may require outgassing again. It is generally good practice to replace them whenever the chamber is opened however, as they have a finite lifespan.

3.3 Shim Coils

Coils centred along three Cartesian axes (axes illustrated in figure 3.9) provide magnetic fields for (i)nulling any stray fields, (ii) provision of bias to create traps and (iii) for modification of the position of the atom cloud inside the MOT. These coils are hand wound (radius of 14cm and 100 turns) on aluminium formers and can carry currents of up to around to 10A. In the vertical (z) direction, a pair of coils around the top flange and below the lower viewport create fields of approximately 6G/A of current at the centre of the chamber with an associated gradient of 0.1G/cm per amp applied. In the y direction, a pair of coils either side of the chamber provide 2G per amp of current, and a single coil in the x direction provides 1G per amp of current (only a single coil was used as there is no space for a second coil due to the ion pump). The currents in these coils is provided by three independent current supplies, and controlled by a series of MOSFET circuits. The control voltage for all the FETs is provided by the computer DAQ interface, described below in section 3.5.

3.4 Imaging

Two cameras have been used throughout the experiment to detect trapped atoms by imaging their fluorescence. Primarily we use a monochromatic 14bit AVT Pike F-033B, but we have also used an 8bit AVT Marlin F-032B which apart from the inferior bit depth is of a similar quality to the Pike.

3.4.1 Camera Calibration

In order to determine the number of atoms in a trap, it is necessary to turn a digital CCD count (in ADUs) into a detected number of photons. The camera response was calibrated against a reference signal at the wavelength of interest. A weak 780nm laser beam was shone onto the bare CCD through the protective glass window, and a sum of the collected pixel values was obtained for a number of different exposure times. The power in the beam was measured with an Ophir power meter (Nova with PD300 head) and the process repeated for a number of different incident powers.

As can be seen in figure 3.11, the response of the CCDs with increased exposure time is highly linear, as is the response with increased power. The fit to the graph on the right determines that the camera will produce a count of $(2.92 \pm 0.01) \times 10^{11}$ counts per incident μW. This means that 13.4 photons (at 780nm) are required to produce one ADU count.

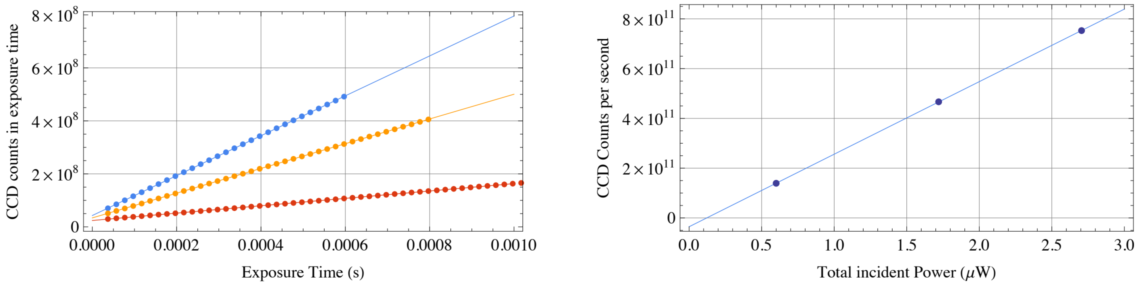

Figure 3.11: Calibration of AVT Pike at 780nm. Left: Response of camera to increased exposure time and linear fits. Right: Fitted response vs. incident power

Theoretical Response We can compare this to the expected efficiency of the CCD at 780nm, which is given by the manufacturer to be approximately 0.10, or 10 photons per photoelectron. We can estimate the conversion factor from photoelectrons to a CCD count from the knowledge that the CCD capacity is around $20\,000e$, corresponding to the maximum CCD count of 2^{14}ADU, thus giving $\eta_{AD} = 0.82$ADU per e. Dividing the factor of 10 photoelectrons/photon by this factor we obtain an estimate of 12.2 photons/ADU which is reasonable agreement with the measured sensitivity.

3.4.2 Camera Imaging Speed

The acquisition of a large number of images in a short duration is of critical importance to the imaging techniques we use. Typically the length of exposure is limited by the saturation of the pixels and to obtain the necessary signal to noise ratio many images have to be taken sequentially. The AVT Pike has an electronic shutter, which is much faster than a mechanical one so the maximum achievable frame rate is limited by a number of other factors. Figure 3.12 shows the maximum frame rate achievable as a function of the size of the region of interest (portion of CCD read out) and the exposure time. The limiting process for a given regime is indicated on the figure. Typically we observe with a shutter speed of 2ms and a region of interest 100 pixels wide. This leads to a maximum frame rate of approximately 450Hz.

3.4.3 Noise on imaging

The attainment of a good signal to noise requires a thorough understanding of all the noise sources arising in the experiment.

Dark noise in the AVT Pike

In order to determine the magnitude of the dark and readout noise for the AVT Pike, a series of exposures were taken with the CCD element blocked from all incoming light. By varying the length of the exposure, the effects of dark noise could be separated from the readout noise (which is independent of exposure time).

The mean value of a particular pixel was recorded from 50 exposures of lengths up to 20ms. This is plotted as a function of the exposure time in figure 3.13. As can be seen there

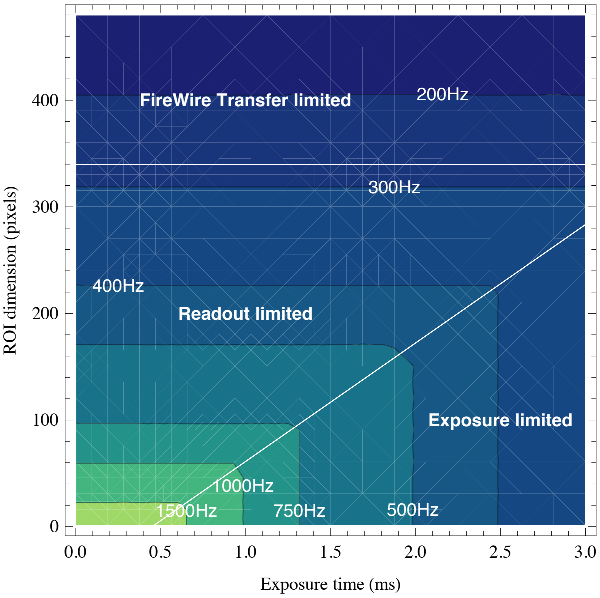

Figure 3.12: Maximum achievable camera rate as a function of exposure time and ROI size (assuming a square ROI). White lines show the boundaries between regimes where different factors limit the rate.

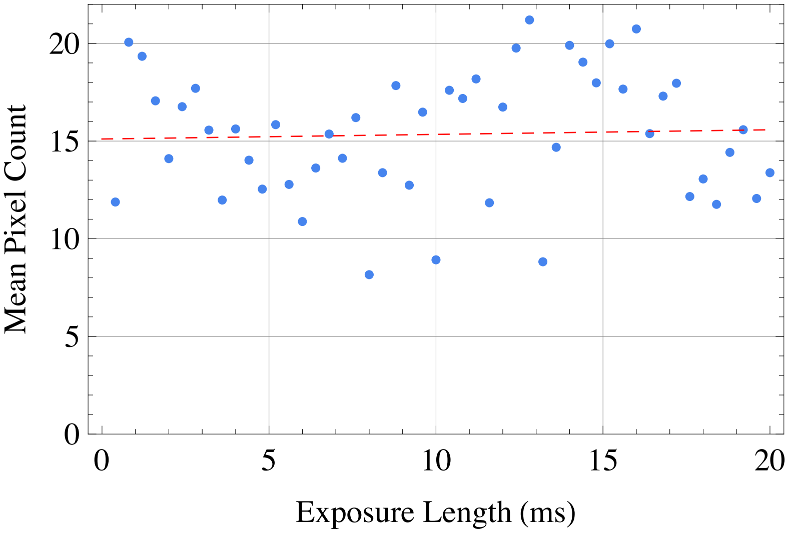

Figure 3.13: Mean noise levels from a dark frame as a function of exposure time (fit shown as a dotted red line)

is no strong linear correlation with exposure time, a linear fit gives a gradient of 0.02 ± 0.08 counts/ms which is consistent with the manufacturers specification for the CCD which quotes $< 1000e/s$ but longer exposures would be needed to confirm this.

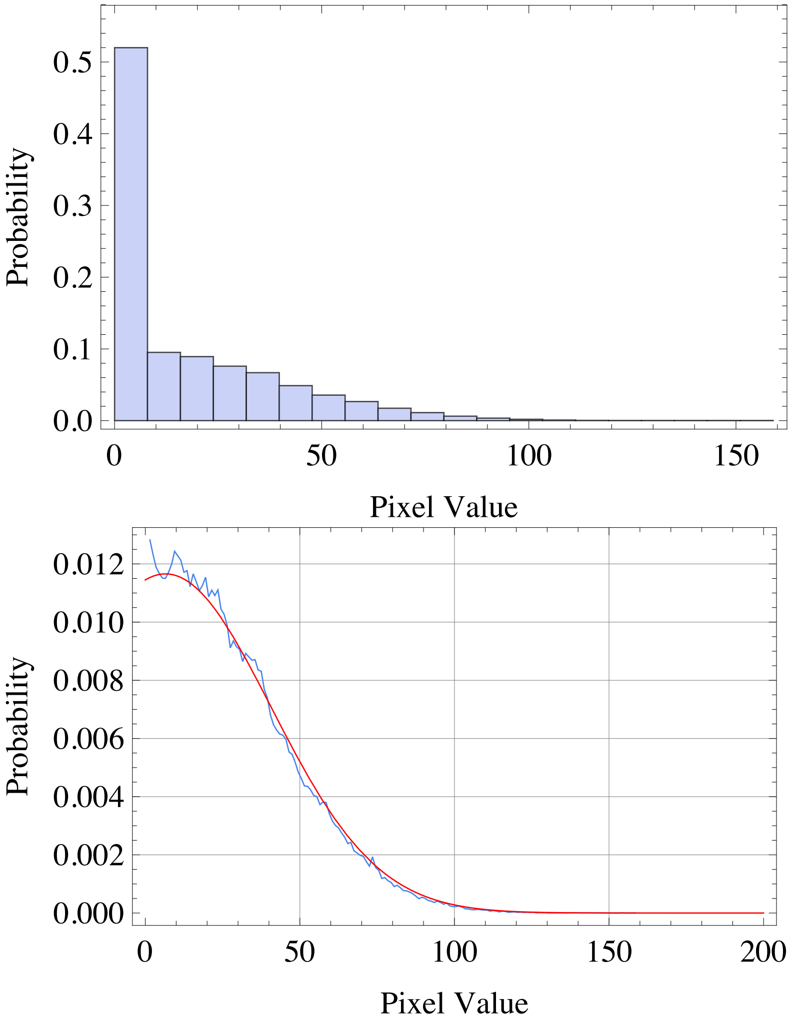

Figure 3.14: Distribution of pixel values occurring in a large number of dark frames. Top:As originally sampled. Bottom: After discarding zero values, with fitted normal distribution (red line)

Read noise in the AVT Pike

If the pixel values for a region of 30 by 30 pixels are recorded for 400 dark observations (so shot noise is negligible), and the pixel values binned according to the frequency of occurrence we get the histogram shown in the left of figure 3.14. This shows a strong skew to lower pixel counts, and if the data is observed more closely it is apparent that any given pixel has an approximately 44% chance of recording a zero count. If we ignore all zero pixel values, we obtain the distribution shown in the right of figure 3.14 which is well described by a normal distribution with $\overline{x} = 6.5 \pm 0.7$ and $\sigma = 34.2 \pm 0.5$. This suggests that the camera truncates negative values to zero. This is verified by the fact that the normal distribution predicts 45% pixel values would be negative, and 44% of the obtained pixel values are zero. In summary we neglect dark noise in our studies and set read noise as $\sigma_{read} = 34.2$ADU per pixel per exposure.

3.4.4 Experiment Noise

In order to judge the effect of noise upon the overall image acquisition process, a background similar to the one observed in the experiment was required. As will be discussed later in the thesis (for example section 4.4) we are observing on a background of scattered laser light, and hence a representative surface was illuminated with laser light and imaged (choosing an exposure time similar to the one used in experiment). The noise in each pixel could then be monitored in order to judge the typical noise performance of the entire system. The photoelectron number is subject to Poissonian statistics, giving a standard deviation of photoelectron number as $\sqrt{N_e}$. Using the predicted value of $\eta_{AD} = 0.82$ to transform into ADU this means the shot noise will manifest as $\sigma_{shot} = \sqrt{0.82}\sqrt{N}$ where N is the number of CCD counts. Using our previous discussions we can henceforth define a total camera noise level as

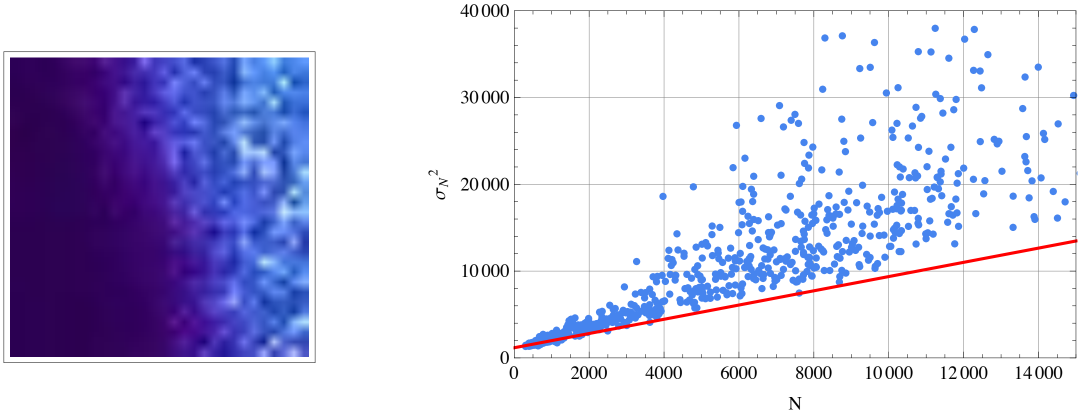

Figure 3.15: Left: 30x30 Image used as representative image to judge noise levels. Right: Plot of mean pixel value (N) against σ_N^2 for 100 exposures at 400Hz. Red line is $\sigma_N^2 = \sigma_{shot}^2 + \sigma_{read}^2 = N + 34.2^2$

.

$$\sigma_{total}^2 = \sigma_{shot}^2 + \sigma_{read}^2 = 0.82N + 34.2^2. \tag{3.2}$$

In figure 3.15, the control image is displayed along with a graph showing the variance per pixel as a function of the mean pixel intensity. It is clear that equation (3.2) defines a lower limit to the noise, but generally the noise level is above this theoretical target, indicating that technical noise is playing a significant part as well.

Noise Spectrum

It is possible to extract temporal information from the observations and identify the dominant sources of technical noise, and it is easiest to simply perform Fourier transformations on the time series of each pixel value. This was performed with a rubidium vapour present and the lasers locked in order to best match typical experimental conditions, results are shown in figure 3.16

It is clear to see that the noise spectrum exhibits excess power below $\sim$ 5Hz. Above this frequency the spectrum appears flat up to the maximum sampled frequency at 400Hz.

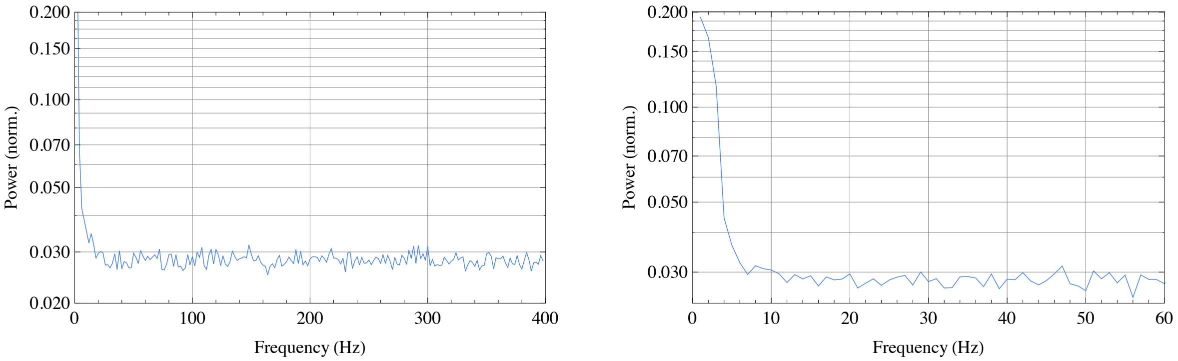

Figure 3.16: Left: Fourier spectrum of camera noise. Right: Showing portion of spectrum below 60Hz

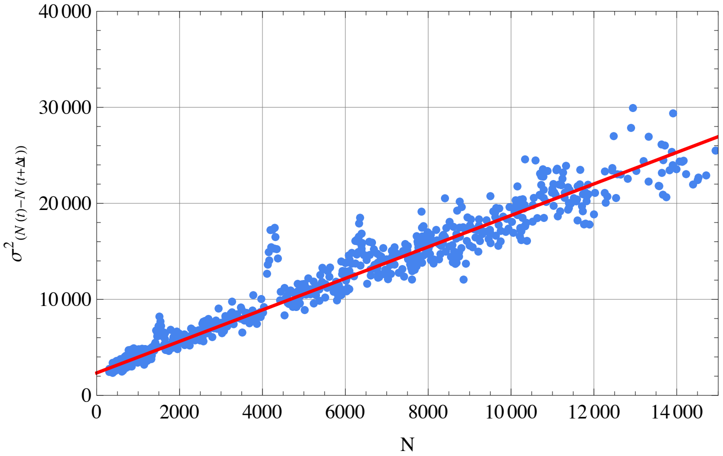

Figure 3.17: Plot of mean pixel value (N) against $\sigma^2_{N(t)-N(t+\Delta t)}$ for 100 exposures at 400Hz. Red line is $\sigma_N^2 = 2(\sigma_{shot}^2 + \sigma_{read}^2) = 2(\eta_{AD}N + 34.2^2)$.

Unlocking the lasers did not give an appreciable reduction in noise, but it is worth stating that the lasers are frequency and not intensity stabilised.

Making the hypothesis that the majority of experimental noise is contained within the low frequency band, the data plotted in figure 3.15 is re-evaluated by subtracting each pixel value from the subsequent measurement $N(t) - N(t + \Delta t)$ and looking at the noise on this measurement, where the time between successive measurements is $\Delta t = 2.5$ms. The results are shown in 3.17 and it appears that the remaining noise is consistent with equation (3.2) multiplied by a factor of 2 due to the subtraction.

3.4.5 Analysing Images

In astronomy 'annulus subtraction' is used, where a target radius is set to encompass the star under observation, and a concentric ring directly surrounding this is used to estimate the magnitude of the sky background. The inner radius of this annulus is chosen so as to be large enough not to overlap any signal, but small enough that it is still representative of the background on the signal.

In the case of our MOT images, we have used a rectangular annulus (a rectangle chosen to simplify the computational steps) surrounding the MOT to estimate the background. The size of the border is chosen to encompass a suitably typically region of scatter, and a function is fit to the pixels in the border region in order to estimate the background the MOT is superimposed upon. For the small regions that the MOT covers in an image (typically less than 100 × 100 pixels) a simple linear planar fit is sufficient for satisfactory background modelling.

This planar fit is subtracted from the MOT region to leave an improved estimate of counts arising from the trapped atoms alone. The counts in this region can then be summed to give a total photon flux (see next section). Alternately the MOT region can be summed across both the horizontal and vertical axes to produce 1D profiles of the MOT. These profiles typically have a Gaussian profile (as the MOT itself has a Gaussian density distribution), and fitting a Gaussian function to these profiles can yield information on the width of the MOT, as well as a precise measurement of the centre of mass of the MOT in the image.

In many situations, reflections of the MOT in the pyramid surfaces was more visible than the MOT itself and thus preferable for fitting (often the MOT lay over a region of bright scatter arising from a mirror edge), and estimates for photon flux were corrected using the reflectivity of the metallic coating.

Estimating the number of atoms

The CCD itself is an instrument which allows for a quantitative measurement of the photon flux, using our previous calibration of the device. We use the principle that if the photon flux from the atoms can be measured at the position of the detector, the number of atoms can be deduced by dividing the flux by the theoretical photon scattering rate per atom. There will be a geometrical factor due to the fact that the imaging optics only captures a small fraction of the scattered photons. The following relation determines the number of atoms (N_A) from the observed number of CCD counts (N_c).

$$N_A = \frac{4\pi\nu}{\delta t R d\Omega} N_c \tag{3.3}$$

where ν is the number of photons required to produce one count, δt is the exposure time, $d\Omega$ is the fraction of the solid angle captured by the imaging optics and R is the scattering rate of a single atom in n beams

$$R = \frac{\gamma}{2} \frac{nI/I_{sat}}{1 + nI/I_{sat} + 4\delta^2/\gamma^2} \tag{3.4}$$

3.5 Experiment Control

The experiment is controlled through a number of computer interfaces by the Labview software package. The interfaces are a pair of National Instrument Data acquisition (DAQ) cards (PCI-6713 and PCI-6221) which allow for digital and analogue input and output.

The analogue outputs are used primarily to control the current supplied to the shim coils and the VCO voltages. Uses of the digital outputs include triggering shutters, cameras, controlling solid state relays to turn current to parts of the experiment on and off and triggering function generators. The Labview software is detailed in appendix B.

3.5.1 Camera Implementation

The camera interface uses the National Instruments IMAQ library. It is almost exclusively run in triggered mode, where either a rising or falling edge on the trigger channel initiates an exposure.

Frames come off the CCD as arrays of 16-bit words, and are converted into floating point arrays for the process of averaging and subtraction into a subtracted image. The floating point precision is maintained throughout, and resultant data files are saved as tab-separated text files with 4 significant figures after the decimal point to ensure little information is lost. Text files were used as it was found that negative values and floating point numbers were handled inconsistently in some image formats. The ‘raw‘ final-image (i.e. the final-image without any form of spatial filtering) from a data set is often saved in addition to any filtered images, along with a PNG preview for quick searching of data sets.

Chapter 4

The route to Silicon MOTs

Unlike previous pyramid MOTs [49], the pyramids formed in silicon by anisotropic etching (see chapter 5) have an apex angle of $\theta = \arccos(1/3) \approx 70.53^\circ$, meaning that the light will not form the 3 pairs of mutually orthogonal reflected beams, typical of a 6-beam MOT. We discuss the implications of this below.

4.1 The 70° geometry

The reflections inside a 70° pyramid angle are more complicated than in the 90° geometry. There are however still only a small number of reflection types possible from an incoming beam directed along the pyramid axis. These are illustrated in figure 4.1.

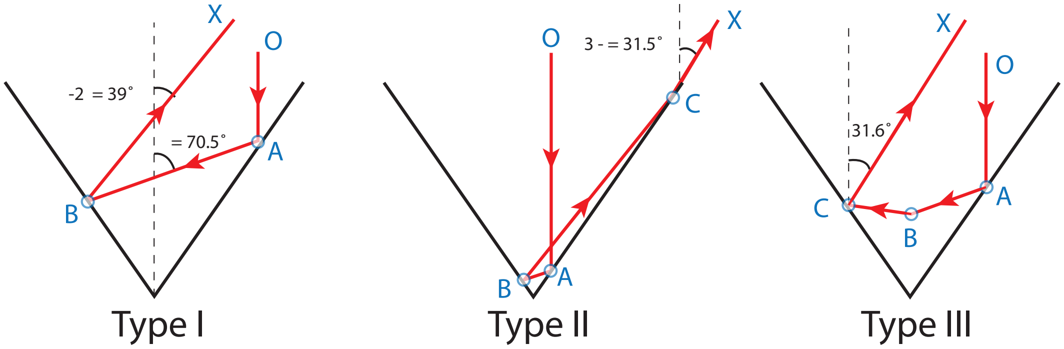

Figure 4.1: Reflections in the 70° pyramid.

Primary beams, responsible for trapping and cooling, shall be referred to as Type I beams. An incoming beam intersects a face at point **A** and is reflected across the pyramid at an angle of $\pi - \theta \approx 109.47^\circ$ to the vertical axis. This beam then crosses the pyramid before being reflected from the opposite face (point **B**) and leaving the pyramid at an angle of $\pi - 2\theta \approx 38.94^\circ$. If the beam is originally incident close to the apex of the pyramid, the exiting beam can intersect a face before leaving the pyramid, reflecting for a third time (point **C**) and exiting at an angle of $3\theta - \pi \approx 31.5^\circ$. These beams, referred to as Type II beams, arise from a central square region $1/9$ of the side length in size. The final type, Type III, arise when a beam following the first reflection is intercepted by an adjacent face (point **B**), before intersecting the opposite face (point **C**) and finally leaving the pyramid at an angle of 31.6° to the vertical. These rays arise from the corner regions, forming triangular shapes with top dimension 21.6% of the pyramid side length. Unlike Type I & II reflections, Type III rays reflect out of plane, and all 8 Type III rays intersect at the centre of the pyramid volume.

If the reflections from the pyramid are viewed from a distance, they form a distinct geometric pattern, and this is shown in figure 4.2.

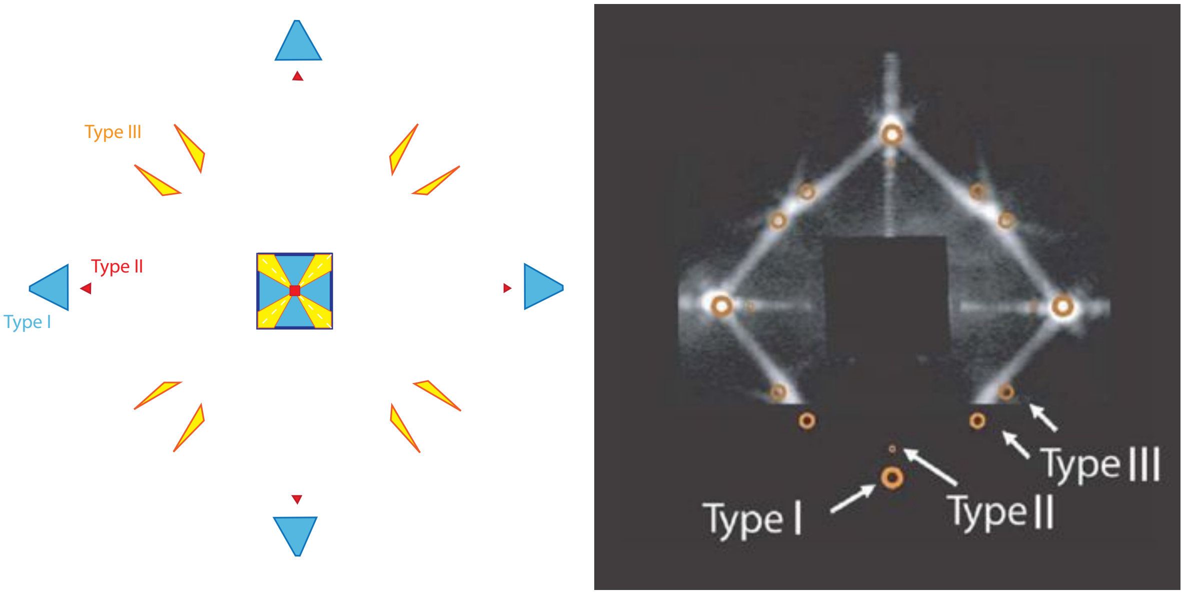

Figure 4.2: Left: Simulated far field reflection pattern as seen on a screen far away from a single pyramid (the regions contributing to each type of reflection inside the pyramid are colour coded) Right: As seen from a real array of micropyramids, illuminated at normal incidence with a HeNe laser. These pyramids feature a degree of rounding at the corners which contributes to a more disperse reflection pattern. [48]

4.1.1 Modification of polarisation upon reflection

If the pyramid were able to be coated with a perfect metal, the helicity of circularly polarised light on reflection from a surface would always result in the reversal of the helicity of the polarisation (in the lab frame). As we are dealing with real metallic coatings (with a less than infinite conductivity) however the polarisation components of light which are parallel and normal to the reflecting plane will undergo phase changes that depend on the angle of incidence, and hence the relative phase between the two orthogonal components $\phi = \phi_{\parallel} - \phi_{\perp}$ will also be angle dependent. For many reflections in the 70° geometry the angle of reflection is far from normal, and hence significantly elliptical polarisations can result.

The theory of phase changes arising in specular reflection from a real metallic surface, as discussed and experimentally confirmed in the previous thesis [79], allows calculation of the ellipticity of the reflections making up the Type I and Type III rays (Type II rays do not intersect the MOT trapping region and so we do not consider them). As can be seen in figure 4.3 the beams remain relatively circular throughout, with the maximum phase shift of $-29.9°$ for gold coatings. The Type III rays, illustrated on the right in figure 4.3 undergo a near grazing incidence reflection becoming nearly linear.

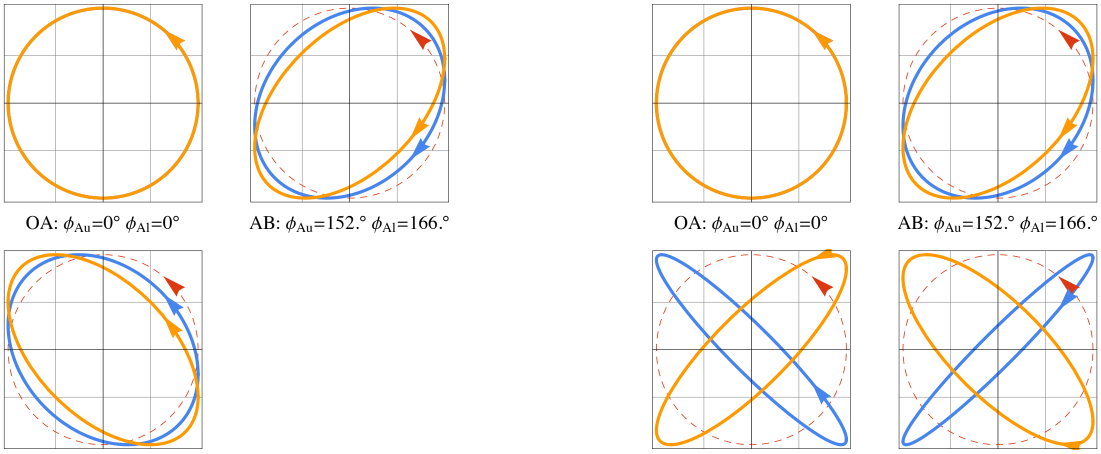

OA: $\phi_{\mathrm{Au}}=0°$ $\phi_{\mathrm{Al}}=0°$ AB: $\phi_{\mathrm{Au}}=152.°$ $\phi_{\mathrm{Al}}=166.°$ OA: $\phi_{\mathrm{Au}}=0°$ $\phi_{\mathrm{Al}}=0°$ AB: $\phi_{\mathrm{Au}}=152.°$ $\phi_{\mathrm{Al}}=166.°$

BX: $\phi_{\mathrm{Au}}=-29.9°$ $\phi_{\mathrm{Al}}=-14.9°$ BC: $\phi_{\mathrm{Au}}=58.°$ $\phi_{\mathrm{Al}}=-70.8°$ CX: $\phi_{\mathrm{Au}}=-131.°$ $\phi_{\mathrm{Al}}=104.8°$

Figure 4.3: Change of polarisation for Type I(l) and III(r) reflections. Orange=gold, Blue=aluminium, Red=original polarisation

4.1.2 Axial imbalance in forces due to mismatched beams

Considering the 70° arrangement in comparison to the 90° geometry, it is quite apparent from the diagram on the left of figure 4.4 that we have significant geometrical deviation from the original 6-beam MOT configuration. This should not preclude the existence of a MOT as long as the forces on an atom balance so as to effect stable trapping and cooling. Every beam with a horizontal component is matched with a reflection from the opposite face. Hence assuming the illumination of the pyramid is uniform on all four faces, symmetry demands that the forces will balance horizontally. In the axial (vertical) direction however the beams do not typically balance, the downwards (towards the apex) components arise from the incoming beam and the four first reflections which make an angle of approximately 70.5° to the vertical direction. These reflections are opposed by the four secondary reflections of the Type I beams, which

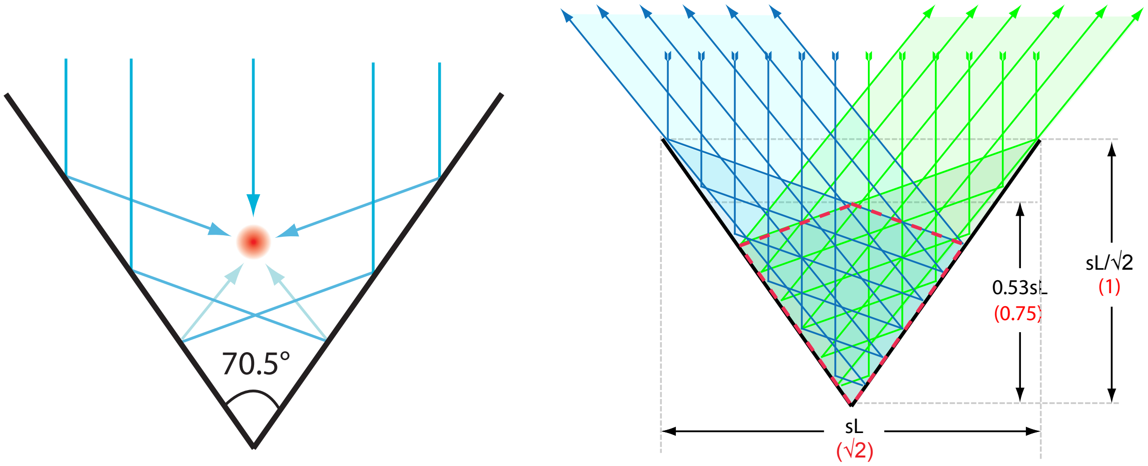

Figure 4.4: Left:Type I rays forming the trapping force. Right:The region that a MOT can potentially form is outlined with a dotted red line

are at 39°, and the potential eight Type III rays at 31.6°. Type II beams do not intersect the trapping region and are hence not considered. As each beam has undergone different numbers of reflections, the intensities are also modified due to the non-unity reflectivity of the mirror coating.

Imbalance from Type I beams

In order to consider this analytically, we define an imbalance parameter [63] which is given as follows,

$$w = \frac{i_+ - i_-}{i_+ + i_-} \tag{4.1}$$

where $i_\pm$ is the sum of beam intensities direction into/out of the pyramid.

Any small imbalance in the intensity is compensated for by the spatial component of the MOT force, and hence for $w \neq 0$ the MOT will be centred at a non-zero value of the magnetic field. In the original consideration of force imbalances in the magneto optical trap [63], the imbalance was studied in a standard 6-beam perpendicular magneto optical trap where the theory concerning sub-Doppler effects in counter propagating σ_+/σ_- beams could be extrapolated from work by Dalibard et al. [80]. In our case we cannot satisfy many of the assumptions required for the sub-Doppler approximation, so we shall restrict ourselves to Doppler theory giving

$$F_{tot} = F_{down} - F_{up} \approx -\alpha v - \kappa z + \frac{\hbar k \gamma w}{2} \frac{i_+ + i_-}{1 + nI/I_{sat} + 4(\delta/\gamma)^2}. \tag{4.2}$$

Equation (2.22) can be used to derive the position of the trapped cloud

$$z = \frac{\hbar\gamma}{\frac{dB}{dz}\mu'} \left(\frac{\delta^2 + (1 + nI/I_{sat})\gamma^2/4}{2\delta\gamma} \right) w, \tag{4.3}$$

which differs from Eq.(70) in ref [63] by the factor of $(1 + nI/I_{sat})$ in front of the linewidth to take account of saturation.

In the 70° pyramid, the force into the pyramid from the Type I beams consists of the incident beam and the four first reflections from each of the faces at 70° to the vertical. The force out of the pyramid is made of the four second reflections at 39° to the vertical. We do not consider the non-ideal polarisation of the light in this simple analytical model, but as will be shown below this does not adversely affect the results.

Substituting the values for the Type I beams into equation (4.3) we obtain the resultant graphs displayed in figure 4.5, showing the displacement (in terms of magnetic field) as a function of reflectivity. Explicit numerical simulations (which take account of the non-ideal polarisation orientation with respect to the magnetic field) agree excellently with our approximate model. This also shows that a stable trap is possible for a large range of reflectivities assuming a suitable vertical magnetic field bias can be applied. It is also worth mentioning that there is no imbalance for a mirror reflectivity of $\rho = 82\%$.

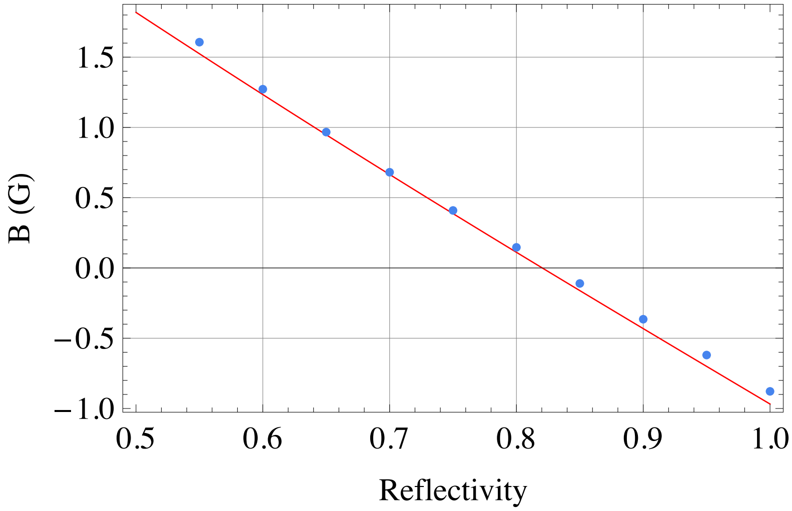

Figure 4.5: Bias field value where the imbalance in the beams is cancelled for $\delta = -2\gamma$ as calculated numerically (blue points) and using equation (4.3) (red line).

Imbalance from Type III beams

The same analysis cannot be used to treat the Type III beams, which are effectively linearly polarised. As a result we numerically simulate the forces. There are eight reflections making up the Type III beams, giving a potential maximum intensity of $8\rho^3$, but this only happens at the very centre of the pyramid where they all intersect. A short distance away and the number of intersecting beams has reduced. As a result in our consideration of these beams we include a parameter η in the intensity of the Type III beams which can vary between 0 and 1. As previously stated small imbalances in the forces can be corrected for by the MOT restoring force, but this cannot increase indefinitely and at some point the imbalance is so great that there is no stable potential minimum and the MOT cannot form.

Using the numeric simulation of the total scattering forces (including the non-ideal polarisations) we can work out whether or not an atom will be trapped in a MOT as a function of the reflectivity ρ and the Type III contribution factor η. Results are shown in figure 4.6, where we plot the value of F_{up}/F_{down} at the position of the magnetic field minimum for the given parameters, and colour the regions where an atom is trapped blue, and the regions where an atom is untrapped red. It is clear that the MOT can potentially correct for quite significant force imbalances, up to just above a 50% difference between the two directions.

We can derive from this condition that a pyramid coated in aluminium ($\rho = 0.78$) can tolerate values of η close to 0.5, while a pyramid coated in gold (with $\rho = 0.986$) requires that $\eta < 0.1$ in order to create a stable MOT. This suggests that gold coated pyramids will not form a MOT without appropriate suppression of the Type III beams.

4.2 Macroscopic Glass Pyramid

In order to test the principle of the 70° pyramid MOT, a macroscopic model of the 70° geometry was created in glass. This was made from four glass cubes (side length of 13mm), each with

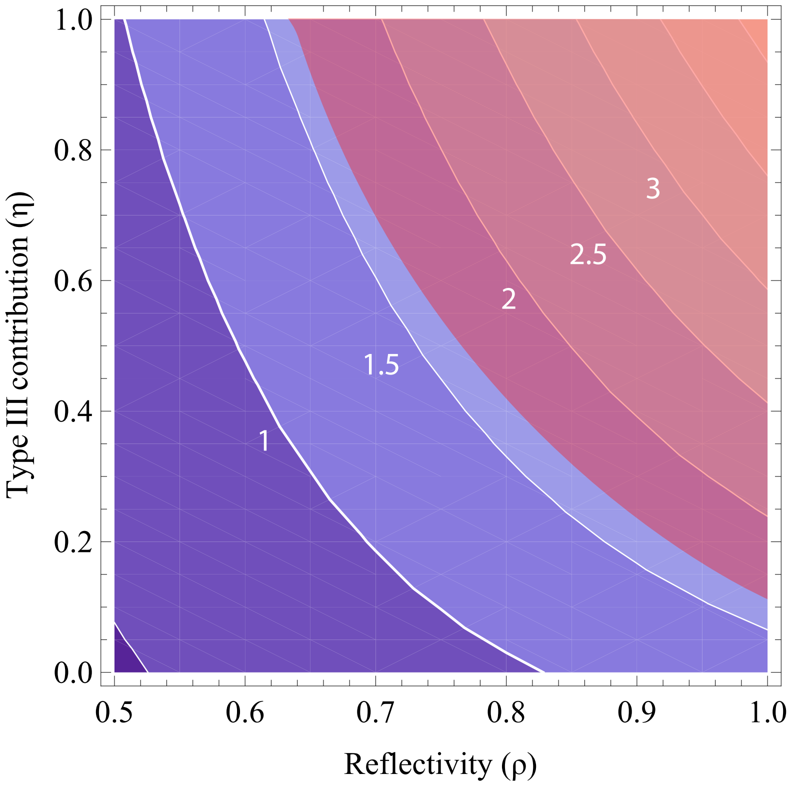

Figure 4.6: Plot showing the force imbalance at the position of the magnetic field minimum as a function of reflectivity and the intensity of the Type III beams (numbers show value of F_+/F_-). Blue regions are where atoms can be trapped, and red regions are where trapping is impossible.

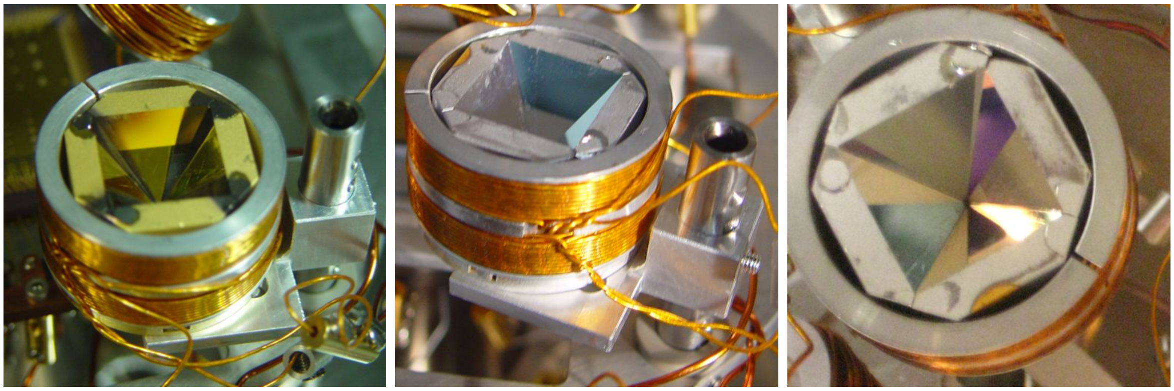

Figure 4.7: Macroscopic Glass Pyramid

the corners cut off (along the plane parallel to the body diagonal), polished and then glued together to create the pyramid hollow structure (with a pyramid side length of approximately 16.3mm). The assembled pyramid was glued onto a metal base, which was screwed onto a set of coils wound in the anti-Helmholtz configuration. These coils centred the quadrupole field at 3mm vertically above the apex of the pyramid.

4.2.1 Making a MOT in the 70° glass pyramid

The mirror surfaces of the pyramid were coated with a metallic coating and mounted in the vacuum chamber. This was illuminated with a circularly polarised light beam (typically composed of MOT intensity 5mW/cm^2 and repump intensity 0.5mW/cm^2) and field gradients of above 15G/cm were used. This has been demonstrated to successfully form MOTs with similar numbers of atoms as an equivalent sized 90° pyramid [48]. A typical fluorescence image of the MOT, compromising fluorescence from the MOT and reflections in the mirror surfaces is shown in figure 4.8.

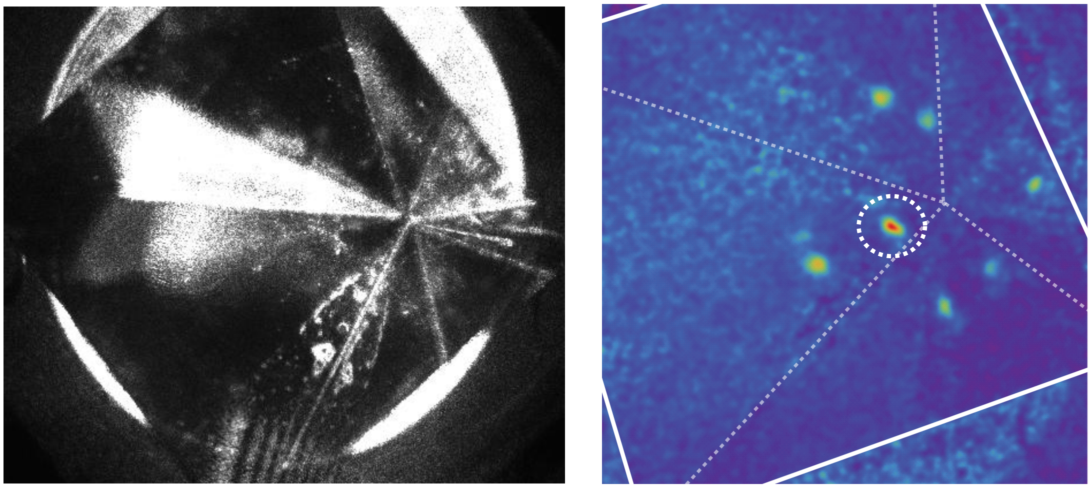

Figure 4.8: Left: Image of the glass pyramid under illumination Right: False colour image of trapped atoms inside the pyramid. The real MOT is circled but a number of reflections of the MOT also visible.

4.2.2 Choice of coatings

The pyramid was initially tested with an aluminium coating (middle and right pictures in figure 4.7), with a thin MgF protective layer and this gave a reflectivity of 78% (determined from the intensity of Type I reflections leaving the pyramid), forming a stable MOT comparable in size to a reference 90° pyramid. The pyramid was then recoated in gold, giving a reflectivity of 95%. No trapped atoms were observed in this pyramid, which is consistent with the predictions made in section 4.1.2. As a consequence it was necessary to suppress the formation of the Type III beams in the gold coated pyramid. This was accomplished with a 'flower' mask (see the left of figure 4.9) placed into the incident beam to block the corner areas giving rise to the Type III beams. Subsequently this was accomplished by only coating gold onto the centre part of the face (see left most picture in figure 4.7), leaving bare glass on the corners. Platinum and chromium were also previously tested with a reflectivity of 70.6% and 63.5% respectively, and these produced MOTs without the need for masking but with reduced atom numbers due to the low reflectivity[79]. All three pyramid coatings (platinum, chromium and aluminium) that worked unmasked also worked with the masking in the beam but showed no appreciable change in atom number trapped when the mask was introduced.

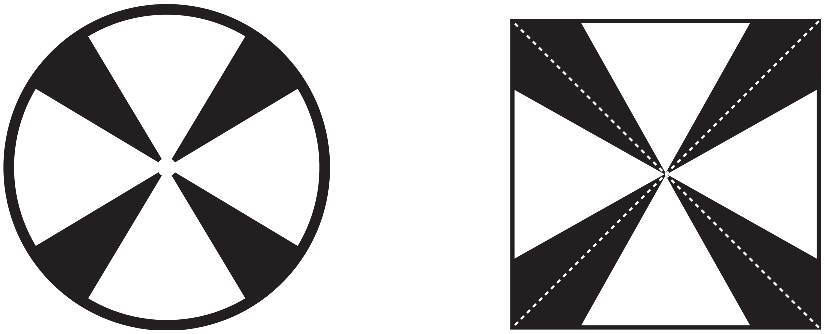

Figure 4.9: Left: Mask inserted into beam to suppress Type III reflection formation. Right: Areas of the pyramid faces coloured white should be coated with metal while the black areas should be left uncoated to prevent Type III reflections

4.2.3 Movement of MOT

An experiment was carried out to determine how the MOT moved inside the pyramid volume if the magnetic zero was displaced using external vertical shim fields. With knowledge of the applied field and measurement of the distance of the MOT from the apex obtained from calibrated CCD images, a measurement of the displacement of the MOT can be made. The distance can be judged in two ways, either by the projection of the displacement of the MOT centre from the apex by the viewing angle or by measuring the distance between two opposite reflections which through simple geometry can be related to the MOT height. We compare the position of the MOT inside the aluminium coated pyramid with and without the flower mask placed in in the incident beam. Displacements are measured from the centre of the quadrupole field created by the anti-Helmholtz pair, which is located approximately 3mm from the apex. Results are shown in figure 4.10, and show a linear translation with applied field. The gradient of a linear fit to the data is consistent with the applied magnetic quadrupole field gradient of 20G/cm in both cases.

The second interesting point is the constant offset between the two positions. For the two cases with and without masking, the positions at zero field are 3.52 ± 0.07 and 4.07 ± 0.08mm respectively, extracted from the linear fits. This corresponds to an approximately 0.6mm or 1.2G difference between the two cases, which we can compare with a value given by the numerical model used in section 4.1.2. The model indicates that a value of $\eta \approx 0.25$ would match the observed displacement. This is surprisingly small, and indicates that the effect of the Type III beams is not as severe as anticipated which could be due to the possibility that the MOT spring constant in the physical MOT is greater than that predicted by Doppler theory.

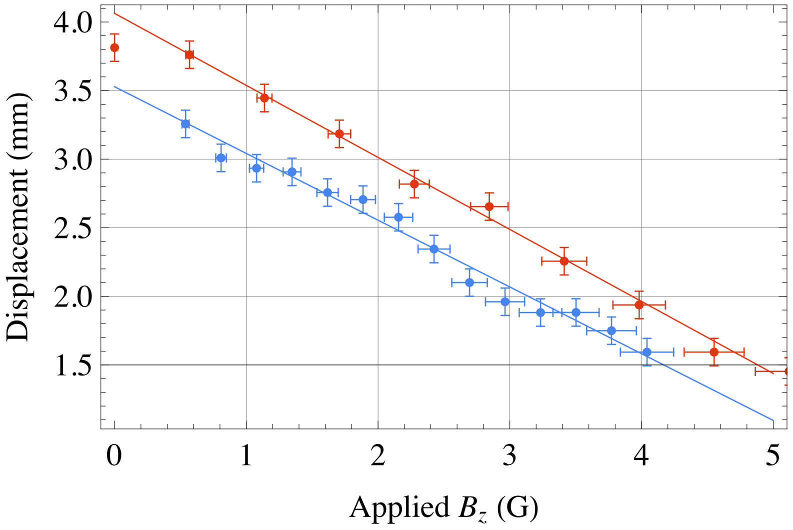

Figure 4.10: Physical position of the MOT as a function of applied shim in an aluminium coated pyramid with (blue) and without (red) masking of Type 3 beams

4.3 The first prototype pyramid chip

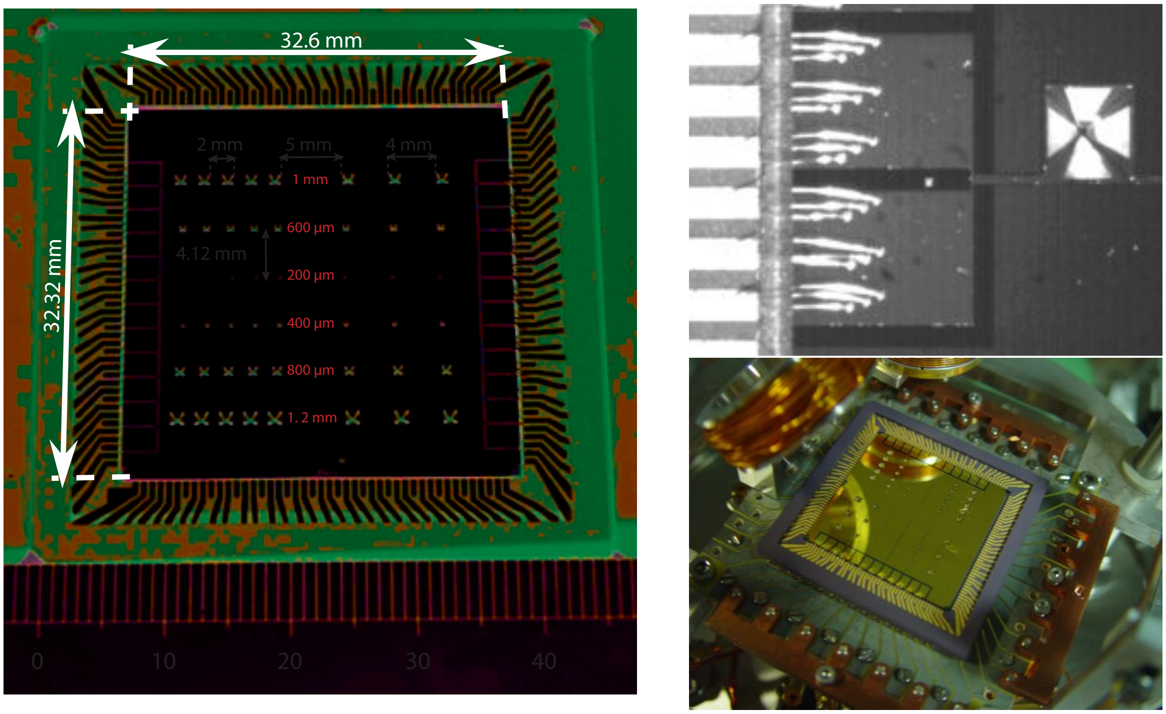

Figure 4.11: Left: The original prototype chip, Right Top: Microscope image of bondwires making connections to chip from CPGA. Right bottom: Assembled chip package mounted in FR4 circuit board

The main goal of this project was the miniaturization of the pyramid MOT. The results above obtained in a macroscopic model were used as a preparatory study for microscopic devices. Figure 4.11 shows the very first integrated micropyramid MOT chip designed and fabricated in collaboration with the Nanoscale Systems Integration Group at Southampton [81]. This device was a fully integrated chip with microfabricated wires and pyramids all contained in a chip package. It featured arrays of silicon pyramids from 140 to 800 microns in depth, all surrounded by a square loop of wire 50μm wide capable of carrying currents in excess of 1A. The chip package was a ceramic pin-grid array (PGA) type connector, which sat in a bed of sockets in a FR4 circuit board, making connections to the outside world. In this way the chip could be hot-swapped instantly without having to remake any connections. The development of this chip & package was covered extensively in a previous thesis [79] and hence I will not cover it in much detail here. This chip was created with gold coated pyramids, with the masking on the corners for Type III suppression. This was typically illuminated with 4mW/cm^2 of circularly polarised light at detunings of 10MHz, found to be optimal for the glass pyramid. The magnetic fields were formed from the square wire loops on the chip and an appropriate vertical bias field created by external shim coils, creating a quadrupole field of 20G/cm and above. The micropyramids were imaged onto a CCD and attempts to detect atoms in the largest 1.2mm pyramids began in January 2007. The most critical aspect for operation of the micropyramids was positioning of the magnetic field at the centre of the pyramid. The assembled chip package had been found to have insignificant stray magnetic fields and field gradients within the region of the micropyramids and hence it was just a question of locating the centre of the quadrupole field within the pyramid volume. Based on studies of the stray magnetic fields present above the chip a region 500mG×500mG×1G in $B_{x,y,z}$ was searched. At a gradient of 20G/cm the magnetic field needed to be accurately positioned to within an estimated 200mG (100mG in x/y) meaning at least 125 combinations of observations were required for a given set of parameters. The number of combinations could be reduced for higher gradients, but for gradients above 40G/cm (requiring 0.5A) the chip wires could not be run continuously due to risk of thermal failure.

Despite an intensive search of trap parameters no atoms were observed in the micropyramids. We had expected to observe numbers in excess of 10 000 atoms in the 1.2mm pyramids, based on equation (2.31) and (2.34). An alternate argument can be made by applying the $L^{3.6}$ scaling we expected from past study of 'normal MOTs' to extrapolate from the atom numbers seen the large glass pyramid, which can trap up to 10^7 atoms with a 12mm beam. If we extrapolate from this we expect to see 2 500 atoms in the 1.2mm pyramid.

Evidently improvements were needed in detection in order to see smaller numbers of atoms.

4.4 Imaging

4.4.1 Absorption vs. Fluorescence Imaging

Two techniques can be used to detect atoms inside the pyramids. Fluorescence imaging where the photons scattered by atoms illuminated by laser light are collected to create an image, or absorption imaging where the atoms are observed by detecting the shadow cast in probe beam due to absorption.

The simplest way to compare the two processes is to compare the ratio of signal to background, $\rho = S/B$, for each case. For absorption imaging the ratio is just $1 - T(x, y)$ where

$T(x, y)$ is the transmittance through an atom cloud of low density $n(x, y, z)$ giving

$$\rho_{Abs} = 1 - T(x, y) = e^{-\sigma_L \int n(x,y,z)dz} \approx \sigma_L \int_{-\infty}^{\infty} n(x, y, z)dz = \frac{3N_0\lambda^2}{4\pi^2\sigma_L^2} \tag{4.4}$$

where the probe beam is traversing in the z direction and imaging onto the xy plane. The absorption cross section σ_L is given by $\frac{3\lambda^2}{2\pi} = 2.9 \times 10^{-13}\text{m}^2$ on resonance for rubidium [55].

For fluorescence imaging we consider the rate of photon scattering per atom R (see equation (2.4)) and multiply by the solid angle covered by the collection optics $d\Omega$, the area per pixel A, and the energy per photon $^{hc}/_{\lambda}$ to give total peak collected power

$$F = \frac{d\Omega Rhc}{4\pi\lambda} \int n(x, y, z)dxdydz = \frac{AfN_0}{8\pi^2\sigma_L^2} \tag{4.5}$$

where $f = \frac{d\Omega Rhc}{4\pi\lambda}$ is the flux per atom at the camera position. The background we detect in our case comes predominately from light scattered by the pyramid surface, and we define the background power per pixel as b giving

$$\rho_{Fluo} = \frac{F}{b} \tag{4.6}$$

If we divide the two ratios, several terms cancel out leaving a ratio independent of atom number and MOT size

$$\frac{\rho_{Fluo}}{\rho_{Abs}} = \frac{Af}{\sigma_L b}. \tag{4.7}$$

Assuming that the atoms in the MOT scatter at close to saturation, i.e. $R \sim \gamma/2$, and using typical experimental values $d\Omega/4\pi = 0.004$ and $A=7.4\mu\text{m}\times7.4\mu\text{m}$ gives the condition that fluorescence detection will be superior than absorption imaging if $b < 3.6$pW per pixel.

The scatter in our micropyramids is around this level, but fluorescence imaging is technically simpler, as it does not require additional beams and images can be acquired continuously from the MOT without having to reload the MOT each shot unlike the destructive absorption imaging process[1]. As a result we will pursue fluorescence imaging for the detection of small atom numbers in the MOT.

4.4.2 Background Subtraction

The primary issue surrounding the imaging of small numbers of atoms trapped in the micropyramids is the issue of resolving the fluorescence from the atoms against the scatter background. Unlike other atom chip experiments, the atoms are trapped inside the pyramids and so we cannot image along a perpendicular axis to the atom chip face. Likewise there is no chance of spectral discrimination as the light from the atoms will be of the same frequency as the scattered laser light.

The basic principle of the imaging, is that two images of the micropyramid are taken, a foreground image with the MOT present, and a background image without. Subtracting the two images from each other should yield an image of the MOT alone.

---

[1] It is possible to make absorption images in-situ, i.e. non-destructive, but the absorption cross section for the probe beam decreases as the atoms are being saturated by the MOT laser beams

Poissonian Noise From the discussion in section 3.4.4 we known that the statistical error on a large CCD count is $\sqrt{0.82N}$, and thus when two identical images are subtracted the standard deviation of the difference is $\sqrt{2 \times 0.82N}$. Considering a foreground image $N_{fg} = N_{sig} + N_{back}$ and a background image $N_{bg} = N_{back}$, subtraction of these two images will yield an image with mean value N_{sig} and a signal to noise ratio of

$$SNR = \frac{N_{sig}}{\sqrt{0.82(N_{sig} + N_{back}) + 0.82N_{back}}} \approx \frac{N_{sig}}{\sqrt{1.64N_{back}}} \tag{4.8}$$

where we have assumed the signal to be much smaller than the background. If we now repeat the process several times, taking n foreground and n background images then subtracting them all we get the following relation

$$SNR = \frac{n \times N_{sig}}{\sqrt{n \times 1.64N_{back}}} = \sqrt{n}\frac{N_{sig}}{\sqrt{1.64N_{back}}} \tag{4.9}$$

thus showing an increase in the signal to noise as $\sqrt{n}$. To give an order of magnitude estimate, it would be desirable to see a signal of 1 count, on a ‘full-well’ background (assumed to be 10 000 counts to maintain sensor linearity). To obtain a SNR of 1 requires 20 000 images (10 000 each foreground and background) which is quite feasible for the camera in use which can potentially run at continuous speeds up to 1kHz.

Crucially foreground and background images need to be taken as close together as possible in order to avoid sampling the low frequency technical noise discussed in section 3.4.4.

4.4.3 Switching the MOT

A factor of some importance is that we need a reliable ‘switch’ for the MOT. Aside from all the noise sources mentioned above, the most critical factor is ensuring the background and the foreground frames differ only by the presence (or absence) of the signal from the MOT. Hence it is has to be possible to switch the MOT off and on without affecting the background at a level comparable to the MOT signal. To use the numbers from above, the background should ideally remain constant to at least 1 part in 10000. In addition to scattered light from the pyramid faces, the background signal contains fluorescence from the rubidium vapour which can also cause attenuation of scattered light. This component adds a sensitivity of the background to changes in laser frequency and magnetic fields, placing constraints on the types of process that can be used to enable and disable the MOT.

The separation of background rubidium signal and MOT signal is also a reason why we use a CCD instead of a single-pixel photodetector (which would offer advantages in noise and bandwidth), as a CCD allows spatial discrimination of a ‘MOT-like’ object as opposed to a uniform background.

Temporal Dynamics

The most successful technique proposed and implemented relies on the fact that the MOT takes a finite time to load. Background and foreground images are taken at the start and the end of the loading process, thus eliminating any background fluorescence change as parameters remain identical. This does limit the number of frames that can be taken however, as background frames can only be taken as long as too few atoms are present in the MOT to be undetectable.

The process is illustrated in figure 4.12. We start with a full MOT, and take n foreground frames $\{f_1, f_2..f_n\}$. We then empty the trap by shutting off the MOT quadrupole field for 50ms, before turning it back on again and taking the background frames $\{b_1, b_2..b_n\}$. For a loading time of ∼1s, and setting the constraint that during the background frames the atom number has to remain less than 5% of the steady state number, there is a time window of around 50ms to take background images. The two sets of frames are then averaged and subtracted to make what we refer to as a 'subtracted image'

$$S_i = \frac{1}{n}\sum_j^n f_j - \frac{1}{n}\sum_j^n b_j. \tag{4.10}$$

The whole process is repeated N times and all the subtracted images are averaged to form the final image R,

$$R = \frac{1}{N}\sum_i^N S_i. \tag{4.11}$$

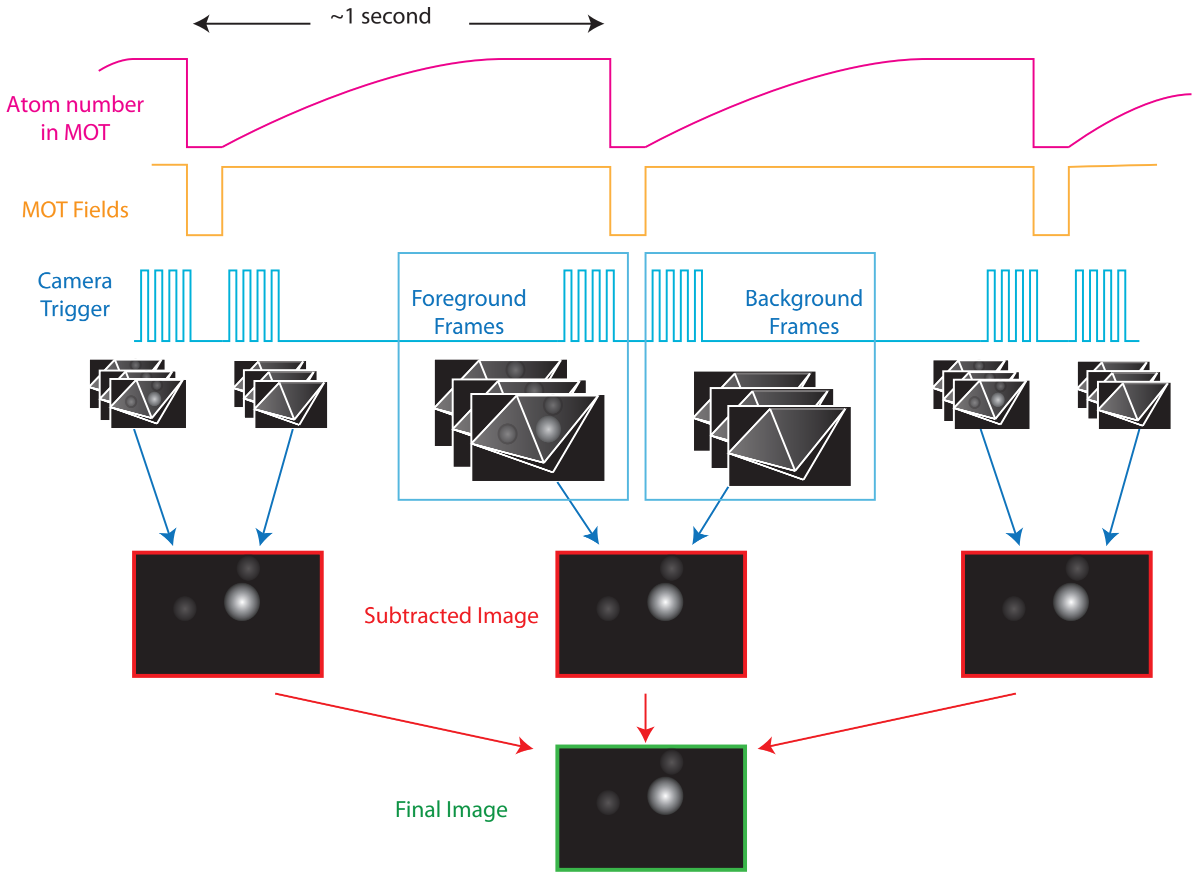

Figure 4.12: How images of the MOT are taken by taking frames at the start and end of the loading process

Shim Coil Switching

Another successful technique was the use of shim fields to control the position of the MOT within the pyramid. This offers the an attractive prospect for the control of the MOT as only small fields need be used to displace the centre of the MOT appreciably. For example a micropyramid with a magnetic field gradient of 10G/cm in x, y, a change of 1G corresponds to a displacement of the MOT by of 1mm. This is excessive for our needs as we only need move the MOT by its own radius in order to see the movement on the imaging plane. Small temperature-limited MOTs formed in a large pyramid using the same parameters typically had $\sigma = 150\mu$m, corresponding to a required step of only 150mG in the shim field.

We can use this 'switch' in two ways. Either to destroy the MOT for the background frames by moving it into a wall, or just to move it by a small amount and keep it intact. The former (destruction of the MOT) offers the best contrast but obviously requires the MOT to be reloaded and hence the rate at which integration can be carried out is limited by the loading time (which is similar to the previous imaging technique). If the latter non-destructive technique is used, there is no need to wait for reloading and large numbers of images can be taken quickly.

For the non-destructive process the images have a characteristic MOT/AntiMOT signature (see figure 4.13 which shows the images formed as the amplitude of the field step was varied) which is a good unambiguous indicator of the presence of a MOT.

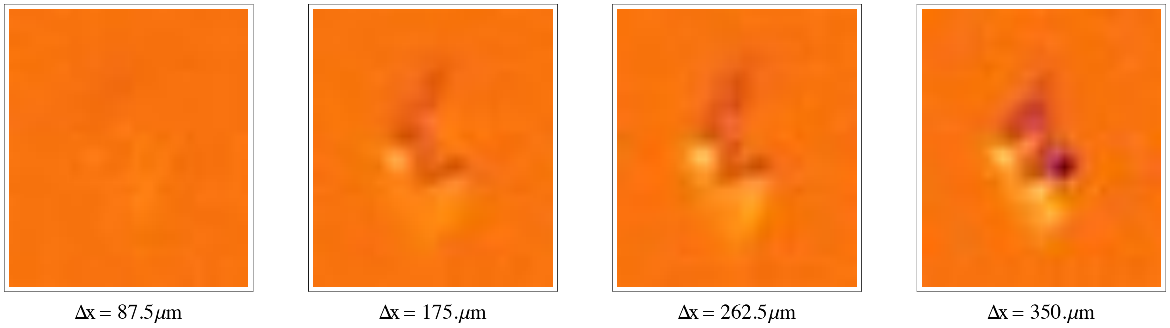

Figure 4.13: Subtracted image as a function of modulation amplitude using non-destructive shim switching. The width of the MOT $\sigma_{MOT} = 120\mu$m

Summary Overall we found that the using the shim fields to turn the MOT on and off (i.e. destructive) offered the best contrast, but at the cost of introducing a slight bias in the final images due to modulation of the background fluorescence. As a result when precise quantitative data was required we used the loading time imaging which eliminates this bias. The non-destructive imaging has the worst contrast, but has the advantage that it is significantly (a factor of 20 times) faster than the other two methods. It is also often troublesome to find an axis of movement for the MOT where the atoms can be moved intact, and parallel to the CCD plane.

4.4.4 Choice of n and N

In order to optimise the imaging procedures a brief study of the optimal values of n (the number of frames per subtracted image) and N(the total number of subtracted images averaged) was performed. The left figure in 4.14 shows the standard deviation of pixel values in a subtracted image initially decreases with increasing numbers of frames n, but after a certain point begins to increase as the low-frequency noise regime is sampled. The optimal number of frames per cycle is found to be 14 from frames spaced by $\Delta t = 10$ms and 80 frames for frames spaced by $\Delta t = 1.7$ms, indicating that the low frequency technical noise is not sampled until $n\Delta t > 135$ms.

The number of subtracted images N is not limited by any physical process apart from available time, and from the right of figure 4.14, it is seen that at values of N above 10 the noise follows a clear $N^{-1/2}$ dependence. This data was taken in far from optimal situations, where the lasers were not well locked and the lab was noisy. This resulted in an overall increase in the noise level by a factor of approximately 3 from typical values.

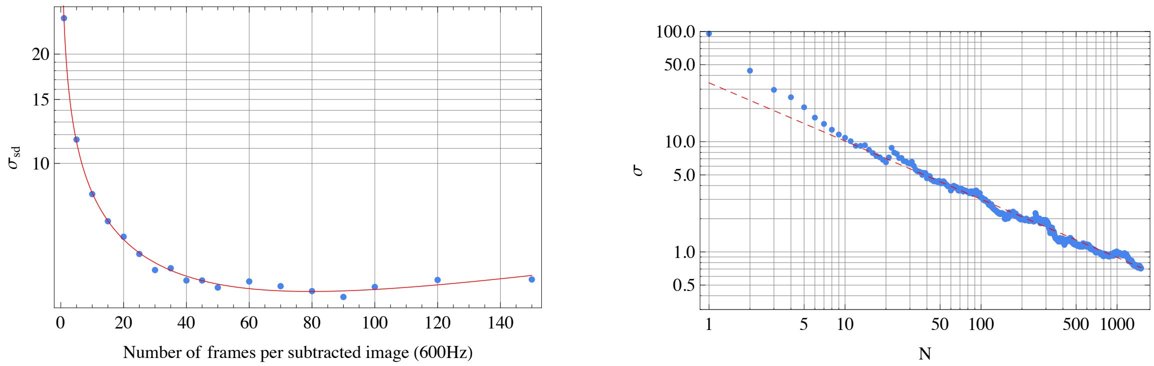

Figure 4.14: Left: Noise level as a function of n at 600Hz after $N = 10$ images. Right: Standard deviation of pixel values in an integrated subtracted image level as a function of N

4.5 Detection of a signal

The ultimate test of our imaging is the ability to accurately identify a MOT in an image. A typical experiment will produce up to 200 images for a variety of parameters, and crucially we want to know if any of these parameters produced a successful MOT in a micropyramid on the prototype chip, confirmed by the visibility of a cold trapped atom cloud in any of these images.

It was found that the time required to scrutinise batches of images manually is not so laborious that it warrants the effort required in creating a potentially quite complex image recognition program to study the data automatically. The main disadvantage of using the human visual system to analyse data is that occasionally it behaves a little too efficiently, and noise artefacts can be misidentified as a real signal leading to wasted time and effort chasing after a red herring. We will consider the issue of these 'false-positives' later.

Our first question is, how prominent does a signal have to be to be unambiguously detected above a noise background? We can try a brief experiment to assess this, where we take a an artificial MOT signal (of variable brightness) and superimpose it on a noisy background. Some images are displayed in figure 4.15, where we have superimposed the artificial MOT upon two types of background, a uniform white-noise background and an inhomogeneous background similar to that obtained in the micropyramids. The results are very subjective, and depend a lot on factors such as the color scheme used, the distance of the eye from the page and the optimism of the person viewing the images. The author finds that he prefers a SNR of at least 1.4. At this level however, the occurrence of 'false positives' is still quite significant. Setting a guideline of SNR=2 has been observed to reduce the occurrence of false positives to an insignificant level and hence we adopt this as our criteria for det

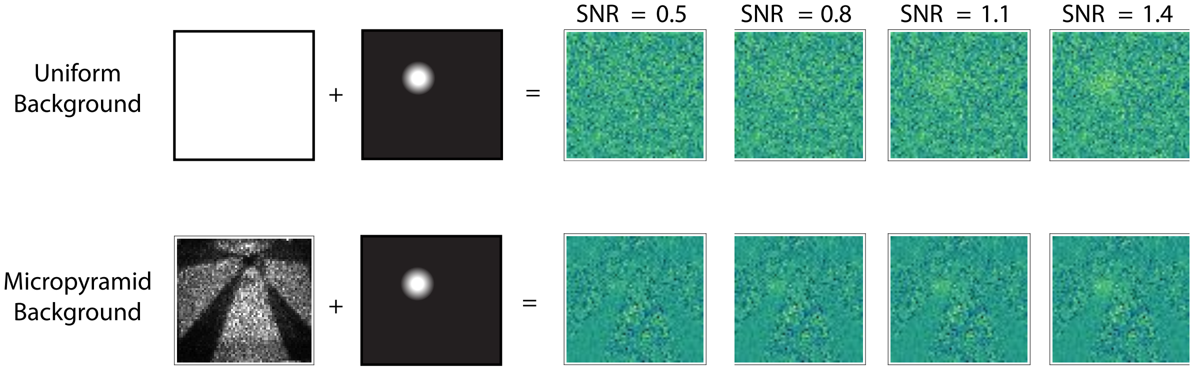

Figure 4.15: Image of a MOT superimposed upon a noisy background for various SNR levels. Top row is a homogeneous background. Bottom row is an inhomogeneous background similar to that inside the micropyramids.

4.5.1 Image Filters

Significant improvements can be made by applying image filtering techniques to the resultant image, and the correct choice of filter can dramatically enhance the obtained SNR and produce usable results from much shorter integration times. Our subtracted images are dominated by shot noise, which in the absence of a signal in an image, gives a range of pixel values distributed about a mean of zero and spatially uncorrelated. The technical term for this kind of noise is 'salt-and-pepper' which nicely describes the appearance of the image when viewed with a monochrome palette.

Kappa-Sigma Filtering Our first filtering step is to discard any outliers in the subtracted image, typically these points arise from features on the pyramid which scatter substantially more light than the norm and hence produce higher counts (and hence higher variation from shot to shot). These anomalous points, if left in the image for further processing, can be artificially enhanced and increase the possibility of a 'false positive'. As a result immediately following subtraction of the background and foreground frames, all points which lie more than three standard deviations away from the mean in the subtracted image are replaced with 0.

This may sound drastic, but the key is only MOT signals with an initial SNR significantly above 1 (which should be readily distinguishable with no need for image filtering) will be removed from the image.

Gaussian Filtering As the majority of the noise is uncorrelated from pixel to pixel large gains can be made by use of spatial filtering. The first approach tried was replacing each pixel value with the median value of the neighbouring pixels (typically 3 to 5 pixels). This is a much more robust technique than purely taking the mean value as isolated anomalous pixels in a neighborhood do not significantly affect the median value.

A better approach was to use a Gaussian filter, which effectively takes a weighted average of the neighboring pixels, with the weighting given by the distribution of a 2D Gaussian centred on the pixel of interest. This was carried out by convoluting the 2D array of pixel values, with a discrete 2D Gaussian kernel truncated three standard deviations from the centre (as it is effectively zero after this point).

This filter acts as a low pass, and hence is very effective in reduction of the salt and pepper noise dominating the image.

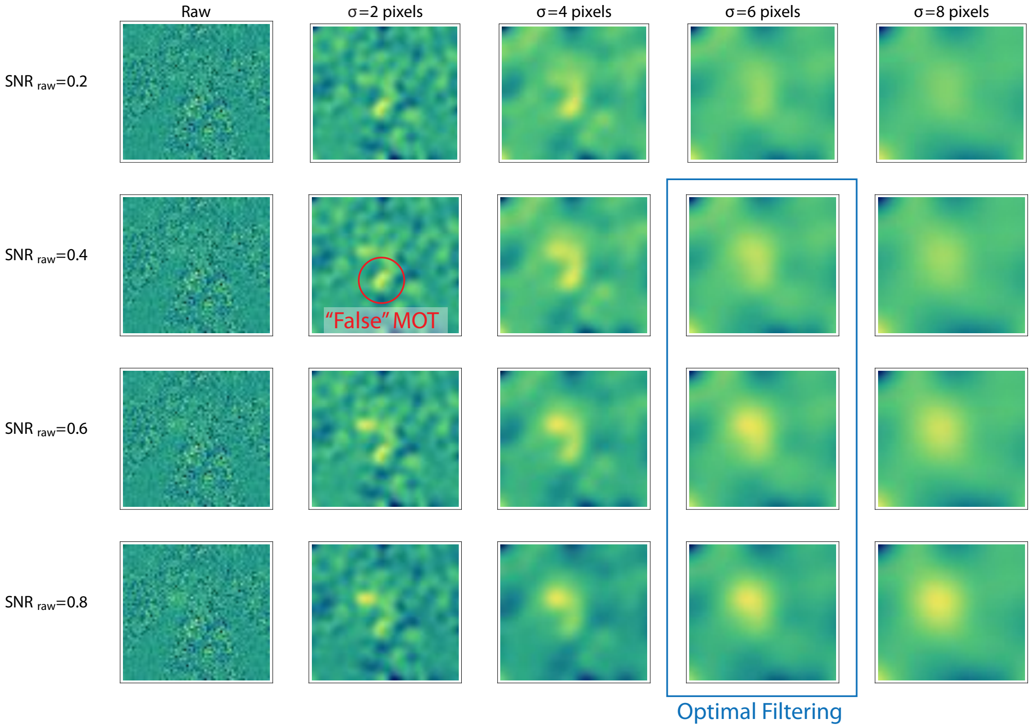

Figure 4.16: Performance of Gaussian Filter as a function of noise level and filter parameters. Optimal filter parameters outlined in blue, and the appearance of a 'fake MOT' is indicated in red. The dynamic range shown in each image is the entire range of the pixel values in each case.

Figure 4.16 shows the effect of Gaussian filter parameters (namely the width of the Gaussian σ_k) on an image of increasing signal to noise ratios (note: the dynamic range of the image is the full range of pixel values). Interestingly the filter for some values ($\sigma_k = 2$) initially enhances the appearance of a noise artefact, but this is suppressed at higher widths. It appears that a

Gaussian width equivalent to that of the original MOT size (i.e. $\sigma_k \approx 5$ pixels) is optimal to suppress the initial noise & artefacts while not adversely affecting the MOT signal by smearing it out as is apparent at larger widths.

The perceptibility of the signal is much improved, with the MOT image much more distinguishable from the noise at an initial level of 0.4ADU.

More advanced filters We trialled two more filtering methods. These were a nonlinear diffusion process[82] and a method based on recreating the background from a basis set of 'eigenframes'[83]. Both offered some slight improvements in effective noise reduction, but at the cost of increased complexity and computational requirements. There was also the fact that non-optimal performance due to poorly chosen filter parameters inhibited detection of a signal and hence overall it was found they did not offer much advantage over our chosen filtering methods.

Implementation of filters

Filtering is applied continuously throughout the acquisition process, with the Kappa-Sigma filter (K_σ) applied to each individual subtracted image, and the Gaussian filter applied at the end of the experiment by convolution of the filter kernel (G) with the final image.

$$S_i = K_\sigma \left[\frac{1}{n} \sum_j^n f_j - \frac{1}{n} \sum_j^n b_j \right] \tag{4.12}$$

$$R = \left(\frac{1}{N} \sum_i^N S_i \right) * G \tag{4.13}$$

$$\tag{4.14}$$

All filters are implemented using the IMAQ toolkit for Labview, allowing for complete integration with the control software. Crucially we also saved the unfiltered data to allow for retrospective filtering of the original image.

4.5.2 Validation of image filtering

Looking at small number MOTs

In order to check the filtering with real atom clouds, small MOTs were created in the macroscopic glass pyramid by limiting the vapour pressure of rubidium in the vacuum chamber. This meant MOTs of only a few 1000 atoms could be created on demand to test the detection. Results are shown in figure 4.17. The initial image of the pyramid as seen by the camera is shown in 4.17(a). There were 30 foreground/background frames per subtracted image, and the resulting image at various stages of the integration are shown in 4.17(b). The Gaussian filter was found to give up to one order of magnitude improvement in the SNR and the MOTs were not distinguishable without the filtering. In total 50 subtracted images were taken (result shown in lower part of 4.17(a)), taking approximately 1 minute in total. There were $\sim 2\ 000$ atoms in the MOT, giving a peak count of 1.5ADU. The increase in the signal to noise is shown in 4.17(c) and proceeds as $N^{1/2}$ as expected. In practice only 10 subtracted images would be needed to detect this MOT, indicating that 100 subtracted images (taking approximately 90 seconds worth of integration time) could detect MOTs with a third of the atoms detected here. As previously discussed the limiting factor to this imaging process is the loading time of the MOT, and if this time could be reduced or eliminated (potentially with the non-destructive shim switching), this time could be reduced from 90 seconds to approximately 15 seconds with the camera running at 400Hz.

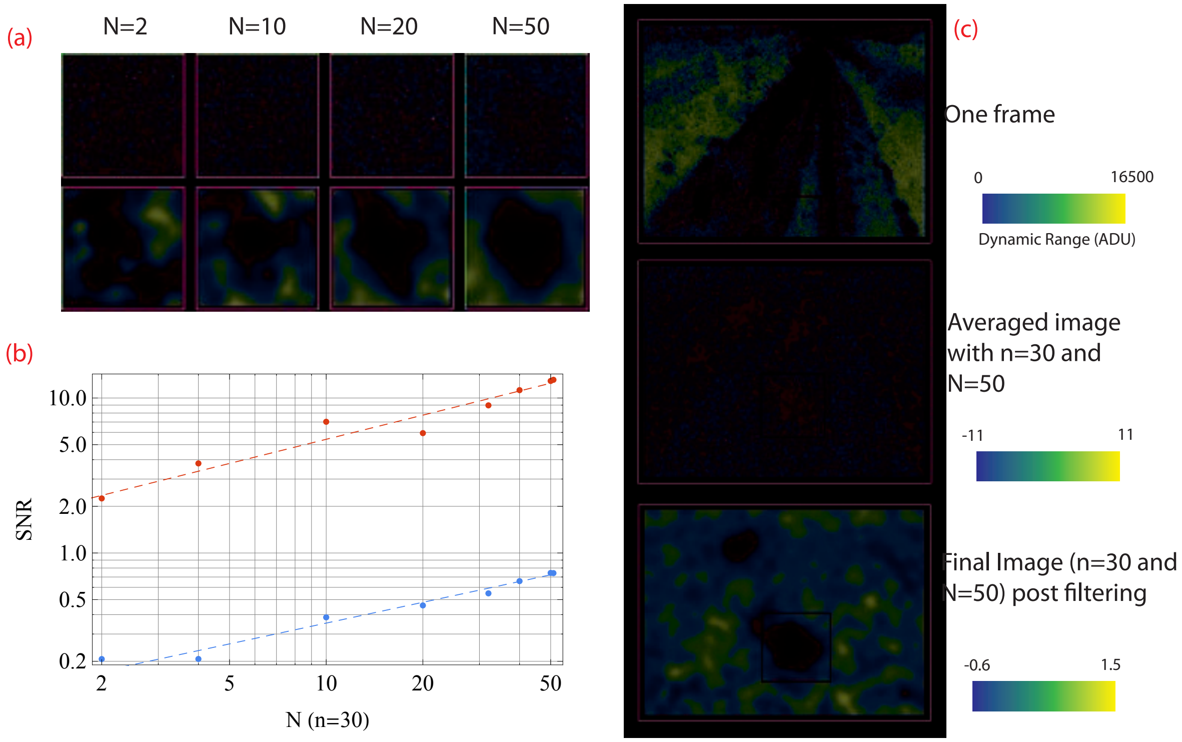

Figure 4.17: Imaging a MOT of approximately 3000 atoms. (a) The images (raw and Gaussian filtered) at various stages in the integration (b) SNR of the MOT as a function of the number of subtracted images N (c) The initial frame, and the final images, requiring 1500 frames in total. White square indicates region of image displayed in (a). Dynamic ranges of false color maps are the full range of pixel values in each image.

How low can we go?

A primary result of some importance is to derive the number of atoms that we can detect on a given background of scatter. Based on time requirements we define an upper limit to integration of $N \leq 100$. At that point, the standard deviation of the unfiltered final image is approximately 1ADU. Following Gaussian convolution this reduces to approximately 0.25ADU. At this level a signal of at least 0.5ADU needs to be obtained to give a good degree of distinguishability.

At typical atomic scattering rates and using our imaging geometry, the camera will collect 16fW of scattered light from each rubidium atom (approximately 20 CCD counts for a exposure time of 4ms). This sets the threshold of 320pW (approximately 400 000 CCD counts) of scatter

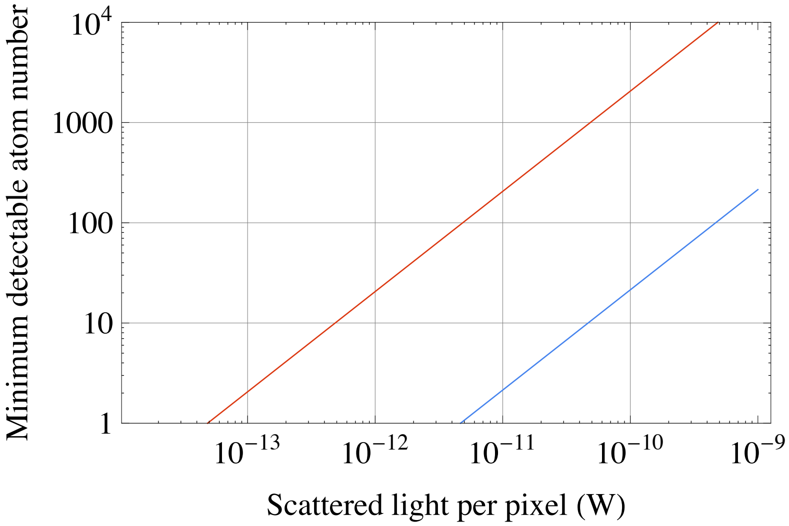

Figure 4.18: Threshold for detectability for a MOT of radius 10μm (blue) and 100μm (red)

per pixel to resolve a single atom assuming all its fluorescence collected at this solid angle is localized on a given pixel. Typically however MOT atoms are not localised to one point, but thermal motion distributes the fluorescence over an area of several pixels. As we know the typical spatial distribution of the MOT, we can work out the number of atoms that are needed to give a distinguishable signal on a given background level. Figure 4.18 shows the threshold for detection given a particular scatter level. This makes the assumption that imaging is 1:1 on the Pike CCD with the solid angle used in the experiment, and that the MOT has a 3D Gaussian profile with $\sigma = 10$ or 100μm. This shows that great benefits can be made with denser MOTs, localising the atomic fluorescence onto fewer pixels.

4.6 Implications for the prototype chip

At the end of the search in the micropyramids the typical noise level, following an imaging process consisting $n = 20$ frames, $N = 100$ subtracted images and filtering with the Gaussian filter, was $\sigma = 0.2$ADU for a scatter-limited exposure time of 4ms. This limited our sensitivity to detecting MOTs with a peak count of 0.4ADU and above. Despite these improvements in detection, increasing sensitivity by several orders of magnitude [2] no consistent signal was detected in any of the pyramids. Assuming typical MOT parameters, and the MOT taking the form of a Gaussian with width 100μm, a peak count of 0.4ADU corresponds to approximately 100 atoms in the MOT. We can hence exclude MOTs above this number and thus conclude that our 1.2mm pyramid captured fewer than 100 atoms. This was not consistent with the extrapolation from the glass pyramid based on the predicted $L^{3.6}$ scaling discussed in section 4.3, and hence a study was undertaken into how the number of atoms in the MOT scaled at small dimensions.

---

[2] The initial sensitivity before improvements to imaging was estimated at 10000 atoms

4.7 Measurement of Scaling Law in Macroscopic Glass Pyramid

A 70° pyramid MOT of variable dimension was achieved with the macroscopic glass pyramid by simultaneously translating the MOT towards the apex and reducing the size of the trapping beam. The scaling of atom number was measured with three different configurations: gold coated with flowers masked on the mirrors; aluminium coating without flowers; and aluminium coating with a flower mask placed into the beam.

The most challenging aspect of these measurements was using aperture sizes below 7mm, where the alignment of the aperture on the central axis of the pyramid became important. A circular aperture was used throughout, square apertures were initially tried but an adjustable circular aperture offered more flexibility and greater ease of alignment. The aperture was mounted on a x-y translation stage so it could be precisely centred on the pyramid axis. For a given aperture diameter, the vertical magnetic field bias was coarsely scanned until a MOT became visible, and then fine adjusted to maximise the atom number. Results of the observed scaling can be seen in figure 4.19. The results are plotted on a log-log scale and the apparent trend is consistent across all four data sets. Starting at large aperture diameters, the rate of decrease in atom number is initially slow, but below 10mm appears to increase significantly. The linear trend below 10mm is characteristic of a power law relation between the aperture size (recorded as $\sqrt{area}$ to allow for comparison between square and circular apertures). Fitting a power law to data points of $L < 9$mm (and excluding the last data points, which were generally difficult to obtain and highly sensitive to aperture alignment) gives values of the exponent listed in table 4.1. If the data sets are all combined, it is apparent that apart from variations that can be explained in terms of variation of rubidium density and background pressure they all have the same trend. Fitting to the points of $L < 9$mm of the combined data set gives a exponent of 6.13 ± 0.27.

Pyramid	**Exponent of scaling**
Masked Aluminium	6.74 ± 0.80
Masked Aluminium Set 2	5.84 ± 0.56
Masked Gold	6.32 ± 0.78
Unmasked Aluminium	5.20 ± 0.40
Combined	6.13 ± 0.27

Table 4.1: Scaling of atom number observed in different types of pyramid

Three conclusions can be initially drawn from this process. Firstly it has little effect whether the mirror coating is gold or aluminum, the exponent of the scaling being 6.32±0.8 and 6.74±0.8 for gold and aluminium respectively. Secondly, the presence or absence of Type III reflections in aluminium coated mirrors has no significant effect on the scaling of the atom number. Finally, the measured exponent of 6.13 ± 0.27 from the combined data sets is completely in agreement with the L^6 scaling expected. If a linear fit is made to the points of aperture size > 9mm, the

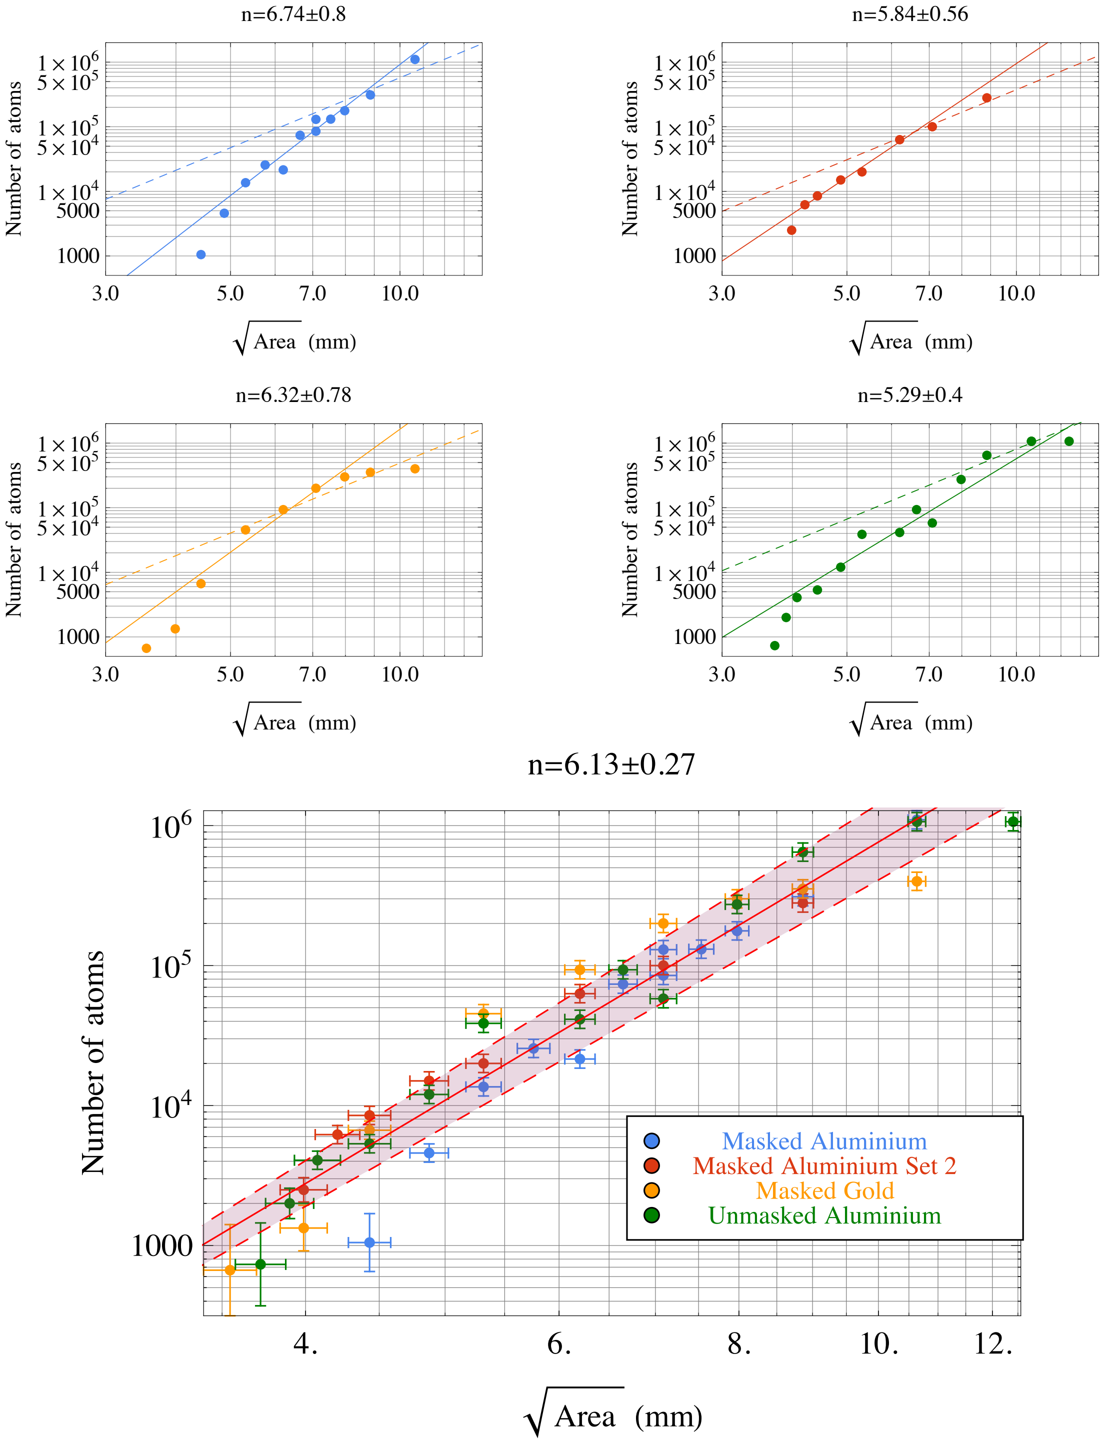

Figure 4.19: Scaling of atom number in the glass pyramid. Top graphs show the individual data sets along with a $L^{3.6}$ trend line (dashed) and a best fit power law (solid line). The exponent of the best-fit power law is shown at the top of the graph. Bottom graph shows all data combined and the combined best-fit model, with the shaded region indicating the range of uncertainty in the fit.

exponent is 4.2 ± 1.2 which is consistent with the expected value of 3.6 for larger MOTs. We were unable to extend our observation to many more larger sizes as we were limited by the size of the trapping beam and available power. Overall we make the initial hypothesis that at aperture sizes below $\sim$ 10mm, the nature of the trapping process begins to change from scattering-rate limited to size limited.

We are obviously still a long way from the 1.2mm pyramids on the prototype silicon chip, as the smallest aperture that gave an observable atom number was 4mm in diameter ($\sqrt{area} =$ 3.72mm). The sharp roll off in atom number as it approaches 4mm (visible in most data sets in figure 4.19) could result from the fact that the glass macroscopic model was not perfect all the way down to the apex. It features a significant 'apex-defect', on the order of several hundred microns in size which will restrict the distance of closest approach to the apex as the second Type I reflections will be absent in a region directly above the defect. From figure 4.20, the size of this region can be estimated to be $1.3a$, where a is the size of the defect, and by estimating the size of the defect in the glass pyramid to be 400μm, trapped atoms will be lost when they get closer than 520μm to the apex.

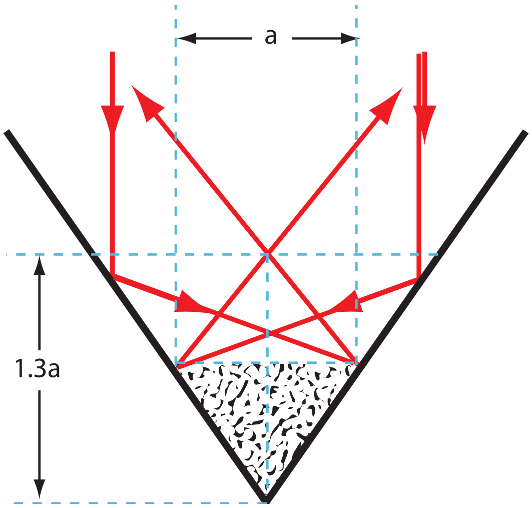

Figure 4.20: The closest distance to the apex in the pyramid is limited by the defect (size a) in the apex. This defect prevents the second reflection (providing the vertical restoring beams) from forming below a height of $1.3w$

Despite this fact, if we extrapolate the observed L^6 trend down to the 1.2mm scale, the graph shown in figure 4.21 implies that a 1.2mm micropyramid will at most capture 4 atoms from the background vapour. This is consistent with the predicted lower limit to detection discussed in 4.6 and hence it is likely that this is the cause of the failure to observe a MOT in the prototype chip. The chip also featured a number of smaller pyramids and these were also tested to but they did not show a positive result either. The absence of a visible MOT was also confirmed in a number of smaller pyramids present on the same chip (down to 200μm in size) to exclude the possibility that the 1.2mm pyramid was defective.

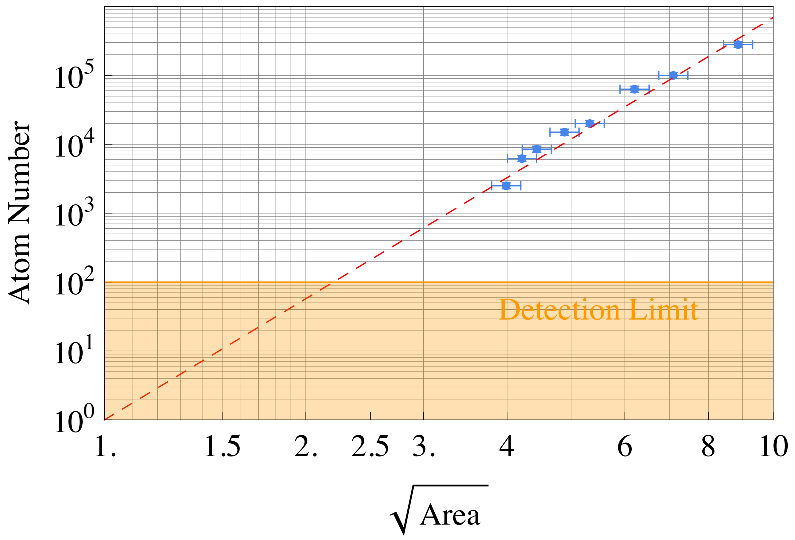

Figure 4.21: Data from section 4.7 (blue) and extrapolated fitted power law (red dashed line). The orange region shows the limit to the imaging obtained and indicates the upper limit to the number of atoms possible in a 1.2mm pyramid.

4.8 Beginning to understand the scaling with pyramid size

It has been shown that the scaling of atom number in the large glass pyramid is apertured scales as L^6, which it is hypothesised arises from the linear dependence of capture velocity to stopping distance. It is important that we are able to confirm this fact, by excluding the possibility that the drop in atom number results from a loss process. As a result a series of further experiments were carried out in the glass pyramid to determine the loading and loss rates of the MOT as a function of distance from the apex. This time we do not aperture the light, so the measured behaviour is solely resultant from the proximity of the MOT centre to the pyramid walls.

For reference, the data contained in figures 4.22 and 4.23, which was obtained experimentally by this author, has appeared previously in the thesis preceding this one[79]. The data is reproduced here as the analysis and conclusions drawn are different.

4.8.1 Loading and Loss rates

As we are in the regime of 'small' MOTs, where density dependent loss rates are negligible the loss rate is determined from the loading curve (from equation (2.36)). This is measured by taking a sequence of CCD images at appropriate time intervals. Since the steady state atom number can be determined from quantitative CCD measurements (section 3.4.5), then both the loading and loss rates at a given position in the pyramid can be determined.

Further information regarding the loading and loss processes can be obtained by varying the background rubidium density and measuring the effect on the loading and loss rates. From the theory in section 2.5.1, the loading rate $R = n_{Rb}\frac{A}{\pi}\left(\frac{v_c^4}{\overline{v}_{Rb}^3}\right)$, and the loss rate due to collisions with hot untrapped rubidium atoms $\Gamma_{Rb} = n_{Rb}\sigma_{Rb}\overline{v}_{Rb}$ both have a linear dependence on the rubidium density n_{Rb} and hence the total number of atoms as a function of n_{Rb} can be written as

$$N = \frac{R}{\Gamma_{tot}} = \frac{\alpha_l n_{Rb}}{\beta_l n_{Rb} + \gamma_l} \tag{4.15}$$

where

$$\alpha_l = \sqrt{\frac{6}{\pi}} A \left(\frac{v_c^4}{v_{rms}^3} \right) \text{ and } \beta_l = \sigma_{Rb} \bar{v}_{Rb} \tag{4.16}$$

and γ_l is the sum of all loss processes independent of rubidium pressure (including background gas collisions, and any hypothetical wall losses). Note that the factor α_l concerning the capture rate is not the same as the MOT damping constant α. The results of α_l, β_l and γ_l are displayed in figure 4.22 as a function of distance from the apex.[3].

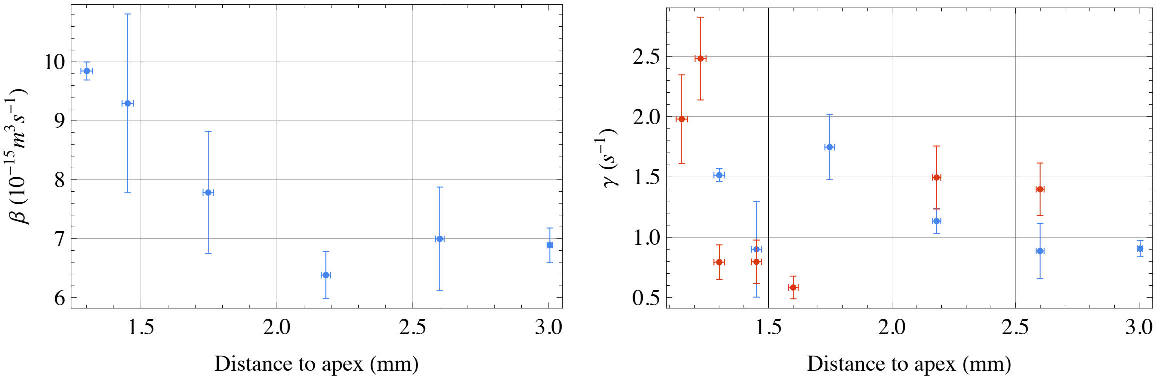

(a) Rb-Rb loss coefficient (β_l)

(b) Other loss processes (γ_l)

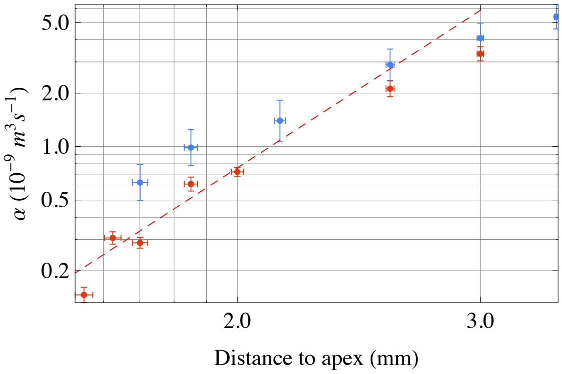

(c) Capture coefficient (α_l) and L^5 fit (dashed red line)

Figure 4.22: The variation of α_l, β_l and γ_l, as the MOT is moved closer to the pyramid apex. Blue points were determined with MOT loading times, and red points were determined from steady state number of atoms

Background rubidium loss (β_l) - We see that the rate at which trapped atoms in the MOT are lost from collisions with hot untrapped rubidium atoms *increases* as the MOT gets closer to the apex (see figure 4.22(a)). The cross section for Rb-Rb collisions, assuming resonant dipole-dipole interactions [63], can be approximated by the following relation

$$\sigma_{Rb} = \pi \left(\frac{4C_3}{m v_{esc} \bar{v}_{Rb}} \right)^{2/3} . \tag{4.17}$$

---

[3] As we are not aperturing the light I have used this distance measurement instead of 'effective pyramid size' which doesn't really mean anything in this context.

where $C_3 = 5.8 \times 10^{-48}$ Jm^3 for rubidium and v_{esc} is the escape velocity from the trap. We therefore surmise that the increased loss rate is due to a decrease in escape velocity as the cold atoms move closer to the walls.

Other loss processes (γ_l) - From figure 4.22(b) it is apparent that there is a small trend to higher γ_l at smaller distances. This can be expected as at a distance of 1.25mm to the apex, the nearest wall is 0.7mm away indicating the walls are beginning to impinge on the extremities of the cold atom cloud($\sim$ 0.2mm $1/e$ radius). The value of γ_l is however highly dependent on the condition of the vacuum in the chamber when the measurements are taken and we expect an amount of variation due to outgassing of components. Over the entire range of motion the increase in γ_l is only a factor of 2, which does not explain the several orders of magnitude drop in the atom number observed at this position.

Capture coefficient (α_l) - By contrast the capture coefficient, α_l, shows a strong drop (see figure 4.22(c)) as the MOT approaches the apex indicating that the drop in atom number is almost solely a consequence of a reduced capture rate. By fitting a power law to the observed data (the data set in red was used as it has more points at smaller distances) a scaling exponent of 5.0 ± 1.0 with distance from the apex was measured, consistent with the L^6 scaling expected.

4.8.2 Translation in horizontal plane

Further evidence for the insignificance of the effect of the walls on trapped atoms was obtained by looking at the behaviour of the MOT when it was translated. We moved the MOT in the xy plane at a constant z height above the apex, measuring the loading and loss rates in a similar manner to that of the previous experiment. As can be seen from the results in figure 4.23, the loading rate remains constant over the range explored, indicating that the capture process is independent of the transverse position of the MOT. At distances comparable to the spatial extent of the MOT we observe an approximately exponential increase in the loss rate, increasing up by one order of magnitude over distances as small as 200μm. It was difficult to extrapolate the precise distance of the MOT from the surfaces, the displacement was inferred from the magnitude of the applied shim and hence a further experiment was carried out in a situation where the MOT displacement could be measured directly.

4.8.3 An improved measurement of the loss rate from a MOT as a function of distance from the surface

A natural candidate for an improved measurement is the reflection MOT (covered in 1.5.1) where the 'surface' we consider is the mirror making up the MOT. Using the gold coated surface of the prototype chip as our mirror, a pair of coils was added aligned on an axis at 45° to the chip face, capable of creating a quadrupole field (typical gradient of 5G/cm/A at the surface). These coils are displayed in figure 4.24 along with the image of some trapped atoms in the resulting reflection MOT above the chip.

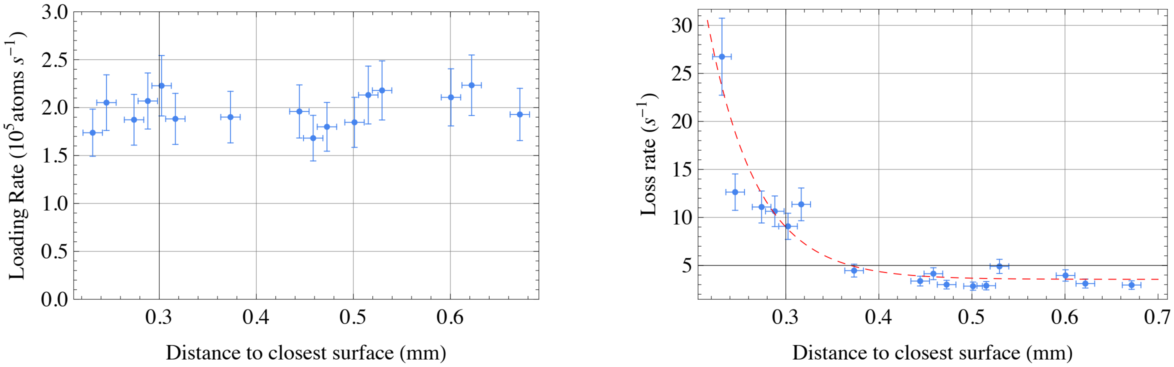

Figure 4.23: Left: Capture rate as a function of distance from the pyramid surface when moving in the radial plane. Right: Loss rate as a function of distance from a pyramid surface when moving radially, dashed line is a guide to the eye.

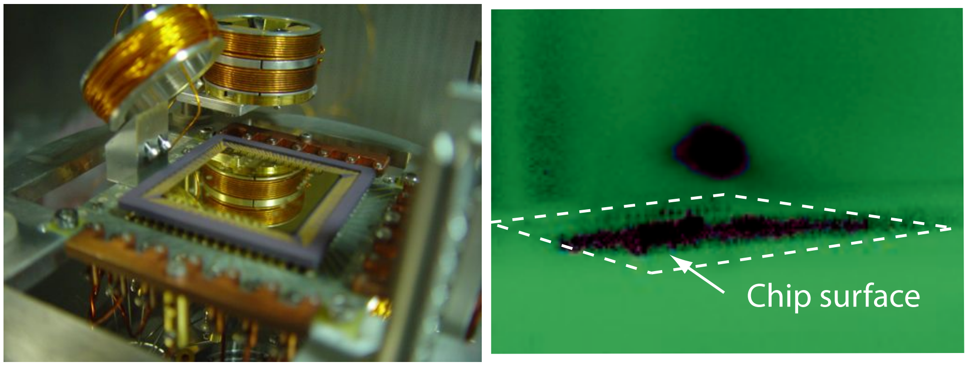

Figure 4.24: Left:MOT coils mounted around prototype micropyramid chip. Right: Image of cold trapped atoms in a reflection MOT above chip surface

A proposed theory for the loss rate of atoms from a temperature limited MOT to a room temperature surface was derived in the thesis of Ramirez-Martinez [79] by assuming that the loss rate was proportional to the fraction of the MOT density that overlaps the surface. This gives the rate of atom loss as a function of distance to the surface, d, as

$$\Gamma_{walls}(d) = \frac{\overline{v}e^{-\frac{1}{2}\left(\frac{d}{\sigma}\right)^2}}{\sqrt{8\pi}\sigma\left(1+\text{Erf}\left(\frac{d}{\sqrt{2}\sigma}\right)\right)} \tag{4.18}$$

where $\overline{v}$ is the mean velocity of atoms in the MOT and σ is the $1/e$ cloud radius.

We were able to test this using our reflection MOT. The loss rate was measured by determining the time constant of the MOT loading (equation (2.36)), and this was done for a variety of shim fields. Simultaneously the MOT was imaged from two perpendicular directions and the position relative to the surface determined. Results are shown in figure 4.25.

The magnitude of the loss rate shown in the data is higher than that of the data in figure 4.23 for comparable distances. This is due to the fact that in the reflection MOT, the atom cloud was much larger as very low field gradients were used to ensure the MOT density was low enough to reach the temperature limited regime.

Equation (4.18) was fitted to the data in figure 4.25, and gave a MOT radius of 0.31 ± 0.1mm and a mean velocity of 17 ± 2cms^{-1} which is consistent with a temperature of $100 \pm 25\mu$K.

Using the CCD images it was able to determine that the radius of the MOT is $316 \pm 12\mu$m which is in excellent agreement with the predicted value. The Doppler temperature of rubidium is known to be 140μK, and hence the predicted values of 100$\pm$25μK suggests slight sub-Doppler temperatures, which is consistent with a typical rubidium MOT. As a result it appears that equation (4.18) very satisfactorily describes the observed behaviour. Conversely this suggests that with suitable measurements, equation (4.18) can be used to determine the temperature of trapped atoms.

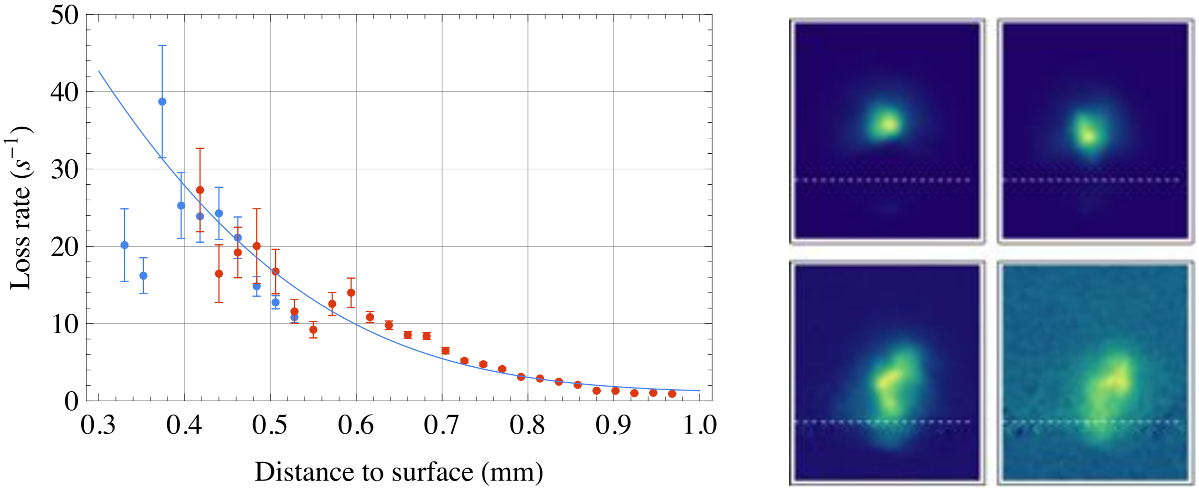

Figure 4.25: MOT surface experiments. Left: Loss rate as a function of distance to the surface. Red and blue points correspond to two different data sets. Right: images of MOT as it is moved towards the surface (indicated with the white dotted line)

4.9 Implications for the Prototype Chip

The observations made above confirm that the scaling of atom number with size results from the reduced capture rate alone. Loss processes only contribute significantly when a pyramid wall is closer than $\sim 3\sigma$. This adds further weight to the probability of the prediction that the number of atoms in the silicon 1.2mm micropyramid will be determined by the L^6 scaling and hence trap too few atoms to be detectable (on the basis of image 4.21). Nevertheless it is possible that loading the 1.2mm pyramids with a source of precooled atoms could result in an observable MOT, and this can be potentially accomplished with the reflection MOT introduced in section 4.8.3.

The reflection MOT was created above a region of the chip near the micropyramids. It was loaded far away ($>$ 1mm) from the surface away so that the trapped atom number could be maximised. When fully loaded, the fields were changed to move the MOT to a region directly beside the micropyramids (figure 4.26). At this point the reflection MOT light was swept to far red detuning to cool the atoms further, before the light was switched off and the light illuminating the chip perpendicularly (the light to create the micropyramid MOT) was switched on at optimal MOT detuning. It was observed that cold atoms were retained, and that they were swept into the chip, but none were observed trapped in the micropyramids.

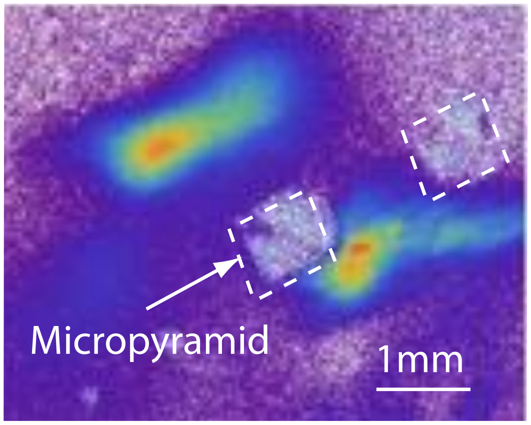

Figure 4.26: False color image of the MOT (And reflection) next to a 1.2mm micropyramid.

4.10 Conclusions

The primary conclusion drawn from this chapter is that the 70° geometry is capable of forming stable MOTs, and that the scaling of the number of atoms trapped can be approximated by a power law L^6 where L is the side length of the pyramid. This scaling is not unique to the 70° geometry and has also been observed in 90° cones (see appendix C). We can detect numbers of atoms down to MOTs of a few hundred, and with this technique we have been able to exclude MOTs of this number from a microfabricated 1.2mm silicon pyramid. To this end, silicon pyramid MOTs of above 1.2mm are likely candidates to demonstrate successful trapping with 4mm pyramids almost certain to work. It was hence decided to revisit the fabrication of these devices to see if several mm scale pyramids could be created.

4.10.1 Roughness

The performance of these devices will depend on the surface roughness, with low scatter levels desirable for imaging and these constraints should be considered during fabrication. We can use the study in section 4.5.2 to estimate what limits this places on the roughness of any microfabricated pyramid. For the threshold of detecting 1000 atoms in a 100μm MOT, we require less than 43pW of scattered light incident per pixel. From an estimate of the total scattered light, the Total Integrated Scatter (TIS) [84] relation can be used to estimate, assuming a random noise spectrum, a threshold roughness for any mirror,

$$TIS = \frac{P_s}{RP_i} = \left(\frac{4\pi\sigma\cos\theta_i}{\lambda}\right)^2 \tag{4.19}$$

where P_s and P_i are the scattered and incident powers respectively, σ is the rms roughness,R is the reflectivity and θ_i is the angle of incidence. For the requirement that $P_s < 10^{-14}$W, and assuming $P_i = 2.5$nW (intensity of 5mW/cm^2 over the area of a single pixel) this gives the constraint that the rms roughness has to be less than approximately 0.1μm in magnitude. Thus as well as adapting the fabrication techniques to produce larger devices, we must ensure that they do not become appreciably rough to inhibit the detection of trapped atoms.

Chapter 5

Microfabrication of a silicon pyramid MOT

The processes by which sub-mm scale pyramids could be produced were well known [81], however proceeding beyond this to the larger scale of several mm has required the development of new processes. This chapter is a brief field guide to some of the processes attempted (successful and otherwise). I highlight the successful processes but I also mention unsuccessful processes that might possibly yield successful results, but for reasons of time or unavailability of facilities I was prevented from following.

5.1 Building with Silicon

Silicon is the second most abundant element in the Earth's crust, and was identified by Antoine Lavoisier in 1787 in flint (named *silex* in Latin) samples. Henri Étienne Sainte-Claire Deville was the first to prepare silicon in its crystalline form in 1854 [1], and purification techniques to allow its use in a semiconductor device began in 1919 for components used in radar. There followed widespread adoption of silicon in integrated circuits and the use of silicon in microfabrication (for example MEMS) began to grow in the mid 1990s.

Silicon Structure Silicon in its crystalline form makes four covalent bonds, leading to tetrahedral structure, the same as diamond. The diamond-cubic structure can be represented as two interpenetrating face-centred cubic lattices, displaced by $(\frac{1}{4},\frac{1}{4},\frac{1}{4})$ of a basis vector relative to each other. If a slice is taken through the {111} planes, it can be seen that the packing density of atoms is highest through these planes, and three of the bonds to each atom lie below the plane.

---

[1] Interestingly he was also responsible for the development of the industrial production of aluminium, the metallic coating of choice for the pyramid mirrors in this experiment

5.2 Anisotropic Etching in Silicon

5.2.1 Anisotropy

Since different crystal planes in silicon have different chemical reactivities some etchants etch faster in one direction than another. Alkali hydroxides (for example KOH, NaOH etc.) of sufficient pH (typically ≥ 12) or organic aqueous solutions such as ethlyenedineamine, hydrazine or choline [85] exhibit an etch rate significantly higher for the (100) and (110) directions than for the (111) direction. For potassium hydroxide (KOH) etches the (100) plane 400 times faster than the (111) plane, while the (110) plane is etched 600 times faster. These rates are highly dependent on the composition of the etchant and factors such as temperature and concentration but in practice the anisotropy is so large that it is easy to reveal the {111} planes[86].

5.2.2 [100] orientated wafers

Silicon wafers can be provided in any orientation desired, but the three most common (and useful) orientations are [100], [111] and [110]. The [100] orientated wafers are often cleaved with a straight edge on the bottom to indicate the [110] direction (manufacturers typically quote a tolerance of ±1° for this primary orientation flat). The {111} planes make an angle of $\cos^{-1}(\frac{1}{\sqrt{3}}) = 54.74°$ with the wafer surface and as can be seen in figure 5.1, the intersections of {111} planes are mutually perpendicular along the $< 110 >$ orientations. If the surface of the [100] wafer is masked so etching occurs in an exposed rectangular region, with edges parallel to the $< 110 >$ orientations, the {111} planes will be revealed to form the beginnings of a pyramid. As the etch proceeds, it will etch downwards, but not outwards and will eventually reach the point of intersection of the {111} planes, whereupon the etching process will effectively stop as there is no further exposed (100) plane to etch. These pyramidal shaped pits can be extremely smooth, as fundamentally this etching is a layer by layer atomic process[87].

5.2.3 Mechanism for anisotropic etching

The processes governing anisotropic etching in KOH begin with the reaction of surface silicon atoms with hydroxyl ions [86]

$$\mathrm{Si{+}2OH^-} \rightarrow \mathrm{Si(OH)_2^{2+}} + 4e^-.$$

The four electrons liberated by each silicon atom are injected into the conduction band, allowing for the reduction of water

$$\mathrm{4H_2O}{+}4e^- \rightarrow \mathrm{4OH^-} + \mathrm{2H_2}.$$

The complexed silicon reacts to form a soluble silicon complex with the hydroxyl ions

$$\mathrm{Si(OH)_2^{2+}{+}4OH^-} \rightarrow \mathrm{SiO_2(OH)_2^{2-}} + \mathrm{2H_2O}$$

The (100) and (110) planes present two dangling bonds at the surface, whereas in the (111) plane only one dangling bond exists (and so only one hydroxide can bond to a surface atom).

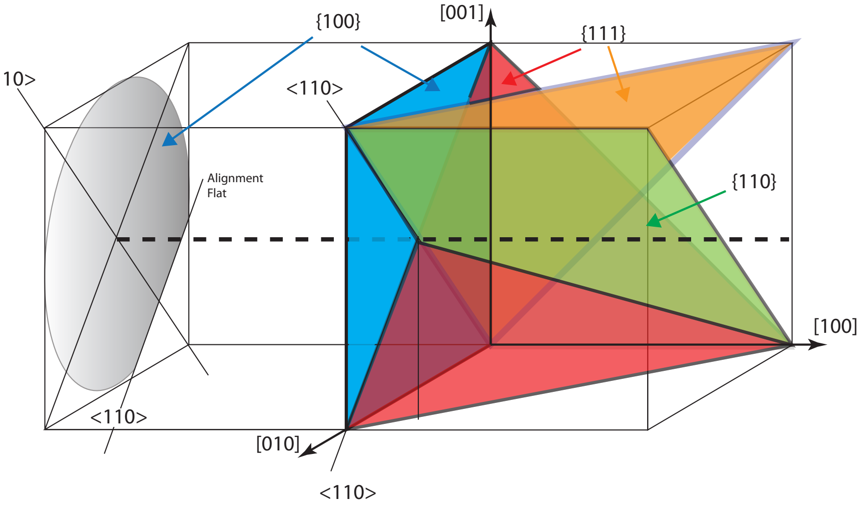

Figure 5.1: Relevant planes in silicon on the unity cube. Adapted from [88]

Simultaneously atoms on the surface of the (111) plane have three bonds with the main crystal surface which need to be broken to liberate the silicon atom, as opposed to only two bonds for the (100) and (110) planes. It is thought that these differences lead to the anisotropy of the etching, though the details remain unclear [86].

5.3 Producing the pyramids

5.3.1 Choice of wafers

Our wafers were sourced from Silicon Valley Microelectronics in California, and they were able to supply a number of $3000 \pm 50\mu$m thick wafers, polished on one side and cut on the $< 100 >$ plane. The wafers were CZ grown (i.e. formed using the Czochralski process), with p-type boron doping to a $1 - 20\,\Omega$cm level.

5.3.2 Choice of reactants

Several factors dictate the choice of reaction conditions in this situation. Having to etch to depths of up to 3000μm dictates that conditions should maximise the etch rate, but not at the expense of producing poor quality surfaces. Figure 5.2 shows the dependence of rate on the two most accessible parameters, temperature and concentration taken from [85]. It is known for the 100 and 110) planes, that roughness reduces as concentration increases[89] [88] and hence a trade off needs to be made between rate and roughness.

A KOH solution of 40% by weight, and a temperature of 80° was chosen to give an etch rate that would allow for a 3mm depth to be etched in ~ 50 hours (i.e. one working week), and this was used for our initial etch.

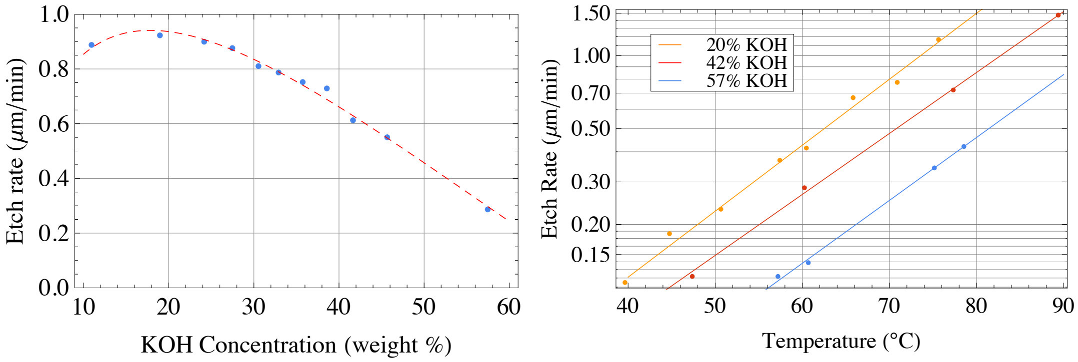

Figure 5.2: Left: Etch rates as a function of concentration at 72°C Right: Etch rates of (100) plane as a function of temperature. Reproduced from [85]

5.3.3 Masking

In order to define a square opening on the (100) plane we use a thin film masking layer. A number of materials are suitable, but low pressure chemical vapor deposited (LPCVD) silicon nitride is chosen because of its low residual stress making it suitable for relatively thick layers [88]. This has an extremely low etch rate of less than 0.1nm/min in KOH, and hence based on a predicted etch rate of the (100) planes of $\sim 1\mu$m/min, a nitride layer 350nm thick is deposited on both sides of the wafer. Photo-lithography is used with reactive ion etching (RIE) to open square apertures in one of the layers, thereby exposing the bare silicon (100) plane underneath.

The design for the mask, with the dimensions of each aperture is indicated in figure 5.3. As etches were done with a 3mm thick wafer, the largest pyramids that could be made had a side length of $3\sqrt{2} = 4.2$mm. Larger apertures (of side length l) form a truncated pyramid with a square hole at the base size $l - 3\sqrt{2}$, and hence 8.34mm apertures were used for the potential of making a large ‘double-stack’ pyramid when combined with the 4.1mm pyramids. The intermediate feature sizes were chosen to ensure a logarithmic range of sizes between the 1.2mm and the 4.2mm were fully explored. The mask was made professionally in glass and chrome, with the black features being produced as an opening in the chrome (positive mask). Minimum feature size possible with this mask was 8 microns, this was more than adequate for our needs.

5.3.4 Initial etches - Wafer I

The wafer, following mask exposure and development, was then ready for etching. To prevent etching of the exposed silicon at the edges of the wafer, PTFE tape was wrapped around the circumference of the wafer. The wafer was placed flat in a large beaker, filled with a solution of 800ml KOH (Sigma-Aldrich sol^n 45% wt) and 150ml of de-ionised water. The etching beaker itself was placed in a temperature stabilised water bath at 80°C. On introducing the wafer to the etchant, the reaction was observed as a vigorous effervescence (most likely hydrogen) arising from the exposed silicon areas. Early in the etch it became apparent that where the dividing lines on the mask met in a corner (for illustration, some are highlighted in green in figure 5.3), unwanted etching was taking place. It was expected that the long alignment lines

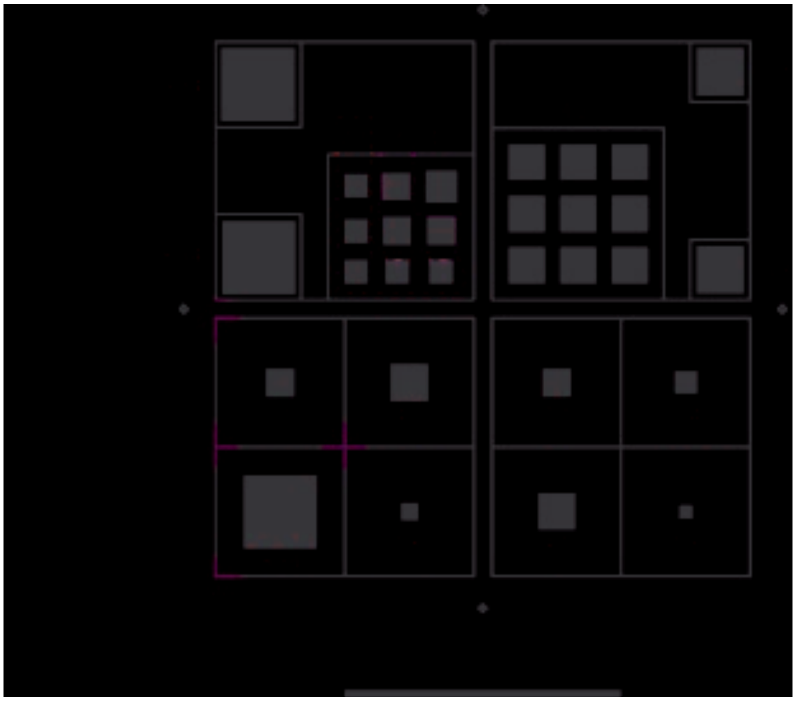

Figure 5.3: Design for the mask. Black features create openings in the nitride layer. The flat line at the bottom is for alignment with the wafer alignment flat. Annotations in red/green are not included on the mask, and all dimensions are in mm. Regions where two lines met in a corner (some are highlighted in green) led to underetching in the first few etches and these corner points were removed in future etches.

would etch into a thin v-groove, but it was apparent that other planes were exposed at the intersections (shown as the green highlights in figure 5.3) and underetching of the mask was taking place. Figure 5.4 shows a computational model of what was happening.

Figure 5.4: Simulation of observed underetching effects - along with the pyramid hollow etching in the centre of the block, we observe underetching effects from the corners of the alignment marks.

As a result a number of the pyramids were lost from the wafer as it was impossible to retrospectively mask off the underetching areas. A similar problem occurred on the back face, where chuck and tweezer damage created small perforations in the rear nitride face causing etching of the back which evolved into further significant damage to a number of pyramids.

The resultant wafer can be seen in the photographs in figure 5.5, we see that despite the unusual topology, there are well defined pyramid structures. Etching was completed in a total of 50 hours, giving a (100) etch rate of 1μm/min which is slightly higher than the times indicated in figure 5.2.

Visual observation of the pyramid faces shows that the pyramid surfaces are far from

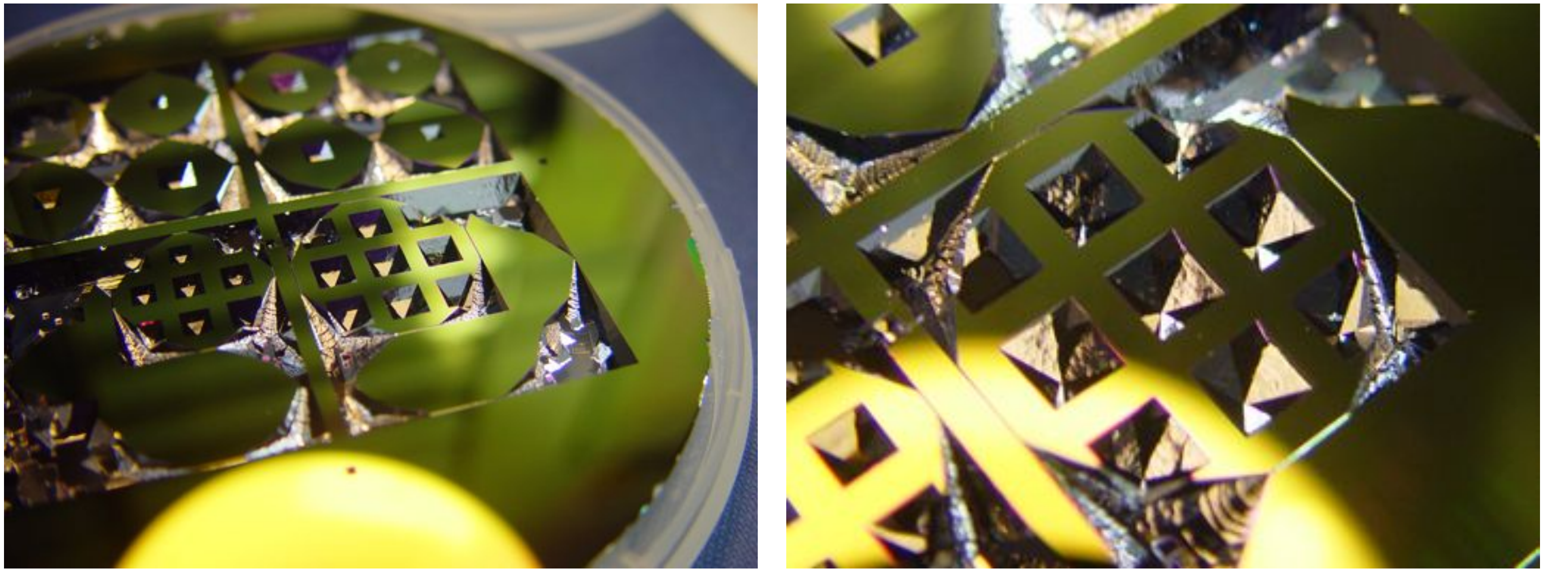

(a) Photograph of entire wafer

(b) Close up on one corner of wafer showing {111} patterning

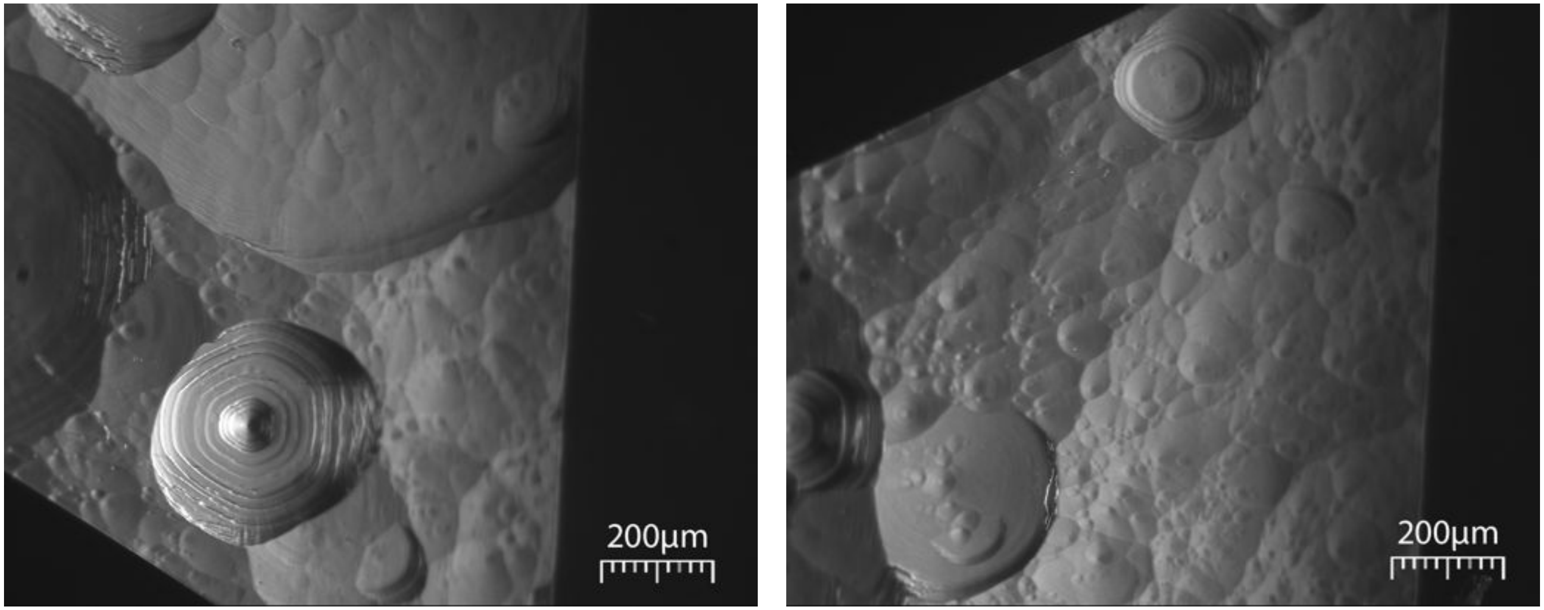

(c) Microscope image showing 111 plane (lower half of face)

(d) Microscope image showing 111 plane

Figure 5.5: Wafer I - this wafer was anisotropically etched in a temperature stabilised KOH solution without any form of agitation.

smooth, with a distinct texture, visible to the naked eye covering the surfaces. On inspection under an optical microscope (lower half figure 5.5) it is clear that the structure is composed of a number of crater like structures[2]. Some of these have a stepped profile, as in figure 5.5(c). Despite the appearance of these features, it is apparent that locally (for example the crater in the bottom left corner of figure 5.5(d)) the surface is smooth suggesting that on a small scale the etching process is uniform.

When viewed under an optical microscope, the edges of the pyramid faces are extremely sharp and in every instance no defect was visible at the apex (down to the resolution of the microscope $\approx 1\mu$m). If a partially complete pyramid is observed looking down towards the base (figure 5.6), the remaining (100) unetched plane appears as a near perfect square.

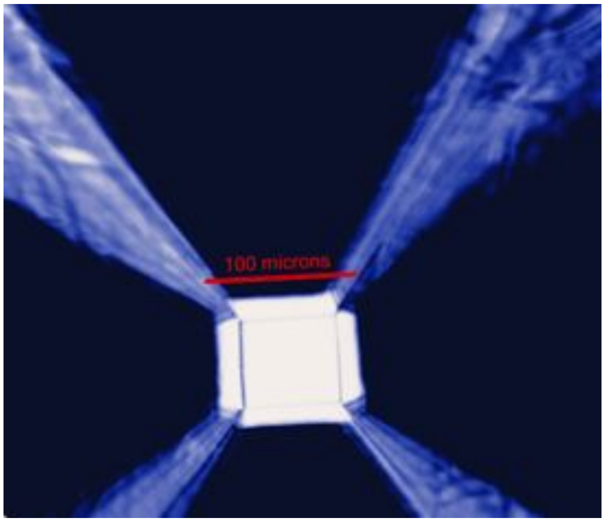

Figure 5.6: Image of the unfinished apex of a pyramid showing the near perfect square of unetched (100) plane

Potential roughness mechanisms

As covered previously it was observed during the etching process that a large number of gas bubbles (most likely hydrogen) were released in the process. These bubbles had a tendency to adhere to the surface, and it is probable that these could be the main cause of the crater-like structures. Indeed the process has been observed in-situ[90], whereupon the hydrogen bubble causes a local variation in the etch rate, either suppressing etching by preventing the flow of fresh solution to the face or by accelerating the reaction in a process known as micromasking. The stepped craters have been attributed to lattice stress created from oxygen defects present in the silicon wafer [91].

5.3.5 Pyramid Angle

The relevant optical properties of the pyramid structure can be measured by illuminating them with light parallel to the pyramid axis, as they are illuminated when used as a MOT.

---

[2] It is not clear from the images whether the artefacts are indentations (craters) or protrusions (hillocks), but using profilometry these were confirmed to be craters

By intercepting the reflected rays exiting the pyramid, the properties of the mirrors can be inferred. A screen with a small hole in it is placed between the incoming beam and a silicon pyramid to image the reflection pattern on the screen, and a typical image is shown in figure 5.7. The pattern formed is essentially identical to that anticipated, as seen in section 4.1. The angle of propagation of the Type I rays, θ_{T1}, is found by measuring the separation between the Type I reflections as a function of distance. From this we infer the pyramid apex angle α using $\theta_{T1} = \pi - 2\alpha$. This gives the silicon pyramid angle as $70.9 \pm 0.6°$ which is consistent with the theoretical value of $\arccos(1/3) \approx 70.5°$.

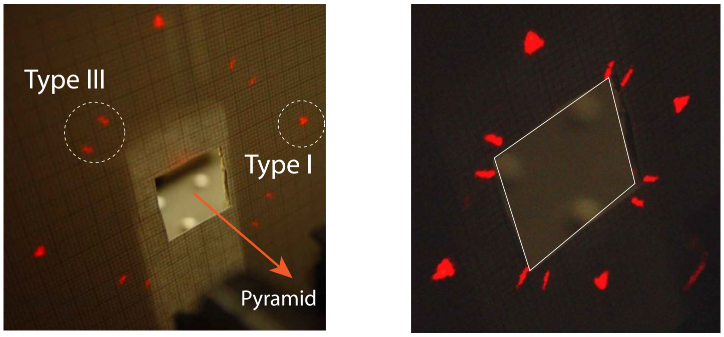

Figure 5.7: Left:Image of far-field reflection pattern on a screen 30mm away from pyramid exit. Right:Same image with background lights turned off

5.3.6 Refinement of etching process - Wafers II and III

As the roughness was primarily attributed to the adhesion of hydrogen bubbles a number of methods to reduce the effect of the bubbles were tried. Addition of a surfactant, in this case isopropanol (IPA), was tested. It is suggested the inclusion of this additive reduces the adhesion of the bubbles and in shorter etches it has been shown to reduce roughness values significantly[88]. However, we found the addition of isopropanol had no obvious effect on the roughness and left a milky white residue on the wafer surface after a long etch session in excess of 10 hours.

We were able to use a combined hot plate and magnetic stirrer for another etch to produce Wafer II. The agitation of the solution by the stirrer was sufficient to set up currents which dislodged the bubbles and resulted in a noticeably better etch to the naked eye as displayed in the upper left of figure 5.8. This observation is backed up by optical microscope images (lower half figure 5.8) which show a much improved surface finish. Cratering is still a problem but the effect is much less marked than before.

For our most recent etch, Wafer III, a number of other changes were made to improve the process further. Most significantly, the mask was modified so that the division lines did not underetch at the corners. The wafer was also protected by a PTFE sheet, which was attached to the backside using PTFE tape to prevent etching of the back. Agitation was provided by bubbling air through the solution. The wafer was laid flat, etch face up on a stainless steel pipe bent into an ‘S’ shape, with a series of holes along its length. Air was bubbled through

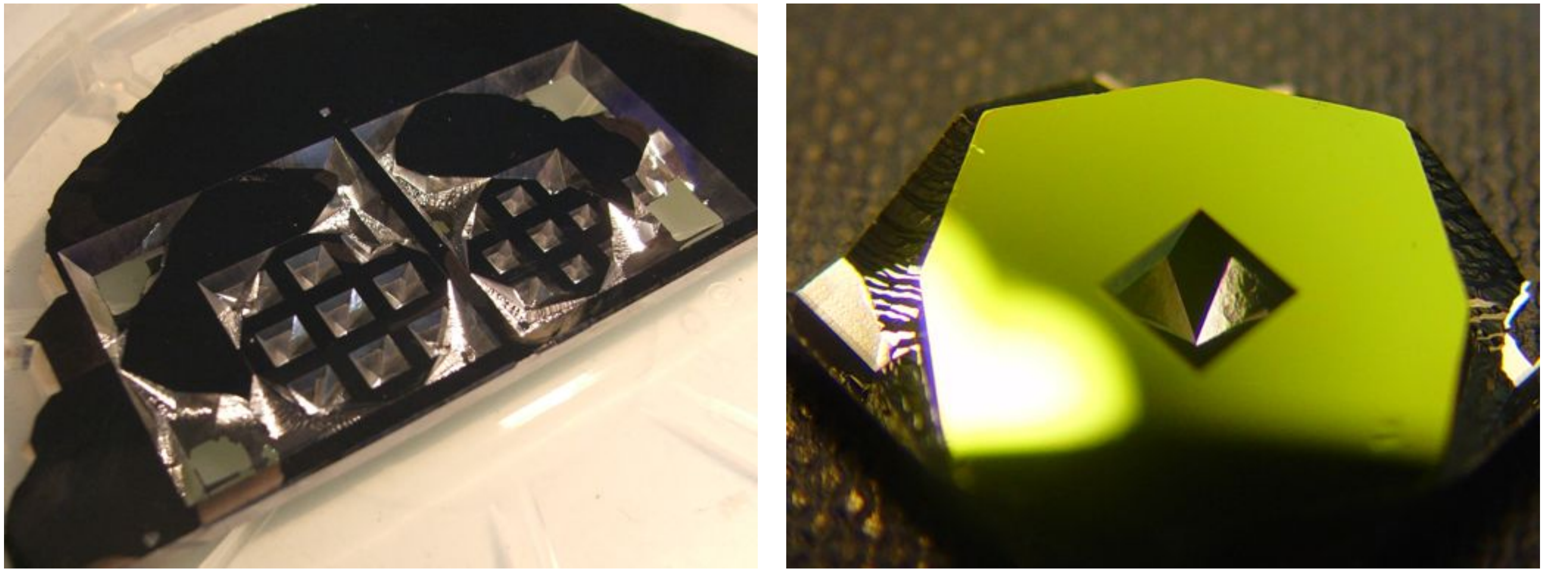

(a) Photograph of half of wafer

(b) Close up on one corner of wafer showing {111} patterning

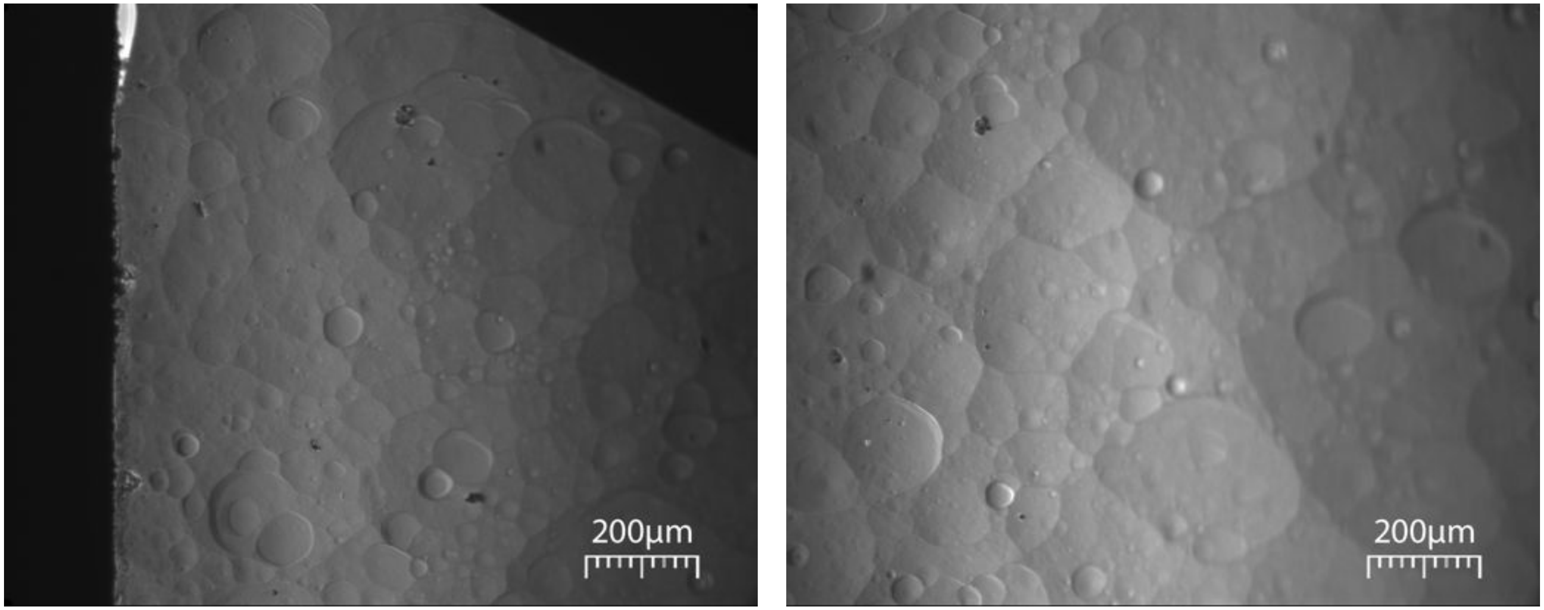

(c) Microscope image showing 111 plane

(d) Microscope image showing 111 plane (lower half of face)

Figure 5.8: Wafer II - This wafer was etched without temperature stabilisation with the KOH solution continuously agitated with a magnetic stirrer

the solution at a rate of $1l\,min^{-1}$. It was hypothesised that as well as agitating the wafer, the oxygen in the air would help to neutralise the excess hydrogen produced.

To the naked eye (see figure 5.9) the wafer appears excellent, with seemingly perfect pyramids and no apparent roughness. The majority of the largest truncated pyramid features had survived with back-etching damaging a few. The etch was terminated slightly prematurely to leave a thin layer at the base of these pyramids as we were not confident that the aperture at the pyramid base would be successfully opened without damage if it was allowed to etch all the way through. The PTFE plate had reduced the amount of damage to the back plate, although some solution did leak into gap between the PTFE and the back face causing a small amount of back-etching.

5.4 Smoothing Techniques

Despite the improvements in etching, the pyramid surfaces were still noticeably rough especially when viewed under the microscope. In order to improve the finish of the pyramid faces a number of techniques were tried. These fell into two main categories, either covering the roughness, or etching the faces isotropically to a depth where the roughness was reduced.

5.4.1 Covering Techniques

Oxidation & Reflow

Silicon naturally oxidises at room temperature forming a layer approximately 2nm thick, but by heating the wafer to high temperatures (between 600 and 1250°C) in a stream of oxygen (or water vapour) the diffusion of oxidant through this initial thin oxide layer can be accelerated to form thick oxide layers. As the oxide layer grows, the rate of diffusion falls and hence after a time the process becomes impossibly slow. Typical thermal oxides are limited at 1-2μm in thickness.

An initial oxidation was made for 48 hours in a dry furnace at 1150°. The resulting surfaces showed interference effects demonstrating that a thin transparent film was present on the surface, and this was coated with a thin aluminium film to ensure the top layer of oxide was visible. Following the coating it was apparent there was little initial reduction in the roughness. Another set of samples were oxidised in a wet furnace at 1000° for a period of 28 days in order to obtain a several micron thick layer. Following this there was still no significant roughness reduction evident so we abandoned this process. There are two possible reasons that this process did not work, firstly not enough oxide was deposited to mask the features underneath, and secondly perhaps the oxide did not reflow sufficiently. This could be potentially improved by doping the glass layer to increase its mobility at 1000°.

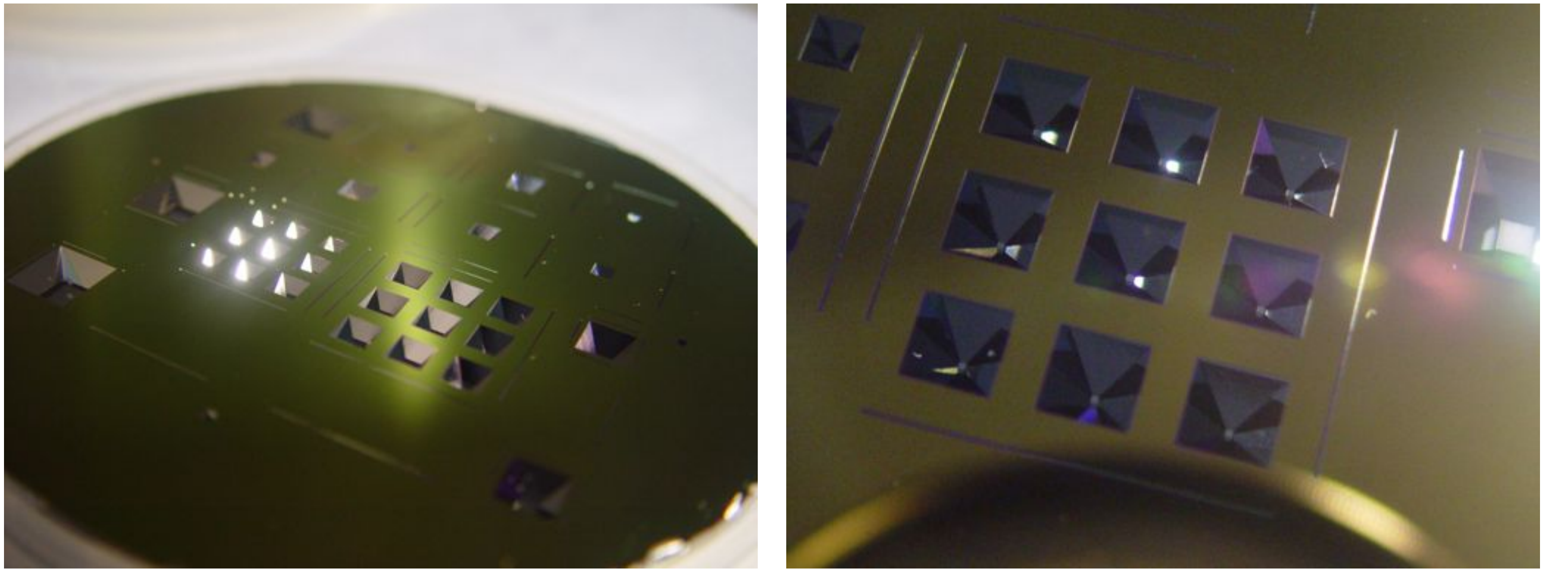

(a) Photograph of entire wafer

(b) Close up on pyramid array (etching not quite complete).

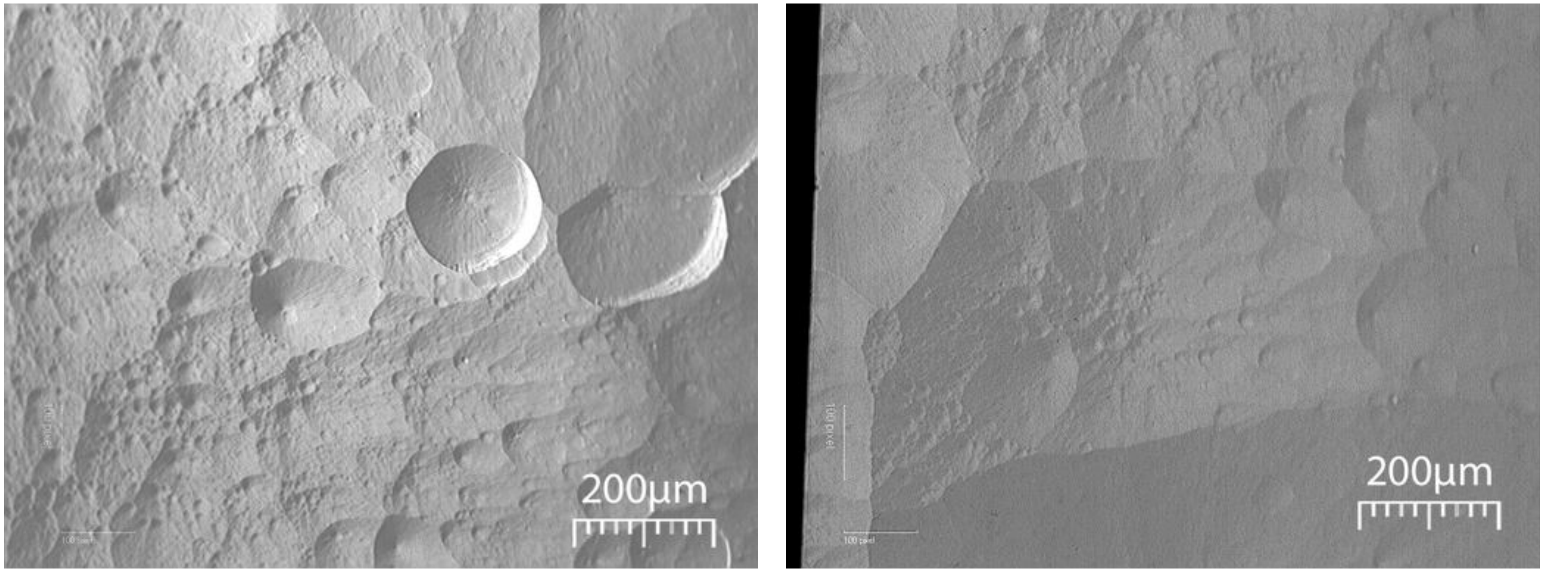

(c) Microscope image showing 111 plane

(d) Microscope image showing 111 plane (lower half of face)

Figure 5.9: Wafer III - This wafer was etched with both temperature stabilisation and agitation provided by bubbling air through the solution.

Monoxide deposition

Rather than oxidising the existing silicon, a film of silicon monoxide can be deposited in evaporation. Silicon monoxide begins to sublime in vacuum at around 1100° and hence can be deposited by the resistively heated evaporation boats in our in-house coating plant.

A layer of approximately 300nm was deposited in 45 minutes. A layer of 20nm of aluminum was sputtered on top to make the top surface visible. Under the microscope (see figure 5.10(a)), comparing the same regions before and after coating there appeared to be little change in the features. All that can be taken from the images is that there appears to be some reduction in the contrast of the features, and some smaller features have disappeared. The sample was put back in for a further 300nm of monoxide and another layer of aluminium but disappointingly there is no significant reduction apparent in the surface patterning.

Finally, in order to attempt to deposit a significantly thicker layer (as a large amount of silicon monoxide was being used to deposit only 300nm), the sample was moved from its original position (approx 60cm from the evaporation boat) to 10cm away. Unfortunately having the sample so close to the source led to an uncontrolled and highly erratic coating. The pyramid faces appear crazed and discoloured. Despite this, areas of even coating give the impression that in some places the thick coating had successfully masked the disorder underneath suggesting that if a way could be found to deposit the same thickness in a controlled manner, without excessive use of raw materials, that the surfaces could perhaps be smoothed in this way.

Resist Deposition

Photoresist is the light sensitive material used to transfer a pattern to a substrate for masking, for example in photolithography. The vast majority of photoresists are organic polymers in solution, but there are some epoxy based photoresists such as SU-8 which has been used in atom chips as a structural component[36]. Crucially relatively thick (several microns) layers can be deposited easily and also ‘reflowed’ when the temperature of the photoresist is raised above its glass temperature allowing for redistribution under surface tension. This seemed a good prospect for smoothing the surface, so several methods of depositing the resist were tried.

Spin Coating Spin coating is a process used to apply uniform thin films of photoresist to wafers. A surplus of photoresist solution is placed in the centre of the wafer and then the wafer is spun at high speed in order to create a uniform layer by centrifugal forces. This was attempted with one of our completed wafers, using a photoresist SPR 220-7.0, and spin parameters (viscosity of solution and speed of rotation) chosen to give a layer approximately 9μm thick. Despite a disruption of the flow by the pyramid pits (causing pooling and striation), the regions that were covered uniformly had a good surface finish. The resist was baked at 150°C in order to reflow the resist and evaporate the solvent. Following this the baked photoresist was sputtered with aluminium[3] and the resultant faces imaged under a microscope

---

[3] A point worth mentioning here is that the baked SPR220 did not adversely affect the base pressure in the coating plant which reached its typical base pressure of 10^{-7}mbar with the resist coated sample in the chamber.

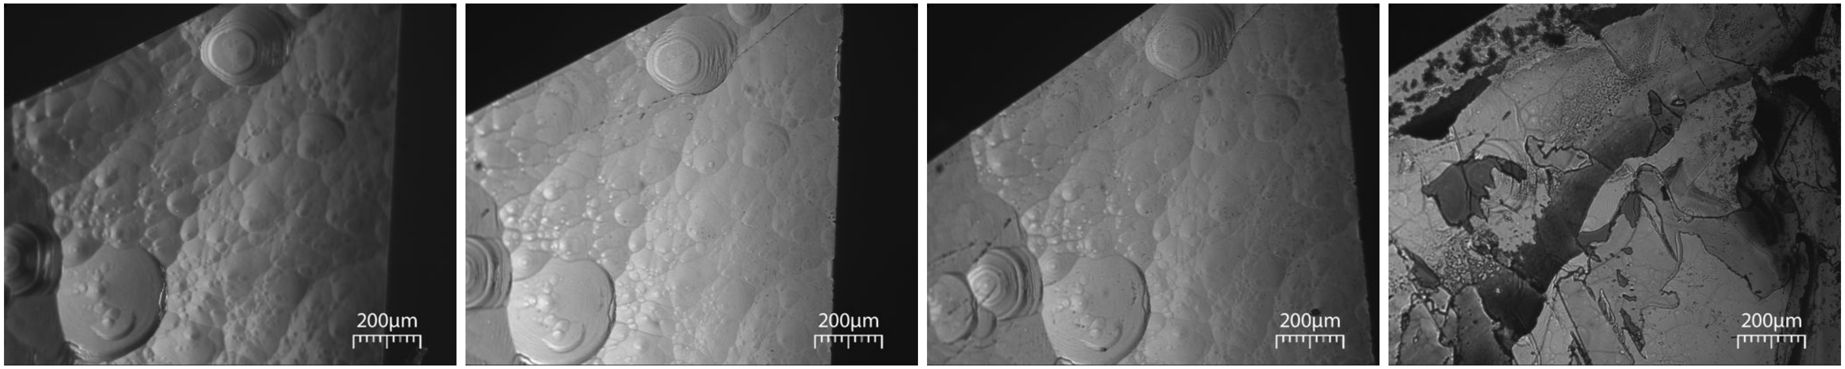

(a) Pyramid faces following (l to r) no coating, 300nm, 600nm and a ‘thick’ layer of silicon monoxide

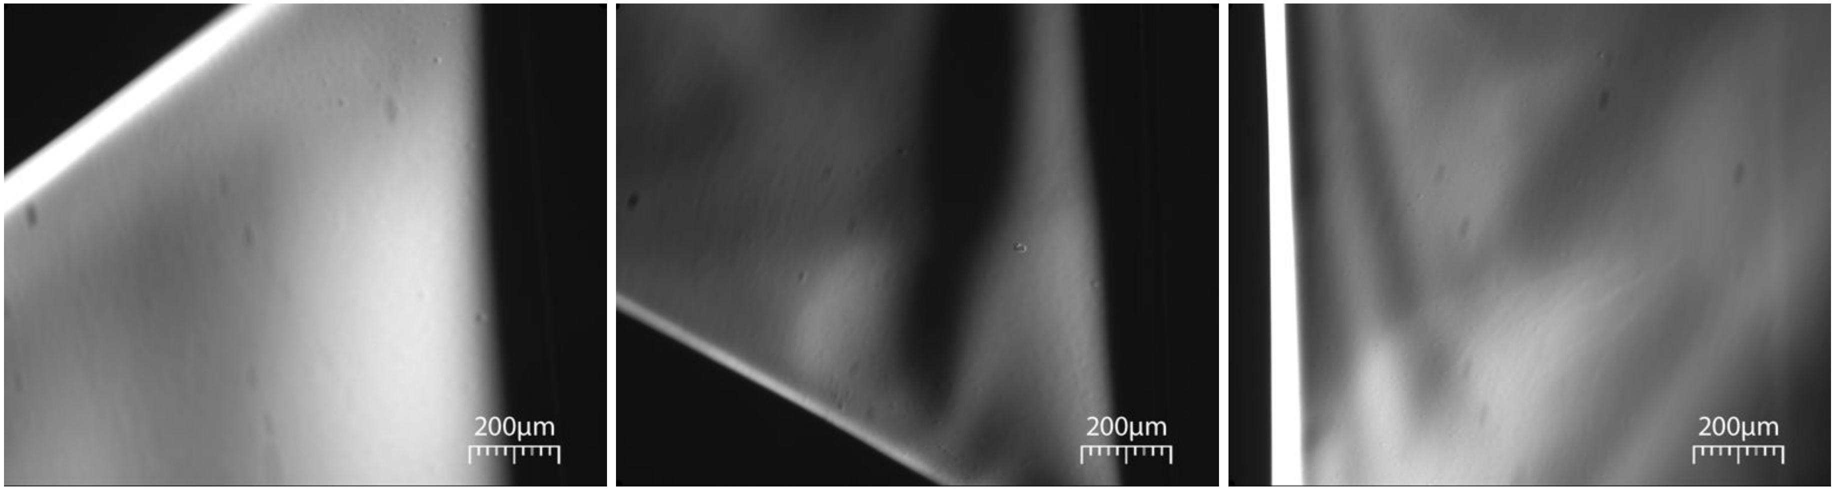

(b) Pyramid faces following spin coating with photoresist

Figure 5.10: Anisotropically etched pyramids coated with top:Spin Coated photoresist and bottom:Evaporated Silicon Monoxide

(figure 5.10(b). Although the surfaces were excellent in some regions, the nonuniformity of coating, coupled with the fact that surfaces were slightly curved suggests that this technique is not particularly suited to the pyramids.

Dip Coating It was suggested that repeated dipping in dilute photoresist, to build up thin layers may have a beneficial effect. We used SU8-2002, which is already a fairly nonviscous photoresist (7.5cSt compared to 1cSt for water), and diluted it further in trichloroethylene. The sample was repeatedly immersed in the solution, and then taken out to let solvent pooled in the pyramids pour away. It was then softbaked at 95° on a hotplate for 2 minutes. The cycle of dip, drain and bake was repeated 10 times. After this, it was looked at under the microscope, and we found that the resist was not forming a smooth layer. Instead it formed layers only in flat regions, with noticeable absences at edges and discontinuities. This attempt was rather crude, but has confirmed the initial thought that making a thin solution flow evenly across the pyramid boundaries is extremely difficult.

Spray Coating In spray coating, the photoresist is sprayed onto the surface from an ultrasonic spray nozzle generating droplets of micrometre size. The surface is scanned in multiple passes to buildup the total film thickness required. Control of the resist deposition is again controlled by the viscosity of the solution, determining the degree of atomization during spraying and the flow of the solution into deep features when deposited. This process is ideally suited to MEMs with the ever increasing use of nonplanar wafers and three dimensional microstructures. [92]

A partial array of pyramids was coated with three layers of resist, masking off a row each time to produce three different thicknesses. The sample was baked after each coating. This process resulted in total thicknesses of 12μm, 19μm and 26μm being deposited. From the photographs in figure 5.11 it is apparent that the results are somewhat unremarkable. The resist appears to have not fully covered the larger crater features on the surface for the thinnest layers, and poorly covers the lower part of the pyramids. The thicker layers appear to hide the underlying roughness well but has developed wrinkles, perhaps as a consequence of the baking which gives the surfaces a distinctly ‘plastic’ appearance.

It is highly likely that with refinement of the deposition, material composition and process, satisfactory coverage could be obtained, but without access to a spray coating machine it would be costly to perform a systematic and exhaustive study using commercial suppliers.

Untried Coating techniques

The oxidation of silicon did not yield successful results, but we propose that artificial deposition of a glass layer (up to several microns using PECVD) could prove beneficial. As the composition of the glass can be tailored with doping, layers of glass could be deposited that would reflow for effective smoothing . Another untried technique which could possibly yield results is the deposition of a thick metallic layer by electroplating (also acting as the mirror

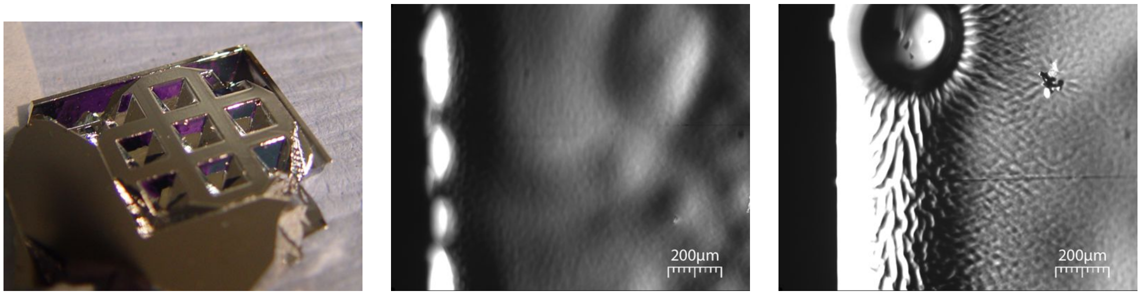

(a) Photograph of sample following spray coating (b) Microscope of sample with 12μm resist (c) Microscope of sample with 19μm resist

Figure 5.11: Spray Coating - Anisotropically etched pyramids are coated with a layer of photoresist deposited by spray coating

coating). Although electroplated surfaces do not normally have as high a quality as sputter deposited surfaces, a thick layer can be deposited faster with less wastage of raw materials.

5.4.2 Isotropic Etches

Wet isotropic etchants are frequently used for polishing silicon. As the name suggests, they etch uniformly in all directions. For silicon, they are typically acidic and are frequently faster than similar dry etching processes.

HNA

The most common wet silicon isotropic etchant is a mixture of $HF/HNO_3/CH_3COOH$ acids which is known as HNA. The etching process is a multi-step process, where the silicon is oxidised, with the nitric acid acting as the oxidant (other oxidants such as hydrogen peroxide or bromine can be used). The resulting silicon oxide products are then dissolved in the hydrofluoric acid. This cycle of oxidation and then stripping of the oxide in situ makes for a very effective etchant. The acetic acid is used as a dilutant instead of water as it is less polar and helps prevent the dissociation of the nitric acid, preserving the rate of the oxidation. [88]

HNA etching properties have been extensively studied and the rate and finish of the etch depend upon whether the reaction is limited by either:

- Reactant product transport rate to the surface
- Surface reaction rate
- Reaction product transport rate away from the surface

Etches are also highly sensitive to agitation and temperature and as a result small changes in conditions can lead to dramatic changes in the etching. At low HF and high HNO_3 concentrations, the reaction is limited by the rate at which HF can remove the oxide layer. This regime is characterised by highly isotropic and smooth etches producing bright finishes.

A solution of composition 8:66:26 HF(49.23% conc.):HNO_3(69.51% conc.):CH_3COOH(99.5% conc.) by volume was initially tried to find its effect on the pyramid faces. A large incomplete

8.34mm truncated pyramid was the test piece as its open structure allows good access to the faces for profiling.

Figures 5.12(a) and 5.12(c) show the results of a 30 minute etch (without agitation or temperature stabilisation). The faces are noticeably smoother, with cratering less visible, however there is a slight rounding on the corners and edges. Figures 5.12(b) and 5.12(d) show the result of a further 30 minute etch on the same sample. The surfaces (both from inspection by eye and under the microscope) have lost most of the features characteristic of the deep KOH etch, but instead possess a long range 'bowing' to the face suggesting a departure from the {111} planar structure.

Two further samples were etched (see figures 5.12(e) and 5.12(f)), this time fully formed pyramids, for durations of 60 and 90 minutes. The 60 minute sample looks very smooth and might be usable, but the degree of rounding of the 90 minute sample significantly changes the original pyramid geometry. The 90 minute sample was also observed when viewing to the apex under a microscope to have some strange planar irregularities, which may either be due to over rounding causing exposure of other planes or possibly due to 'pooling' of reaction byproducts for the extended period of the etch.

The most undesirable aspect of this etchant is the need to handle concentrated HF. Having stared death in the face on numerous occasions while performing the process the writer would like to recommend HF as a sobering reminder of ones own mortality.

ICP

Inductively Coupled Plasma (ICP) etching is a dry etching process, a definition which covers a range of methods where a solid surface is etched in the gas vapour phase. The etching can either proceed by ion bombardment (e.g. sputtering) or through chemical reaction at the surface. Many of the techniques are plasma-assisted processes. Etching can be either isotropic or anisotropic processes through dynamics controlled by plasma conditions rather than chemical selectivity.

ICP is a reactive ion etching process where the plasma is driven not by electric fields but the magnetic fields from a 13.56MHz RF field thus creating high-density, low-pressure and low energy plasmas. The high plasma densities lead to considerably higher etching rates (on the order of several microns per minute) than traditional plasma etchers which cannot achieve such high plasma densities at low pressure.

A schematic of a typical ICP setup is shown in figure 5.13. In this case, the plasma is generated by RF in the coil around the chamber. There is only weak coupling between the ion energy and the plasma power, and hence there is isotropic angular ion distribution in the plasma. Unlike other (e.g. RIE) plasma etches, ICP has separate coil and platen power supplies, allowing for fully independent control of ion energy. If an AC voltage is applied to the platen, electrons and ions are alternately attracted. The electrons are much more mobile than the heavier ions so the platen acquires a negative charge giving a net time averaged negative potential that is sufficient to accelerate the ions towards the substrate (at energies up

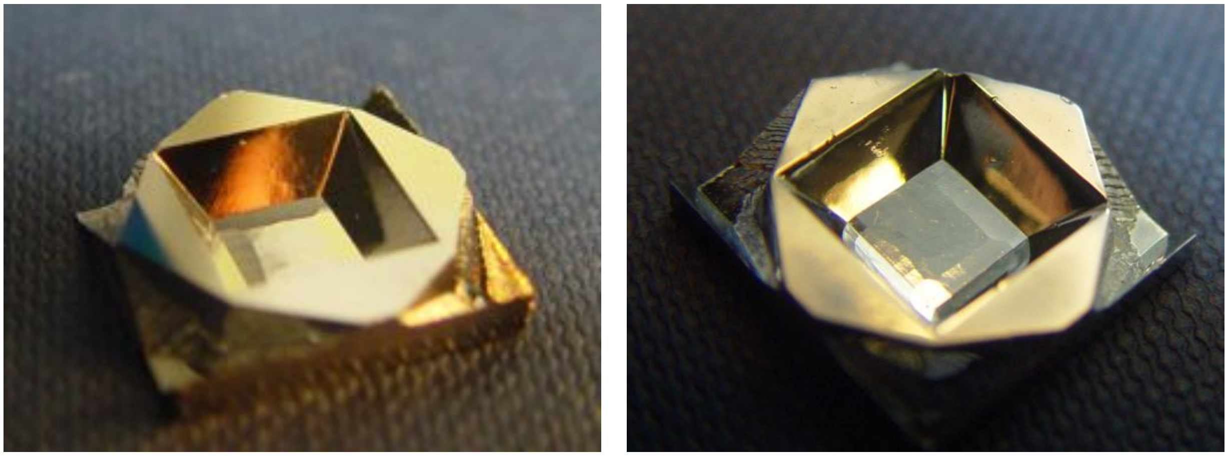

(a) Photograph of sample post 30 minutes HNA etching (b) Photograph of sample post 60 minutes HNA etching

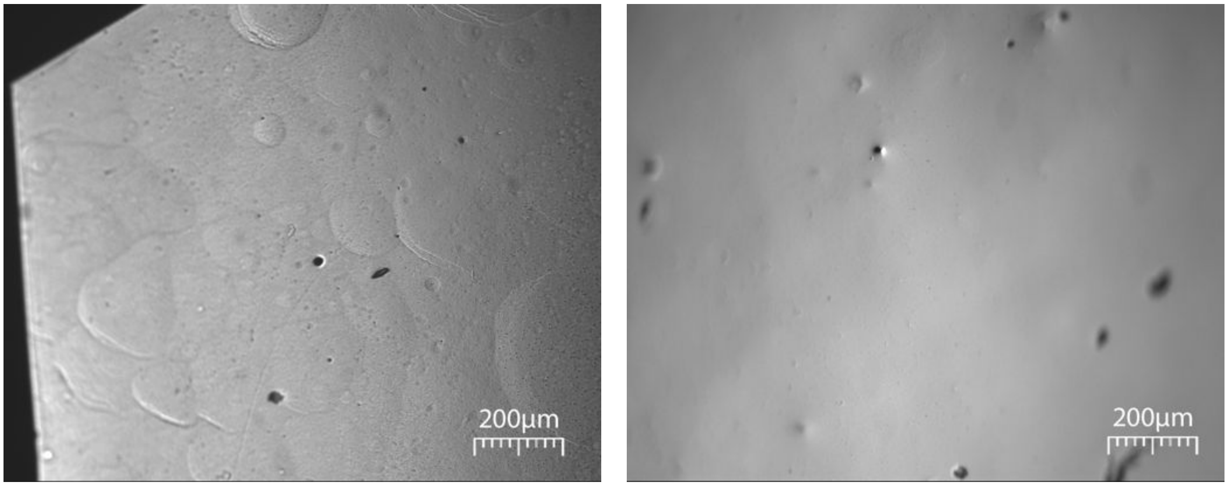

(c) Microscope of sample post 30 minutes HNA etching (d) Microscope of sample post 60 minutes HNA etching

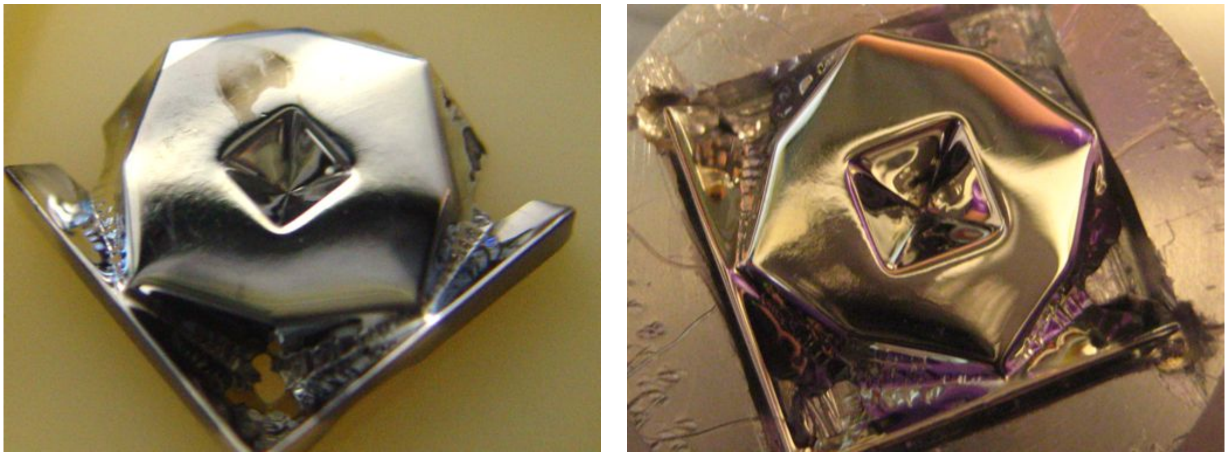

(e) Pyramid etched in HNA for 60 minutes (f) Pyramid etched in HNA for 90 minutes

Figure 5.12: HNA Etching - Anisotropically etched pyramids are polished with an isotropic wet etch in HNA

to several hundred eV). This allows the process to switch between a highly isotropic etch (low plasma potential) and a highly directional etch (high plasma potential for high ion energies) also allowing for good reactant penetration into high aspect ratio structures [93].

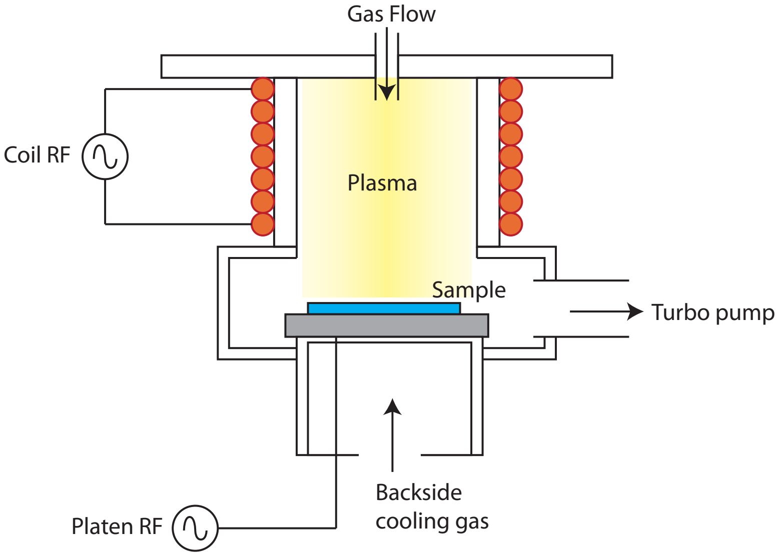

Figure 5.13: Schematic of a typical ICP etcher

The source for the plasma was the inert sulphur hexaflouride (SF_6) which disassociates in the RF field to form fluorine free radicals. These fluorine radicals react with the silicon to form volatile SiF_4. A recipe developed for isotropic etches of silicon microlens molds [93], detailed in table 5.1, was tested to see its effectiveness in removing the {111} plane roughness.

Isotropic ICP etch recipe	
Pressure (mTorr)	10
Coil Power (W)	3000
Platen Power (W)	3000
SF_6 flowrate (sccm)	200
O_2 flowrate (sccm)	0
Chuck Temperature (°C)	20

Table 5.1: ICP recipe for the isotropic smoothing of silicon hollows. Taken from [93].

Using a STS etcher at LCN, only 1500W of coil power was possible rather than the required 3000W, but all other parameters from the recipe could be reproduced faithfully. Etching diagnostics reported that the peak-to-peak voltage was 134V and the bias voltage (between the platen and the chamber) was reported to be 17V. As well as pyramid samples, a portion of a thin wafer was placed in the chamber, partially masked so as to give a calibration of the depth etched. Initially a large wafer was used to mount the samples (using Crystalbond as a temporary mounting adhesive), and this was covered in a thick layer of photoresist to protect the silicon underneath.

Three samples from wafer LCN1 (see figure 5.5 for the pre-etch surface quality) were etched,

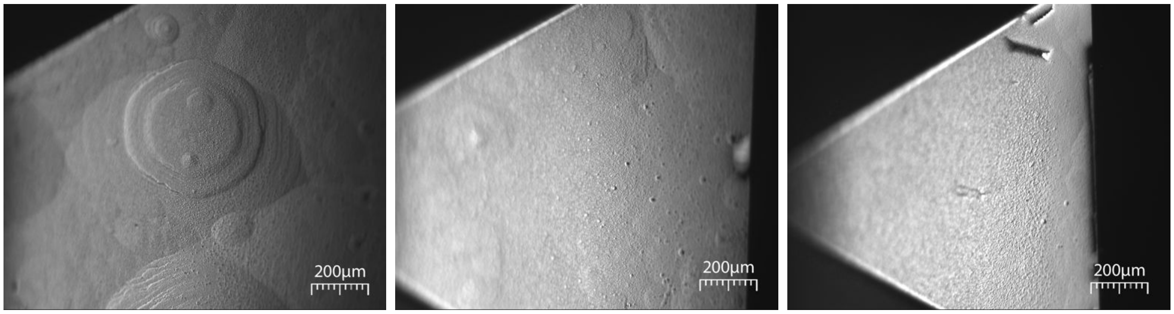

(a) Wafer II samples etched for (l to r) 10,20 and 30 minutes

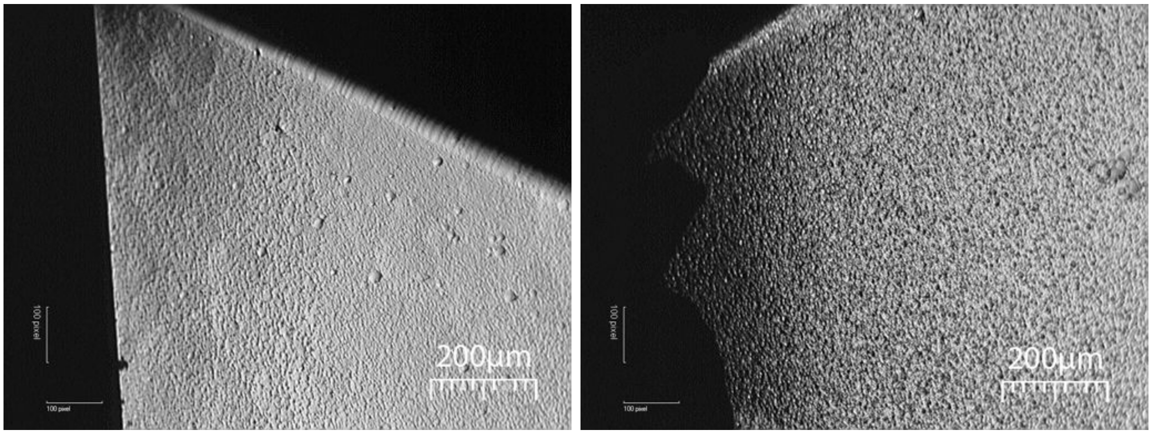

(b) Wafer III samples etched for (l to r) 40 and 60 minutes

Figure 5.14: ICP etches - Anisotropically etched pyramids are polished with an isotropic plasma etch.

and the recorded etch depths (as measured on the calibration piece) are listed below in table 5.2. Variations in the etch rate were due to changing the carrier the samples were mounted on.

Sample (pyramid size)	Etch Time (min)	Etch Depth (μm)
A (3.1mm)	10	34
B (2.4mm)	20	87
C (1.3mm)	32	100

Table 5.2: Initial SF_6 ICP trial

As can be seen from the microscope images (figure 5.14) the crater features are reduced in prominence with increased etching time, with the features almost completely suppressed after 30 minutes (100 μm depth). Despite the large amount of material removed, the faces remain visibly flat and the edges are still sharp, unlike the isotropic etch with HNA.

The reduction in feature prominence is accompanied by the appearance of a fuzziness to the faces. This noise does not appear on any (100) planes and is more apparent on the parts of the faces which are closest to the apex. The mechanism that produces this texture is not known, but it could be a result of sputtering damage from reactant products, which will be more concentrated inside the pyramid.

The three samples etched were of differing size, so in order to remove the possibility that geometric effects contributed to the different surface finishes, a batch of identical pyramids from Wafer III were etched for progressively longer periods of time. Similar etch rates were observed, and extended etches (up to 60 minutes in duration) were tried to see if the high frequency noise would decrease. These pyramids were initially 4mm in size, and hence 2.82mm deep, so holes appeared at the apexes for etches exceeding 20 minutes and increased in size with duration of etch. These holes appeared clean to the eye but had a jagged appearance under the microscope, as seen on the edge of righthand figure 5.14(b). We conclude that between 30 and 40 minutes of ICP etching is sufficient to completely remove the unwanted features from the {111} faces, and that this produced no degradation (apart from the aforementioned apex holes) in the shape or planar nature of the faces.

Combined ICP & HNA

The two etching processes what we tried had relative strengths and weaknesses. The HNA produces a very smooth face, but introduces unacceptable rounding. ICP removes surface features whilst maintaining planar structure but introduces high frequency noise to the faces. It was therefore decided to try a hybrid etch, where ICP would be used to remove the cratering, and a short HNA etch would be used to remove the high frequency remainders.

Samples etched by using the ICP previous process (see figure 5.14(b)) were used as a starting point. With a aggressive HNA composition (9:13:8) only 1 minute of etching was needed to remove the high frequency component of the surface profile without rounding the faces significantly. Images of the resulting faces before and after the HNA polish are seen in figure 5.15(b), some slight rounding can be seen to the edges but the clarity of the faces is quite apparent.

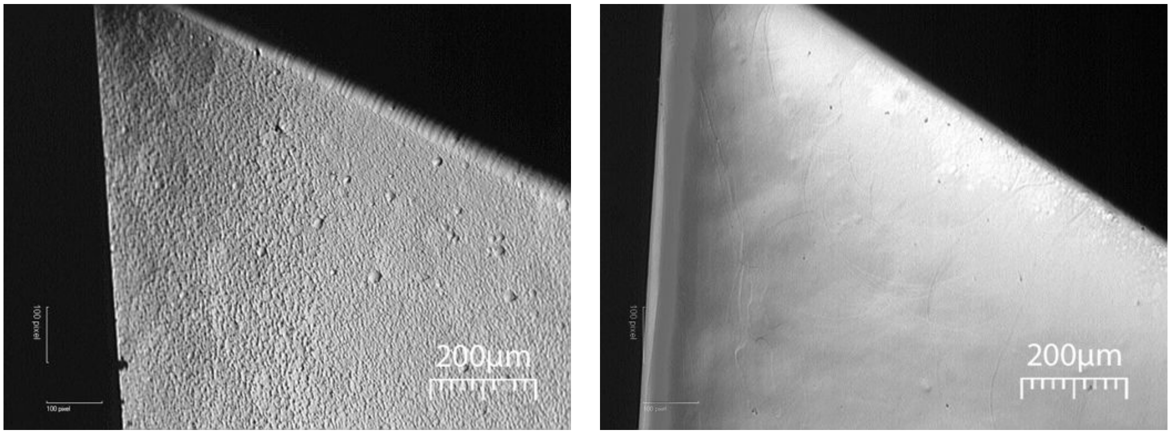

(a) (b)

Figure 5.15: Anisotropically etched pyramids polished with a)40 min ICP b)combined 40 min ICP & 1 min HNA

5.5 Profiling Surfaces

In order to achieve some quantitative measures of the surface properties and the effect of the various postprocessing techniques, the surfaces were profiled using a scanning white light interferometer (SWLI). This offers many advantages over other profilometry techniques, as it is a non-contacting technique and much faster than other optical techniques as it only requires scanning in the vertical direction. It is limited by the Rayleigh diffraction limit in the lateral direction to around 0.5μm for large numerical apertures. In the vertical direction the limit is the precision to which the phase of the reflected light can be measured, and this is around half a nanometre.

There are some notable problems with SWLI, less prevalent with modern systems but still a matter for some concern. The most prevalent ones encountered are stepped artefacts corresponding to a 2π phase jump where the fringe orders have been misclassified, and these are often referred to as 'ghost-steps' as they do not correspond to a physical step in the material. There is also a general overestimate of surface roughness, arising from multiple scattering but this arises from length scales below the diffraction limit of the observation aperture [94].

Unlike AFM/SEM measurements, the instrument used (Zygo NewView 200), can scan regions up to 1mm by 1mm in size and is thus ideally suited to viewing the cratering on the pyramid surfaces. Unfortunately the areas of interest were contained inside the pyramid volume and hence it was impossible to get the SWLI objective within a usable distance and at a usable angle for most of the pyramids. However with some of the larger flat bottomed pyramids the faces are sufficiently far apart to position the objective to make an observation. In other cases the pyramid sample was diced to provide access to the face. The raw data from the SWLI is provided in the form of a 2D array of values (typically spaced by approximately 1μm) recording the height measured at that position.

5.5.1 Form, Waviness and Roughness

To analyse and discuss the profiles we define three spatial regimes, **form**, **waviness** and **roughness**. Form refers to the general shape of the surface, and correspondingly can be known as the 'shape' or 'geometry'. Waviness can be considered intermediate scale surface profile typical of the craters and roughness is the finest grade of texture such as we see following the ICP etching. There is no fixed definition for what separates these regimes and the cutoff for each can vary according to the surface or process under study. Filtering of the images can isolate each of these spatial regimes, allowing independent measurements of the noise at each level to be made[95]. For the pyramid faces the following definitions have been chosen.

- **Form** - The shape of the pyramid face, $\lambda_c \geq 500\mu$m
- **Waviness** - The mid-range undulations in the face such as cratering $50\mu\text{m} \leq \lambda_c \leq 500\mu$m
- **Roughness** - Non-geometric deviations $\lambda_c \leq 50\mu$m

When we process the data, form is removed by subtracting a quadratic surface fit. The coefficients of the x^2 and y^2 terms will be used assess the degree of curvature of the surface.

Assuming the quadratic form was appropriate, the resulting profile is then the sum of roughness and waviness. We use a lowpass filter to isolate the waviness[4]. The most widely used lowpass filter is a Gaussian filter

$$S(r) = \frac{1}{\alpha\lambda_c} e^{-\pi\left(\frac{r}{\alpha\lambda_c}\right)^2} \tag{5.1}$$

$$\tag{5.2}$$

where $\alpha = \sqrt{\ln 2/\pi}$, λ_c is the cutoff frequency (the filter has transmission of 1/2 at the cutoff) and r is the distance of a given profile point from the centre of the weighting function[95]. We apply this filter by 2D convolution to the form-subtracted profile data to obtain the waviness. The roughness is obtained by subtracting the obtained waviness from the original form-subtracted profile.

Surface profiles have been obtained for a number of different samples

- Top flat nitride coated 100 plane surrounding a pyramid, this surface should act as a reference as when coated with a metallic coating this surface is indistinguishable from a high quality commercial mirror
- 111 face of an unpolished pyramid from Wafer II (as shown in figure 5.8)
- 111 pyramid face from a pyramid polished with 30min HNA or 30min ICP (as shown in figures 5.12 and 5.14 respectively)

Illustrations of the obtained profiles, the extracted waviness and the roughness of the surfaces are shown in figure 5.16. To allow comparison the colour mapping of the waviness and roughness plots are consistent between all images (spanning ±400nm for waviness, and ±15nm for roughness). The exception to this is the roughness of the ICP samples which had to be plotted with a larger dynamic range.

In order to quantify the texture at each level, two measures of roughness were used, root-mean-square (R_{RMS}) and average (R_a), defined as follows

$$R_{RMS} = \sqrt{\frac{1}{N}\sum_{i}^{N}(z-\overline{z})^2} \qquad R_a = \frac{1}{N}\sum_{i}^{N}|(z-\overline{z})| \tag{5.3}$$

If the surface is dominated by large features such as craters or hills then the rms roughness is generally higher than the average roughness. When calculating these values the data was filtered to remove 2π steps, and for the roughness profile a region 100μm square from each sample was extracted so regions of dirt and wafer damage could be excluded. Results can be seen in table 5.3.

---

[4] Care must be taken that form is correctly subtracted as use of the lowpass here means any residual form would be included in the waviness

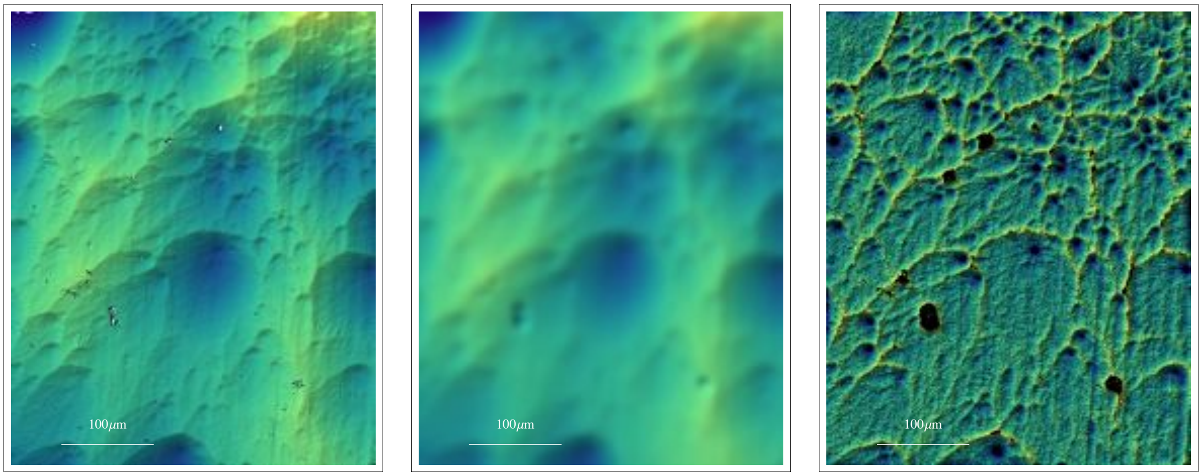

(a) Unprocessed LCN1

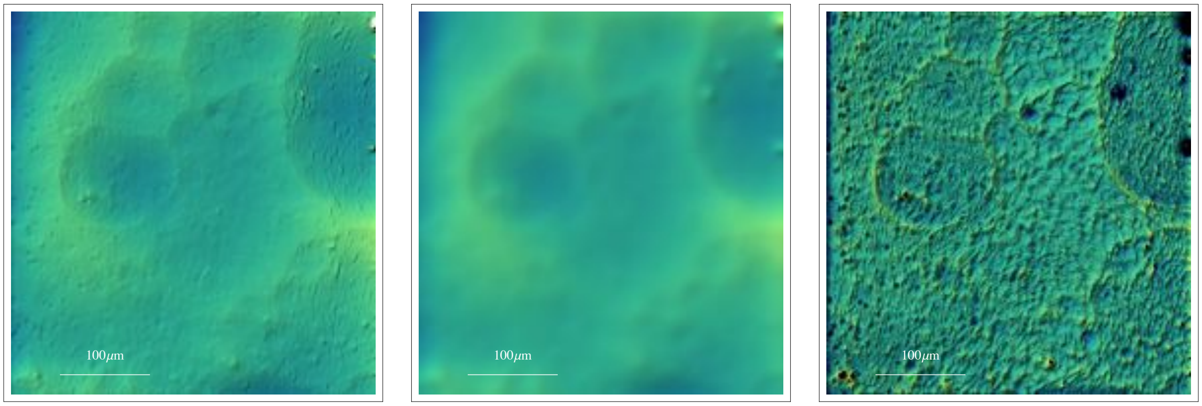

(b) HNA 30min

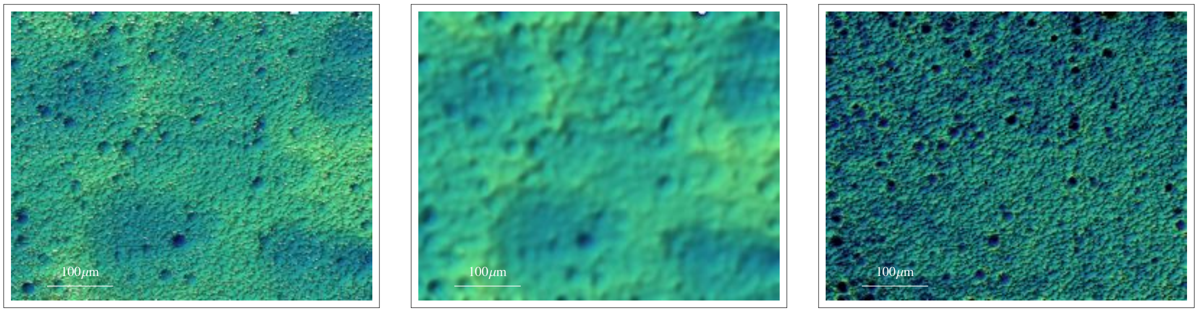

(c) ICP 30min

Figure 5.16: Profile data. Left:Raw data, Centre:Waviness, Right:Roughness

Sample	Roughness (nm)		Waviness (nm)		Form	
	R_{RMS}	R_a	R_{RMS}	R_a	$\frac{\partial}{\partial x^2}$	$\frac{\partial}{\partial y^2}$
(100) plane	0.83*	0.6*	4.6	1.3*	0.01	0.01
Wafer II (unpolished)	4.0	3.1	102	84	-0.7	2.7
HNA 30min	3.0	2.2	58	44	6.7	-29.6
ICP	22.7	16.9	61	48	1.0	2.1

* These values are below the manufacturers quoted resolution of 2nm

Table 5.3: SWLI Profiling results

There we see that the roughness of the HNA polished sample is close to that of the unpolished sample from Wafer II at only a few nm, almost as small as the untreated (100) plane. The HNA etch therefore has little effect on the roughness but reduces the waviness by around a factor of two. However the HNA increases the curvature by a factor of 10. The ICP etched sample shows a marked increase in roughness, but little change in the form and a similar reduction in the waviness as the HNA. The rms values are consistently higher than the average roughness for all surfaces studied, which matches the observation that the surface is dominated by cratering.

5.6 Surfaces as mirrors

5.6.1 Mirror coatings

As silicon is only 33.1% reflective at 780nm the surfaces require a reflective coating. Following the trials in the macroscopic glass pyramid, aluminium was selected as the preferred coating. Aluminium coatings were sputtered onto glass slides, and a reflectivity of $82.5 \pm 0.5\%$ was measured, which is slightly below the expected value of 87.9% at 780nm. These slides were heated in air to several hundred degrees to see if the lack of a protective overcoat would lead to degradation, but no change in reflectivity was observed. Typically layers of 25nm were deposited, with deposition rates around $0.1\text{Å}s^{-1}$ at a substrate temperature of $100°$. We have also successfully sputtered other metals onto the silicon pyramids, namely chromium, platinum, copper and gold (with a 5nm chromium adhesion layer used for gold).

When the silicon pyramids are coated, there is a slight difference in reflectivity between the (100) plane and the {111} surfaces, consistent with a loss of some light to diffuse scattering on the {111} surfaces. The reflectivity of the ICP etched mirrors was measured to be identical to an unpolished mirror.

5.6.2 Ray Tracing

In order to judge how the light reflected from the surfaces appears to an atom close to the mirrors, the SWLI profile data can be used to obtain a map of surface normals and hence simulating with raytracing the effect on a uniform plane wave incident on the face. The simulated rays were incident on the face, which is inclined at $70.5^\circ/2 = 35.75^\circ$ and the imaging plane was initially placed 1mm away to observe what the intensity pattern would be at the position of the MOT in the centre of the pyramid. We also simulated the far-field intensity pattern (at 20mm) to see if any difference could be observed in the intensity distribution outside of the pyramid. In total 300 000 rays were traced (one per point in the profile data), and the intensity pattern in the image plane constructed by counting the number of rays intersecting in a 5μm square areas.

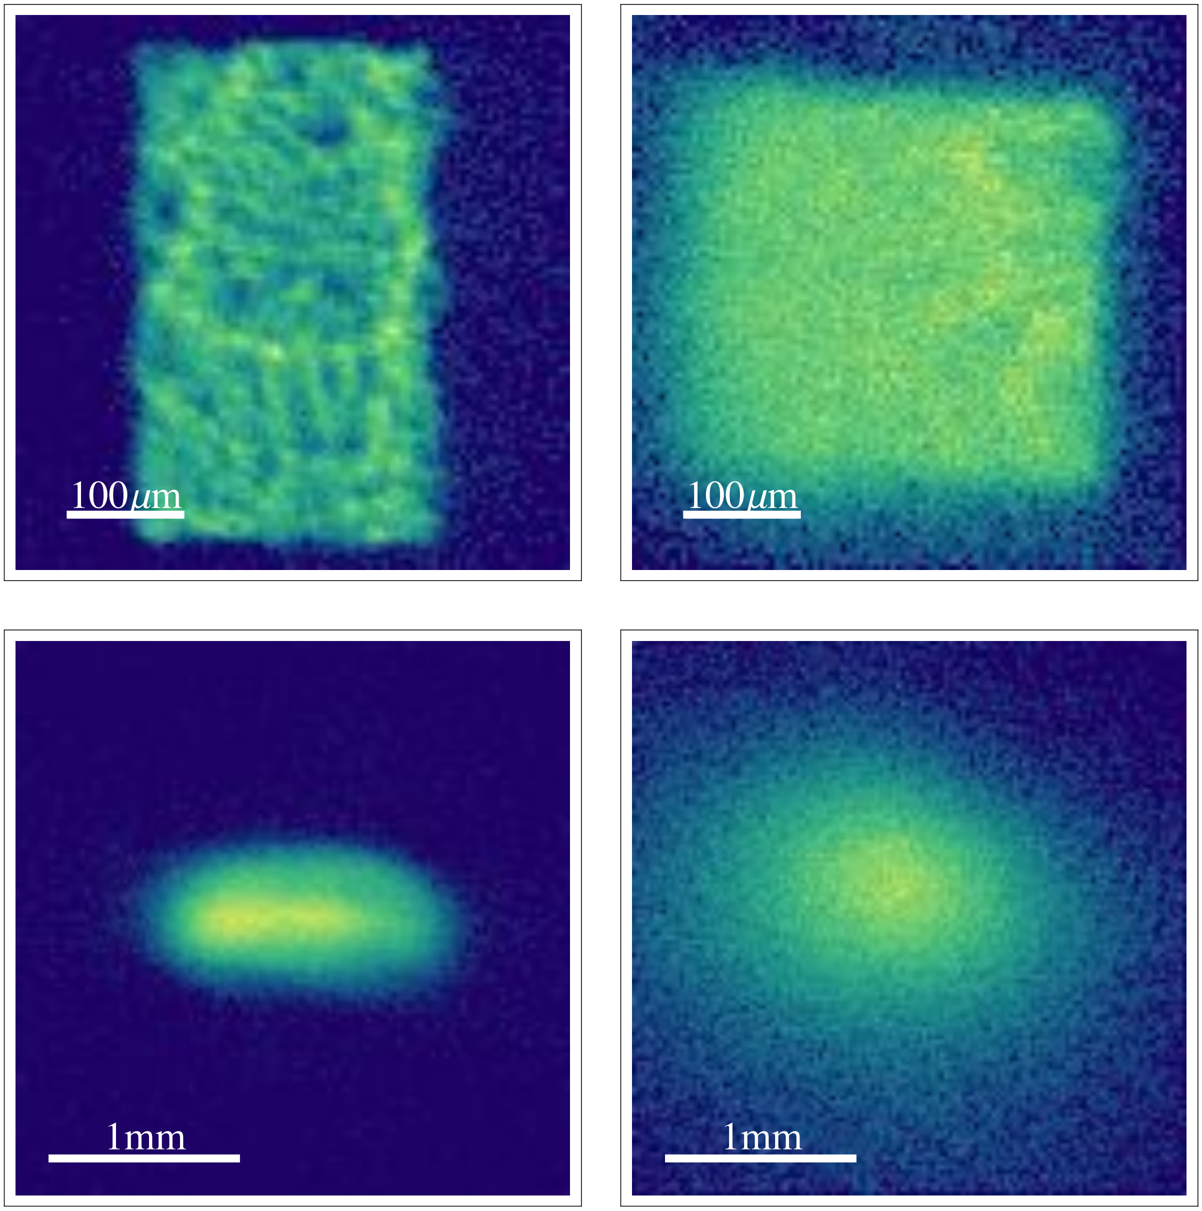

Figure 5.17: Raytracing results. Left:Unpolished Right:30min ICP. Top:Near field (1mm) Bottom:Far field (20mm)

As can be seen from the top images in figure 5.17 the intensity pattern from the unpolished surface features a large amount of intensity variation and structure while the ICP etched face produces essentially uniform illumination in the MOT region. Moving the image plane further

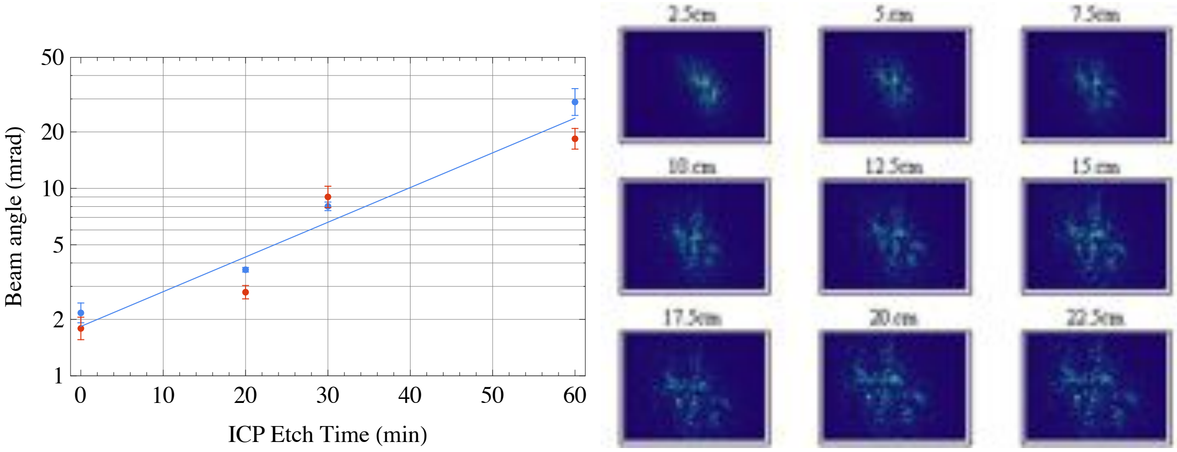

Figure 5.18: Left:Reflected beam divergence as a function of etching time. Right:Images of beam at 2.5cm intervals.

away, the reflected beams have increased in size(with more size increase present in the the ICP sample) and no major difference can be seen in the homogeneity of the illumination.

This observation is confirmed by measurements made with laser illumination of pyramid faces (by a collimated beam with size comparable to the face). The reflected beam was captured by a CCD after approximately 3cm of propagation from the face and shows slightly more small-angle scatter from the ICP sample, but no significant differences between the two.

The intensity patterning predicted on the unprocessed samples suggests that the trapping force produced will be highly nonuniform, and the coalescence of a stable MOT in unprocessed pyramids is quite likely to be inhibited by the spatial imbalance between opposite beams.

5.6.3 Angular Spread

Any remaining roughness in the pyramids will cause light to be scattered diffusively into a solid angle around the direction of specular reflection. The increased roughness of the faces etched with ICP implies that the reflected beams will spread in a larger angle. If the width of the reflected intensity pattern is measured as a function of distance from the face, then the angle of scattering of the beam can be determined. This was measured for a number of samples etched with ICP for different durations and the results are illustrated in figure 5.18, showing an approximately exponential increase in divergence with etching time.

5.6.4 Scattering Properties

The scattered light from the pyramid faces is the background which must be subtracted to obtain only light arising from atom fluorescence. In order to judge the magnitude of this component, an unprocessed silicon pyramid, sputter coated with 50nm of gold was illuminated at normal incidence with 3.6mW/cm^2 of 780nm light, and a CCD camera was used to collect the scattered light at a variety of angles around the pyramid. The camera was placed 12cm away, as this is the typical working distance for the imaging lens used in the experiment.

Each image can be considered as an array of intensity measurements, and if we extract these values (I) and bin the data we obtain an intensity distribution for each image. Figure

5.19 shows the distribution of intensity values ($p(I)$) collected from the scatter in each image. This intensity distribution, which compromises speckle due to the coherent illumination and detected using a finite aperture detector (the CCD pixels) is expected to be described using the Gamma distribution [96]

$$p(I) \approx \frac{I^{k-1} e^{-\frac{I}{\eta}}}{\Gamma(k)\eta^k} \tag{5.4}$$

where $\eta =< I > /k$ and k is the number of speckle correlation cells within the measurement aperture[5]. It can be seen in 5.19 that the Gamma distribution (plotted as red lines) is a good fit for the observed data in each case. If the mean, $< I >= \eta k$ is calculated for each angle, and plotted on the right of figure 5.19 we can see a roll off in the mean value as the angle of viewing increases. It was anticipated that the roll off would follow a $\cos^2\theta$ dependence (plotted as the red dashed line), but the morphology of the pyramid surface (for example corrugation effects) leads to some deviation from this trend.

Taking a peak scatter level at close to incident detection of 0.3pW/pixel, and referring to figure 4.18 this places the threshold for detection of a 100μm MOT at well below 100 atoms above this surface.

Scatter in postprocessed pyramids

The postprocessing discussed in this chapter changes the morphology of the surface and results in different scattering properties. Various pyramids, processed with either ICP, HNA or combined ICP/HNA and coated in aluminium were observed under illumination with 5mW/cm^2 of laser light. It was observed that the pyramids processed with HNA and the combined ICP/HNA process are comparable in scatter to the unprocessed pyramid, with only regions at the corners and apex showing increased scatter levels as we would expect for the increased small scale roughness.

The pyramid processed with ICP alone shows a significant increase in the amount of light scattered. The level of scatter is estimated at 10pW per pixel, increasing the threshold of detection in this pyramid to MOTs of 200 atoms.

5.6.5 Polarisation

An important characteristic of the coated pyramid faces is how they alter the polarisation of the incoming circular light, which has implications for how they behave as mirrors for a MOT. In section 4.1.1 the non-normal incidence of the light on the surfaces introduces ellipticity into the initially circularly polarised light. If the degree of ellipticity is too severe the MOT will not operate as expected. This is studied in the silicon pyramids by illuminating them at normal incidence with circularly polarised light, passing a reflection (Type I or III) from the pyramid through a quarter wave plate in a rotatable mount rand then through a fixed linear polariser.

---

[5] Note that for a measurement aperture which is small compared to the correlation area of the speckle $k = 1$, the probability density function approaches a negative exponential, or for large measurement apertures where k is large the pdf approaches the form of a Gaussian

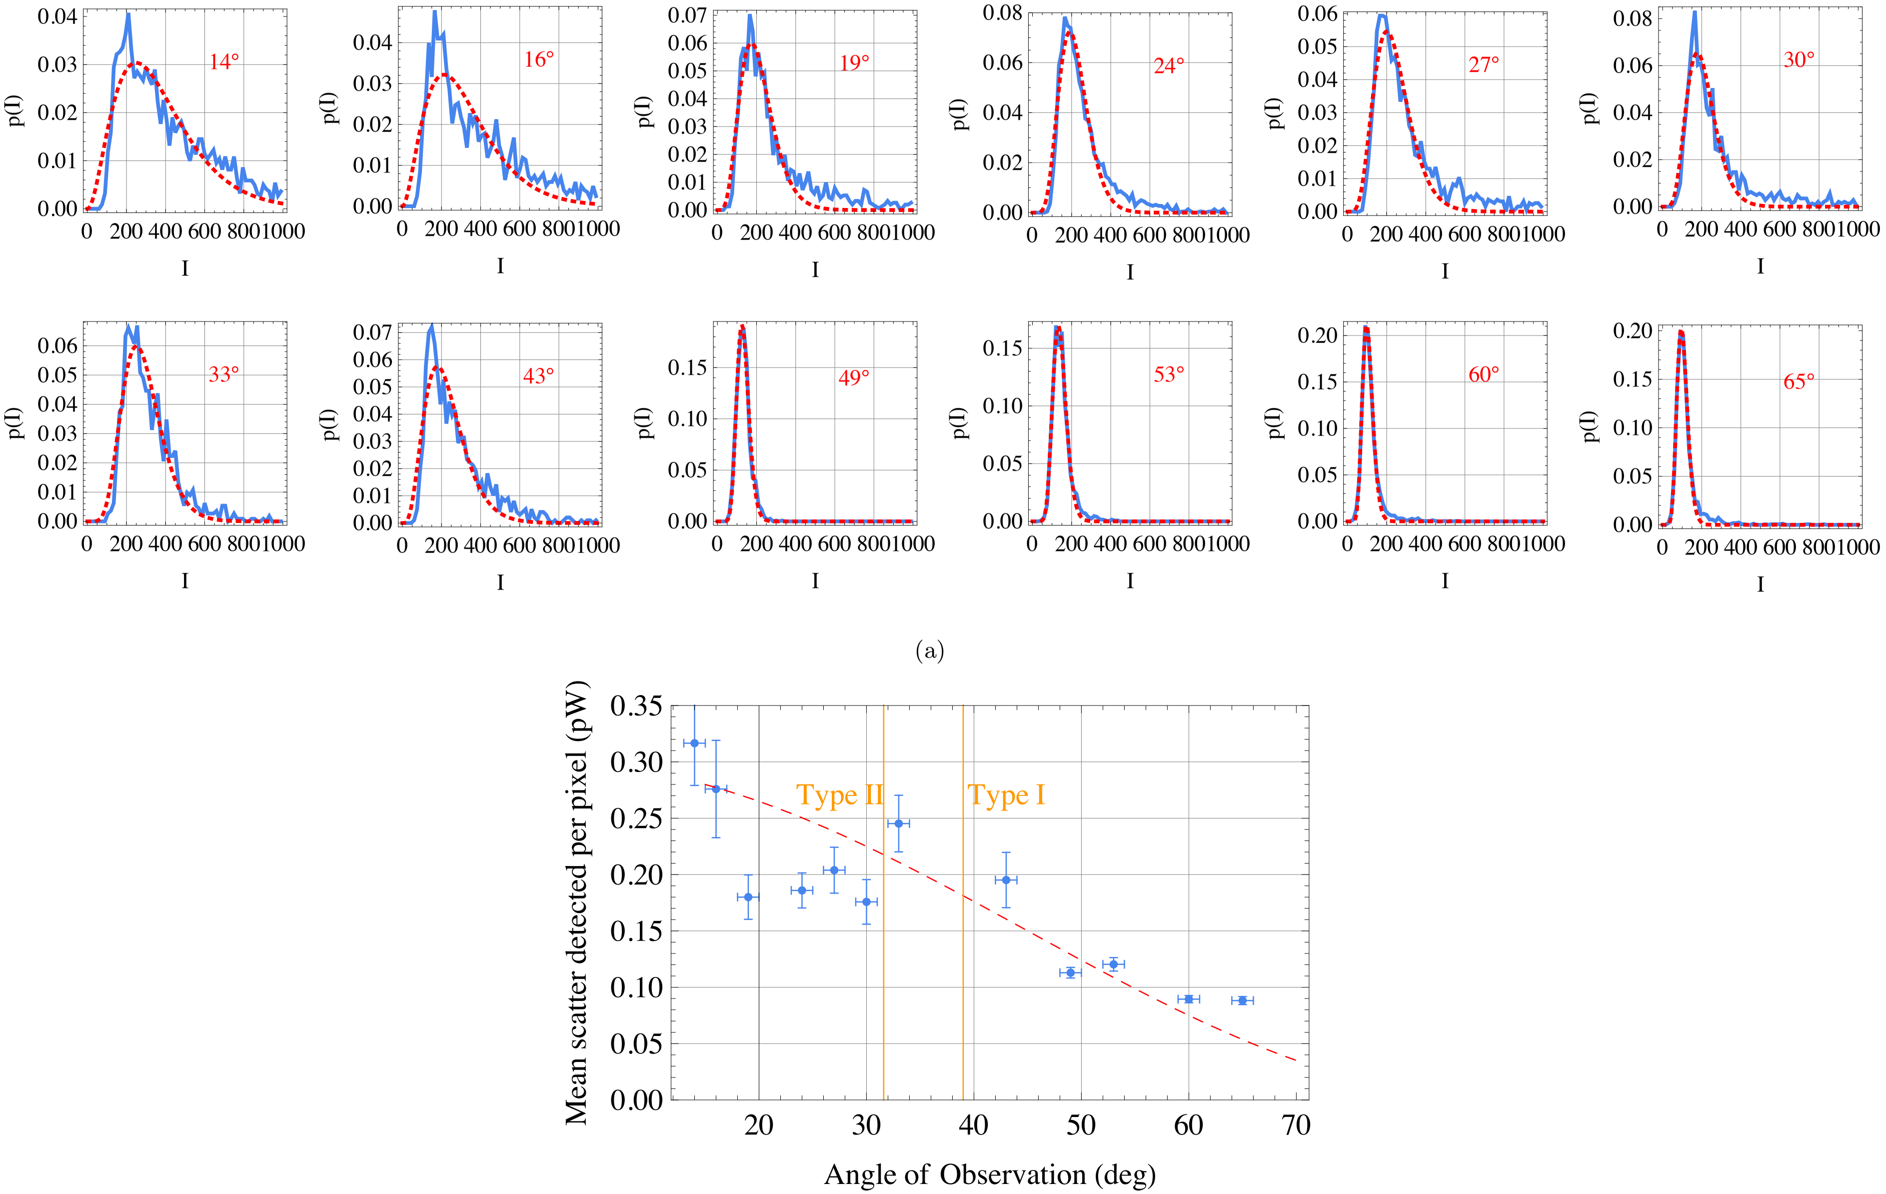

(a)

(b)

Figure 5.19: (a) Distribution of pixel intensities and associated Gamma distribution fits in images taken at a variety of angles around a silicon pyramid. (b) Extracted mean scatter levels as a function of detection angle. Orange indicators indicate positions of type I and II reflections. Red line is a $\cos^2(\theta)$ trend.

The intensity of light transmitted through the linear polariser is recorded for a range of quarter wave plate angles and by comparing this data to a theoretical model the polarisation of the reflected light can be inferred. The theory used is Jones calculus, where polarisation of light is represented in a vector, $\begin{bmatrix} E_x(t) \\ E_y(t) \end{bmatrix}$. Optical elements are represented as matrix operators, and thus the expected output field, $\vec{E}_{out}$, can be modelled as a sequence of operators on the input light $\vec{E}_{in}$.

$$\vec{E}_{out} = [\boldsymbol{J_{VLP}} \boldsymbol{R}(-\theta) \boldsymbol{J_{QWP}} \boldsymbol{R}(\theta)] \vec{E}_{in} \tag{5.5}$$

where $R(\theta)$ is the rotation matrix, and the linear polariser $\boldsymbol{J_{VLP}}$ and quarter wave plate $\boldsymbol{J_{QWP}}$ are given by

$$\boldsymbol{J_{LP}} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad \boldsymbol{J_{QWP}} = e^{i\pi/4} \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \tag{5.6}$$

The measured value is an intensity, given by

$$I = \vec{E}_{out}^{\dagger} \vec{E}_{out} \tag{5.7}$$

The input light is modelled as $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ ie^{-i\delta} \end{bmatrix}$, where δ is the phase shift introduced on reflection. By fitting the function $I(\theta) = (\boldsymbol{J_{tot}}.E_{in})^{\dagger}.(\boldsymbol{J_{tot}}.E_{in})$ to the observed data, the value of δ can be determined. In addition we need to include the offset angle θ_0 (fast axis of the QWP in its rotation mount). To find this we make a control measurement with a known polarisation. Results are shown in figure 5.20, and show that as expected, reflection from the gold coated silicon flat at close to normal incidence introduces very little ellipticity to the beam. The values of phase introduced into the Type I and III beams are $-30.5 \pm 1.8^\circ$ and $-91.4 \pm 1.5^\circ$ respectively. The value for the Type I beam is certainly consistent with the expected -29.9° (see section 4.1.1), but the Type III beams are much closer to -90° (hence being linear) than the anticipated -131° for gold. Similar results are expected for aluminium pyramids.

5.6.6 Testing Performance as a MOT

Throughout the fabrication process, various pyramids were tested to see if they could operate as a MOT. They were installed in the chamber, which typically required just over a week to return to suitably low pressures, and attempts were made to detect trapped atoms.

5.6.7 Initial test silicon micropyramid

The mm scale silicon pyramids were initially tested inside a set of coils identical in size to those used for the large glass pyramid. As these coils are 25mm in diameter, great care had to be taken to site the small pyramid exactly at the centre of the field. Subsequently a smaller arrangement was constructed, with the scale of the coils reduced to a more suitable size. This assembly is illustrated in figure 5.21. Four tapped holes were made in an aluminium plate on a 8mm square, and into these four brass long studs were screwed. To create the lower coil

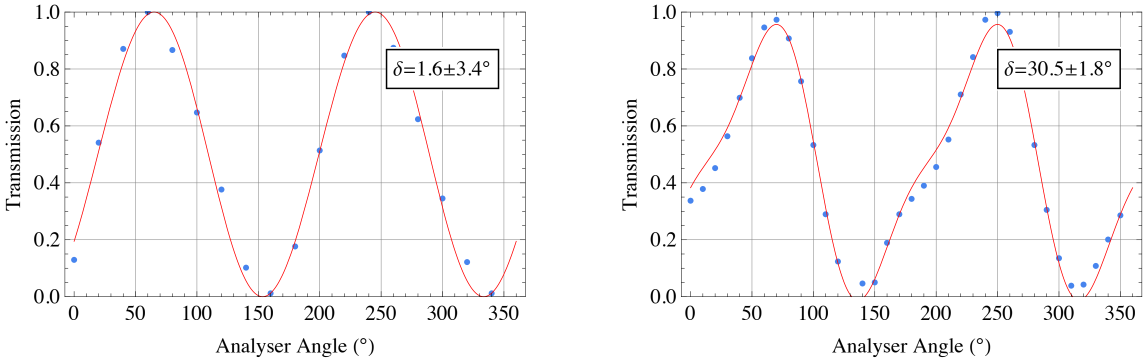

(a) Left:Reflection from top silicon flat at close to normal incidence, Right:Type I reflections on exiting the pyramid

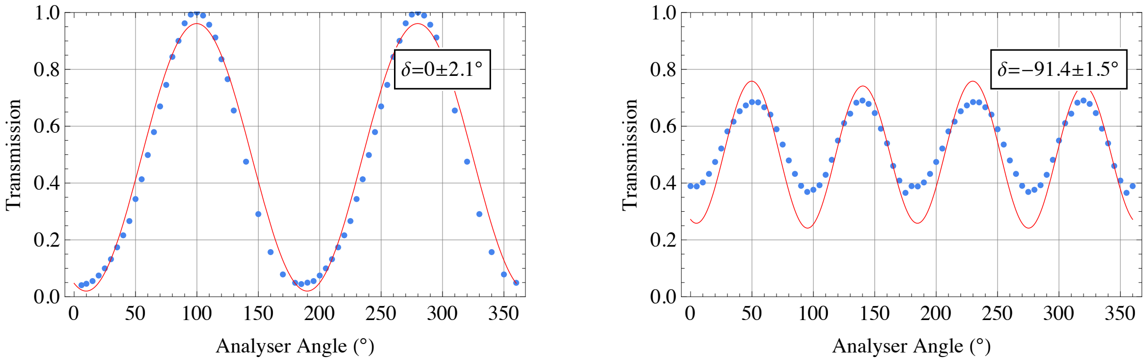

(b) Left:Reflection from top silicon flat at close to normal incidence, Right:Type III reflections

Figure 5.20: Polarisation modification in silicon pyramids

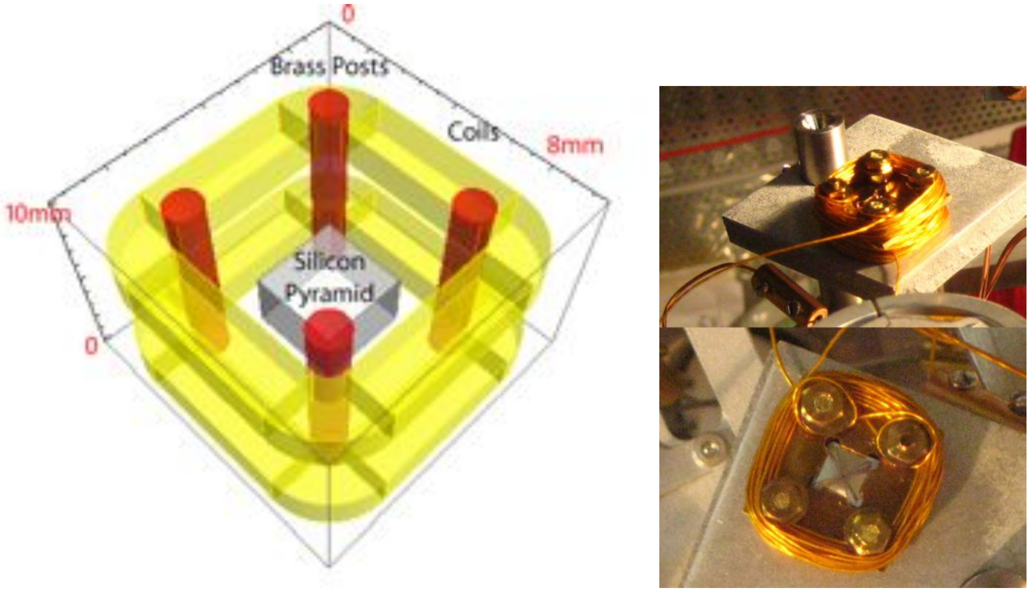

Figure 5.21: Initial Test Pyramid

25 turns of 0.5mm wire were made around the studs to create a square loop, and this was fastened into place with nuts threaded down the brass studs. A brass plate with 4 holes in each corner was placed over the top of the coil on which the silicon pyramid device sat. Another brass plate, with another 4 holes in the corners for the post and a 4mm square hole in the centre, similar to a picture frame was used to lock the pyramid into place. This was fastened down with a further 4 nuts and another 25 turns of wire made around the top to complete the anti-Helmholtz configuration. Although not the most aesthetically pleasing arrangement, it proved suitable for our requirements.

In order to understand the fields generated inside this package, a model was constructed using the *Radia* libraries for Mathematica. The model showed that at the position of the field zero a gradient of $\frac{\partial B}{\partial z} = 42\text{G/cm}$ was obtained for a current of 1A. In vacuum the coils were able to be able to support continuous currents of up to 1A without reaching the temperature at which the polyimide coating would have burnt (above 120°C).

Wafer	Size	Post Processing	Coating	Notes
Wafer II	4.1 & 3.1	None	Evaporated Al/MgF	Coating did not give uniform reflectivity
Wafer II	4.1 & 3.1	None	25nm Sputtered Al	No MOT $N > 250$ observed
Wafer I	4.1	60min HNA (recipe A)	25nm Sputtered Al	No MOT $N > 250$ observed
Wafer III	4	30min ICP/1min HNA (recipe B)	25nm Sputtered Al	Apex epoxy damage. No MOT > 250 observed
Wafer III	4 x 2	30min ICP	25nm Sputtered Al	Successful operation

Table 5.4: Pyramids tested

Testing procedure

Typically a range of parameters would be have to be explored, principally magnetic field as there was some uncertainty in positioning the field zero in the optimal place inside the pyramid volume. The magnetic field zero was stepped using the 3-axis bias coils, through a 3D volume typically 2mm cubed in step sizes of around 200μm. At each position the image was integrated until noise levels fell to a point where we could conclusively confirm the absence of MOTs more than a few hundred atoms in size (requiring generally around 1500 pairs of images). The pyramids tested are listed in table 5.4.

The first pyramids tested in May 2008 were an unpolished batch from wafer II compromising 4.1 and 3.1mm pyramids. They were evaporatively coated with a layer of aluminium and a protective MgF overcoat. No successful MOT was observed in either pyramid, we had observed degradation of the mirror finish on some other samples coated at the same time (taking on a milky appearance for the more damaged samples) and thus decided to recoat the pyramids without a protective layer. These pyramids were more uniform in finish but still did not form a MOT. We then embarked upon the investigation into the postprocessing techniques. Initially we tried two pyramids, one pyramid polished solely with HNA, and one pyramid polished with a combination ICP/HNA etch. A lot of hope was pinned on the pyramid polished with the combination etch, but unfortunately epoxy leaked through the hole in the apex during baking and partially filled the pyramid spoiling the mirrors. Neither pyramid successfully trapped atoms at a detectable level.

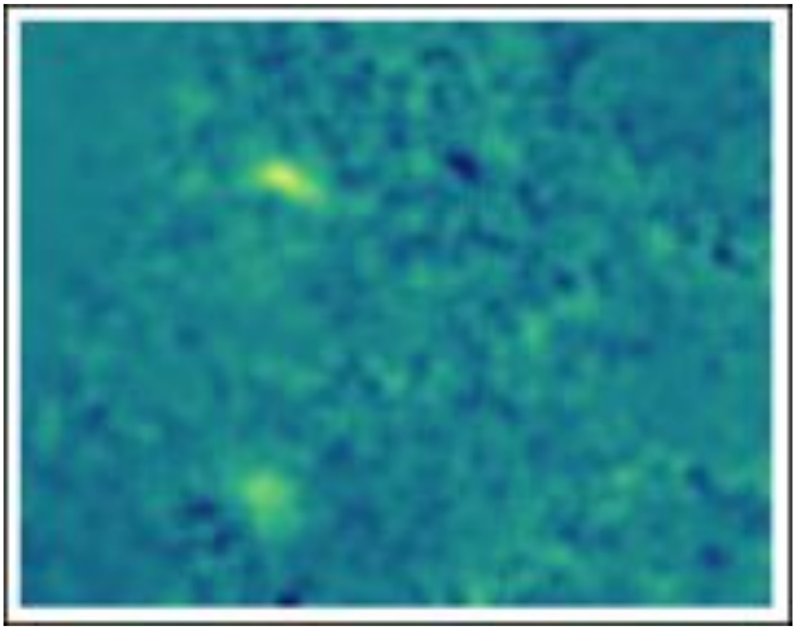

Figure 5.22: The first MOT observed in a silicon pyramid

Finally, we tried a 4.2mm 30min ICP polished sample from Wafer III. This showed high levels of scattered light as predicted but we persevered with observations. On the 12th December 2008, at about half past two in the afternoon, the image shown in figure 5.22 was observed.

Chapter 6

Trapping atoms in a silicon pyramid MOT

This chapter covers studies of MOTs formed in a 3.0mm deep silicon pyramid (4.2mm wide), with some discussion at the end about early studies in smaller pyramids. Some initial observations such as population and size of the MOT are given, followed by an exploration of the effect of laser and magnetic field parameters on the behaviour of the trap. Finally the temperature of the trapped atoms is measured and studies on the physical parameters of the trapping force are made.

6.1 First observations

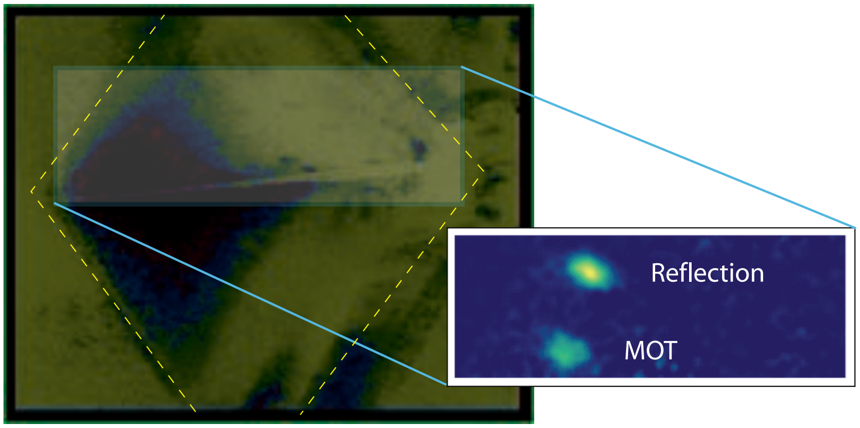

Figure 6.1: Large image shows the 4.2mm pyramid (opening outlined with the dashed line) as seen with the detection CCD, the MOT is much weaker than the background and hence invisible. If we remove the background using the techniques from section 4.4 we obtain the inset image showing the MOT and its reflection.

Figure 6.1 shows in the large image what the CCD sees when it looks at the 4.2mm pyramid. The CCD was positioned so as to minimise the amount of scatter, but also maximise the viewing angle into the pyramid volume. As previously mentioned in section 4.4.2, the MOT cannot be seen against this scatter, so we use the background subtraction from section 4.4 to obtain the inset image and see the trapped atoms in the silicon pyramid. There is a reflection of the MOT visible against a relatively homogeneous and weak region of scatter, we can also see the MOT but as it sits above a region of bright scatter from the intersection of two faces it has a poor SNR compared to the reflection.

We observed up to a maximum of 7000 atoms in the MOT when we used parameters of $\delta = 10\text{MHz}$, $I = 6\text{mW/cm}^2$ and a gradient of 30G/cm.

The image formed on the CCD is a projection of the 3D distribution onto the imaging plane and the size of the cloud can be measured by fitting to the image a 2D Gaussian function of the form

$$A \exp\left[\frac{-((x_1 - r_1)\cos\theta - (x_2 - r_2)\sin\theta)^2}{2\sigma_1^2} + \frac{-((x_1 - r_1)\sin\theta + (x_2 - r_2)\cos\theta)^2}{2\sigma_2^2}\right] + B \tag{6.1}$$

where A is the peak value at the centre of the distribution, (r_1, r_2) is the position of the MOT centre on the image plane, θ is the angle between the x, y directions in the image and the principal axes of the cloud, B is a constant (to take account of imperfect background subtraction) and σ_1 and σ_2 are the rms radii. We observe a size of $103 \pm 10\mu\text{m}$ by $140 \pm 12\mu\text{m}$ in the shortest and longest directions respectively. This shows that the radial is larger by a factor of $\sqrt{2}$ as predicted in section 2.5.6 due to the radial gradient being half that of the axial gradient.

Since this optimal MOT contains 7000 atoms, the peak density $(N/((2\pi)^{3/2}\sigma_1\sigma_2\sigma_3))$ is $2.3 \times 10^8\ \text{cm}^3$. A laser beam on resonance with the cooling transition in ^{85}Rb (hence $\sigma_L = \frac{3\lambda^2}{2\pi}$) would be attenuated by only 1.6%, indicating the cloud is optically thin and therefore the MOT is in the temperature limited regime.

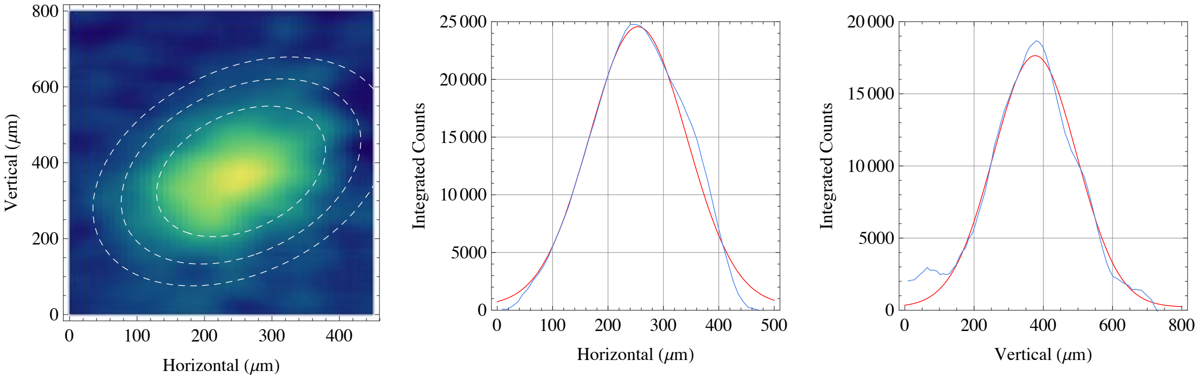

Figure 6.2: Left:Image of the MOT, with 1,2 and 3σ radii of the Gaussian fit indicated. Middle and Right: Blue line shows integrated profile of the MOT density, red shows the Gaussian fit.

6.2 Laser Parameters

We study the fluorescence collected from the MOT as a function of intensity and detuning of the primary laser, and the intensity of the repumping laser.

6.2.1 Detuning

If we scan the detuning of the MOT laser using the VCO to scan the offset lock, and measure the amount of fluorescence collected from the MOT we obtain the data shown in the left of figure 6.3. This was collected from multiple scans (each in a different colour) and shows that the shape is repeatable, with a sharp edge at low detunings beyond which no MOT is observed. As the detuning is increased it rolls off until at about -20MHz where very few atoms are detected. Averaging all data sets produces the graph on the right of figure 6.3 which show the shape of the curve with a lot more clarity. The shape in the graph can be understood by considering that the fluorescence yield is a product of the atom number $N(\delta)$ and the total scattering rate $R(\delta)$. The scattering rate as a function of detuning is given in equation (2.4) but the relation of detuning to the number of atoms has to be estimated. Let us assume the force on the atom is given solely by $m\ddot{x}(t) = -\alpha\dot{x}(t)$ where α is the damping coefficient defined in equation (2.17). Solving for an initial velocity v_0 we find

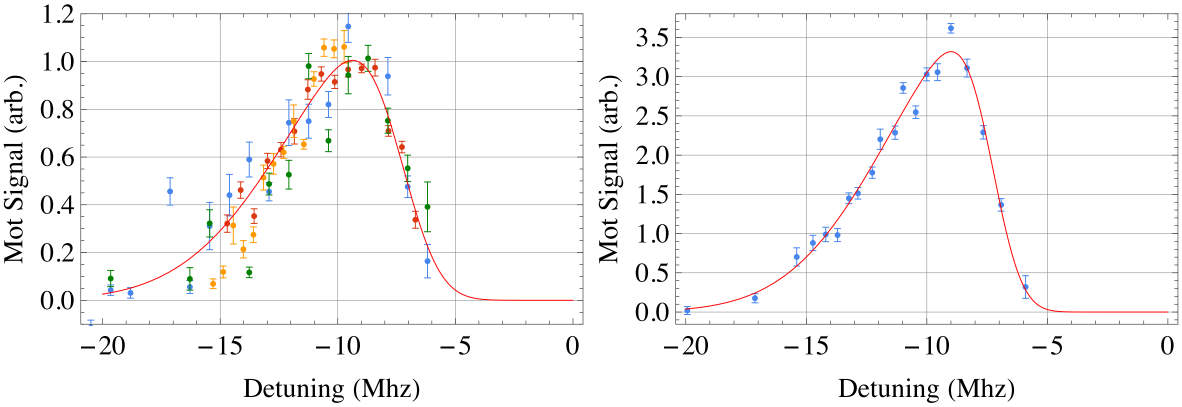

Figure 6.3: Left: Amount of fluorescence observed from the MOT as a function of trapping laser detuning taken from a number of experiments. Right: Multiple data sets averaged to give good coverage of the full range. Line is a guide to the eye.

$$x(t) = \frac{mv_0}{\alpha}(1 - e^{-\alpha t/m}) \tag{6.2}$$

Assuming a short stopping distance (in this case the depth of the pyramid), the maximum trappable velocity v_c is proportional to α. Hence with knowledge of α (equation (2.17)) and the scattering rate (equation (3.3)) predictions for the fluorescence yield detected Ω as a function of detuning δ can be given by

$$\Omega(\delta) = N(\delta)R(\delta) = A\alpha(\delta)^4 \frac{ns}{1 + ns + 4(\frac{\delta}{\gamma})^2} \tag{6.3}$$

where A is a normalisation constant which contains constants independent of the detuning (e.g. imaging geometry). Both the formulae for α and R contain the term n, which takes account of the number of beams interacting with the atom. The magnitude of n is not immediately obvious, but we can choose a number of potential values which make themselves immediately apparent

- $n = 3$: This lower limit can be derived by only considering the beams contributing to the vertical trapping, the incoming beam and the four second reflections of the Type I beams (the first reflections have the wrong polarisation in the vertical direction) giving $n \approx 3$ for a reflectivity of 75%.
- $n = 6$: If we consider the incoming beam, the four first & second reflections of the Type 1 beam we obtain $n \approx 6$ for a reflectivity of 75%.
- $n = 9$: An upper limit arises when we consider all the beams interacting the MOT region, including the 8 Type III reflections.

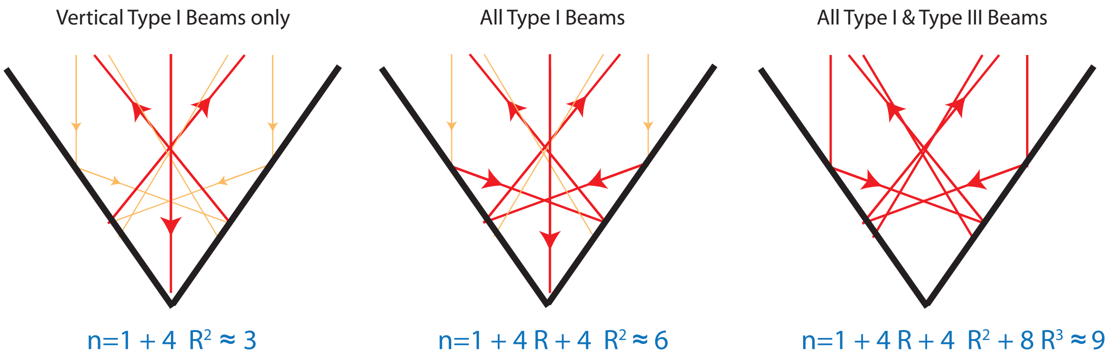

Figure 6.4: Potential values of n - only contributing beams are coloured red, orange beams are not considered. Beams coming into and out of the page are not drawn.

Equation (6.3) is plotted in figure 6.5 for this range of values of n, and it can be seen that the best correspondence to the data occurs for a value of $n = 6$ or 9. If we evaluate the χ^2 probability for all three fits, this gives the result that for $n = 6$, $P(\chi^2) = 99.5\%$, for $n = 9$, $P(\chi^2) = 94.5\%$ and for $n = 3$, $P(\chi^2) = 1\%$. This indicates that it is highly unlikely that $n = 3$ and hence we should assume that n is either 6 or 9. All three curves required a shift in the detuning to best fit the data, for the case of $n = 6$ it is offset by approximately 4MHz to the red. This offset could arise from not locking the lasers to the precise centre of the transition or more likely, from the fact that the MOT does not sit at the minimum of the magnetic field due to the force imbalance but at a field of $\sim$ 1G shifting the absolute detuning to the blue.

The good correspondence between the proposed theory and the model indicates that we are most likely on the right track with the capture mechanism and thus additional dynamics are most likely not necessary to explain the behaviour of the silicon micropyramids.

6.2.2 Intensity

The intensity of the trapping light was varied by rotating a half wave plate and monitoring the total power in the beam. Knowing the beam dimensions, the peak intensity of the beam could hence be determined from $2P/\pi w^2$ where w is the beam waist. Figure 6.6 shows the fluorescence yield from the trapped atoms as a function of the intensity of incident trapping intensity. Again, the model given in equation (6.3) can be used to try and explain the shape

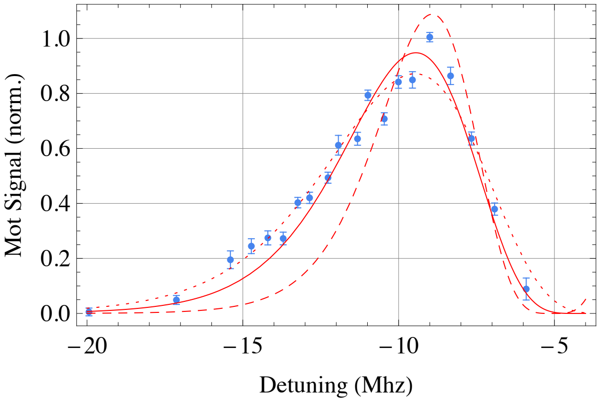

Figure 6.5: Proposed theory for detuning dependence of MOT number. Red Lines: Theoretical model displayed with the detuning data for $n = 3$(dashed), 6(solid) and 9(dotted).

of the data, however this requires us to convert our physical intensities into the parameter s, for which we need the saturation intensity. Intuition tells us that a MOT should be optimised for $I \approx I_{sat}$, which from our data appears at $I \approx 5\text{mW/cm}^2$. This indicates that the isotropic saturation intensity (3.9mW/cm^2) is more appropriate than the value for the $\langle F = 3, m_F = \pm 3| \rightarrow \langle F = 4, m_F = \pm 4|$ cycling transition (1.6mW/cm^2).

The model is plotted as a function of intensity for the three proposed values of n in figure 6.6, all three curves are normalised to the same peak value. In each case the optimal detuning was used. Here it appears that $n = 6$ shows the most reasonable match to the data overall. In all cases is appears that at lower intensities the models diverges from the data, suggesting that our approximations are no longer valid in the region where the transition is not saturated.

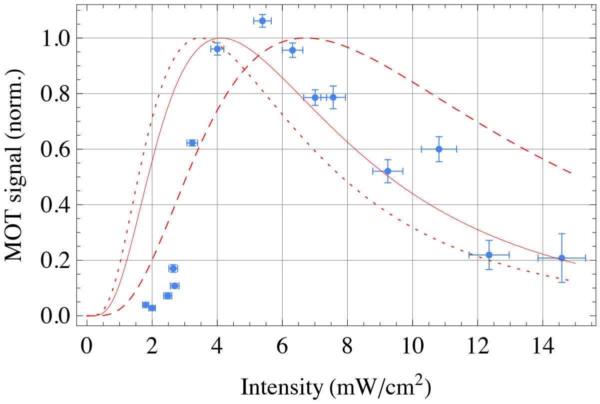

Figure 6.6: Amount of fluorescence observed from the MOT as a function of trapping light intensity. Red lines are theoretical curve with $\delta = 1.5\gamma$ and $n = 3$(dashed), 6(solid) and 9(dotted).

In conclusion it appears that the same model can give a reasonable prediction of the dependence of the MOT fluorescence on both detuning and intensity (at intensities above saturation) if we assume that $n = 6$, i.e effectively six beams contribute to saturation. This value of n seems to indicate that only the Type I beams contribute to the trapping and cooling operation of the MOT, and that the Type III reflections only introduce a constant perturbation to the position of the MOT. This fact could be verified by repeating the experiments in a silicon pyramid with masked corners to suppress the Type III reflections.

6.2.3 Position of the MOT as a function of trapping light intensity

It was observed that the vertical position of the trapped atoms varied as the intensity of the trapping light was changed. Figure 6.7 shows that the MOT moves approximately linearly with the variation in intensity with a gradient of 23±3μm per mW/cm^2. As discussed previously in section 4.1.2, the imbalance in the light fields created by the 70° mirror geometry causes a displacement of the centre of the MOT. In the coil arrangement used for the test silicon pyramid, due to their rudimentary nature, we cannot be sure of the exact position of the magnetic field centre. If equation (4.3) is differentiated with respect to I, we can see that it predicts that displacement is also linearly related to intensity

$$\frac{\partial z}{\partial I} = \frac{nw\gamma^2\hbar}{8I_s\delta\mu\frac{\partial B}{\partial z}} \tag{6.4}$$

where as previously explored we assume $n = 6$. If equation (6.4) is evaluated for the laser parameters used we obtain 18μm per mW/cm^2, in reasonable agreement with the measured value.

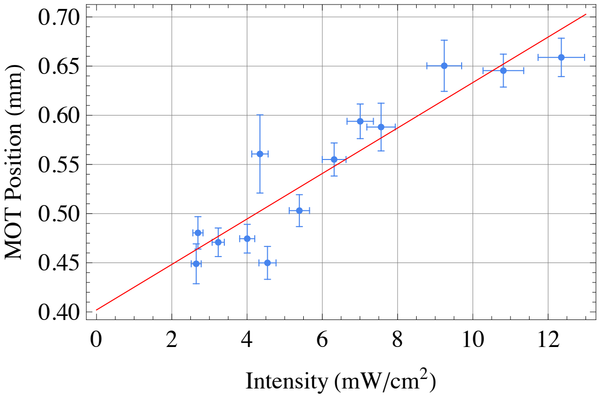

Figure 6.7: Position of the MOT as a function of trapping light intensity, the linear fit has gradient 23±3μm/(mW/cm^2). Error bars are 10% of the MOT radius.

6.2.4 Repump Intensity

Figure 6.8 shows how the MOT fluorescence varies as the amount of repump light is increased. It is apparent that intensities in excess of 0.75mW/cm^2 are sufficient for optimal MOT operation. The repump laser pumps atoms in the dark F=2 ground state back to the F=3 ground state and hence into the cycling transition. As a result the only requirement is that when an atom decays into the dark state, that it always sees enough light to repump it before it leaves the trapping region. It is preferable to limit the repump intensity, as surplus repump light only adds to the scattered light making imaging more difficult.

Initial conclusions

In conclusion the behaviour of the MOT as a function of the trapping light parameters has been explored and it has been found that the signal from the MOT is optimised for a detuning of

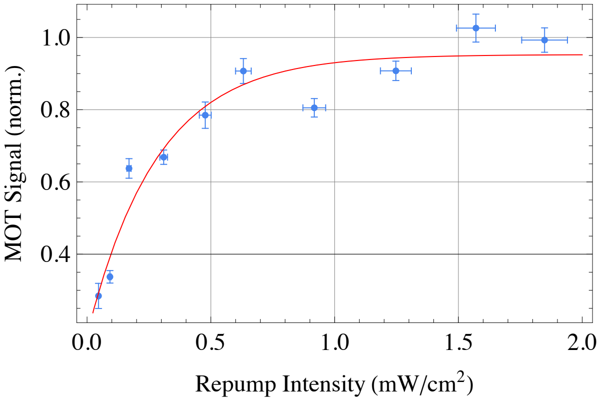

Figure 6.8: Amount of fluorescence observed from the MOT as a function of repump light intensity.

10MHz, a trapping laser intensity of 5mW/cm^2 and a repump laser intensity of 0.75mW/cm^2. The behaviour of the MOT from the variation in trap light parameters is found to correspond reasonably well to a model deriving from a basic Doppler theory capture argument, if we assume that the saturation parameter n takes on a value close to 6.

6.3 MOT loading properties

The signal from the MOT as a function of time following the turn on of the trap is shown in figure 6.9. This was achieved in the lab by synchronising a series of camera trigger pulses (up to 500 images can be taken at up to 1kHz) to coincide with the TTL signal which turns the MOT coils on. A number of integrations of this sequence (typically 100) were required to obtain data of the quality shown in figure 6.9. This figure shows that the loading of the MOT corresponds well to equation (2.36) for a value of $\Gamma = 5.46 \pm 0.28s^{-1}$ which is plotted as a red line.

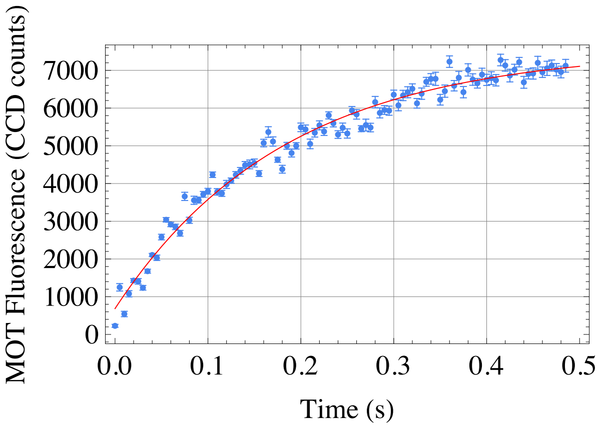

Figure 6.9: Typical data showing the amount of fluorescence observed from the MOT as a function of time after the fields are turned on. Red line is a fit of $N(t) = N_0(1 - e^{-\Gamma t})$ giving $\Gamma = 5.46 \pm 0.28s^{-1}$.

We would like to know what the loss processes are in our silicon pyramid, so we use these loading curves, which as previously covered in section 2.5.4, allow us to measure the sum of the one-body loss rates (Γ) from the loading time of the MOT. We expect that if the silicon pyramids do not introduce significant additional loss processes, the trap loss will be dominated by collisions with hot rubidium atoms. As a result we record the MOT loss rate as a function of rubidium density in the chamber. Similarly, we record the maximum number of atoms in the trap for a given rubidium pressure and we can determine the capture rate of atoms into the trap.

The output of the rubidium dispensers as the supply current is increased was calibrated by measuring the attenuation of a resonant probe beam passing through the chamber, and at the same time loading curves were recorded for MOTs in the silicon pyramid. Figure 6.10 shows the behaviour of the loading curve as the rubidium density is increased.

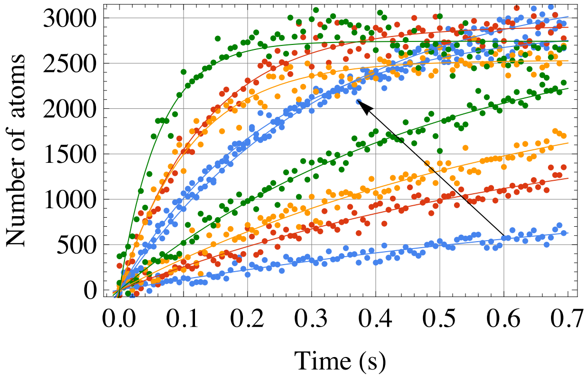

Figure 6.10: Loading curves for the MOT with increasing rubidium pressures (arrow indicates increasing pressures).

Looking at the figure we see as the Rb density increases, the number of trapped atoms initially increases (eventually saturating), but the time taken for the MOT to approach its maximum population decreases indicating an increase in loss rate as the pressure rises. By fitting equation (2.36) to each observation the loss rate, and the steady state number can be extracted from images of fully loaded MOTs. These two numbers can be used to derive the capture and loss rates as a function of rubidium density, and these values are shown in figure 6.11.

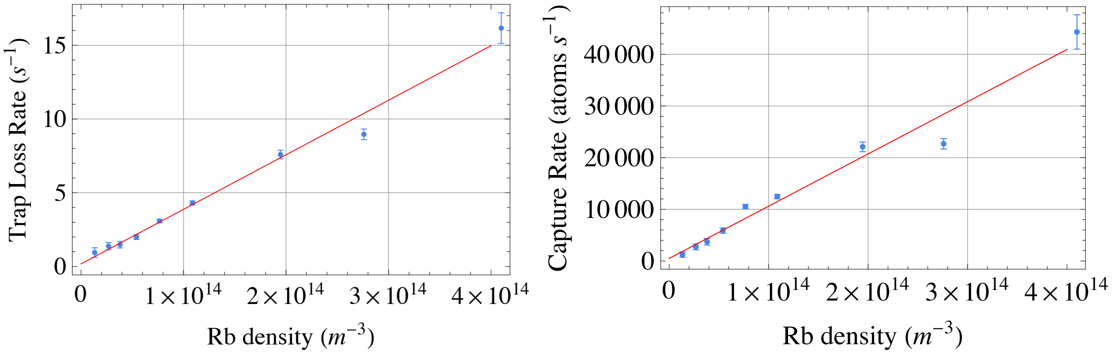

Figure 6.11: Left: Loss rate as a function of rubidium density. Right: Capture rate as a function of rubidium density.

Both show a good linear correlation to rubidium density and from this we extract factors α_l, β_l and γ_l, defined in equation (4.15), from the linear fits. A fit to the data shown in the left graph in figure 6.11 gives

$$\Gamma_{tot} = (3.7 \pm 0.1) \times 10^{-14} n_{Rb} + (0.2 \pm 0.1), \tag{6.5}$$

giving $\beta_l = (3.7 \pm 0.1) \times 10^{-14}\text{m}^3\text{s}^{-1}$ and $\gamma_l = (0.2 \pm 0.1)\text{s}^{-1}$. The value of γ_l is consistent with the base pressure of our chamber (indicating a lifetime of around 5 seconds) and shows that in this 3mm deep pyramid that the proximity of the walls does not increase the loss rate of atoms from the trap.

The coefficient $\beta = \sigma_{Rb-Rb} v_{Rb}$ is the product of the cross-section for Rb-Rb collisions (for ground/excited state collisions) and the rms velocity of hot untrapped atoms. Assuming that the hot untrapped atoms have thermalised with the chamber they will have an rms velocity of 286ms^{-1}. This implies that $\sigma_{Rb-Rb} = (13.0 \pm 0.3) \times 10^{-13}\text{cm}^2$, a value which is approximately 3 times higher than the values quoted for a free space MOT of 3.3×10^{-13} [97]. We attribute the higher value to a decrease in the escape velocity of atoms from the trap due to the proximity of the walls, which was originally proposed in section 4.8.1.

A linear fit to the right hand graph in figure 6.11 gives the loading rate coefficient $\alpha_l = (1.01 \pm 0.05) \times 10^{-10}\text{m}^3\text{s}^{-1}$. Using equation (2.31), this gives an effective capture velocity of $\sim 4.0\text{ms}^{-1}$. Assuming the atoms are stopped as they reach the MOT (so stopping in 1.5mm) this gives an overall average acceleration of $540g$, a tenth of the maximum possible laser-cooling deceleration for an ^{85}Rb atom.

6.4 Habitable zone

It is interesting to study the volume of the pyramid that is accessible to the trapped atoms. At each point inside the pyramid the number of atoms will be determined by the rate of atom capture at that position, and the proximity of the four walls which could cause atom loss from the MOT. It is interesting to study the size of this 'habitable zone' (i.e. the area where a MOT can survive) in order to gain further insight into how we anticipate the behaviour of these pyramids to perform on smaller scales, but also to consider the process of extracting the atoms which may require moving them closer to a surface for magnetic trapping. Firstly we consider motion in the vertical direction (referred to as the axial direction), and then the x-y plane (referred to as the radial direction).

6.4.1 Axial Direction

The behaviour of the MOT in the axial direction was tested by moving the central position of the MOT vertically using external magnetic bias fields.

The first stage is to determine accurately the position of the atoms in the pyramid. Where possible the position of the MOT on the CCD image is used to give an exact physical measurement of displacement. In some observations the reflections of the MOT are often the only visible objects thus requiring some simple geometrical transformations to resolve the position of the actual MOT. One critical factor required is the angle of observation of the pyramid and this is determined below.

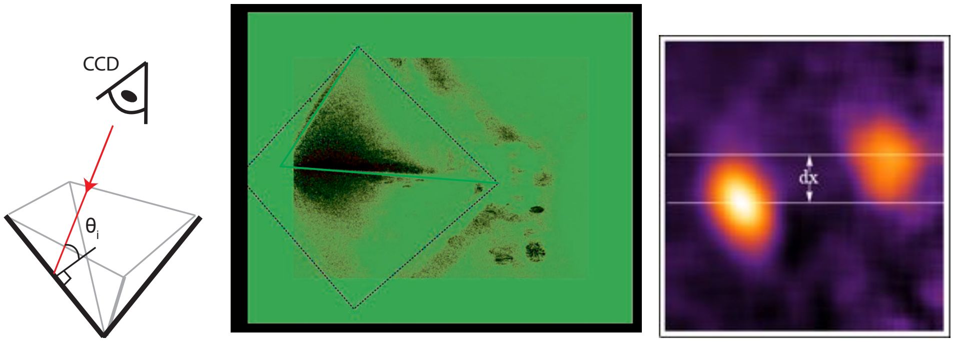

Figure 6.12: Left:Schematic of arrangement of CCD and pyramid indicating angle of interest. Centre: Image of silicon pyramid as seen by CCD. Dotted white line shows pyramid aperture and green lines show edge of faces. Right: Image of MOT (right) and MOT reflection (left) in pyramid face, orientated with the opening of the pyramid to the bottom, and the apex of the pyramid to the top.

Determining the angle of viewing Using the image in the right of figure 6.12, and by measuring the separation of the MOT to the reflection image we can infer θ_i, the angle the camera's optical axis to the diagonal edge of the pyramid. The separation along the axis of interest (vertical in the image, which corresponds to the axial direction in the pyramid) is measured to be $210 \pm 10\mu$m, giving an angle of incidence of $\theta_i = 7.7 \pm 2°$. This value corresponds well with an estimate of the same angle to be $\approx 6°$ from measurements made of the projection of the square pyramid opening onto the CCD plane (see the central figure in 6.12).

6.4.2 Displacement of the MOT with applied bias

With this geometric conversion factor we could obtain data on displacements as a function of applied bias field, for a variety of trap gradients (defined by the current in the coils), and this is displayed in figure 6.13. In order to judge the absolute position of the MOT relative to the apex (which was not able to be included in the field of view), two lines were extrapolated along the edge of the face (shown as green lines in the central image in figure 6.12), and the point of convergence defined to be the apex.

On the left of figure 6.13 are shown graphs of MOT position as a function of applied field for a number of different quadrupole coil currents. As one would expect, the displacement is linear with applied field, and linear fits to these points tell us what the gradient of the trap is. The range of data shown is the full range of visibility for the MOT, beyond the extremities of each graph, no trapped atoms were detected for the integration time chosen.

The graph on the right of figure 6.13 shows the field gradient $\frac{\partial B}{\partial z}$ derived from each linear fit as a function of the current in the MOT coils, and shows a linear trend of

$$\frac{dB}{dz} = (3.55 \pm 0.25)I + (-0.17 \pm 0.15). \tag{6.6}$$

The gradient of this fit, (3.55 ± 0.25)(G/mm)/A, is slightly below the theoretical model of the MOT coils which predicted a peak gradient of 4.2G/mm/A.

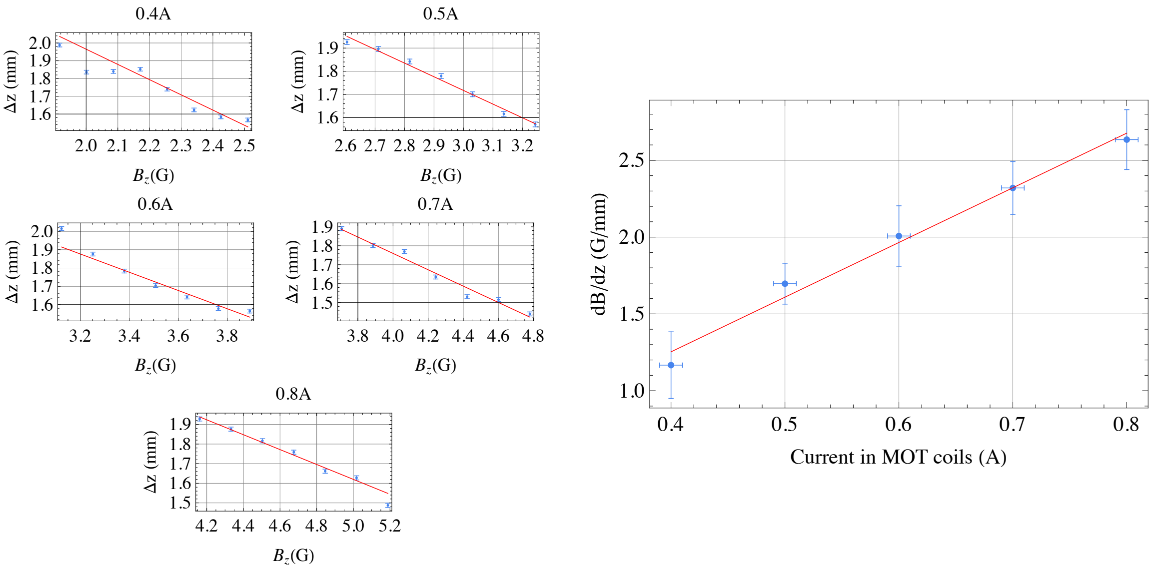

Figure 6.13: Left: Position as a function of B_z for a variety of gradients Right: Observed field gradients vs. predicted gradients.

Optimal Position

The number of atoms trapped in the MOT does not remain constant over the range traversed in the axial direction. The images on the left of figure 6.14 show the MOT fluorescence changing as it is moved in the axial direction, it initially increases to a peak, before decreasing and disappearing. This is due to the fact that as the cloud approaches the apex, the capture rate decreases because the stopping distance is reduced. At the top end of the range, the unmatched transverse beams cause loss of atoms from the MOT.

As a result the number of atoms we observe will be a competition between the increased capture rate as we move closer to the pyramid opening and an increased loss rate as we get closer to the top of the trapping region (which is at a distance $\frac{3}{4\sqrt{2}}L$ from the apex). We have previous studied the dependence on loss-rate to a 'wall' in section 4.8.3, and use approximation that the loss rate to the top of the trapping region can be approximated by equation (4.18). We know that the rate of atom capture with pyramid size scales as L^6, and hence we assume that approaching the apex is akin to trapping atoms in a smaller pyramid. Thus we make the prediction that the number of atoms as a function of distance from the apex can be given by

$$N(z) = \frac{R(z)}{\Gamma_{bg} + \Gamma_{wall}(z)} \propto \frac{z^6}{\Gamma_{bg} + \frac{\overline{v}e^{-\frac{u^2}{2}}}{\sqrt{8\pi}\sigma(1+\mathrm{Erf}(u))}} \tag{6.7}$$

where u is $\left(\frac{z-\frac{3}{4\sqrt{2}}L}{\sigma}\right)$, z is the axial distance to the apex and Γ_{bg} is the loss rate due to collisions with hot background atoms.

On the right of figure 6.14, the blue points show how the fluorescence varies as a function of distance from the apex for the case of the pictured MOT. The red points show translation over the same physical distance but at a lower field gradient. Equation (6.7) is plotted for the condition that the trapped atoms are at the Doppler temperature i.e. $\overline{v} = 0.2\text{ms}^{-1}$. For each case σ and $\Gamma_{bg} = 1$ were left as free parameters and adjusted to match each data set. It was observed that a value of $\sigma = 100\mu$m shows reasonable correlation with the observed trends, and this value is consistent with the measured value of σ. Similarly the model was optimised for a value of $\Gamma_{bg} = 3$, which is close to the value measured in the section above. This suggests that our approximate model is a good match for the experimental observations. We define the boundaries of the habitable region as the points where the atom number drops by one order of magnitude from its maximum value, which from our model we can estimate to give a total traversable distance of 0.8mm. The smaller field gradient, represented by the red points, has a slightly more prominent roll-off as it approaches the apex. This data set does not fit the model particularly well, even if other parameters are varied and hence we assume that in this case the MOT maybe slightly off-centre leading to a premature collision with a wall.

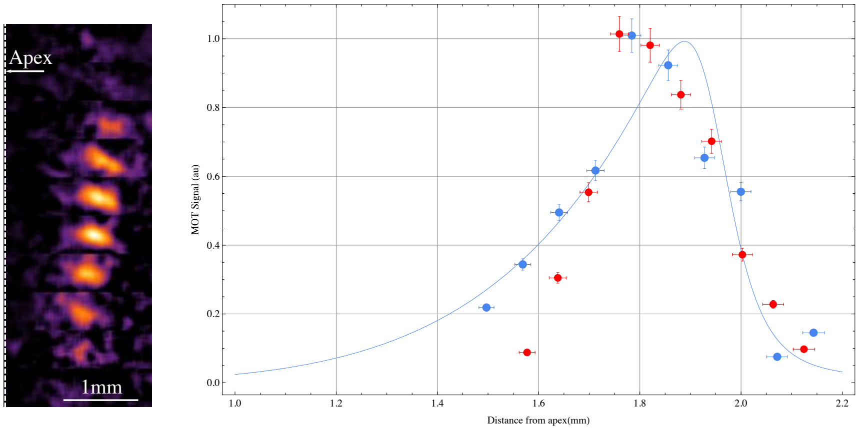

Figure 6.14: N vs. z for two currents (Red=500mA in MOT coils (18G/cm), Blue=700mA in MOT coils (25G/cm)). The solid line is equation (6.7) for the condition that the atoms are at the Doppler temperature and $\sigma = 100\mu$m

From figure 6.14, it is also clear that there is a bias field where the capture process is optimised for both data sets. This was studied for a range of trapping gradients and the results are shown in figure 6.15. It can be seen that the point of optimal operation is approximately fixed in space as the gradient is varied. This point can be calculated to be 1.84 ± 0.10mm from the apex, which is approximately 0.4mm below the boundary of the first Type I reflections.

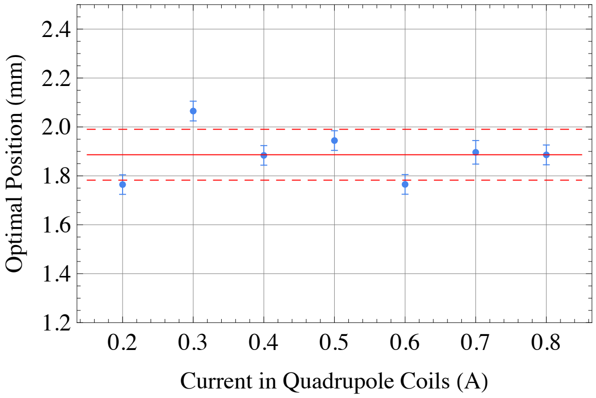

Figure 6.15: Optimal position of the MOT as a function of applied current in the quadrupole coils.

6.4.3 Radial habitable region

Similarly we can study the behaviour of the number of atoms in the MOT as it is translated in the x-y plane at a constant height of 1.9mm above the apex. Images of the MOT were taken for a particular pair of values of $\{B_x, B_y\}$, tracing out a 2D grid. The number of atoms observed at each position is indicated graphically in figure 6.16, brighter colours indicate more atoms were trapped at that point. As expected the most atoms are observed in the centre, and the number of atoms trapped decreases as the MOT approaches the walls which are indicated as a solid white line. The rotation of the axes is due to a slight misorientation of the pyramid to the shim coil axes.

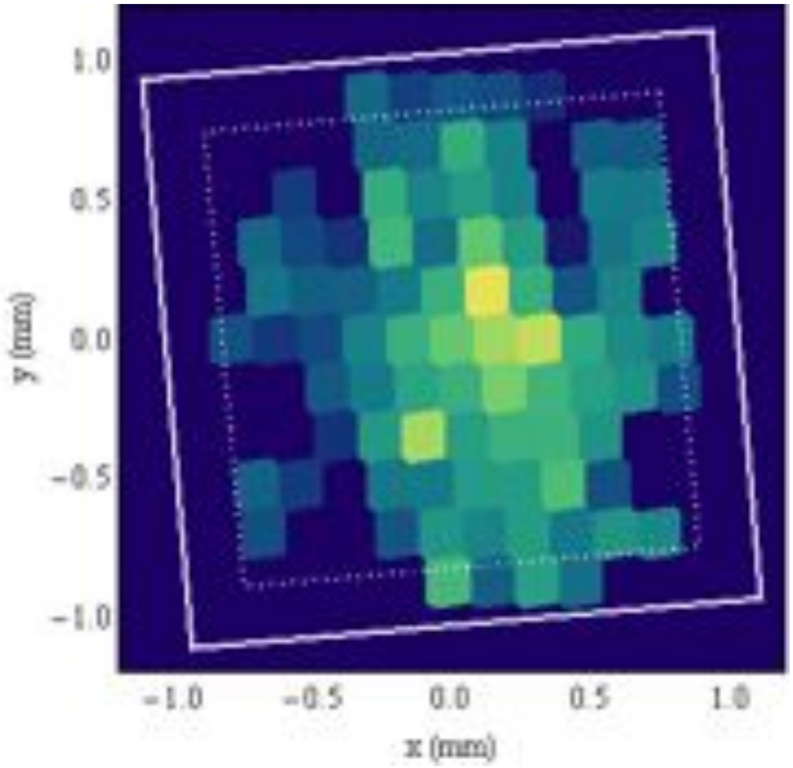

Figure 6.16: MOT population as a function of x and y shim fields. Solid lines are the position of the pyramid walls, while the dotted line indicates the border of the habitable region predicted. More atoms are represented as brighter colours. The magnetic field gradient was approximately 11G/cm in the x and y directions.

We can use the same methodology proposed in the previous section to understand the size of this habitable region. In this case however the rate of atom capture (R) remains constant as it is translated radially and only the losses to the pyramid faces are considered. This gives the number of atoms as a function of the radial position as

$$N(x,y) = \frac{R}{\Gamma_{bg} + \Gamma(x,y)_{walls}} = \frac{1}{\Gamma_{bg} + \sum \Gamma_{wall}(d_i)} \tag{6.8}$$

where $\Gamma_{wall}(d_i)$ is the loss rate (equation (4.18)) from a given surface at distance d_i. This is summed over all four pyramid faces as loss can occur to all faces simultaneously. Using the same parameters as the previous section (assumed Doppler temperature, $\sigma_{x,y} = \sqrt{2}\sigma_z = 140\mu\text{m}$) we again define the boundaries of the habitable region as the point where the peak atom number falls by 10, i.e. the loss rate to all surfaces $\Gamma_{wall} = 10\Gamma_{bg}$ where Γ_{bg} is the background collisional loss. This gives a square region 1.64mm in size, which when plotted as the dotted white line on figure 6.16 appears in good agreement with the region where the MOT signal is visible.

In summary we can observe that MOTs can withstand significant translations across the radial plane in the pyramid, exploring a square region 1.64mm in dimension and can access a volume approximately 0.8mm in height in the axial direction. We can predict the size of this habitable region with some accuracy using a model based upon the capture rate, and loss rate to surfaces as a function of position.

6.5 Temperature

A crucial factor of interest to determine is what temperature the trapped atoms reach when captured and cooled in our silicon MOTs. From the theory discussed in section 2.2.1 we expect the equilibrium temperature to be reached when the rate of heating from spontaneous emission is balanced by the rate of cooling, as detailed in equation (2.12). If we make the assumption that the atoms are scattering at close to saturation, and that we use equation (2.17) to derive α from Doppler theory to be $\alpha = 3 \times 10^{-22}\text{Nm}^{-1}\text{s}$ we obtain an equilibrium Doppler temperature of

$$T_{Doppler} = 550\mu K \tag{6.9}$$

This is higher than the rubidium Doppler limit of 140μK, but in our MOT we have optimised laser parameters for capture rather than cooling. There are a number of ways to experimentally estimate the temperature of the trapped atoms. The most common technique is the ‘time-of flight’ where the atoms are allowed to fall freely. After a delay the spatial distribution is studied, either by absorption imaging onto a CCD or via attenuation of a probe beam. After ballistic expansion of the cloud size for a time t, the radius of the cloud is

$$r(t) = \sqrt{r_0^2 + \alpha^2 t^2} \tag{6.10}$$

where r_0 is the initial size of the cloud and $\alpha = \sqrt{\frac{k_B T}{m}}$ is the radial velocity. Since the micropyramids are mounted upside down, so the apex points upwards, the minimum distance the atoms have to fall before clearing the pyramid opening and coils is about 3mm, which takes approximately 25ms. After this time an atomic sample, initially 100μm in size and at temperature 100μK will be 1.7mm in size, approximately 5000 times less dense than before and thus not easily made visible with any standard detection scheme.

A more suitable technique is therefore ‘release-and-recapture’ in which the trap is switched off for a short period before being turned back on. The fraction of atoms remaining in the trapping region after the period of free evolution indicates the temperature of the sample. We define the recapture region to be a sphere of radius R_{cap} surrounding the MOT centre[98]. After release, the Gaussian distribution of velocities evolves into a Gaussian distribution of positions under the transformation $v \rightarrow r/t$. This gives the fraction of recaptured atoms as

$$\frac{N_{cap}}{N_0} = \frac{4}{\sqrt{\pi} r(t)^3} \int_0^{R_{cap}} r^2 e^{\left[\frac{-r^2}{r(t)^2}\right]} dr \tag{6.11}$$

where $r(t)$ is as described above in equation (6.10) and t is the time the atoms are released for.

6.5.1 Experimental procedure

In order to make the release and recapture measurement, it is necessary to turn the MOT light on and off quickly - in times short compared with 10ms. Unfortunately the optimal solution, optically switching with an Acousto-Optical Modulator was not possible as a suitable device was not available at the time. An alternate solution had to be found.

We discounted detuning the MOT or repump lasers, as they could not be detuned far enough from resonance fast enough to effectively turn off the trapping force. Moreover, the sweep to the far red for the MOT laser would also act as an artificial cooling stage.

We therefore decided on a fast shutter acting at a focus of the trapping beam, where the waist was approximately 20μm. The first device we used was a latching solenoid, where a sprung bolt is alternately driven out by an electromagnet and returned by a combination of the electromagnet in reverse and the release of the spring. The return motion is extremely quick at the midpoint of its motion and the range of travel is on the order of several cm. A thin ($\approx$ 500μm) opaque flag was attached to the solenoid and this was positioned so that it cut the beam at the focus at the midpoint of its travel. By adjusting the position (and size) of the flag towards and away from the focus, transit times of 2ms up to 15ms could be achieved with a reasonably square profile. The solenoid is controlled by a shutter driver which puts out ±30V from a TTL input.

The second solution was attaching a small drill bit ($\approx$ 300μm in diameter) to the outer rim of a wheel, and driving the wheel at a constant rate with a motor. Driving the wheel at around 1 revolution per second the beam was interrupted for durations from 500μs to 2ms, with a full second to allow for the MOT to reload between measurements.

Both methods were used to make a range of measurements for delays of up to 10ms, the wheel was found to be optimal for delays up to 2ms and the solenoid useful to cover the longer delays.

In order to judge the fraction remaining, the fluorescence of the MOT immediately following the recapture was monitored. Using equation (2.36), and using the initial condition that $N(0) = \nu N_0$, the MOT will behave as

$$N(t) = N_0(1 - (1 - \nu)e^{-t\Gamma}) \tag{6.12}$$

Γ and N_0 are known since these can be found by taking a ‘control’ loading curve immediately before or after the release and recapture measurement. It is therefore a simple matter to derive the factor ν by fitting equation (6.12) to the observed data. The experimental sequence is illustrated in figure 6.17. A large degree of variation of the timing is found with the solenoid and the wheel. In the case of the solenoid the time between the triggering pulse and the light being extinguished varies by up to tens of ms. For a one shot measurement this is not a problem, but to obtain a good signal, the measurement had to be repeated up to 100 times

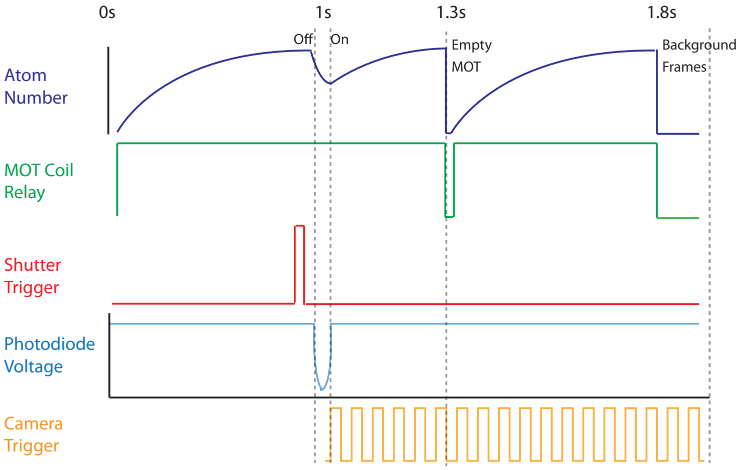

Figure 6.17: Experimental sequence for release and recapture.

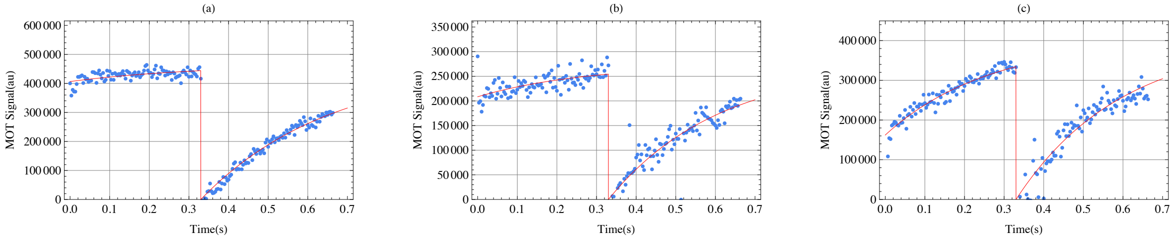

Figure 6.18: Graphs showing the loading process following release and recapture. $t = 0$ is the point at which the trap is restored, and at $t = 0.33$s the trap is fully emptied to measure the loading time. The release times are (a)0.7ms (b)0.5ms (c)1.8ms.

and hence when integrating over this many measurements the signal quickly becomes blurred out and it is tough to produce a reliable measurement of ν.

As a result the main part of the experiment was triggered from a photodiode signal, which monitored the laser intensity. The rising edge of the photodiode (as the light returned) was used as a trigger to start the digital pattern output, thus eliminating any timing drift. As a result of this triggering requirement, the 'control' measurement of the MOT (to determine N_0 and Γ) was made after the release and recapture measurement and not before. The final step in the sequence was to reset all parameters so the entire sequence would restart for the next integration with a fully loaded MOT (allowing for a time delay of 1s between integrations).

6.5.2 Results

Some typical results from the experiment sequence are shown in figure 6.18. They all show a fraction of recaptured atoms remaining in the trapping region immediately following the release, and the loading curve obtained from an empty trap at 0.33s.

Equations (6.12) and (2.36) were fit to both parts of the data set using non-linear regression, and the obtained value of ν used to determine the recaptured fraction. Release times were measured during the experiment by observing the monitor photodiode signal on an oscilloscope and recording both the FWHM and the total extinction time. The results are shown in figure 6.19 with the error bars given by the parameter uncertainties in the fits and the variation of the extinction time as measured on the oscilloscope.

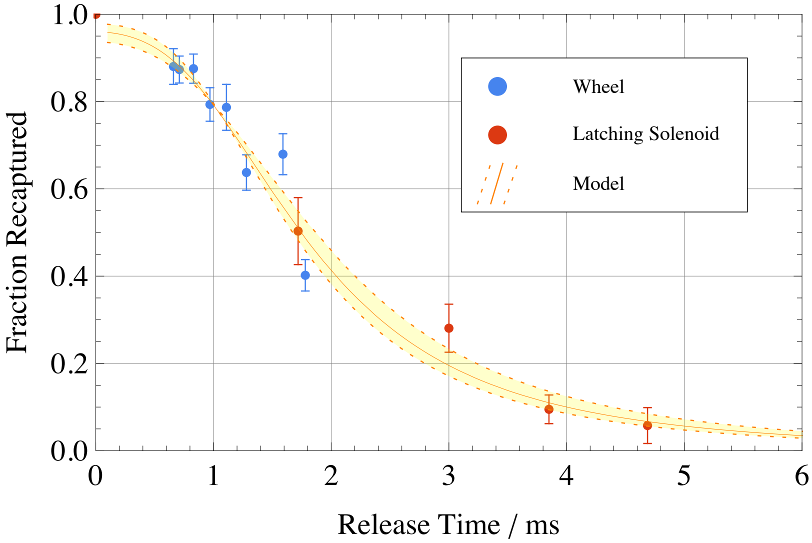

Figure 6.19: Release and Recapture results.

In order to determine a temperature from the data, we use equation (6.11). The original radius, r_0 is determined from CCD images, but the recapture radius of the trap in this case is not known. As a result we attempt to determine optimal R_{cap} and T values simultaneously by a process of Chi-Squared minimisation. Figure 6.20 shows the magnitude of χ^2 from our experimental results for equation (6.11) as a function of T and R_{cap}. It can be found that the minimum value of χ^2 (i.e. best correspondence between the model and the data) occurs for the values of $T = 170\mu$K and $R_{cap} = 290\mu$m. The value of χ^2 for these parameters is 5.42, which gives a probability of accepting the model of 0.94. Indicated on figure 6.20 by the dotted line is the region where $P(\chi^2) > 0.9$, i.e. there is a less than 10% chance that the obtained value of χ^2 arises from a random data set. This region indicates that $127 < T < 200$, and $270 < R_{cap} < 300\mu$m at this 90% significance level. Equation (6.11) is plotted on figure 6.19 as the orange region for the possible range of values of T and R_{cap} and it is clear that there is good correspondance, but does this mean that they are the actual physical values giving rise to the experimental data? Obviously the spherical recapture region model is a simplification of the actual geometry, but the optimal value of R_{cap} is not too dissimilar from the expected value. The closest 'barrier' to the MOT is the top of the region of trapping beams, at a distance of approximately 400μm, and allowing for a stopping distance of 100μm for the atoms to be recaptured (they will all be moving 'outwards' at the point where the trap is restored) the recapture radius of $\approx 300\mu$m seems feasible. As a result it is highly probable that the temperature of the atoms is

$$T = (170 \pm 40)\mu\text{K} \tag{6.13}$$

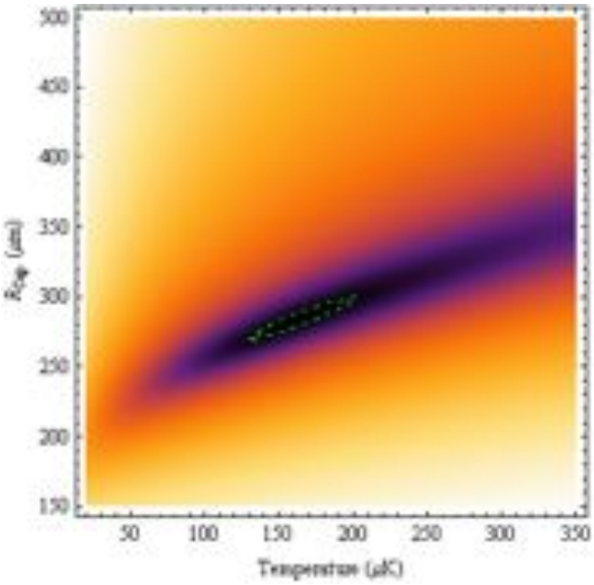

Figure 6.20: Value of χ^2 as a function of T and R_{cap}. Darker colours indicate lower values of χ^2 and the dotted green line indicate values significant at a 90% level.

The switching process used in the experiment will undoubtedly perturb the MOT, most likely due to the fact that the light is not instantaneously switched, but has a transition time of approximately 200μs as the shutter moves into or out of the focus. As a result it is likely that the actual MOT temperature is somewhat below the measured value.

In conclusion the experimental results have allowed us to determine the temperature of atoms trapped in the MOT is $(170 \pm 40)\mu$K, which is significantly below the predicted Doppler limit of 550μK. This indicates that there is the possibility of sub-Doppler cooling mechanisms at work inside our MOT. These results could be improved if the switching process could be made less perturbative, perhaps with an AOM which could in theory switch the light off faster.

6.6 Trap spring constant

The equipartition theorem predicts that due to thermal motion the radius of the distribution of trapped atoms is given by

$$\kappa\sigma^2 = k_B T. \tag{6.14}$$

The above release and recapture experiment gave us value for the temperature and we can use the observed value of $\sigma = 100 \pm 10\mu$m to estimate that the trap spring constant is

$$\kappa = (2.3 \pm 0.7) \times 10^{-19}\mathrm{Nm}^{-1}. \tag{6.15}$$

A comparison can be made to the value predicted by Doppler theory by numerically evaluating equation (2.28) for all the beams and taking the derivative of the force with respect to z,. For the values of $\delta = 1.5\gamma$, $I/I_{sat} = 1.3$, $\frac{\partial B}{\partial z} = 0.3$ and $R = 0.77$ this predicts

$$\kappa_{Dopp} = 4.4 \times 10^{-19}\mathrm{Nm}^{-1}. \tag{6.16}$$

This prediction is slightly larger than the upper bound obtained above, and this value implies that to obtain the observed cloud size the temperature of that atoms would have to be approximately twice times the experimentally determined value. The overestimation of κ by Doppler theory is not unexpected and has been observed in many other experiments, for example Wallace et al. [68].

6.7 Damping Constant

In order to measure the damping constant of the MOT a series of studies were performed where the motion of the atom cloud following an instantaneous displacement was monitored. By observing the form of the motion as the atoms moved back towards the equilibrium point, information relating to the MOT damping constant α and its relative magnitude compared to the spring constant κ could be derived.

Assuming the atoms are displaced by small amount about the trap centre we can make the approximations

$$m\ddot{z}(t) = -\alpha\dot{z}(t) - \kappa z(t) \tag{6.17}$$

At time zero the atoms are released at rest from a non-equilibrium point $z(0) = z_0$ and depending on the relative values of α and κ the atoms will undergo simple harmonic motion with under damped, critical or overdamped motion around the trap centre. It can be shown, for example by Steane et al.[63], that in the overdamped limit the collective motion of the atoms in the trap can be described as an exponential decay to the equilibrium position with decay rate

$$\Gamma_{decay} = \frac{\alpha - \sqrt{\alpha^2 - 4m\kappa}}{2m} \approx \frac{\kappa}{\alpha}. \tag{6.18}$$

In Doppler theory, as the MOT spring constant κ is proportional to the MOT damping constant α this decay rate can be written as

$$\Gamma_{decay}^{Doppler} \approx \frac{\kappa}{\alpha} = \frac{\frac{\partial B}{\partial z}\mu}{\hbar k} = 5310\frac{\partial B}{\partial z}\mathrm{s}^{-1}. \tag{6.19}$$

Hence for typical MOT gradients we expect to see motion on the sub-ms timescale. Throughout this we have neglected the effect of stochastic forces arising from absorption, spontaneous emission and polarisation gradients in order to keep the analysis simple.

6.7.1 Experimental Technique

In order to displace the atoms, similar experiments[99] have used a weak chopped ‘pushing’ beam to displace the atoms from the centre. In the micropyramids there is not the optical access to introduce a beam without it undergoing reflection and potentially traversing the cloud twice. Therefore the cloud is moved by adding or removing a small bias magnetic field to displace the magnetic centre as discussed in Kohns et al[100].

The current used to generate the bias field has a finite response time because of the gain/bandwidth limit of the electronics that control it. This response can be measured by monitoring the voltage across a sense resistor when the current is stepped, as shown in figure 6.21. The response is exponential with a time constant of $t_c = 0.86 \pm 0.02$ms which limits observations to frequencies of atomic motions slower than this.

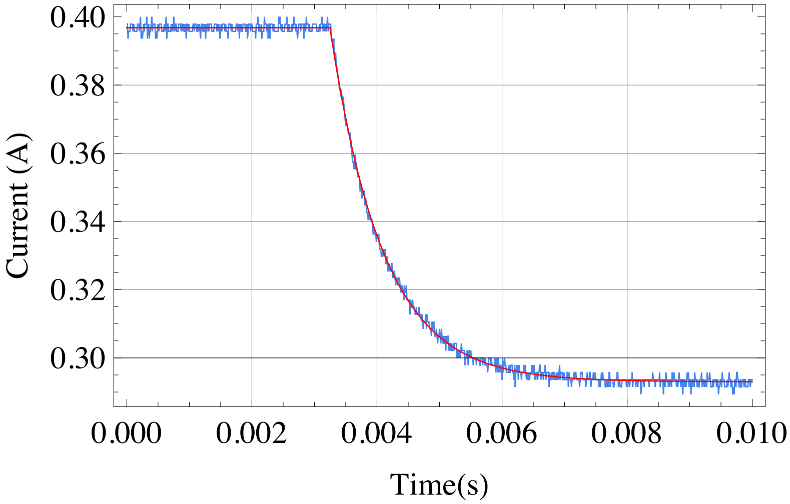

Figure 6.21: Response of current in the shim coils to a step function in the control voltage. The measured response time is 0.86ms. Blue line is measured data and red line is fitted exponential decay model.

6.7.2 Measurement

The MOT was initially loaded at optimal detuning and z position, and then subject to up to 8 translations back and forth, with sufficient time between translations to let the motion fully decay before it was moved again. The camera was run at rates up to 600Hz allowing it to map the motion of the cloud with a good temporal resolution. The entire sequence was repeated up to 100 times to obtain a good SNR.

The results of the experiment are a series of images $I_1, I_2....I_N$, to each a 2D Gaussian is fitted to determine the position of the MOT, $z_1, z_2...z_N$. The camera was run at a constant rate, assigning a time to each image with the centre of the exposure chosen as the observation time. Each experiment contained up to 8 translations of the MOT back and forth, and the data was broken up to look at each translation individually. Equation (6.18) was fit to the MOT motion in each translation using non-linear regression to give a value of Γ_{decay}. All the fitted values of Γ_{decay} for each translation were then averaged to give an overall measure of the decay rate for that experiment. The decay time of the motion was found to be independent of displacement distance for displacements up to 200μm, but typically displacements on the order of 100μm were used to minimise atom number loss.

6.7.3 Initial estimation of the damping constant

Figure 6.22 shows the results from an experiment taken at typical MOT operating parameters. The top part of the figure shows the integrated vertical profile of the image as a function of time, and the lower part shows the position of the MOT centre determined from the 2D Gaussian fit. It can be seen that a decaying exponential model (green line in the lower half

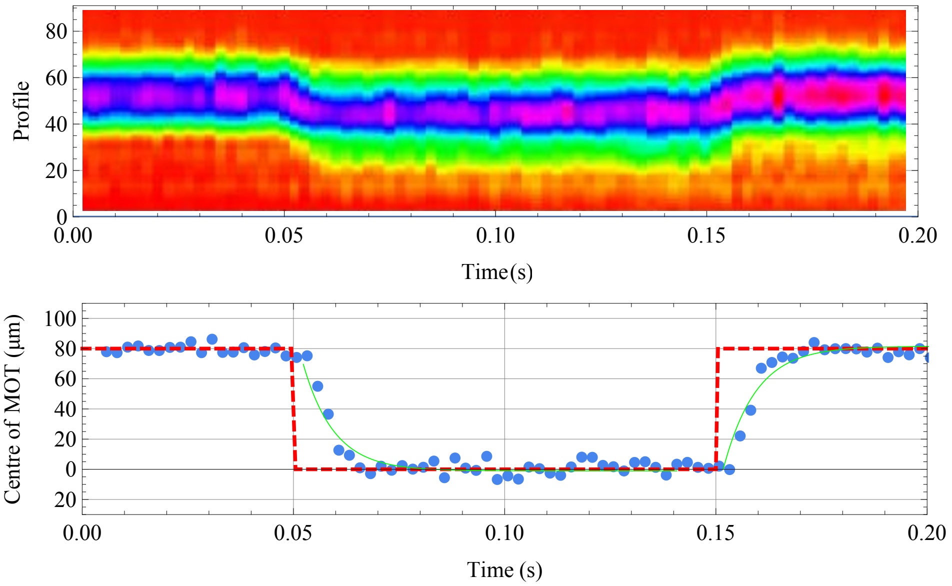

Figure 6.22: Top:A vertical slice taken from each image shows the overdamped motion of the cloud. Bottom: The centre of the MOT as a function of time (blue markers) responding to a shift in field (dashed red line) can be fit using a decaying exponential(green line).

of the figure) gives a good fit to the observed motion following each translation, and we can determine the average decay rate to be

$$\Gamma_{decay} = 138 \pm 7\text{s}^{-1}. \tag{6.20}$$

We expect from equation (6.19) that Doppler theory predicts a decay rate of $\sim 1500\text{s}^{-1}$ at a gradient of 30G/cm, which is a factor of 11 larger than observed. We know from the temperature and size of the MOT that κ is not significantly different from the value predicted from Doppler theory so therefore it looks likely that

$$\alpha \approx 11\alpha_{Doppler} = (3.0 \pm 1.0) \times 10^{-21}\text{Nm}^{-1}\text{s} \tag{6.21}$$

to obtain the observed damping rate. This is consistent with our expectations from the temperature that sub-Doppler processes such as Sisyphus cooling are enhancing the damping constant of the MOT.

6.7.4 Dependence on detuning

Doppler theory predicts that the decay rate will be independent of detuning (see equation (6.18)) and so the decay rate of the atomic motion was studied as a function of detuning in order to test the hypothesis that sub-Doppler processes are contributing to the velocity damping in the MOT. The results are shown in figure 6.23 and show that there is significant dependence of the decay rate on detuning.

The dependence of the damping coefficient arising from Sysiphus cooling is given by

$$\alpha_{Sis} \approx \hbar k^2 \frac{3\delta}{20\gamma}\left(\frac{3\delta^2 + 14\gamma^2}{\delta^2 + \gamma^2}\right) \tag{6.22}$$

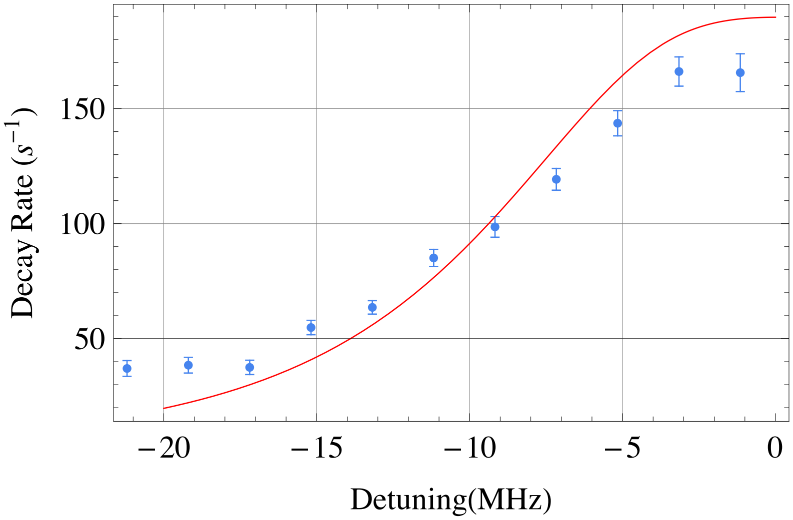

Figure 6.23: Left: Damping time as a function of detuning. Predicted damping rates using Sysiphus (red solid line) sub-Doppler theory is shown.

for three dimensional lin ⊥ lin molasses[101]. This is not the situation in our 70° pyramid MOT, but a useful starting place to make some initial observations. If we assume that the decay rate is given by

$$\Gamma_{decay} = \frac{\alpha_{Sis}}{\kappa_{Doppler}}, \tag{6.23}$$

the decay rate can be derived for a given detuning. This is shown as the red line in figure 6.23 and it can be seen that the Sisyphus mechanism is a good match for the observed magnitude and behaviour of the data.

6.7.5 Dependence on Intensity

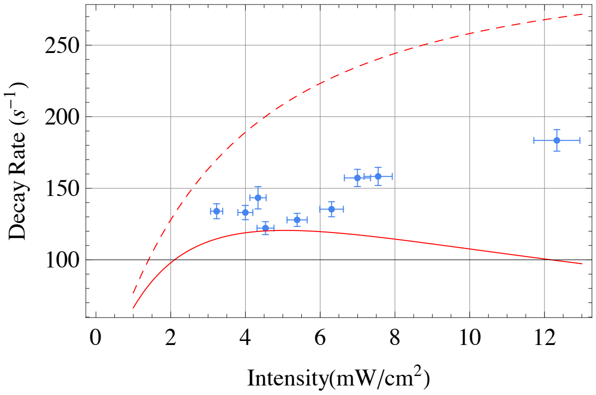

Figure 6.24: Damping time as a function of intensity. Solid line is predictions from Sysiphus cooling, and dashed line is modified Sysiphus cooling to take account of line broadening.

As shown in figure 6.24, variation of the decay time is also observed as the intensity varies. The theory used for α_{Sis} is independent of intensity and thus only the variation of κ should affect the decay rate. It is observed though that in this case equation (6.23), plotted as the solid line in figure 6.24, is a poor reproduction of the trend observed. The increase in the decay time observed in the data is not matched by the predictions, and the proposition is made that the efficiency of the sub-Doppler cooling reduces as the intensity increases (which would lead to an increase in the decay time as required). One potential mechanism for the reduction of α_{Sis} could be line-broadening, and indeed if we make the transformation $\gamma \rightarrow \gamma\sqrt{1+s}$ to account for this in equation (6.22) we obtain the red dotted line in figure 6.24 which shows an increase in the damping time as the intensity is reduced (albeit at the expense of overestimating the magnitude of the damping rate).

6.7.6 Decoupling α and κ

One final point of interest was to try and isolate the effect of κ on the damping process, whereupon the damping rate was measured for a variety of magnetic field gradients. From both Doppler and sub-Doppler theory α should be invariant with gradient and hence should allow estimation of the relationship between the two trap parameters. Results are shown in figure 6.25 and interestingly show little variation of the damping rate with the gradient. We can reject the hypothesis that κ does not depend on gradient, as the size of the cloud has been observed to vary in the expected way so we can only assume that α also shows a dependence on magnetic field gradient. The invariance of the damping rate with gradient was also observed in work by Hertog [102].

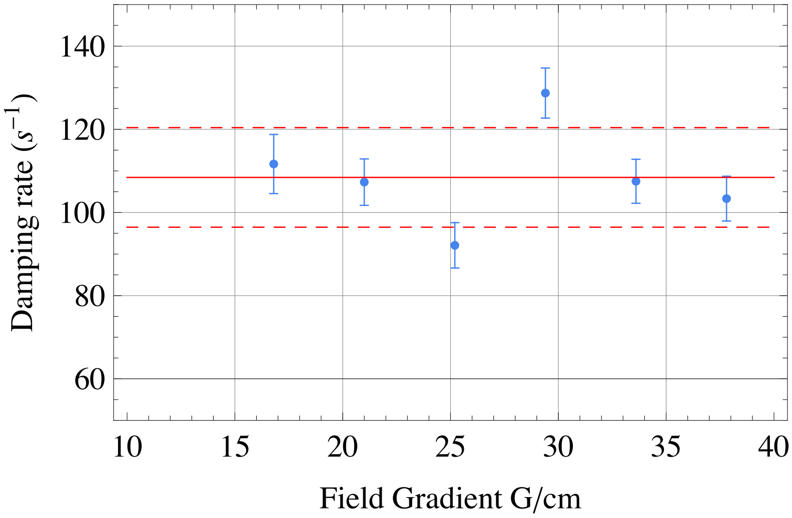

Figure 6.25: Damping rate as a function of magnetic field gradient.

In conclusion we have shown some evidence for the existence of sub-Doppler damping effects at work in the magneto optical trap. There are however a number of anomalous results which could benefit from further studies and it is proposed that in future rather than displacing the MOT with a step function, the amplitude and phase of the MOT motion in response to a sinusoidal modulation of the bias fields, as demonstrated in work by Kohns. et al [100], could yield more information about the behaviour of the atoms inside the trap. It is worth mentioning here that the damping constant we have measured is the force experienced by cold atoms near the centre of the trap. The effective damping constant felt by atoms being captured from the background thermal vapour is likely to be closer to the value predicted by Doppler theory.

6.8 Smaller Silicon Pyramids

With the extensive knowledge gained about the properties of trapped atoms in a 4.2mm wide (3.0mm deep) pyramid, the next logical step is to take the study to ever smaller pyramids. The etched wafers featured a three by three array of pyramids, ranging in pyramid size from side lengths of 3.5 to 2.5mm (depths of 2.5 to 1.75mm) and this array was processed with ICP in a similar manner to the previous pyramid and sputter coated with aluminium. In order to test the array it was mounted in the package illustrated in figure 6.26. This features a PEEK chip carrier in which the silicon device sits, and an array of copper current carrying structures. The wires were created on a wire-eroder from a 0.7mm thick sheet of copper, and are typically 0.75mm wide. Three sets of wires form the grid structure on the top allowing creation of a quadrupole inside each pyramid with an appropriate external vertical bias field. The middle wire is mirrored by an identical set on the back face so the central diagonal row can have quadrupole fields created without the external bias field. Beryllium-Copper barrels are used to make the connections from the copper wires to wires connected to the feedthroughs.

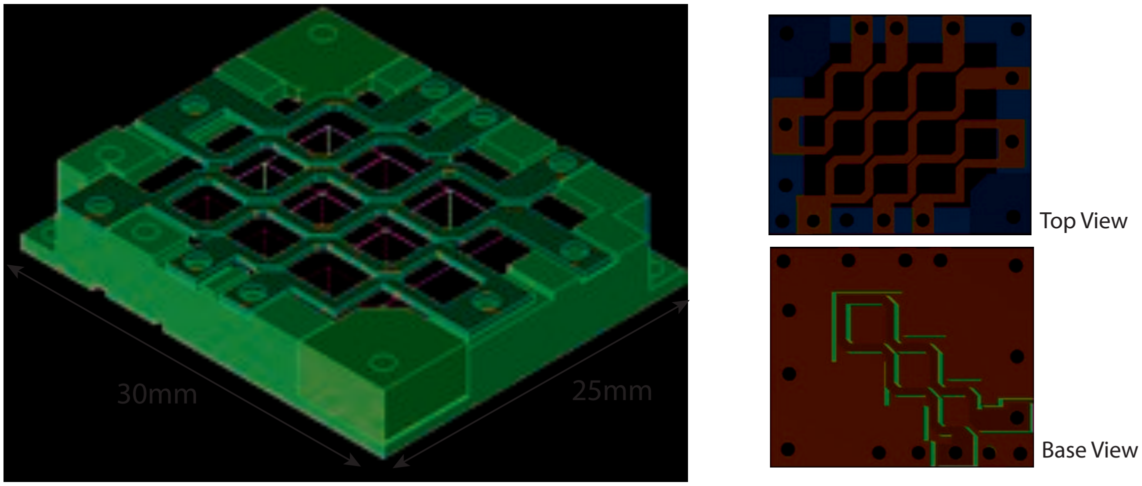

Figure 6.26: Mount for array of silicon pyramids

Fields

The fields generated by the copper structures were modelled in Radia, and to test the accuracy of this model a Hall probe was used to map out the fields above the central structure on the chip package. The probe was fixed at a distance of 1.6 ± 0.2mm above the surface of the chip package, and the package was translated on an x-y translation stage which offered μm precision in translation. Comparison between the model and the measured fields are shown in figure 6.27 and show good consistency within the specified bounds of uncertainty.

It was found that we could run currents of up to 4A continuously through these structures under UHV conditions. Above this current, undesirable levels of outgassing were observed, and it is not known whether the outgassing arises from wires making the connections to the feedthroughs or overheating of the package itself.

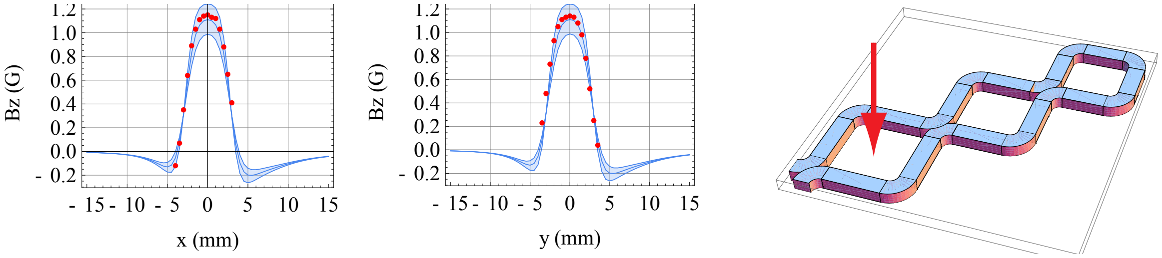

Figure 6.27: Fields generated by the central copper structure on the pyramid array chip package when running 1A. Measured fields are plotted in red, and calculated fields are in blue. A range of calculated values are shown (shaded region) to take account of the ±0.2mm uncertainty in the vertical position of the probe when making measurements.

6.8.1 3.5mm pyramid

Using the pair of central copper structures above and below the central diagonal column of pyramids, a current of 3A generated a field gradient of 25.5G/cm at the centre of the pair of coils (which corresponded to a point 0.4mm above the apex of the 3.5mm pyramid). Without any need for any supplementary bias field a MOT was observed in this 3.5mm pyramid, containing up to 1800 atoms in a MOT.

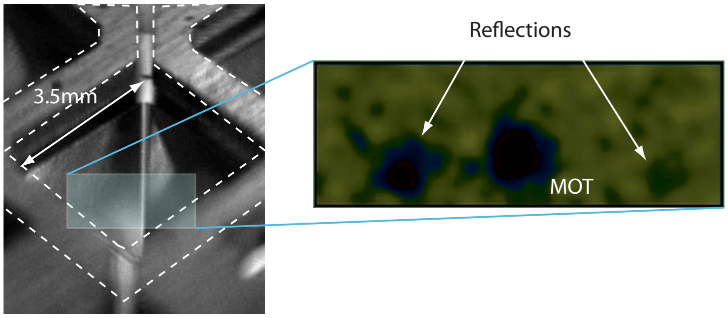

Figure 6.28: A MOT in the 3.5mm wide pyramid.

6.8.2 Size and density

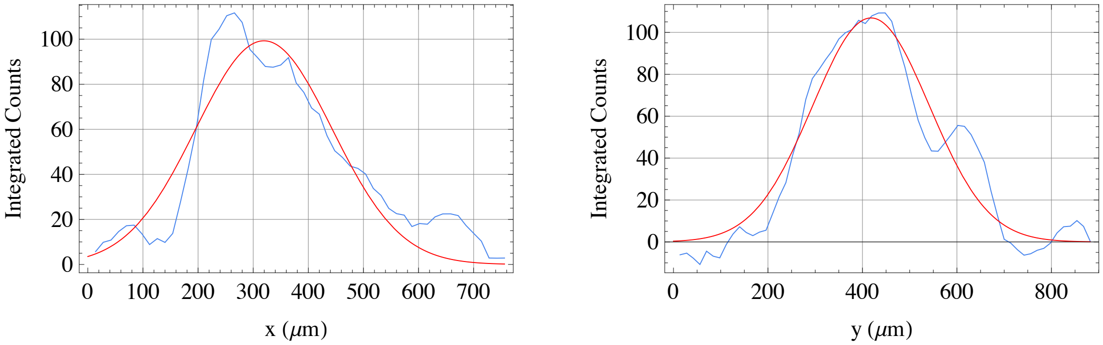

Figure 6.29: Horizontal and Vertical Profiles of the MOT in the 3.5mm pyramid.

The physical dimensions of the MOT were estimated to be $123 \pm 6\mu$m by $124 \pm 4\mu$m (figure 6.29), which is similar to that in the 4.2mm pyramid. The angle of viewing suggests that we are almost observing the MOT from top down hence viewing primarily the radial axes. The peak density of $1.5 \times 10^{-8}\text{cm}^{-3}$ is similar to the previous measurement. Unlike the larger pyramid, we were able to find a viewing angle which made the MOT more prominent than the reflections and this is illustrated in figure 6.28.

Position and habitable region

Unlike the coil arrangement used for the previous pyramid, the position of the magnetic field minimum relative to the pyramid is well known in this setup. We observe MOTs, for a given integration time, over a vertical shim range of -0.75G to 0.50G, which at the gradient used corresponds to displacements of the magnetic field from 0.15 to 0.65mm above the apex. The actual position of the MOT can be determined from images in the same way as discussed previously (see section 6.4.1) and we find a consistent offset of 0.8 ± 0.1mm between the position of the field zero and the centre of the MOT. This is consistent with the predictions and observations made in the large glass pyramid.

In the larger 4.2mm pyramid we could traverse approximately 0.80mm, and the MOT sat at an optimal height of 1.8mm (see section 6.4.2), in the smaller pyramid the optimal position is approximately 1.40mm from the apex and can translate approximately 0.5mm which is slightly smaller proportionally than that the 4.2mm pyramid.

Similarly in the 4.2mm pyramid, the MOT could traverse a region 1.64mm in size radially (see section 6.4.3), indicating the region should be 1.37mm in size in the 3.5mm pyramid. We observe MOTs forming over field ranges corresponding to displacements of 0.91 ± 0.15mm in the x direction and 0.75 ± 0.12mm in the y direction which is significantly smaller. This is probably due to the fact that in the 3.5mm pyramid the MOT sits proportionally lower than in the 4.2mm pyramid.

6.8.3 3mm Pyramid

The next size down was a pyramid 3mm in side length. A signal was observed, at a statistically significant level of SNR=2.5, and this signal also appeared at the expected value of vertical shim field and detuning. This signal however was difficult to reproduce with any fidelity, as it seemed to show an extremely reduced range of operation in the axial direction and unless laser parameters were well optimised, the signal was swamped by the scatter. The poor reproducibility of this signal was attributed to instability in the experiment, and a number of tests were performed to try and isolate the factor which was most detrimental to stable trapping in the smaller pyramid. Drifts in the magnetic field were investigated, and found to be insignificant, with variation in the ambient magnetic field less than 100mG. Low levels of noise and vibration were found to be crucial for low-noise image acquisition, as well as noise adversely affecting stability of laser frequency. Attempts were made to acoustically damp the experiment using thick curtains around the table, but even this is not sufficient during working hours in a lab with two other experiments (and the associated personnel). Laser intensity was found to fluctuate significantly (up to 10%), and hence the optics were rebuilt (to the configuration now detailed in 3.8) for improved intensity stability but even with all these improvements we were not able to obtain a sustainable signal in the 3mm pyramid. It is most probable that we are coming up against the limit hypothesised in section 5.6.4 from the high scatter levels due to the ICP etching. This predicted the threshold under optimal imaging conditions to be 200 atoms. The 3.5mm pyramid traps numbers of atoms only one order of magnitude above this lower level, and hence it is entirely possible that we may have reached the limit of our detection with this schema.

As a result of these difficulties we were not able to systematically optimise this signal and hence we associate the observation of 600 atoms with a significant error bar to reflect the difficulty in resolving it against the scatter.

6.9 Scaling law

We can combine the observations of optimal atom numbers made in the 4.2, 3.5 and 3mm silicon pyramids and combine these with the observations made in the glass pyramid. The combined results are shown in figure 6.30.

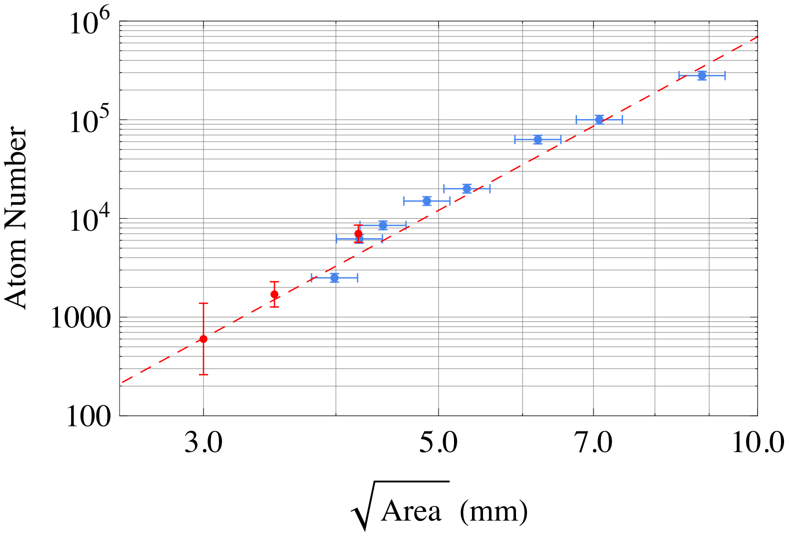

Figure 6.30: Data showing scaling of atom number with size of pyramid from the silicon pyramids and the glass pyramid.

It is apparent to see that the scaling trend which begins in the glass pyramid at scales below 10mm continues across the range of silicon pyramids. We can henceforth exclude any imperfections in the glass pyramid as a mechanism for reduction in atom number. If the data from the two pyramid types is combined, a power law can be determined from simple regression which gives an exponent of 5.74 ± 0.25, and this is shown as the red dotted line in the figure.

Based on the observation of this consistent scaling law, we can extrapolate the trend to predict that a 1mm wide pyramid will trap only a single atom from the background vapour. This adds further weight to the prediction that we would be unable to detect atoms in the 1.2mm pyramid in the original prototype chip with our current imaging scheme without a form of assisted loading from a precooled source.

6.10 Conclusions

The performance of the trap was studied in the largest 4.2mm pyramid and found to exhibit similar dependencies on laser parameters as the large macroscopic model. The loading dynamics of the MOT were observed, and it was found that the loss rate was dominated by collisions with background untrapped atoms and not losses to the walls. The effective capture velocity was estimated to be 4ms^{-1}.

An exploration of the 'habitable region' for the MOT, the volume that the MOT can inhabit inside the silicon pyramid, was presented. It was found that the MOT could traverse a height of up to 1mm and a region 1.6mm square at the optimal trapping height of 1.9mm from the apex.

Using the release and recapture method, the temperature of the atoms in the MOT were estimated to be $T = 170 \pm 40\mu$K, however additional work is suggested to try and verify this measurement.

Dynamics of the MOT were studied when the cloud was subject to an impulse, and this showed damping of the motion potentially arising from the Sysiphus sub-Doppler cooling mechanism.

Trapped atoms have also been observed in 3, and 3.5 wide silicon pyramids. The number of atoms trapped in these pyramids is found to scale consistently with the L^6 scaling predicted in the macroscopic model. The struggle to detect a reproducible signal in the 3mm pyramid suggests that the current limit to detecting trapped atoms in these particular pyramids, and raises issues to be addressed in order to proceed to smaller pyramids and fewer numbers of atoms. These issues, and proposals for solutions are presented in the final chapter.

Chapter 7

Theory of MOT properties on a small scale

7.1 Assembling the evidence so far

In order to simplify the MOT dynamics, we must consider what approximations are reasonable to make. Firstly we know the exponent of the scaling law does not vary significantly whether the surfaces are gold or aluminium coated, and whether the Type 3 reflections are present or suppressed by masking. Secondly a similar scaling consistent with the observed exponent of L^6 has been observed by Vangeleyn et al.[61] in a different MOT geometry, suggesting that the scaling is independent of the beam geometry. Finally the exponent of the scaling arises from a reduction in capture rate, suggesting that an L^2 component arises from the reduction in capture area, and a relation of $v_c \propto L$ of the size to the capture velocity.

It is known that 'large' MOTs scale as $L^{3.6}$, and thus a question of note is at what scale does the transition between the two regimes of scaling occur? We made predictions in section 2.5.1 suggesting that the transition occurred at a velocity of $v_c = \delta/k$ but turning this into a physical stopping distance requires experimental data. In figure 7.1, a data set from the glass pyramid has been used to estimate how the exponent of the scaling varies across the size range explored. A power law was fitted to sets of 5 data points at a time, and this gives a rough indication of how the scaling is proceeding at that size.

It is apparent that amongst the scatter in the points, that above 8mm the trend is approaching the expected $L^{3.6}$. Below this scale the exponent varies between 5 and 7, which is consistent with the expected exponent of 6. From this we can make the proposition that the average capture velocity in an 8mm pyramid is approximately $v_c = \delta/k = 8\text{ms}^{-1}$ (for $\delta = 10\text{MHz}$). Assuming a linear scaling, this would imply that the capture velocity in our 4.2mm pyramid should be $\sim 4\text{ms}^{-1}$ which was exactly the value we measured experimentally in section 6.3.

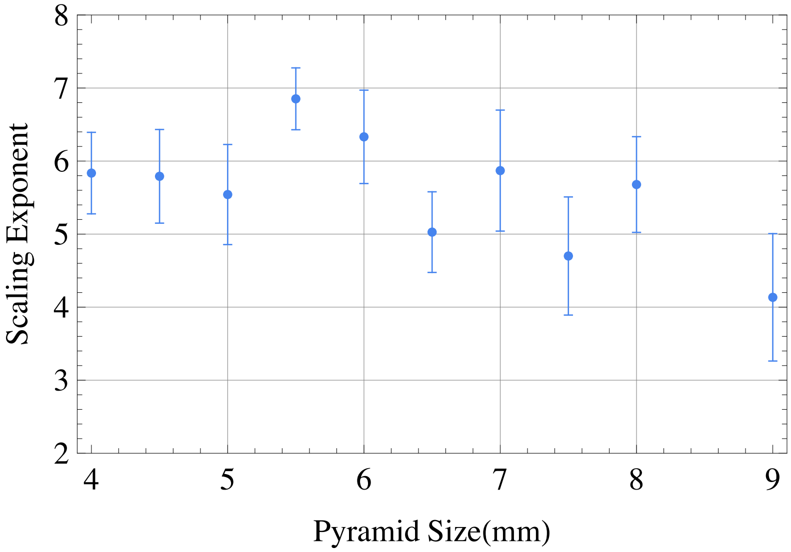

Figure 7.1: Exponent of scaling as a function of size.

7.2 Numerical Simulation

In order to more fully explore the capture properties of the pyramid MOT in 3 dimensions, a full numeric simulation was created, considering capture of atoms from a thermal vapour. It was also beneficial to include the effects of non-orthogonal beam geometries in the 70° geometry.

7.2.1 Methodology

We consider that scattering occurs independently from each beam, in an identical manner to that presented in equation (2.28). Each beam is defined with the properties

- Wavevector - a 3 dimensional vector giving the direction of propagation of the beam.
- Intensity - A scalar giving the intensity of each beam relative to the input beam as determined by the reflectivity of the mirrors.
- Profile - A function determining the intensity profile of the beam. In most models this is just implemented as a top hat function to define the region of space where the beam is present.
- Helicity - The phase angle between the s and p components of the beam.

The force arising from each beam is calculated, considering the relative probability of the three possible transitions σ_+,σ_- and π. The forces arising from all beams are summed to give the total force on the atom, with the saturation parameter in the denominator of equation (2.28) ensuring the atom never scatters at a rate above $\gamma/2$.

The approach chosen has already led to good results in work done by a number of previous researchers [64][103] and hence it is expected to be a reasonable approximation of the capture process. As in most of these previous works, our model does not take account of the heating process from spontaneous emission, and does not consider the fact that the effective magnetic moment of the cooling and trapping transition varies according to the hyperfine sublevel occupation. It also does not consider the effect of optical dipole forces, or the polarisation gradients that lead to sub-Doppler effects. As a result it is not expected to reflect the intricacies of the behaviour of the atoms at the centre of the trap.

Simulation Implementation

The simulation was produced in Mathematica, and makes extensive use of the internal numerical differential equation solving routines. The programme has not been particularly optimised for performance, but will simulate a typical atom capture in approximately 200ms on a reasonably humble PC.

The criterion for trapping was defined to be $v < 0.2\text{ms}^{-1}$ (the Doppler velocity) after an evolution of the trajectory for 1s. In the case of simulations in pyramids, an additional constraint is imposed by collisions with the wall. The position of the atom with respect to the pyramid walls was continually evaluated and if it crossed a wall in the process of being captured then the trajectory was rejected as untrapped.

Monte Carlo or bust?

Two approaches were compared for running the simulation, firstly a systematic study was considered. In this case the simulation would systematically step through all possible positions on the entrance surface, and all possible incident angles. At each point, the initial velocity of the atom would be increased until it was no longer trapped, and the maximum successful velocity recorded for that point. It was found that $\sim$ 3000 points were required for reasonable resolution, and this required about 6 hours to run on a 2GHz PC.

An alternate approach was to consider the use of a Monte-Carlo simulation. In this case, trajectories are picked at random from an appropriate distribution (in our case a Maxwell Boltzmann distribution at 298K) and simulated to see if it was captured or not. By determining the fraction of random incoming atoms that were successfully captured, the capture rate of the MOT could be estimated. The initial sampling distribution was cutoff at a finite velocity $v_{cutoff} \sim 30\text{ms}^{-1}$ in order to ensure that fast atoms that could never be trapped were not needlessly simulated.

By experimentation it was found that 2000 trajectories per simulation, taking approximately half an hour, gave a good trade-off between execution time and estimated precision.

7.2.2 Initial Observations

Results found from preliminary simulations have been illustrated in figure 7.2. The translation of simulation results into physical atom numbers will be discussed in the next section. It can be seen that the scaling at low pyramid sizes is consistent with the expected L^6 scaling, and above $\sim$ 7mm scales as $L^{3.6}$. Above 20mm however the increase in atom number begins to slow down which is not consistent with what we expect. By examining the raw trajectories,

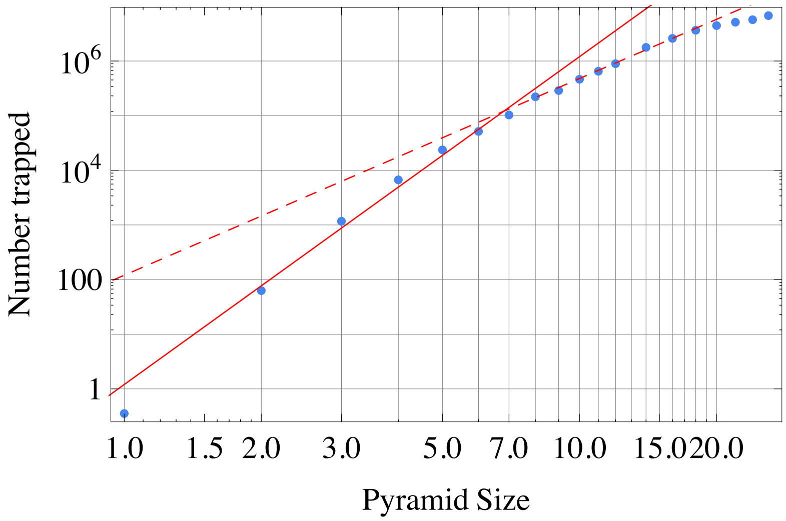

Figure 7.2: Numbers of atoms trapped as a function of pyramid size as determined by simulation. Red solid line is L^6 and red dotted line is $L^{3.6}$.

it becomes clear that atoms are initially accelerated towards the trap centre, which is due to the region of unbalanced beams at the entrance of the pyramid. At large pyramid scales the distance over which the atoms are accelerated by these unbalanced beams is quite significant and leads to a artificial depletion of the low velocity atoms. It is not certain whether this effect would be replicated in real large scale glass pyramids.

7.2.3 Comparison to experimental results and other pyramids

In order to use the simulation results to predict physical atom numbers, we used the equilibrium condition $N(t \to \infty) = R/\Gamma$ (see section 2.5.4). The capture rate is defined in equation (2.31) and in the case of the simulation gives

$$R = \frac{n_{Rb} A \xi v_{\text{cutoff}}^4}{\pi \Gamma \overline{v}^3} \tag{7.1}$$

where v_{cutoff} is the upper velocity limit used in the simulation and ξ is the fraction of simulated trajectories that were successfully trapped. All other parameters are known physical parameters and can be chosen to best match the conditions at the time the experimental measurements were made. Figure 7.3 shows three sets of experimental data, aluminium with and without flower masks, and silicon pyramids. As there is an intrinsic degree of fluctuation in this data arising from variation of background pressure and rubidium density we thus define a range of typical parameters

$$\Gamma = 2 \pm 1\text{s}^{-1} \qquad\qquad n_{Rb} = (8 \pm 2) \times 10^{13}\text{m}^{-3}. \tag{7.2}$$

These values give the simulation results, shown as a shaded orange region in figure 7.3. This region indicates that the simulation is a good match for the data and hence it seems to capture the essential physics needed to explain the scaling process.

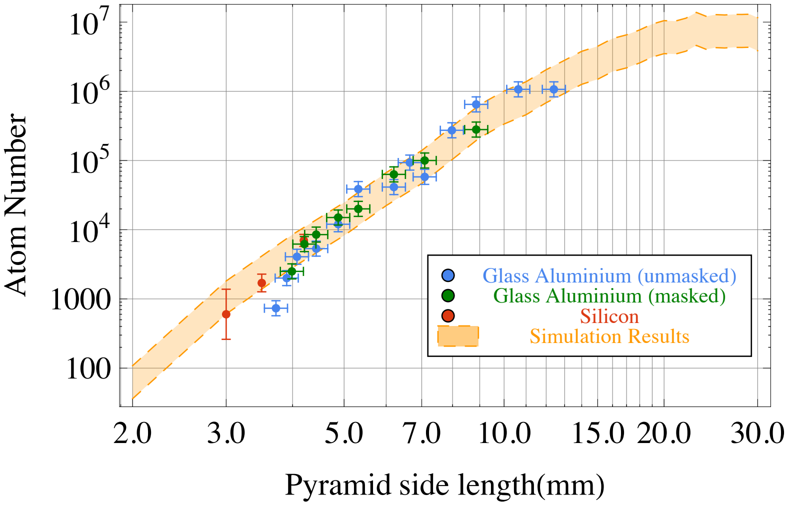

Figure 7.3: Simulation results (yellow banded region) compared to observed values

Comparison to 90 degrees

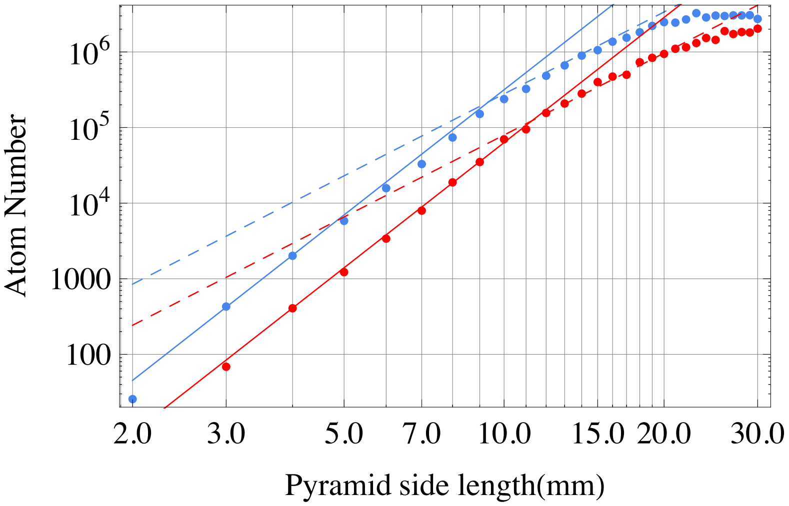

Figure 7.4: Comparison between a 70° (blue) and 90°(red) simulated pyramid. Lines are a L^6(solid) and $L^{3.6}$(dashed) fit.

One big question that has haunted this project throughout has been whether 90° pyramids could be microfabricated, and if so, would they offer more favourable scaling at the mm-scales. The research referred to earlier in a tetrahedral MOT[61] suggests not, but the flexibility of our model allows us to create a 90° pyramid and observe its behaviour. Results from a 90° simulation, with a comparison to the 70° results, are shown in figure 7.4. This indicates that not only does the scaling appear to be the same for both geometries, but for a direct size to size comparison, the 70° geometry appears to outperform the 90° one. This not unexpected as the 70° is deeper by a factor of $\sqrt{2}$, hence increasing the available stopping distances.

7.3 A deeper look at the simulation results

More information can be obtained from the simulations in the 70° pyramid by considering the trajectories of atoms that have been captured. We choose to consider this in terms of the component axial velocities (v_z) and the radial velocities ($v_r = \sqrt{v_x^2 + v_y^2}$).

7.3.1 Axial and radial capture velocities

The maximum and mean axial and radial velocities trapped are plotted as a function of the pyramid size in figure 7.5. They are both initially linear with size, before rolling off above the transition point.

The maximum radial velocities are slightly higher than the maximum axial velocities. This may indicate either an asymmetry in the cooling forces in the two directions or a difference in stopping distances.

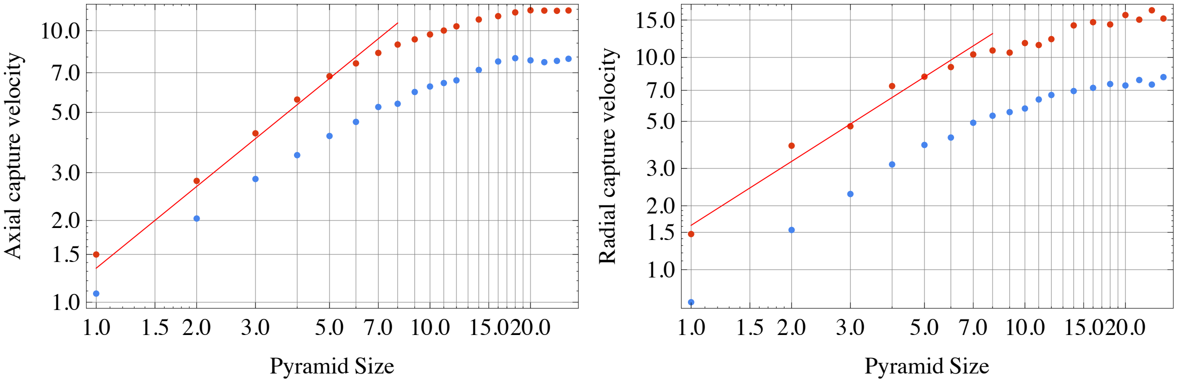

Figure 7.5: Left: Mean (blue) and Maximum (red) axial capture velocities as a function of pyramid size. Right: Mean (blue) and Maximum (red) radial capture velocities as a function of pyramid size.

The dependence of capture velocity on initial trajectory can be considered by studying the mean axial and radial velocities at each point on the pyramid opening as illustrated in figure 7.6. Unsurprisingly, atoms have the highest axial capture velocity when incident close to the centre where the depth is maximised. The radial capture velocity is largest from atoms entering from the corners, again consistent with trajectories maximising the available stopping distance.

7.3.2 Approximating the simulation with an analytic solution

To summarise the findings so far, we know that a full 3D simulation shows that for the small scale pyramids the velocities trapped scale linearly with size as expected. It would be useful if we could use this knowledge to produce a simple relation between pyramid size and the capture velocity, leading to a prediction of capture rate.

We choose to consider the overall picture by looking at the radial velocities. This decision can be justified by considering the flux of incoming atoms through the pyramid aperture, as we never actually observe an atom come in directly on axis (due to a $\sin\theta\cos\theta$ distribution of polar angles[104]). As a result more atoms enter at trajectories closer to 45° than normal to

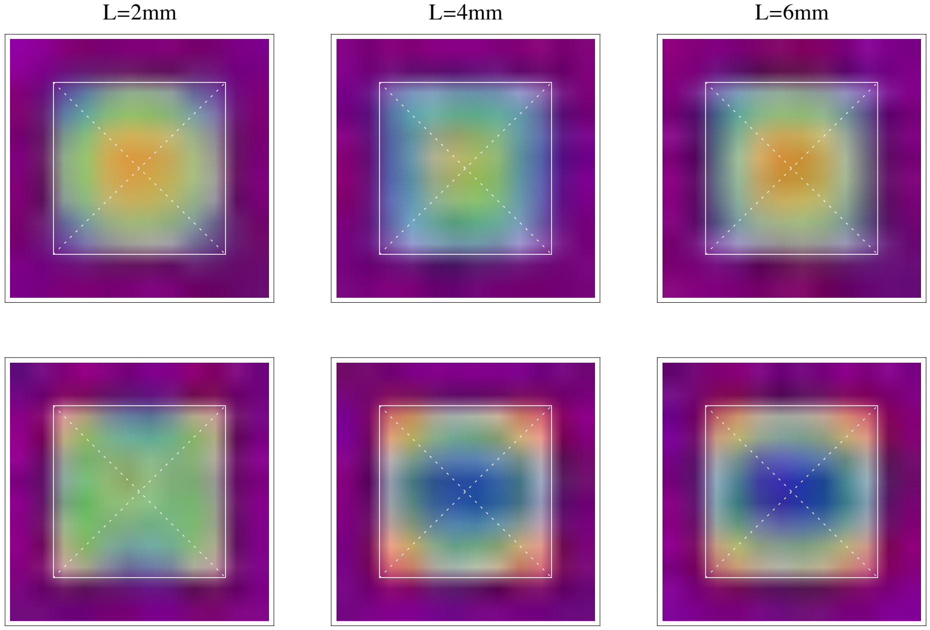

Figure 7.6: Mean velocities (Top row is axial & bottom row is radial) as a function of position on the pyramid aperture, which is shown as the white square. Red colours are higher velocities.

the entrance aperture and hence it is predicted that radial cooling forces dominate the capture process.

If the line of best fit from figure 7.5 for the region $1 < L < 6$mm is obtained we get the following relation

$$v_r^{max} = (1.62 \pm 0.07) \times 10^3 \text{ms}^{-1} L. \tag{7.3}$$

We know that the capture velocity, as a function of stopping distance, can be approximated by equation (2.33) , allowing us to estimate the effective damping constant α. Whichever component we choose, the maximum stopping distance is $\sim L$. Combining this maximum stopping distance with our knowledge of the fastest atoms captured from equation (7.3) we obtain

$$\alpha = \frac{m v_r^{max}}{L} \approx (2.3 \pm 0.1) \times 10^{-22} \text{Nm}^{-1}\text{s}. \tag{7.4}$$

which is consistent with the value obtained if we explicitly calculate α from the force equations in the simulation

$$\frac{\partial F_{sim}}{\partial v_x} = 2.5 \times 10^{-22} \text{Nm}^{-1}\text{s}. \tag{7.5}$$

Seeing that a reasonable approximation to the capture velocity can be made by considering the damping coefficient and the available stopping distance, the final piece of the puzzle is to ask, what is the stopping distance of an average atom that ends up in our MOT? The available space for deceleration depends on the initial position of the atom on the pyramid aperture, and its initial velocity vector. We look at the average distance each trapped atom would have traveled before hitting a face in the absence of a MOT, and this is plotted as a function of pyramid size in 7.7.

This figure shows a linear trend between pyramid size and stopping distance given by

$$\overline{d} = (0.63 \pm 0.01) L. \tag{7.6}$$

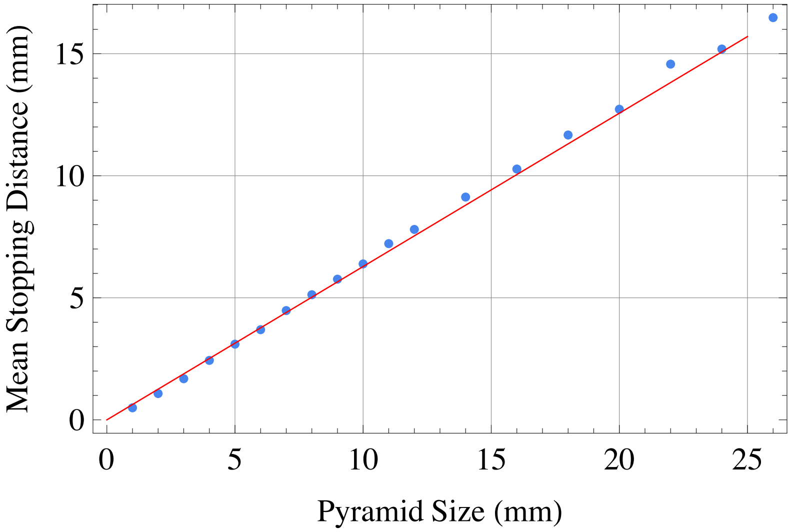

Figure 7.7: Mean stopping distance as a function of pyramid size.

This tells us that in the simulation, the typical stopping distance is roughly two-thirds of the side length of the pyramid. Therefore, the capture velocity for a pyramid of size L, averaged across the entire pyramid volume is given by

$$v_{cap} = \bar{d}\frac{\alpha}{m} \approx \frac{2\alpha L}{3m} \tag{7.7}$$

7.3.3 Assembling all the pieces together

Given we now have the capture velocity, it is a simple matter to derive the total capture rate in a pyramid. Recalling that the capture rate is given by

$$R = \frac{n_{Rb}A}{\pi}\left(\frac{v_c^4}{\bar{v}^3}\right), \tag{7.8}$$

we can insert our prediction for v_c and the pyramid area $A = L^2$ to obtain

$$R(L) = \frac{n_{Rb}\ L^2}{\pi\bar{v}^3}\left(\bar{d}\frac{\alpha}{m}\right)^4 = \frac{0.16\ n_{Rb}\ \alpha^4}{\pi\bar{v}^3\ m^4}\ L^6. \tag{7.9}$$

The best way to validate this is to use the capture rate experimentally measured in a 4.2mm pyramid in section 6.3, where it was observed to be

$$R_{exp} = \alpha_l n_{Rb} = (1.01 \pm 0.05) \times 10^{-10}\mathrm{m}^3\mathrm{s}^{-1}\ n_{Rb}. \tag{7.10}$$

If evaluate expression (7.9) for the case of $L = 4.2$mm and the values listed in section 7.2.3 we obtain

$$R_{sim} = (1.89 \pm 0.34) \times 10^{-10}\mathrm{m}^3\mathrm{s}^{-1}\ n_{Rb}, \tag{7.11}$$

which is a good correspondence, and suggests that the extrapolation we have made from our model is not unjustified in its assumptions.

7.4 Conclusions

We have presented a 3D model of the capture process by simulating the trajectories of incoming atoms when interacting with the complex light arrangements making up the 70° pyramid MOT. We find that the capture process scales for different pyramid sizes in a manner identical to that observed experimentally. This scaling is found to be independent of beam geometry and instead solely a consequence of the reduced stopping distance.

By studying the results of the simulation a theoretical quantitative estimate of the capture rates in small pyramids is made which is in good agreement with the observed empirically obtained values in the silicon pyramids.

Chapter 8

Conclusions & Future Prospects

8.1 Summary of major findings

From studies in a macroscopic glass model of the 70° geometry, it was found that the number of the atoms trapped from the background vapour scales as L^6, where L is the side length of the pyramid. This is expected for a MOT capturing from the velocity regime where the damping force is linear with velocity. Refinements were made to the process of distinguishing the MOT from a background of scattered light. It was estimated that in typical situations as few as 100 atoms could be detected.

The refined imaging techniques were used to place an upper limit on the capture capability of a preexisting set of microfabricated 1.2mm silicon pyramids to 100 atoms, which is consistent with the L^6 scaling. As a result development of larger pyramids (up to several mm in size) was pursued. The anisotropic etching process was investigated and found to be capable of producing mm-scale pyramids, albeit with some residual roughness that appeared to prevent them operating as a MOT. An isotropic ICP etching process was found to reduce the long-range roughness and result in mirrors which worked successfully to produce a MOT.

The performance of the trap was studied in the largest (4.2mm) silicon pyramid and found to exhibit similar dependencies on laser parameters as the large macroscopic model. It was found that the smaller number of atoms in a MOT close to the pyramid walls are due to a reduced capture rate, rather than as a increased loss rate. Indeed, the observed rate of atom loss from the MOT was not dominated by loss to the walls, but by background collisions. The effective capture velocity for the 4.2mm pyramid was estimated to be on the order of 4ms^{-1}.

It was found that the MOT could explore a large proportion of the available volume in the radial direction, but the formation in the axial direction was restricted to a smaller volume than that suggested by the accessible geometry.

The temperature of the atoms in the MOT was estimated to be $170 \pm 40\mu\text{K}$ which indicates the potential presence of sub-Doppler cooling mechanisms. This observation was backed up by measurements of the dynamics of the MOT when subject to an impulse, which showed damping of the motion consistent with the Sysiphus sub-Doppler cooling mechanism.

Further studies were carried out in 3 and 3.5mm wide silicon pyramids. The number of atoms trapped continued to follow the L^6 scaling.

8.2 Immediate improvements for detection

The most obvious area for refinements, particularly if there is a desire to move towards ever smaller scales, is to improve the detectability of atoms trapped in the MOTs.

8.2.1 Reduction of scatter

The current largest impediment to detection is the excessive amount of scatter arising in a ICP etched pyramid. It was predicted that in the 3mm-wide pyramids we would observe 500 atoms, which was on the order of the detection threshold imposed by the amount of light scattered from the ICP-processed faces. Pyramids smoothed with a combination of ICP and HNA (section 5.4.2) are yet to be tested with atoms, but it has been observed that the scattered intensity is about ten times less than in ICP only pyramids. This suggests that MOTs in pyramids as small as 2.0mm, trapping an expected 50 atoms, could be detected with certainty. As a result it is proposed that further refinement of the mirror postprocessing, and optimisation of the coating deposition could offer great improvements on the performance of these devices.

8.2.2 Blue light imaging

An alternative way to improve the detection is to separate the MOT fluorescence from the background scattered light by changing the frequency of the fluorescence. Rubidium has a transition from the $5P_{3/2}$ state to the $5D_{5/2}$ state, which is at 776nm and hence within the reach of the diodes we use to produce 780nm light for the D2 transition. As shown in figure 8.1(a), if an atom is pumped to the $5D_{5/2}$ state, it can decay via the $6P_{3/2}$ state emitting a blue (420nm) photon. The relevant branching ratios dictate that an atom promoted to the $5D_{5/2}$ state will return to the $5S_{1/2}$ state via the 420nm emission route 1 in every 9 times. By using this two-step excitation process, and an appropriate spectral filter on the CCD to eliminate 780nm light we should be able to observe atomic fluorescence with no background. Although the rate of 420nm emission will be less than the rate of 780nm emission, elimination of the background means that the reduction in noise should more than compensate for the loss in signal.

This principle was initially tested in a vapour cell, heated to increase the rubidium density, where laser beams tuned to the $5S_{1/2} \rightarrow 5P_{3/2}$ and $5P_{3/2} \rightarrow 5D_{5/2}$ transitions counter-propagated through the cell. Blue fluorescence was easily visible with the naked eye (see figure 8.1(b)), and the emergence of the blue fluorescence was accompanied by a drop in the 780nm absorption within the cell (monitored with a photodiode). A bandpass filter was used to select only the blue fluorescence (Comar Optics bandpass filter 445GB25, transmitting at 445±50nm), which has a density of 7.5 (i.e. $10^{-7.5}$ power reduction) at 780nm. The cell was imaged with the CCD through the filter, and 420nm fluorescence appeared much weaker than the 780nm fluorescence (by a factor of approximately 1000). Long exposure times and high gains had to

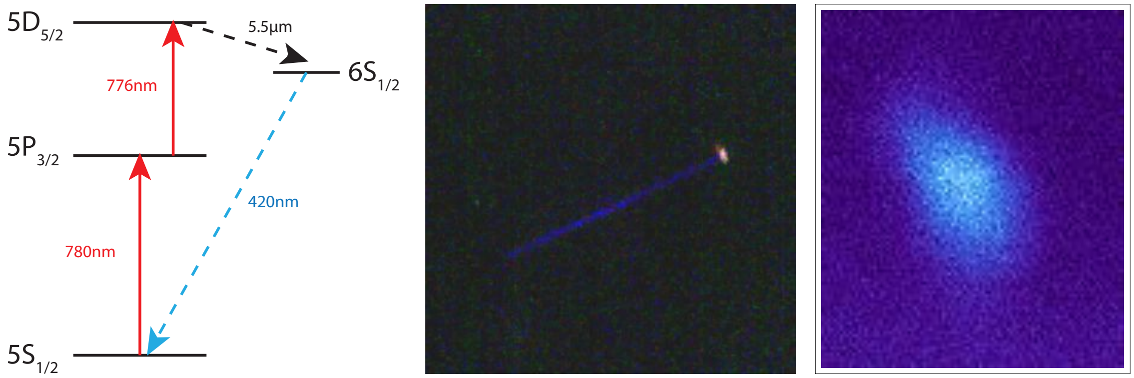

(a) Relevant Energy levels

(b) Blue fluorescence in a vapour cell (colour image)

(c) Blue fluorescence in a MOT (false colour image)

Figure 8.1: Blue fluorescence

be used to resolve it against a background of CCD noise (introduced with the high gain). This seems to be weaker than we would expect, but is consistent with the observations made by Sheludko et al.[105].

The 776nm laser ($P = 1.7$mW) was overlapped with one of the MOT trapping beams and sent inside the vacuum chamber. The atomic cloud of approximately 10^8 ^{87}Rb trapped atoms was monitored with the CCD camera pointing through a vacuum chamber viewport. We observed blue fluorescence in the MOT (see figure 8.1(c)) peak for a similar detuning as used to maximise blue fluorescence in the cell. The yield of blue light is still low, and this is most likely due to the limited power available from the diode used to generate the 776nm light. Using this method to detect small MOTs in the micropyramids is suggested as a highly promising route for further investigation.

8.3 Enhancement of capture

For many applications it is undesirable to have small numbers of atoms. The applicability of the pyramids could therefore be improved by somehow improving the number of atoms trapped in them.

8.3.1 Apexless pyramids

A benefit of the 70° geometry is that the apex region does not contribute to the reflections that intersect the trapping region, and hence if this region is absent a MOT can still form. We could use this property for a number of extensions to the silicon pyramid concept.

Removing the apex

The size of the hole at the apex is given by $a = w - \sqrt{2}d$, where d is the thickness of the wafer and w is the width of the square aperture used to define the pyramid opening. Removing the apex region is equivalent to the problem of the apex defect (see figure 4.20, where we are restricted by the requirement for the second reflection of the trapping light to balance the incoming light. It was previously determined that a defect of width a reduces the vertical trapping region by $1.3a$. With a total trapping height of $3/4d$ in a full pyramid, this means the usable volume in a truncated pyramid is approximately $3/4d - 1.3(w - \sqrt{2}d)$. This limits the size of the hole to a maximum of approximately $a < 1.7d$. Assuming a minimum distance of 300μm from the centre of the MOT to the ends of the trapping region, giving a height of 600μm, this limits the size of the opening aperture for a 3mm wafer to $w < 5.6$mm.

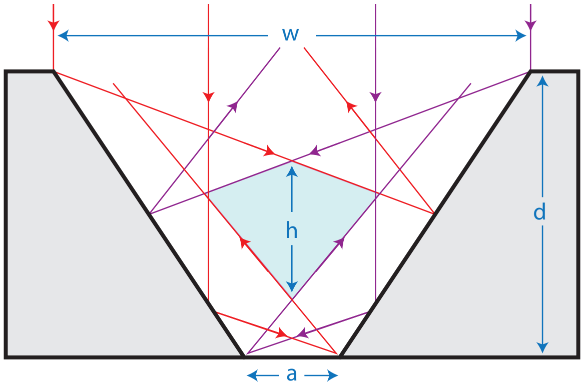

Figure 8.2: Diagram showing how the accessible volume for trapping reduces as the apex hole increases in size

Adding beams

This limit can be circumvented by added a vertical beam coming back through the hole at the apex of the pyramid. This means there is no longer any restriction on the size of the hole, as the vertical component of the trapping and cooling force is completed by the beam coming through the hole. The transverse confinement is still provided by the first reflection of the incoming light. The counterpropagating beam can be easily obtained by means of using a quarter wave plate with an metallic coated backside placed behind the pyramid. This ensures the light is retroreflected with the correct polarisation for the MOT. It is predicted that this beam counterpropagating with the incoming beam will act to increase the effective stopping distance of the trap, hence enhancing the capture rate.

Modified Mount for apexless pyramids

In order to test pyramids fabricated without an apex, a mount was designed and constructed in order to accommodate a half silvered quarter wave plate to allow retroreflection of the beam coming through the apex. This mount also included coils of a small radius in order for high fields and high trapping gradients. The design, illustrated in figure 8.3, is designed for a 5.3mm pyramid with an approx 1mm hole at the apex. The coils produce a field gradient of 50G/cm per amp applied. The holder is constructed from PEEK, and the coil formers are made from aluminium. The quarter wave plate is a standard multilayer waveplate, and it had the rear face coated with aluminium to retroreflect the beam. When initially testing this device we noticed anomalously high levels of outgassing, and it was suspected that the adhesive used to create the multi-layer structure of the waveplate itself is not particularly compatible with UHV. Investigations into this are ongoing.

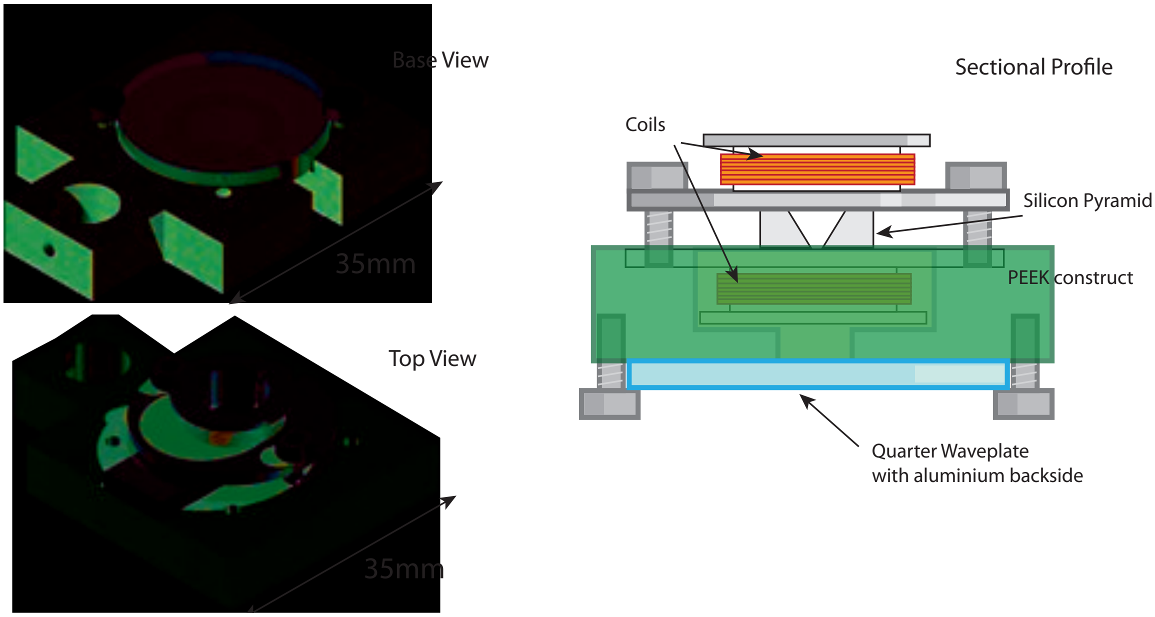

Figure 8.3: Mount for apexless silicon pyramid

8.3.2 Back to Back Pyramids

It is entirely possible to anisotropically etch pyramids on both sides of a silicon wafer, assuming that wafers can be sourced that are polished on both sides. If the masks are chosen appropriately then it is feasible to envision a 'back-to-back' structure with both pyramids connected via a hole at the apex[1]. This opens opportunities for cooling different atomic species in the two magneto optical traps (see figure 8.4), before combining the two species in a magnetic trap (transporting one through the connecting aperture) for heteronuclear cold molecule experiments. Another application of this setup is illustrated in figure 8.5, where two sets of pyramids are used to create a differential pumping between two parts of a small vacuum system. On one side the pyramids could be used to precool atoms from a background rubidium vapour. The cold atoms would then drift through the differential pumping regime to the secondary MOTs at low pressure.

8.3.3 Diamond trap

An interesting arrangement occurs if two wafers of depth d are placed back to back with truncated pyramids of top side length $^5\!/\!\sqrt{2}d$. In this case, illustrated in 8.6, the first reflections pass completely through the aperture and are reflected from the opposite pyramid face. If the device is simultaneously illuminated from both directions with the same circular polarisation

---

[1] This can also be accomplished by etching devices on two wafers separately, and then placing them back to back

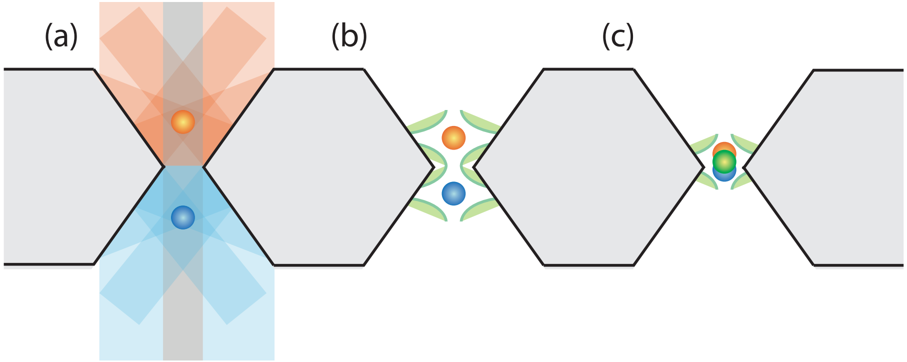

Figure 8.4: Proposed double MOT system. (a) Two back to back pyramids are illuminated with trapping light for two different species. (b) The atoms are magnetically trapped and (c) overlapped

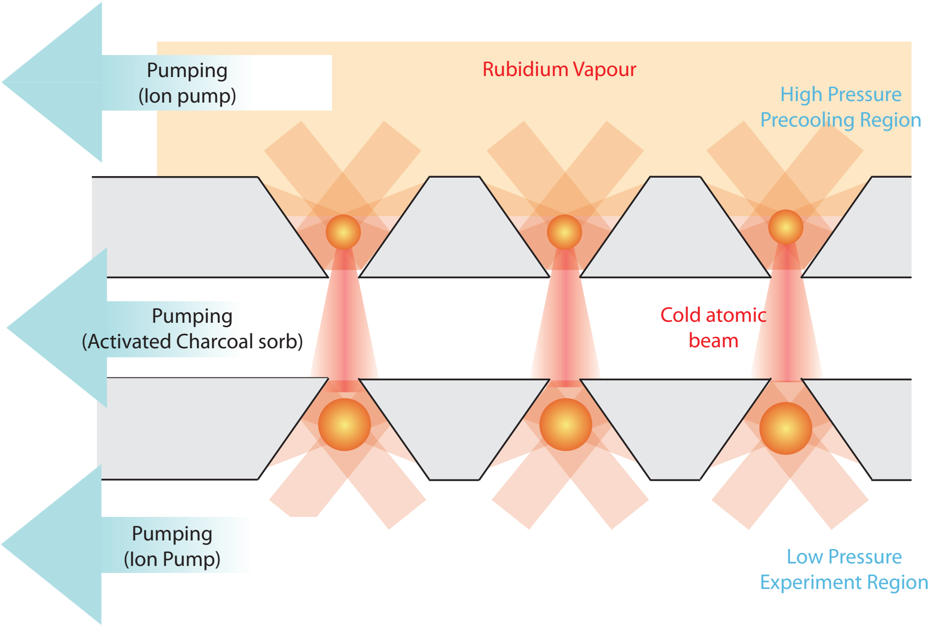

Figure 8.5: Differential Pumping arrangement. Both sides of the setup are illuminated with trapping light, atoms are precooled in the upper pyramids, extracted as a cold beam and used to load the lower traps in a low pressure region.

this forms a naturally balanced magneto-optical trap. This has a very large volume for trapping and good optical access from both sides. If access is only required from one side, then the lower pyramid can be capped with a retroreflecting arrangement (with suitable placement of a quarter wave plate).

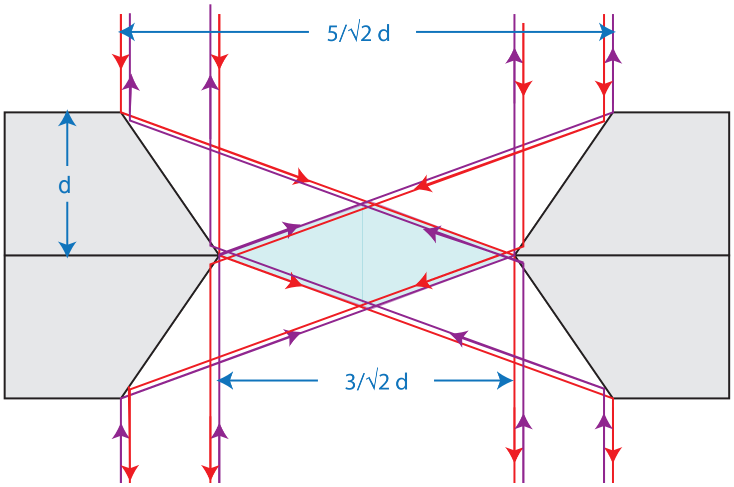

Figure 8.6: Schematic of a proposed diamond trap

Simulations show that a device made from 1mm wafers, i.e. with a width of 3.5mm could trap an atom number as large as 10^7 from the background vapour

Crucially the arrangement of beams means that balanced optical molasses are expected if the illumination of each side of the device is equal.

8.4 Final thoughts

Throughout this thesis I have studied pyramid MOTs spanning 200μm to 2cm in size. I have investigated methods of etching and polishing silicon on the scale of several mm, which is far above the scales previously studied. I successfully demonstrated the operation of a silicon pyramid MOT on scales from 4.2mm down to 3.0mm. Trap parameters such as temperature, damping and spring constants were measured and their behaviour was compared with expectations from both analytical and numerical theories. I have shown that the proximity of a surface reduces the atom number in these pyramids by reducing the capture rate, not by increasing the loss rate as previously expected.

The demonstration of these pyramids as a compact cold-atom source is a significant step towards more fully integrated devices requiring small atom clouds, perhaps in arrays. It is the hope of the writer is that the first steps taken here are used to further develop and improve the concept, and may good luck, hot tea, stable lasers, and low base pressures follow those who choose to take up this challenge.

Appendices

Appendix A

UHV procedures

A.1 Cleaning

Before any kind of pumping can begin, all components (Atom chips, atom chip mounts, magnetic field coils etc.) all must be scrupulously cleaned to ensure the vacuum is not contaminated by outgassing from any kind of grease or oils. The typical cleaning process for any parts required to go under vacuum is as follows:

- Removal of any obvious dirt or grease by manual cleaning in a soap and water solution. If a part arrives from the workshop with an obvious oil residue, it is advisable to make sure the component appears visibly clean before any more stringent cleaning is attempted
- Ultrasonic cleaning in detergent. A dilute solution of a detergent (Decon 90) will remove the vast majority of contaminants. Typically 15 minutes suffices for all but the most stubborn residues.
- Ultrasonic cleaning in distilled water. This will remove any detergent remaining on the part.
- Ultrasonic cleaning in a variety of solvents. Typically isopropanol, followed by methanol, and ethanol for durations of 15 minutes each. In between each solvent, the sample should be rinsed in distilled water for 5 minutes. Use of acetone is advantageous if the sample will not dissolve in it.
- Final ultrasonic cleaning with distilled water. Sample should then be left to dry in a scrupulously clean environment, or blown dry with a clean compressed nitrogen gun if available.

Throughout this process, samples should only be handled with clean gloves and never with bare skin. Where possible, if a device can be disassembled, component parts should be cleaned individually.

A.2 Pumps

1000mbar to 10mbar From room pressure, an oil-free diaphragm pump (model Leybold DIVAC 2.5VT) is used to initially pump out the system. In order to minimise disruption to other experiments, mechanical pumps are where possible kept externally in the pump room and connected via a series of pump lines to the main lab. This will quickly get the pressure down to the point where a turbomolecular pump can be started.

10mbar to 10^{-5}mbar We use a turbomolecular pump (Leybold TurboVac TW70H) attached via a CF flanged bellows to initially reduce the pressure to the high vacuum regime. This should allow pressures of 10^{-5}mbar to be reached in a few hours, if no obvious leaks are present at any of the flanges.

10^{-5} to 10^{-7} mbar After at least 12 hours of pumping it is advantageous to begin selectively heating current carrying components in the vacuum chamber. This helps to drive off any gas/water vapour trapped inside tightly wound coils for example. By passing initially low currents through a component, typically one will observe a sharp rise in pressure followed by a slow drop as the outgassing stops, at which point the current can be increased to further increase the temperature. All components which will carry current should be outgassed up to the maximum level (and if possible, slightly above) at which they will be operated. Typically following this process the base pressure will be slightly better, and if it is below 10^{-5} mbar, the ion pump can be turned on. We typically have to cycle the ion pump voltage a few times, as the initial shock to the electrodes dislodges a lot of detritus accumulated during the exposure to atmosphere raising the pressure by up to an order of magnitude. Turning the ion pump off and waiting for the pressure to drop again before repeating the process a few times will ensure the pump starts successfully without any danger of it tripping. If it fails to start successfully, we have had success in the past with small mechanical shocks to the casing to dislodge whiskers on the electrodes and manually hi-potting the electrodes.

10^{-7}mbar to below 10^{-9} mbar The ion pump is relatively ineffective compared to the turbo pump at pressures above 10^{-7}mbar, but over 24 hours the pressure will drop gradually and the effectiveness of the ion pump will increase. Over 2 days following the ion pump turn on we observe pressures of a few 10^{-8}mbar. To get below this point, the entire chamber has to be baked. We wrap the chamber in an initial layer of aluminium foil, making sure to cover any viewports in a thicker layer before wrapping heater tapes around the chamber. Care must be taken to ensure that heating is even, and no large temperature gradients are created across a viewport to prevent potentially critical damage to the window seal. The whole chamber is then wrapped in a final few layers of aluminum foil.

Unlike many atom chip experiments in our situation the chamber is opened up to atmosphere several times a year. As a result we found that the performance of the ion pump degraded with time, so we started to bake the pump at the same time as the chamber to ensure any liberated contaminant is removed to the turbo pump. We do not remove the magnets in the ion pump and keep it running throughout.

We use variacs to gradually raise the current in the heater tapes, ensuring that we do not raise the temperature at a rate faster than 10° per hour, monitoring the temperature at several key points with thermocouples. Over 2 days the temperature is raised to 120°, the maximum temperature for some materials used in the chamber. Once the maximum temperature is reached, it typically takes at most two days before the pressure begins to drop, and at that point the heaters are turned off and the chamber allowed to cool slowly. By the time the chamber and pump reach room temperature, the pressure is just above 10^{-9}mbar and following activation of the dispensers, the angle valve is closed and the turbo pump detached. Over the course of the first week of operation, there is typically some small residual outgassing, not above the 10^{-8}mbar level, but after a week the pressure as measured on the ion pump is below 10^{-9}mbar (the lower limit of the pressure as measured on the ion pump).

Appendix B

Labview control

The Labview software to control the experiment is developed using the NI-DAQ and NI-IMAQ function libraries which allow for complex routines, including some powerful image processing routines, to be assembled quickly and easily using the Labview graphical programming environment.

The experiment is unified through one master virtual instrument (VI), containing a range of subVIs to operate the experiment and display diagnostics. The benefit of using a master VI, rather than individual programmes is that resources can be shared (for example global settings such as channel definitions and camera settings) and that conflicts such as two programmes accessing the same device can be prevented.

Current subVIs include:

Setup

- Define Channels - Sets up the mapping between physical hardware channels and internal channel definitions, see section B.1 below.
- Set Channels - Allows full control of all analogue and digital outputs.
- Setup Camera - Gives a real time monitor of the camera image, and allows acquisition parameters to be defined for all experiments.
- Experiment Parameters - A list of properties which are saved with every experiment (for example laser power, type of pyramid, dispenser current) for reference purposes

Diagnostics

- Camera Speed Test - This measures the maximum acquisition rate possible for a given set of camera properties. A series of 5 triggered images are repeatedly acquired from the camera, with the rate of acquisition increasing until the camera fails to make all the acquisitions.

- Spectroscopy - While simultaneously ramping the VCO frequency for the offset lock of the primary MOT laser, a photodiode voltage is read from a saturated absorption spectroscopy setup and hyperfine peaks identified in order to calibrate the absolute detuning of the MOT transition.
- Fabry Perot - Reads in the voltage from the Fabry Perot photodiode as the cavity is scanned. Peaks in the cavity transmission are identified and the relative positions and sizes measured in order to ensure that all lasers are single mode and behaving as expected.

Experiments

- Image MOT - A utility to observe and quantify a MOT in a pyramid. Uses the procedure specified in section 4.4
- Iterative Image - A routine to iteratively obtain images of MOT over a range of up to three experiment parameters (for example scanning detuning or magnetic fields for optimisation)
- 1984 Sequencer - The sequencing procedure outlined in section B.3
- Display images - Loads a batch of images from an experiment and allows for further processing or quantification.

B.1 Hardware interface

The interface between the individual programmes and the hardware is separated by a layer of abstraction. This relates hard coded channels in the code (e.g. *dev0/ao1* or *dev0/PI0* for the first analogue output and digital input respectively of the device with hardware ID *dev0*) with simple labels such as 'Laser Detuning' and 'X-Coil Current'. A software interpreter has been created (illustrated in figure B.1) which takes a request for a value on a specific named channel, and looks up the mapping for that particular channel. This also allows the use of calibration factors and conversions, so mappings can be used to turn 'useful values' into voltages output from the DAQ.

B.2 Waveform output

The DAQ cards have a timebase of up to 20MHz, and are capable of outputting arbitrary pattern sequences on up to 32 digital channels and 8 analogue channels. In order to synchronise output from different devices, a counter channel on one of the DAQs is defined as the master timing source, and this is hard wired to a spare digital input of all DAQ devices. When preparing to write data, each device is instructed to synchronise the output to this timing source, and hence all devices operate in step.

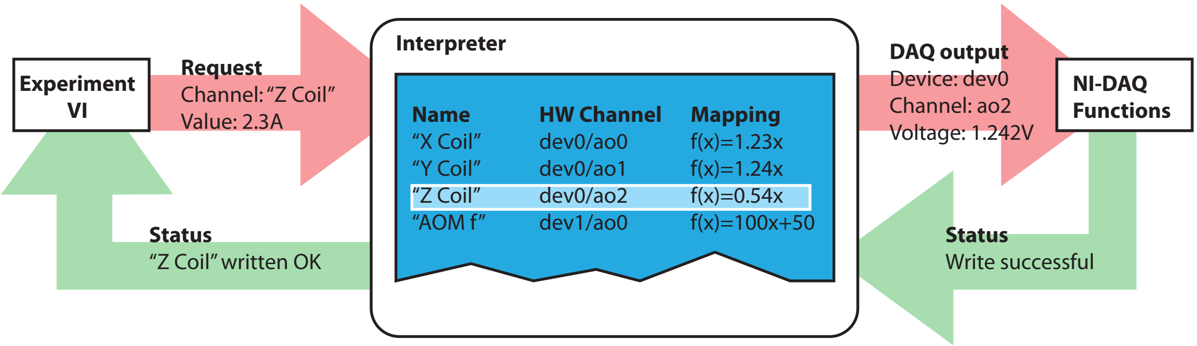

Figure B.1: Process by which the experiment is controlled via a hardware interpreter

B.3 Sequencing

Within the CCM group experiment sequencing has normally been carried out in a 'matrix' style, where the state of the channels is defined at various time positions. This makes sense for digital pattern generation, but makes less sense when using analogue waveforms or when a large number of channels or events are required. Matrix sequences (shown in the left of B.1) often require that if one parameter is to be changed, this change must also be propagated down through all subsequent events.

time/ms	ao0	ao1	ao2	ao3	ao4	do0	do1
5	0	1.5	0	1.2	1	T	F
10	0	0.5	0	1.2	1	T	F
20	0	0	2.6	1.2	1	T	F
50	0	0	2.6	0	0	F	F
100	0	0	0	0	0	F	F
110	0.5	0	0	0	0	F	F
250	0.5	0	0	0	0	F	T

t/ms	Channel	Event Type	Value 1	Value 2
5	All	Default	-	-
10	X Coil	Ramp	0	1
20	AOM	Set Value	50	-
50	Camera	Pulse Sequence	10	
100	X Coil	Ramp	1	0
110	AOM	Set Value	20	-
250	Camera	Pulse Sequence	10	-

Table B.1: Two experimental sequences, Left:the traditional matrix sequencing style, Right:an alternative script-type sequence

The alternative approach used in this experiment is to define events in terms of a 'script'(shown in the right of B.1). In this situation, events are written in a list of actions for individual channels. The actions are then turned into a waveform only at run-time by a script compiler, which creates a blank waveform for all channels with the duration defined by the last event in the script. Each event in the script is then interpreted in turn with the relevant event translated into an output waveform and inserted into the master waveform at the relevant place. Some of the events currently implemented include:

- Single Value [Analogue & Digital] - A single value `[Value 1]` is written to the channel at time `[t]`
- Ramp [Analogue] - Ramps from `[Value 1]` to `[Value 2]` from `[t]` to `[t + Value 3]`
- Pulse Sequence [Digital] - Sends `[Value 1]` TTL pulses out with rate `[Value 2]`

More complicated events could be implemented, such as arbitrary analogue functions or pre-defined digital sequences. As it happens those above events are sufficient to construct most of the required experimental sequences. Scripts are saved as XML files, and not compiled until runtime. A future implementation of this process would allow for multiple scripts to be used in an experiment, thus a complicated experiment could be built up from many simple small sub-scripts.

B.4 Batch data taking

It is often advantageous to perform an experiment many times, varying one or more experimental parameter to measure the response to those parameters. To that end a sequence is often repeated many times but with only small additions making defining a new script for each iteration wasteful. In order to implement this, the script allows events to be given labels, and at run time the events with labels can have parameters modified dynamically or iterated. A typical batch experiment is shown in figure B.2. The output of the experiment (typically the contents of the camera buffer at the end of the sequence) are read out and saved at the start of the subsequent iteration while the script is being recompiled.

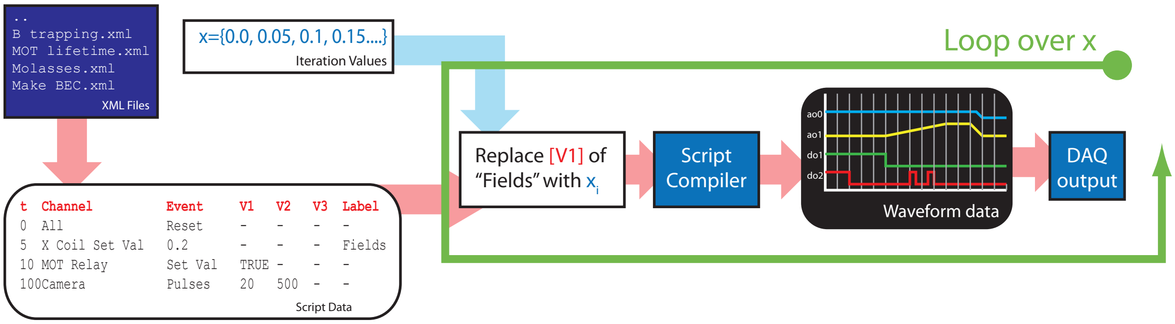

```
..
B trapping.xml
MOT lifetime.xml
Molasses.xml
Make BEC.xml
```

t	Channel	Event	V1	V2	V3	Label
0	All	Reset	-	-	-	-
5	X Coil Set Val	0.2	-	-	-	Fields
10	MOT Relay	Set Val	TRUE	-	-	-
100	Camera	Pulses	20	500	-	-

Figure B.2: Typical Sequence for taking a batch of experiments

Appendix C

Conical MOTs

In order to judge how the performance of a MOT scaled with size, a range of sizes of conical or axicon MOTs were created. It was found that almost perfect 90° hollow cones can be created in aluminium using a process known as Electrical Discharge Machining (EDM). This process, also known as ‘spark eroding’, removes material using repeated electrical arcing from an electrode placed above the aluminium. Cones of diameter 20mm, 15mm, 10mm and 5mm were formed, and following some mechanical polishing, which achieved a reasonable optical finish, these cones were sputtered with gold. The resultant reflectance of the conical mirrors was determined to be 80%. The quadrapole magnetic fields were formed from copper rings mounted above and below each cone, small PEEK clamps held the copper rings in place and isolated them from the metallic surface.

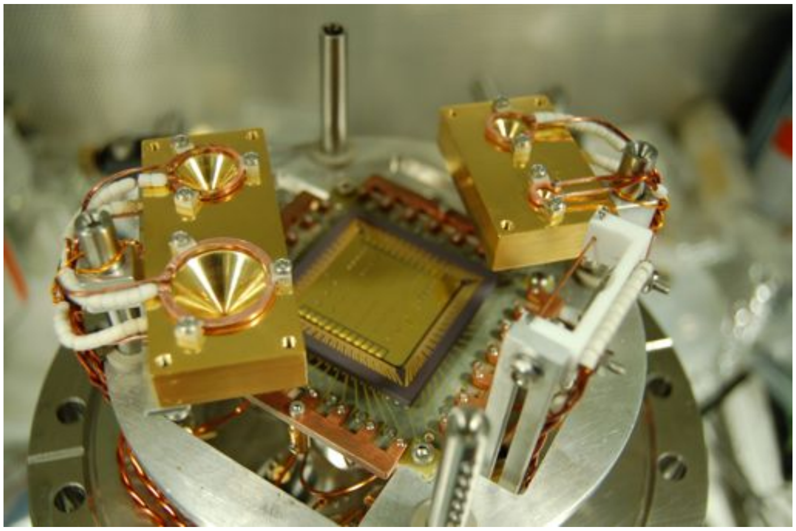

Figure C.1: Hollow conical mirrors made from gold coated aluminium

These cones work in the same way as a 90° pyramid, apart from the fact that trapping and cooling is present along all radial directions as beams are reflected from the inner surface. Vertical confinement comes from doubly reflected beams which counterpropagate parallel to the incoming beam. Four sizes of pyramid, 5, 10 15 and 20mm in diameter were created and all of them were observed to successfully trap atoms. The typical maximum number of trapped atoms for each cone size are displayed in figure C.3 and show scaling for the smallest three cones consistent with the L^6 expected from the large glass pyramid. Clouds formed in a number of configurations, typically a bi-lobal shape (see right figure C.2, but if the magnetic minima was displaced using an external field, ring shapes were clearly visible as predicted by the theory detailed in [71]. The mirrors produced following polishing were obviously satisfactory for MOT operation but the residual roughness caused some difficulty in imaging the trapped atoms, due to the excessive amount of scattered light produced by the region of the cone surface facing the camera. This scattered light was particularly problematic in the largest of the cones (20mm) as the larger curvature of the surface meant more light was reflected to the camera. It is proposed that further mechanical polishing could reduce this.

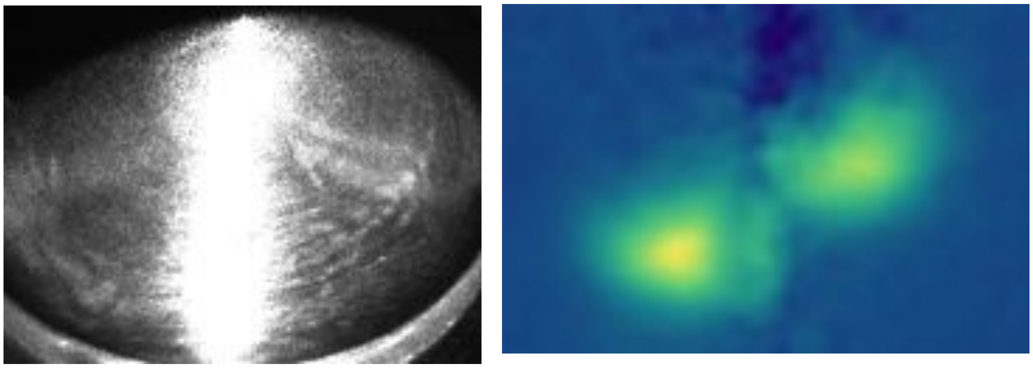

Figure C.2: Left: image of illuminated cone. Right: image of typical bilobal MOT formed inside a 10mm cone

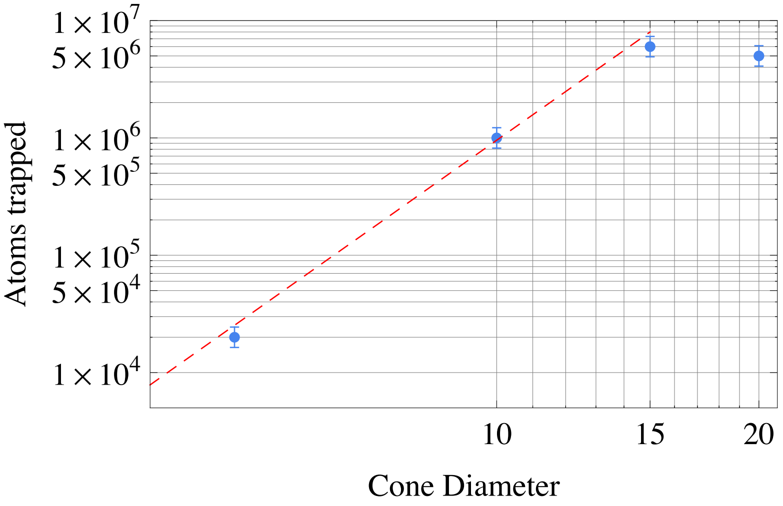

Figure C.3: Scaling of atom number with cone size. Scaling is found to be $x^{5.2\pm0.3}$ (shown as dotted line)

Although the cones do not lend themselves readily to microfabrication or integration with atom chips, they are still remarkably simple to produce with the method we present here. It is proposed that the larger ones we produce could be of some use in providing a quick and simple MOT apparatus solution, or with a hole drilled at the apex, as a convenient source of a cold atomic beam. More crucially the demonstration of trapped atoms in a 5mm cone was further evidence that MOTs can form in small enclosed volumes close to room temperature surfaces

Appendix D

Publications

D.1 Journal Publications

Fabrication of Magneto-Optical Atom Traps on a Chip

G. Lewis, Z. Moktadir, C. Gollasch, M. Kraft

School of Electronics and Computer Science, University of Southampton

S. Pollock, F. Ramirez-Martinez, J. P. Ashmore, A. Laliotis, M. Trupke, E. A. Hinds

Centre for Cold Matter Imperial College

JMEMS (18) 2 347-353

Integrated magneto-optical traps on a chip using silicon pyramid structures

S. Pollock, J. P. Cotter, A. Laliotis, and E. A. Hinds

Centre for Cold Matter Imperial College

Optics Express, Vol. 17, Issue 16, pp. 14109-14114

Experimental study of micro-fabricated silicon pyramid magneto-optical traps

S. Pollock, J. P. Cotter, F.Ramirez-Martinez, A. Laliotis, and E. A. Hinds

Centre for Cold Matter Imperial College

In preparation

D.2 Other publications featuring the silicon pyramids

Atom chip moves on

Belle Dumé - NanoTechWeb (`http://nanotechweb.org/cws/article/tech/39093`)

Physics arXiv blog

The second paper was covered in the MIT Technology Review (`http://www.technologyreview.com/blog/arxiv/23503/`)

Bibliography

[1] T. W. Hansch and A. L. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13:68–+, January 1975.

[2] Wineland D. and Dehmelt H. Proposed $10^{14}\delta\nu/\nu$ Laser Fluorescence Spectroscopy on $T1^{+}$ Mono-Ion Oscillator. *Bulletin of the American Physical Society*, 20:637, 1975.

[3] W. D. Phillips and H. Metcalf. Laser Deceleration of an Atomic Beam. *Physical Review Letters*, 48:596–599, March 1982.

[4] A. L. Migdall, W. D. Phillips, J. V. Prodan, T. H. Bergeman, and H. J. Metcalf. First observation of magnetically trapped neutral atoms. *Physical Review Letters*, 54:2596–2599, June 1985.

[5] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269:198–201, July 1995.

[6] J. D. Weinstein and K. G. Libbrecht. Microscopic magnetic traps for neutral atoms. *Physical Review A*, 52(5):4004–4009, Nov 1995.

[7] J. Fortagh, A. Grossmann, C. Zimmermann, and T. W. Hänsch. Miniaturized wire trap for neutral atoms. *Physical Review Letters*, 81(24):5310–5313, Dec 1998.

[8] J. Reichel, W. Hansell, and T. W. Hansch. Atomic micromanipulation with magnetic surface traps. *Physical Review Letters*, 83(17):3398–3401, 1999.

[9] R. Folman, P. Kruger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer. Controlling cold atoms using nanofabricated surfaces: Atom chips. *Physical Review Letters*, 84(20):4749–4752, 2000.

[10] D. Muller, D. Z. Anderson, R. J. Grow, P. D. D. Schwindt, and E. A. Cornell. Guiding neutral atoms around curves with lithographically patterned current-carrying wires. *Physical Review Letters*, 83(25):5194–5197, 1999.

[11] N. H. Dekker, C. S. Lee, V. Lorent, J. H. Thywissen, S. P. Smith, M. Drndic, R. M. Westervelt, and M. Prentiss. Guiding neutral atoms on a chip. *Physical Review Letters*, 84(6):1124–1127, 2000.

[12] W. Hansel, P. Hommelhoff, T. W. Hansch, and J. Reichel. Bose-Einstein condensation on a microelectronic chip. *Nature*, 413:498–501, 2001.

[13] H. Ott, J. Fortagh, G. Schlotterbeck, A. Grossmann, and Zimmermann A. Bose-Einstein condensation in a surface microtrap. *Physical Review Letters*, 87:230401, 2001.

[14] P. K. Rekdal, S. Scheel, P. L. Knight, and E. A. Hinds. Thermal spin flips in atom chips. *Physical Review A*, 70(1):013811, Jul 2004.

[15] S. Eriksson, F. Ramirez-Martinez, E. A. Curtis, B. E. Sauer, P. W. Nutter, E. W. Hill, and E. A. Hinds. Micron-sized atom traps made from magneto-optical thin films. *Applied Physics B-Lasers and Optics*, 79(7):811–816, 2004.

[16] P. Krüger, L. M. Andersson, S. Wildermuth, S. Hofferberth, E. Haller, S. Aigner, S. Groth, I. Bar-Joseph, and J. Schmiedmayer. Potential roughness near lithographically fabricated atom chips. *Physical Review A*, 76(6):063621, 2007.

[17] 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, 1999.

[18] C. D. J. Sinclair, E. A. Curtis, I. Llorente-Garcia, J. A. Retter, B. V. Hall, S. Eriksson, B. E. Sauer, and E. A. Hinds. Bose-Einstein condensation on a permanent-magnet atom chip. *Physical Review A*, 72:03160(R), 2005.

[19] M Singh, M Volk, A Akulshin, A Sidorov, R McLean, and P Hannaford. One-dimensional lattice of permanent magnetic microtraps for ultracold atoms on an atom chip. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 41(6):065301 (6pp), 2008.

[20] A. Shevchenko, M. Heilio, T. Lindvall, A. Jaakkola, I. Tittonen, M. Kaivola, and T. Pfau. Trapping atoms on a transparent permanent-magnet atom chip. *Physical Review A*, 73(5):–, 2006.

[21] A. Gunther, M. Kemmler, S. Kraft, C. J. Vale, C. Zimmermann, and J. Fortagh. Combined chips for atom optics. *Physical Review A*, 71(6):–, 2005.

[22] M. Fattori, C. D'Errico, G. Roati, M. Zaccanti, M. Jona-Lasinio, M. Modugno, M. Inguscio, and G. Modugno. Atom Interferometry with a Weakly Interacting Bose-Einstein Condensate. *Physical Review Letters*, 100(8):080405, 2008.

[23] B. V. Hall, S. Whitlock, R. Anderson, P. Hannaford, and A. I. Sidorov. Condensate splitting in an asymmetric double well for atom chip based sensors. *Physical Review Letters*, 98(3):–, 2007.

[24] M. Vengalattore and M. Prentiss. A reciprocal magnetic trap for neutral atoms. *European Physical Journal D*, 35(1):69–73, 2005.

[25] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Kruger. Matter-wave interferometry in a double well on an atom chip. *Nature Physics*, 1(1):57–62, 2005.

[26] Y. Colombe, E. Knyazchyan, O. Morizot, B. Mercier, V. Lorent, and H. Perrin. Ultracold atoms confined in RF-induced two-dimensional trapping potentials. *Europhysics Letters*, 67(4):593–599, 2004.

[27] S. Wildermuth, S. Hofferberth, I. Lesanovsky, S. Groth, P. Kruger, J. Schmiedmayer, and I. Bar-Joseph. Sensing electric and magnetic fields with Bose-Einstein condensates. *Applied Physics Letters*, 88(26):–, 2006.

[28] D. M. Harber, J. M. Obrecht, J. M. McGuirk, and E. A. Cornell. Measurement of the Casimir-Polder force through center-of-mass oscillations of a Bose-Einstein condensate. *Physical Review A*, 72(3):033610, Sep 2005.

[29] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hansch, and Immanuel Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.

[30] P. Treutlein, T. Steinmetz, Y. Colombe, B. Lev, P. Hommel-Hoff, J. Reichel, M. Greiner, O. Mandel, A. Widera, T. Rom, I. Bloch, and T. W. Hansch. Quantum information processing in optical lattices and magnetic microtraps. *Fortschritte Der Physik-Progress of Physics*, 54(8-10):702–718, 2006.

[31] A. Grabowski and T. Pfau. A lattice of magneto-optical and magnetic traps for cold atoms. *European Physical Journal D*, 22(3):347–354, 2003.

[32] S. Whitlock, R. Gerritsma, T. Fernholz, and R. J. C. Spreeuw. Two-dimensional array of microtraps with atomic shift register on a chip. *New Journal of Physics*, 11(2):023021–+, February 2009.

[33] S. Eriksson, M. Trupke, H. F. Powell, D. Sahagun, C. D. J. Sinclair, E. A. Curtis, B. E. Sauer, E. A. Hinds, Z. Moktadir, C. O. Gollasch, and M. Kraft. Integrated optical components on atom chips. *European Physical Journal D*, 35(1):135–139, 2005.

[34] Peter Horak, Bruce G. Klappauf, Albrecht Haase, Ron Folman, Jörg Schmiedmayer, Peter Domokos, and E. A. Hinds. Possibility of single-atom detection on a chip. *Physical Review A*, 67(4):043806, Apr 2003.

[35] M. Wilzbach, A. Haase, M. Schwarz, D. Heine, K. Wicker, X. Liu, K. H. Brenner, S. Groth, T. Fernholz, B. Hessmo, and J. Schmiedmayer. Detecting neutral atoms on an atom chip. *Fortschritte Der Physik-Progress of Physics*, 54(8-10):746–764, 2006.

[36] Marco Wilzbach, Dennis Heine, Sönke Groth, Xiyuan Liu, Thomas Raub, Björn Hessmo, and Jörg Schmiedmayer. Simple integrated single-atom detector. *Optics Letters*, 34(3):259–261, 2009.

[37] A. Takamizawa, T. Steinmetz, R. Delhuille, T. W. Hansch, and J. Reichel. Miniature fluorescence detector for single atom observation on a microchip. *Optics Express*, 14(23):10976–10983, 2006.

[38] I. Teper, Y. J. Lin, and V. Vuletic. Resonator-aided single-atom detection on a microfabricated chip. *Physical Review Letters*, 97(2):–, 2006.

[39] M. Trupke, J. Goldwin, B. Darquié, G. Dutier, S. Eriksson, J. Ashmore, and E. A. Hinds. Atom detection and photon production in a scalable, open, optical microcavity. *Physical Review Letters*, 99:063601, 2007.

[40] T. Steinmetz, Y. Colombe, D. Hunger, T. W. Hansch, A. Balocchi, R. J. Warburton, and J. Reichel. Stable fiber-based fabry-perot cavity. *Applied Physics Letters*, 89(11):–, 2006.

[41] Y. Colombe, T. Steinmetz, G. Dubois, F. Linke, D. Hunger, and J. Reichel. Strong atom-field coupling for Bose-Einstein condensates in an optical cavity on a chip. *Nature*, 450:272–276, November 2007.

[42] G. Birkl, F. B. J. Buchkremer, R. Dumke, and W. Ertmer. Atom optics with microfabricated optical elements. *Optics Communications*, 191:67–81, May 2001.

[43] M. Kohnen, M. Succo, P. Petrov, R. Nyman, Trupke M., and E. Hinds. An integrated optical-waveguide chip for measurement of cold atom clouds. *In preparation*, 2009.

[44] P. D. D. Schwindt, S. Knappe, V. Shah, L. Hollberg, J. Kitching, L. A. Liew, and J. Moreland. Chip-scale atomic magnetometer. *Applied Physics Letters*, 85(26):6409–6411, 2004.

[45] S. W. Du, M. B. Squires, Y. Imai, L. Czaia, R. A. Saravanan, V. Bright, J. Reichel, T. W. Hansch, and D. Z. Anderson. Atom-chip Bose-Einstein condensation in a portable vacuum cell. *Physical Review A*, 70(5):–, 2004.

[46] ColdQuanta. Ultracold atom and BEC devices, Instruments, Atom optics and Systems, 2009. `http://www.coldquanta.com`.

[47] J. Reichel. Microchip traps and Bose-Einstein condensation. *Applied Physics B-Lasers and Optics*, 74(6):469–487, 2002.

[48] M. Trupke, F. Ramirez-Martinez, E. A. Curtis, J. P. Ashmore, S. Eriksson, E. A. Hinds, Z. Moktadir, C. Gollasch, M. Kraft, G. Vijaya Prakash, and J. J. Baumberg. Pyramidal micromirrors for microsystems and atom chips. *Applied Physics Letters*, 88:071116, 2006.

[49] K. I. Lee, J. A. Kim, H. R. Noh, and W. Jhe. Single-beam atom trap in a pyramidal and conical hollow mirror. *Optics Letters.*, 21:1177, 1996.

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

[51] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin. Storage of light in atomic vapor. *Physical Review Letters*, 86(5):783–786, Jan 2001.

[52] A. Naik, O. Buu, M. D. Lahaye, A. D. Armour, A. A. Clerk, M. P. Blencowe, and K. C. Schwab. Cooling a nanomechanical resonator with quantum back-action. *Nature*, 443:193–196, September 2006.

[53] J. Verdú, H. Zoubi, Ch. Koller, J. Majer, H. Ritsch, and J. Schmiedmayer. Strong magnetic coupling of an ultracold gas to a superconducting waveguide cavity. *Physical Review Letters*, 103(4):043603, Jul 2009.

[54] Samuel A. Meek, Horst Conrad, and Gerard Meijer. Trapping Molecules on a Chip. *Science*, 324(5935):1699–1702, 2009.

[55] H.J Metcalf and P van der Straten. *Laser Cooling and Trapping*. Graduate Texts in Contemporary Physics. Springer, 1999.

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

[57] C.S Adams and I.G Hughes. *Handbook of Laser Technology and Applications*, volume 3. Taylor & Francis, 2004.

[58] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. *ArXiv e-prints*, February 1999.

[59] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. Optical molasses. *Journal of the Optical Society of America B-Optical Physics*, 6(11):2084–2107, 1989.

[60] W. Wohlleben, F. Chevy, K. Madison, and J. Dalibard. An atom faucet. *European Physical Journal D*, 15:237–244, August 2001.

[61] Matthieu Vangeleyn, Paul F. Griffin, Erling Riis, and Aidan S. Arnold. Single-laser, one beam, tetrahedral magneto-optical trap. *Optics Express*, 17(16):13601–13608, 2009.

[62] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman. Low-velocity intense source of atoms from a magneto-optical trap. *Physical Review Letters*, 77(16):3331–3334, Oct 1996.

[63] A. M. Steane, M. Chowdhury, and C. J. Foot. Radiation force in the magneto-optical trap. *Journal of the Optical Society of America B*, 9(12):2142–2158, 1992.

[64] 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, 1992.

[65] 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 magnetooptical trap. *Physical Review A*, 52(2):1423–1440, 1995.

[66] B. Ueberholz, S. Kuhr, D. Frese, V. Gomer, and D. Meschede. Cold collisions in a high-gradient magneto-optical trap. *Journal of Physics B Atomic Molecular Physics*, 35:4899–4914, December 2002.

[67] Joe Kerckhoff. Mot dynamics. *qwiki.stanford.edu*, 2007.

[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 magnetooptical trap. *Journal of the Optical Society of America B-Optical Physics*, 11(5):703–711, 1994.

[69] P. D. D. Schwindt, E. A. Cornell, T. Kishimoto, Y. J. Wang, and D. Z. Anderson. Efficient loading of a magnetic waveguide on an atom chip. *Physical Review A*, 72(2):–, 2005.

[70] H. R. Noh and W. Jhe. Atom optics with hollow optical systems. *Physics Reports-Review Section of Physics Letters*, 372(3):269–317, 2002.

[71] J. A. Kim, K. I. Lee, H. R. Noh, W. Jhe, and M. Ohtsu. Atom trap in an axicon mirror. *Optics Letters*, 22(2):117–119, 1997.

[72] Fujio Shimizu, Kazuko Shimizu, and Hiroshi Takuma. Four-beam laser trap of neutral atoms. *Optics Letters*, 16(5):339–341, 1991.

[73] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A simple extended-cavity diode laser. *Review of Scientific Instruments*, 69:1236–1239, March 1998.

[74] C P Pearman, C S Adams, S G Cox, P F Griffin, D A Smith, and I G Hughes. Polarization spectroscopy of a closed atomic transition: applications to laser frequency locking. *Journal of Physics B*, 35(24):5141–5151, 2002.

[75] T Stace, A N Luiten, and R P Kovacich. Laser offset-frequency locking using a frequency-to-voltage converter. *Measurement Science and Technology*, 9(9):1635–1637, 1998.

[76] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski. Simple scheme for tunable frequency offset locking of two lasers. *Review of Scientific Instruments*, 70:242–243, January 1999.

[77] G. Ritt, G. Cennini, C. Geckeler, and M. Weitz. Laser frequency offset locking using a side of filter technique. *Applied Physics B-Lasers and Optics*, 79:363–365, 2004.

[78] K. L. Corwin, Z. T. Lu, C. F. Hand, R. J. Epstein, and C. E. Wieman. Frequency-stabilized diode laser with the zeeman shift in an atomic vapor. *Applied Optics*, 37(15):3295–3298, May 1998.

[79] F Ramirez-Martinez. *Integration of Optical Components and Magnetic Field Sources in Atom Chips*. PhD thesis, Imperial College London, July 2008.

[80] J. Dalibard and C. Cohentannoudji. Laser cooling below the doppler limit by polarization gradients - simple theoretical-models. *Journal of the Optical Society of America B-Optical Physics*, 6(11):2023–2045, 1989.

[81] G.N. Lewis, Z. Moktadir, C. Gollasch, M. Kraft, S. Pollock, F. Ramirez-Martinez, J.P. Ashmore, A. Laliotis, M. Trupke, and E.A. Hinds. Fabrication of magnetooptical atom traps on a chip. *Microelectromechanical Systems, Journal of*, 18(2):347–353, April 2009.

[82] J. Weickert, B.M. ter Haar Romeny, and M.A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7(3):390–410, 1998.

[83] Jochen Kronjaeger. *Coherent Dynamics of Spinor Bose-Einstein Condensates.* PhD thesis, University of Hamburg, 2007.

[84] J.M Bennett and L. Matteson. *Introduction to Surface Roughness and Scattering.* Optical Society of America, 1989.

[85] H. Seidel. The mechanism of anisotropic, electrochemical silicon etching in alkaline solutions. In *Solid-State Sensor and Actuator Workshop, 1990. 4th Technical Digest., IEEE*, pages 86–91, Jun 1990.

[86] G.T.A. Kovacs, N.I. Maluf, and K.E. Petersen. Bulk micromachining of silicon. *Proceedings of the IEEE*, 86(8):1536–1551, Aug 1998.

[87] Z. Moktadir and H. Camon. Monte carlo simulation of anisotropic etching of silicon: investigation of $\langle 111 \rangle$ surface properties. *Modell. Simul. Mat. Sci. Eng*, 5:481–488, 1997.

[88] M. J Madou. *Fundamentals of Microfabrication.* CRC-Press, 2nd edition, 2002.

[89] Kazuo Sato, Mitsuhiro Shikida, Takashi Yamashiro, Masaki Tsunekawa, and Shunji Ito. Roughening of single-crystal silicon surface etched by koh water solution. *Sens. Act. A*, 73(1-2):122 – 130, 1999.

[90] A. S. Louro and J. R. Senna. Real-time, in-situ microscopic observation of bubbles and roughening in KOH etching of silicon. In *Micromachining and Microfabrication Process Technology VII*, 2001.

[91] A. J. Nijdam, J. G. E. Gardeniers, C. Gui, and M. Elwenspoek. Etching pits and dislocations in si111. *Sensors and Actuators A: Physical*, 86(3):238 – 247, 2000.

[92] Nga Phuong Pham, E. Boellaard, J.N. Burghartz, and P.M. Sarro. Photoresist coating methods for the integration of novel 3-d RF microstructures. *Microelectromechanical Systems, Journal of*, 13(3):491–499, June 2004.

[93] K. Pontoppidan Larsen, J. Ravnkilde, and O. Hansen. Investigations of the isotropic etch of an icp source for silicon microlens mold fabrication. *J. Micromech. Microeng.*, 15:873–882, 2005.

[94] F. Gao, R. K. Leach, J. Petzing, and J. M. Coupland. Surface measurement errors using commercial scanning white light interferometers. *Measurement Science and Technology*, 19(1):015303–+, January 2008.

[95] J. Raja, B. Muralikrishnan, and Shengyu Fu. Recent advances in separation of roughness, waviness and form. *Precision Engineering*, 26(2):222 – 235, 2002.

[96] J.C. Dainty. *Laser Speckle and related Phenomena.* Springer-Verlag, 1975.

[97] U. D. Rapol, A. Wasan, and V. Natarajan. Loading of a Rb magneto-optic trap from a getter source. *Physical Review A*, 64(2):art. no.–023402, 2001.

[98] R. S. III Williamson. *Magneto-Optical Trapping of Potassium Isotopes.* PhD thesis, University of Wisconsin-Madison, 1997.

[99] Xinye Xu, Thomas H. Loftus, Matthew J. Smith, John L. Hall, Alan Gallagher, and Jun Ye. Dynamics in a two-level atom magneto-optical trap. *Physical Review A*, 66:011401, 2002.

[100] P. Kohns, P. Buch, W. Suptitz, C. Csambal, and W. Ertmer. On-line measurement of sub-doppler temperatures in a Rb magneto-optical trap-by-trap centre oscillations. *Europhysics Letters*, 22(7):517–522, 1993.

[101] Klaus Molmer. Friction and diffusion coefficients for cooling of atoms in laser fields with multidimensional periodicity. *Physical Review A*, 44(9):5820–5832, Nov 1991.

[102] M.I den Hertog. *Measuring the temperature of a Rb MOT using an oscillating magnetic field.* PhD thesis, Utrecht University, 2004.

[103] Y. T. Chough and W. Jhe. A computational analysis of an axiconic magneto-optical trap. *Journal of the Physical Society of Japan*, 69(5):1366–1373, 2000.

[104] F. Reif. *Statistical and Thermal Physics.* McGraw-Hill, 1965.

[105] David V. Sheludko, Simon C. Bell, Russell Anderson, Christoph S. Hofmann, Edgar J. D. Vredenbregt, and Robert E. Scholten. State-selective imaging of cold atoms. *Physical Review A*, 77(3):033401, 2008.