Manipulating Single Atoms with Optical Tweezers

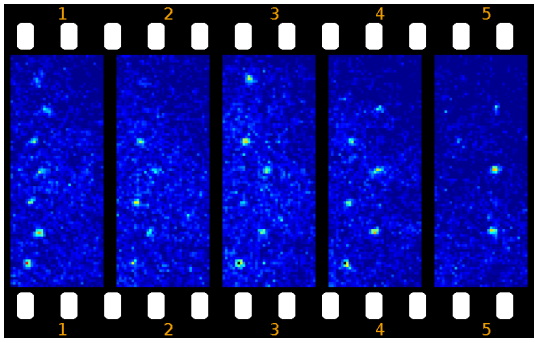

Dustin Stuart

Merton College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2014

This thesis is dedicated to
my parents.

Statement of Originality

This thesis is entirely my own work, except for Chapter 3, which is adapted from the paper by Muldoon et al. [1]. I contributed to these experiments in my first year, but the text itself was written by Cecilia Muldoon. I have included it as it provides an essential background perspective for the current work.

Acknowledgements

This thesis wouldn't have been possible without the generous support and encouragement of many people.

First and foremost, I would like to thank my supervisor, Axel Kuhn, for his unfailing support over the past four years, for his breadth of knowledge of the field, and his disarmingly friendly personality. His uncanny ability to make the most complicated ideas seem completely trivial is something which I can only hope to emulate.

I am deeply grateful to the Rhodes Trust for funding my degree, to the Western Australian selection panel who believed in me, and in particular, to Don Markwell, Mary Eaton, Andrew Graham and Charles Conn, for nurturing such a vibrant community of scholars within Oxford. I also thank the Department of Atomic and Laser Physics, especially Paul Ewart, for their support.

To my college, Merton, of which I am immensely proud to be a member, thanks to Simon Jones, to my friends in the MCR, and to all of the staff who have made my time here so special. To Sukumar at Wadham, Pat Roche at Hertford, and to all the students whom I've had the privilege of teaching. I've learnt more physics from those tutorials than anywhere else. To my lab mates Ollie, Annemarie, Jian, Cici, Lukas, Jerome, Peter, Gunnar, Tom and Marwan: thanks for making Room 301 such a banterous and stimulating place to work. Also, thanks to those of you who kindly proofread this thesis.

Finally, thanks to my beautiful fiancée, Sarah Wilson, for your love and support, and for putting up with my nonsense. To the IVI, especially those who made the trek to visit, and lastly, to Mum, Dad, Riley, Eliza, Gran and Grandad. This is for you.

Abstract

Single atoms are promising candidates for physically implementing quantum bits, the fundamental unit of quantum information. We have built an apparatus for cooling, trapping and imaging single rubidium atoms in microscopic optical tweezers. The traps are formed from a tightly focused off-resonant laser beam, which traps atoms using the optical dipole force. The traps have a diameter of $\sim$1 µm and a depth of $\sim$1 mK.

The novelty of our approach is the use a digital mirror device (DMD) to generate multiple independently movable tweezers from a single laser beam. The DMD consists of an array of micro-mirrors that can be switched on and off, thus acting as a binary amplitude modulator. We use the DMD to imprint a computer-generated hologram on the laser beam, which is converted in to the desired arrangement of traps in the focal plane of a lens. We have developed fast algorithms for calculating binary holograms suitable for the DMD. In addition, we use this method to measure and correct for errors in the phase of the wavefront caused by optical aberrations, which is necessary for producing diffraction-limited focal spots.

Using this apparatus, we have trapped arrays of up to 20 atoms with arbitrary geometrical arrangements. We exploit light-assisted collisions between atoms to ensure there is at most one atom per trapping site. We measure the temperature of the atoms in the traps to be 12 µK, and their lifetime to be 1.4 s. Finally, we demonstrate the ability to select individual atoms from an array and transport them over a distance of 14 µm with laser cooling, and 5 µm without.

Contents

1 Introduction **1**

1.1 Atoms, cavities and optical tweezers 5

1.2 State of the art 7

1.3 Thesis outline 9

2 Theory **11**

2.1 Light-atom interaction 11

2.1.1 The scattering force 13

2.1.2 The dipole force 14

2.1.3 Dipole force for multi-level atoms 15

2.1.4 Doppler cooling 18

2.1.5 Sub-Doppler cooling 20

2.1.6 The magneto-optical trap 21

2.2 Collisional blockade 23

2.3 Diffraction-limited focusing 24

2.3.1 Beyond the paraxial approximation 25

2.3.2 Spin-orbit coupling of light 30

2.4 Conclusions 32

3 Control and Manipulation of Cold Atoms in Optical Tweezers **33**

3.1 Trap design and optical arrangement 34

3.2 Atoms trapped in arbitrary potential landscapes 37

3.3 Deterministic re-arrangement and controlled transport of atoms . . . 41

3.4 Conclusion and Outlook 43

3.5 Methods 43

4 Experimental Design **45**
4.1 MOT . . . 47
4.1.1 Vacuum chamber . . . 48
4.1.2 Magnetic field coils . . . 50
4.1.3 Cooling lasers . . . 51
4.1.4 Sub-Doppler temperature . . . 54
4.1.5 Control of the MOT density . . . 54
4.2 Dipole trap . . . 57
4.2.1 Trapping laser . . . 58
4.2.2 Testing the dipole trap . . . 60
4.2.3 Aspheric lens . . . 61
4.2.4 Lens mount . . . 62
4.2.5 Digital mirror device . . . 63
4.2.6 DMD speed . . . 64
4.2.7 Relay optics . . . 66
4.2.8 Aberrations . . . 67
4.3 Single atom imaging . . . 69
4.3.1 Calibrating the EMCCD camera . . . 72
4.4 Conclusions . . . 73

5 Fast algorithms for generating binary holograms **75**
5.1 Introduction . . . 75
5.2 Computation of holograms . . . 78
5.2.1 Hologram of a single focused trap . . . 79
5.2.2 Binary rounding algorithm . . . 80
5.2.3 Extension to multiple traps . . . 81
5.2.4 Ordered dithering algorithm . . . 82
5.2.5 Weighted Gerchberg-Saxton algorithm . . . 83
5.2.6 Wavefront correction . . . 84
5.3 Comparison of different algorithms . . . 87
5.4 Conclusion . . . 89

6 Moving single atoms **91**
6.1 Introduction . . . 91
6.2 Trapping and imaging single atoms . . . 93
6.2.1 Dipole traps . . . 94
6.2.2 Imaging system . . . 95

6.2.3 Observation of single atoms . . . 96
6.2.4 Lifetime of trapped atoms . . . 100
6.2.5 Temperature of trapped atoms . . . 101
6.2.6 Trap frequencies . . . 102
6.3 Trapping arrays of atoms . . . 103
6.4 Transporting single atoms . . . 106
6.5 Prospects for quantum information processing . . . 109
6.5.1 Scalability of qubits . . . 109
6.5.2 Preparation fidelity . . . 109
6.5.3 Coherence time . . . 110
6.5.4 Quantum gates . . . 111
6.5.5 Qubit readout . . . 112
6.5.6 Coupling via optical cavities . . . 113
6.5.7 Concluding remarks . . . 114

7 Conclusion **115**

Bibliography **118**

List of Figures

1.1 Comparison between different physical implementations of qubits . . 4
1.2 Atom-cavity architecture 5

2.1 Energy levels in a two-level atom 12
2.2 Energy levels of rubidium 16
2.3 Trap depth and scattering rate for different wavelengths 17
2.4 Trap depth with circularly polarised light 18
2.5 Cooling force as a function of velocity 19
2.6 Energy levels of rubidium-87 involved in laser cooling 21
2.7 Molecular potential for ground and excited state of Rb_2 23
2.8 Electric field of a focused Gaussian beam 24
2.9 Geometry for calculation of focal field 26
2.10 Calculated intensity of the trapping light, along with the intensity of the x and z-polarised components, in the focal plane of the aspheric lens 27
2.11 Optimum waist of the trapping beam 28
2.12 Intensity of the E_x, E_y and E_z components of the focused laser beam 29
2.13 Trapping potential taking the m_F sublevels in to account. 30
2.14 Calculated intensity for a right-circularly polarised beam 31

3.1 Tweezers arrangement 35
3.2 Fluorescence images of trapped atoms 37
3.3 Illustration of the concept of the atom transport 40
3.4 Ballistic atom transport 40

4.1 Photograph of the complete vacuum chamber taken with an infra-red camera 45
4.2 Experimental setup 46
4.3 Fluorescence image of the MOT 48
4.4 Perspective view of the vacuum chamber 49
4.5 Diagram of cooling laser setup 51

4.6 Frequency-modulated (FM) spectrum of the rubidium transitions . . 52
4.7 Detailed frequency modulation (FM) spectra of the cooling and repump transitions. The vertical line indicates the locking point. 53
4.8 Relevant energy levels for laser locking 53
4.9 Time-of-flight measurement of the temperature of sub-Doppler cooled atoms 55
4.10 Variation of the diameter and density of atoms in the MOT with magnetic field gradient 56
4.11 Number of atoms in the MOT measured over one month 57
4.12 Dipole trapping laser setup 58
4.13 Normalised power spectral density of the dipole trapping laser 59
4.14 Transmission of MaxLine 785 filter 59
4.15 Absorption image of atoms being loaded into a dipole trap 61
4.16 Fluorescence from the focused dipole trapping beam 62
4.17 Photograph and dimensions of custom lens mount 63
4.18 Images of the DMD 63
4.19 Trajectory of mirrors after reset 65
4.20 Intensity modulation due to DMD settling time 66
4.21 Dipole trapping optics 67
4.22 Surface of the DMD as part of a Michelson interferometer 68
4.23 Wavefront error measured using the DMD 69
4.24 Operating principle of the EMCCD camera 70
4.25 Single atom imaging optics 70
4.26 Background light variation with aperture diameter 71
4.27 Variance versus mean of EMCCD counts 73
4.28 Variance versus mean of photoelectrons 74

5.1 Experimental setup used for testing holograms 77
5.2 Examples images for different hologram algorithms 79
5.3 Measuring wavefront errors 85
5.4 Performance of different algorithms 86

6.1 Experimental setup 93
6.2 Photon count statistics from dipole-trapped atoms 96
6.3 Post-selected images of trapped atoms 97
6.4 Histogram of trap duration 98
6.5 Lifetime of dipole-trapped atoms 100

6.6 Temperature and oscillation frequency of dipole-trapped atoms 102
6.7 Gallery of various arrangements of trapped atoms 105
6.8 Bright transport with continuous laser cooling 106
6.9 Dark transport for different step sizes 107

Chapter 1

Introduction

In 1982, Richard Feynman pointed out that the complexity of quantum mechanical systems means that they are impossible to simulate with classical computers [2]. Furthermore he suggested that it might be possible to build a quantum computer that could take advantage of this complexity rather than being limited by it. In the 30 years since, there has been tremendous progress towards building and controlling quantum systems, however, a universal quantum computer does not yet exist.

To appreciate the inherent difficulty in simulating quantum mechanics, consider a simple quantum system with two eigenstates, $|0\rangle$ and $|1\rangle$. This system forms a quantum bit or *qubit*: the fundamental unit of quantum information. In general, the state of the system $|\psi\rangle$ is fully described by two complex numbers α and β:

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle \tag{1.1}$$

Now consider a large quantum system comprised of N qubits. This composite system has 2^N eigenstates and a complete description requires $2^N - 1$ complex numbers. For just $N = 50$ qubits, storing the quantum state would require more than 9000 TB of memory, far beyond the capability of any classical machine.

If a quantum computer did exist, it would have an immediate and revolutionary

impact on a wide range of challenging problems. For example, Grover's algorithm would allow fast searching of databases [3], which is of immense practical importance in the information age. Shor's algorithm [4] famously allows the factorisation of large numbers faster than any known classical algorithm, enabling one to break the public key encryption [5] that is used for government and military communication. Most importantly, a quantum computer could be used to simulate the physics of many-body quantum systems. This may lead, for example, to the development of new high-temperature superconductors for more energy-efficient electricity distribution.

The theoretical principles of a quantum computer are well established [6]. It is known that any quantum operation can be decomposed into a sequence of one and two-qubit logic gates. In a real computer, the gates will be subject to errors due to experimental imperfections, and decoherence due to interactions with the environment. Fortunately, error correction techniques [7] have been developed which allow a single logical qubit to be encoded robustly in a system of many physical qubits.

The DiVincenzo criteria [8] set out the requirements for building a practical quantum computer. Briefly, they are:

1. a scalable physical implementation of a qubit
2. the ability to initialise the system to the $|00\cdots0\rangle$ state
3. long coherence times (relative to gate times)
4. a universal set of gates
5. the ability to selectively readout individual qubits

There are many examples of physical systems in which qubits have been implemented, including single neutral atoms [9], photons [10], ions [11], superconducting Josephson junctions [12], optomechanical resonators [13], and quantum dots [14]. However,

demonstrations so far have been limited to a handful of qubits [15] and scaling to larger numbers remains a challenge. In addition, there are a variety of schemes for implementing two qubit gates, using vibrational coupling of ions [16], cold controlled collisions between atoms [17], Rydberg states [18], geometric phases [19], and cavity-mediated interactions [20]. Many of these schemes rely on being able to precisely position qubits, for example, to control the interaction strength or to maximise coupling to a cavity. We can compare these vastly different quantum systems (see figure 1.1) by the rate of coherent coupling g between qubits versus the rate incoherent coupling κ to the environment. The ratio g/κ is a figure of merit which roughly corresponds to the number of gate operations that can be performed before an error occurs due to decoherence. Error correction techniques work as long as the gate errors are below a certain threshold. Fowler formally proved that this threshold is $< 7.4 \times 10^{-4}$ [21], and numerical simulations suggest that it is as high as 10^{-2} [22, 23]. As illustrated in figure 1.1, one-qubit rotations are already well below this threshold, and rapid progress is being made towards improving the error rate of two-qubit gates.

In addition to the first five DiVincenzo criteria, it is important to be able to transfer quantum information between distant quantum systems. The resulting quantum network [34] is a hybrid quantum system consisting of stationary qubits (for example, single atoms) and flying qubits (single photons). At each node of the network, stationary qubits are converted into flying qubits using, for example, Fabry-Perot cavities [35]. Such networks provide a general framework for implementing several quantum devices, for example, quantum repeaters [36], surface codes for error correction [37], and even a one-way quantum computer [38]. From a practical standpoint, the hybrid nature of the quantum network allows one to exploit the relative strengths of different quantum systems, for example, using strong interactions between atoms, as well as the high speed and excellent coherence properties of photons.

We have chosen one particular architecture for quantum information processing,

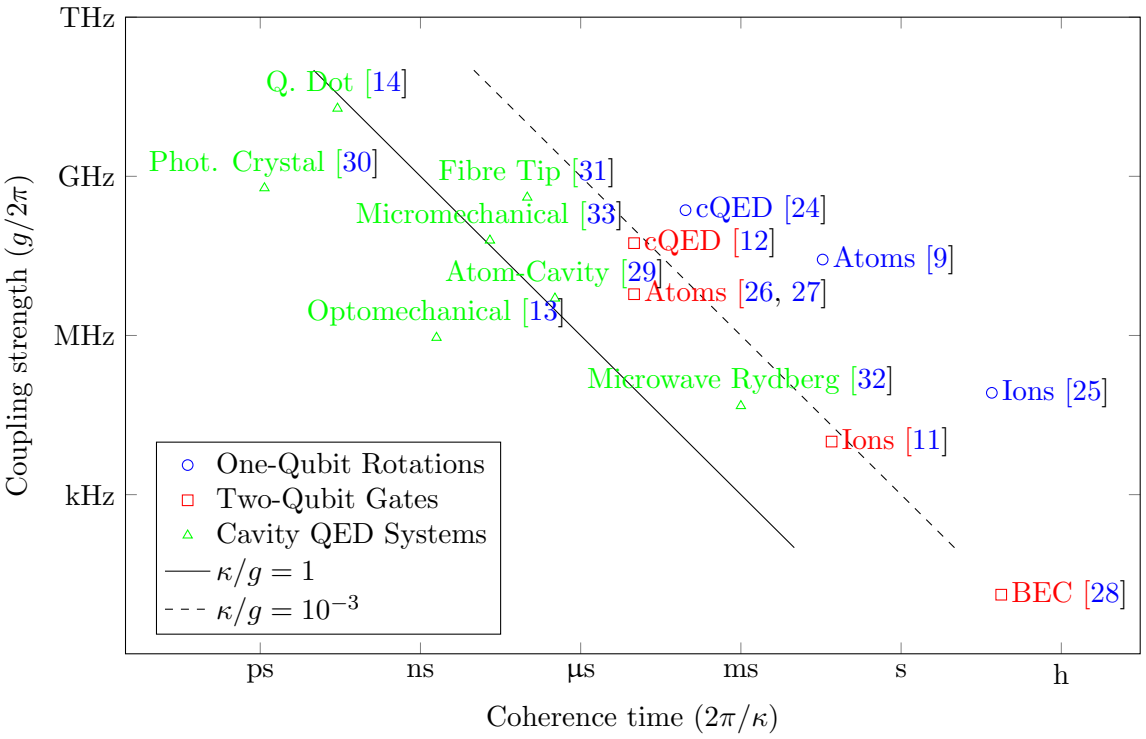

Figure 1.1: Comparison between different implementations of qubits. The coupling strength g has different meanings for the different systems. For one qubit gates, g is the Rabi frequency; for two qubit gates, g is the coupling between two qubits; and for cavity QED systems, g is the vacuum Rabi frequency.

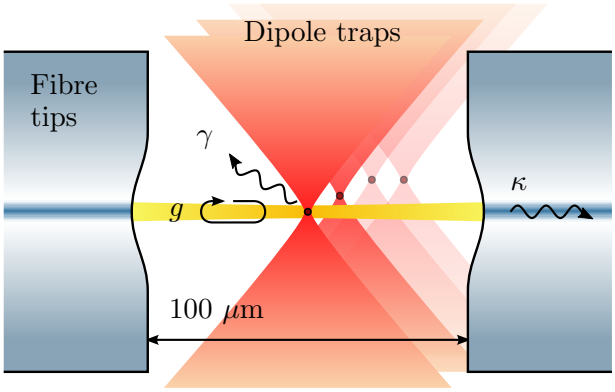

Figure 1.2: Atom-cavity architecture. Single atoms are coupled to single photons inside an optical Fabry-Perot cavity. The cavity is formed from two optical fibre tips with highly reflective mirror coatings. Atoms can be selectively positioned within the optical field mode of the cavity using microscopic dipole traps. The parameters g, κ and γ are the rates of coherent atom-cavity coupling, cavity field decay, and spontaneous emission.

namely, using single neutral rubidium atoms as stationary qubits and single photons as flying qubits. In this thesis, I describe a method for trapping, individually addressing, and moving these single atoms. Our technique addresses three major obstacles in building a computer: firstly, we can trap large numbers of atomic qubits (addressing DiVincenzo criterion #1). We demonstrate the trapping of 20 atoms, with prospects of scaling up further by increasing the laser power. Secondly, our scheme can trap atoms in arbitrary geometries, which is necessary for position-dependent quantum gates. Finally, our method is ideally suited for coupling of atoms to a cavity for building a quantum network.

1.1 Atoms, cavities and optical tweezers

Our architecture for a node in a quantum network is illustrated in 1.2. In rubidium, the $|0\rangle$ and $|1\rangle$ qubit states may be chosen to be two Zeeman sublevels of the hyperfine ground states (which correspond to two different orientations of the nuclear magnetic moment). Extremely long coherence times of 53 seconds have been observed for these states at room temperature [39]. On the other hand, single photons are a good

choice for flying qubits as they are easily transmitted by low loss optical fibres. The conversion of the quantum state from the atom to the photon is accomplished by exciting the atom from the side, after which it emits a photon into the cavity mode. Eventually, the photon leaks out of the cavity due to the imperfect reflectivity of one of the mirrors. Thus the system acts as a deterministic source of single photons, or *photon pistol*. This idea was first introduced by Kimble [40], and developed by Rempe and Kuhn [41] and others.

In a concurrent experiment in our research group, Kuhn et al. [42, 43] have built a photon pistol consisting of a single atom loaded into a centimetre-scale optical cavity with a finesse of 80,000 and a cavity coupling strength $g = 2\pi \times 12\,\mathrm{MHz}$. Rempe et. al. [29], using a similar apparatus, have demonstrated the re-absorption of the photon by a second atom, forming a two-node quantum network. There are two requirements to scale up these demonstrations: a smaller cavity with higher g resulting in higher photon production efficiency; and the ability to trap hundreds of atoms in an individually addressable way. For the first point, we propose using fibre-tip cavities [31]. These cavities are made by melting the end of an optical fibre with a CO_2 laser to form a curved surface, to which a highly reflective mirror coating is applied. Our fibre-tip cavities have a finesse of 120,000 [44] and a projected $g > 100\,\mathrm{MHz}$, with the output directly coupled into an optical fibre. For the second point, we propose to use optical tweezers to trap and transport large numbers of atoms. Optical tweezers are formed from tightly focused laser beams (see fig. 1.2). The beam is red-detuned from the atomic resonant frequency, causing atoms to be trapped due to the dipole force. In this work, we use a digital mirror device (DMD) [45, 46] to generate multiple independent tweezers from a single dipole trapping beam, an idea first proposed by [47]. We use the DMD to display a hologram, which is converted into the desired arrangement of traps in the focal plane of a lens. The main advantage of DMDs is their high frame rate of 20 kHz, which enables high speed transport of atoms.

Another promising candidate for the implementation of stationary qubits are single ions. The main advantage of ions compared with neutral atoms is that they are much easier to trap because of their non-zero electric charge. Ion traps typically have a depth $> 1000\,\mathrm{K}$, while atom traps are on the scale of $\sim 1\,\mathrm{K}$ for a magneto-optical trap or $\sim 1\,\mathrm{mK}$ for a dipole trap. To illustrate the point, the first image of a single trapped ion was made in 1980 [48], while that of a single atom wasn't made until 1994 [49, 50]. The main advantage of neutral atoms is their ability to be positioned close to dielectric (and therefore electrically insulating) surfaces, for example, those that make up the mirrors of an optical cavity. In the case of ions, there is a risk that, during the ionisation process or otherwise, electrons can become stuck on the dielectric surface and gradually build up stray charge. The electric fields from these stray charges can disrupt the trapping fields and prevent the ion trap from working. It was only recently, a single ion was coupled to a microscopic optical cavity [51], despite over a decade of research [52]. This was achieved by first trapping the ion without a cavity to avoid charging the mirrors, and then moving them into position (a distance of several millimetres) with a mechanical translation stage.

1.2 State of the art

The concept of trapping particles with light was pioneered by Arthur Ashkin in 1970 [53, 54]. There are two aspects to the mechanical force exerted by light on matter: the radiation pressure, which acts along the direction of the beam; and the gradient force, where the particle is attracted toward the point of highest intensity. The optical Earnshaw theorem [55], analogous to the Earnshaw theorem from electrostatics, states that no particle can be trapped by radiation pressure alone, however, for a tightly focused beam, the large intensity gradient can be used to trap particles.

With the advent of laser cooling, gradient force traps were used as a versatile

method of trapping cold atoms [56, 57, 58]. Atoms are so small relative to the wavelength of light that only the dipole polarisability is relevant, hence the 'dipole force'. Dalibard and Cohen-Tannoudji described the theory of this dipole force in terms of dressed states of atoms and light - a form which is well-suited to cold atoms [59]. Wieman et al. performed detailed studies of loading cold atoms into these dipole traps [60]: they were able to load 8×10^6 atoms (40% of the atoms from a MOT) into a dipole trap with a waist of 58 µm and a depth of 1 mK.

Proposals for quantum gates with neutral atoms (e.g. collisional gates [17] and Rydberg gates [18]) provide a strong motivation for trapping single atoms. At the Institut d'Optique, Grangier et al. extended the dipole trapping technique to single atoms by using a 1 µm-sized dipole trap. During the course of these experiments, they demonstrated the collisional blockade effect [61] (see section 2.2) in which light-assisted collisions limit the number of atoms per trap to either zero or one. Meanwhile, they developed techniques to measure the temperature [62] of the trapped atoms, while the Regal group have recently implemented Raman cooling to reduce the temperature of the atoms to the motional ground state of the trap [63].

To address the challenge of scaling up these dipole traps, Weiss et al. demonstrated a three-dimensional lattice of atoms [64] with single-site resolution. Single site addressability was demonstrated by Bloch and Kuhr [28] by loading a BEC into an optical lattice. To extend the range of trapping geometries beyond simple periodic lattices, Grangier used a liquid crystal spatial light modulator to holographically generate multiple trapping sites [65]. These results have recently been extended to larger numbers of atoms in more complicated geometries [66]. The performance of DMDs for holographic beam shaping have also been explored by Zupancic [67]. Birkl et al. used a combination of a DMD and a micro-lens array to create large numbers of traps [68]. The group of Mark Saffman has made particularly innovative traps using a combination of polarisation optics and diffractive elements. They have also made

blue-detuned bottle beam [69, 70] traps, which trap atoms at an intensity *minimum* using a repulsive dipole force.

Toward the goal of moving atoms, Grangier was able to transport atoms in two dimensions using a piezo tip-tilt mirror [71]. In parallel, the group of Dieter Meschede trapped and transported single atoms in a running wave dipole trap [72, 73, 74]. Birkl et al. used an acousto-optical deflector (AOD) to transport lines of qubits in microlens arrays [75].

Finally, prospects for processing quantum information with single atoms has been explored through the demonstration of entanglement [27] and a controlled-NOT gate [26] using Rydberg interactions, matter wave Hong-Ou-Mandel interference [76], and coupling to a photonic crystal cavity [30]. The error rate of Rydberg gates scales as

$$E \sim \frac{1}{(B\tau)^{2/3}} \tag{1.2}$$

where B is the frequency shift due to the interaction between neighbouring Rydberg atoms and τ is their natural lifetime, which is realistically estimated to be $E < 10^{-3}$ [77]. There are good prospects for improving current experiments [78] to approach these limits.

1.3 Thesis outline

This project builds on the work of Brandt, Muldoon and Dong on a previous experiment in our group [79, 1]. In this former experiment, multiple dipole traps were created by *directly imaging* the surface of the DMD, as opposed to holographically.

The structure of this thesis is as follows: in chapter 2, I describe the theory of laser cooling and dipole trapping of single atoms. I also present a calculation of the field of the strongly focused laser beam. In chapter 3, I reproduce the paper [1] describing the former experiment for trapping and transporting atoms. This is included as back-

ground to motivate the design of our new apparatus. In chapter 4, I describe this new apparatus in detail, along with preliminary experimental results to characterise its performance. In chapter 5, I describe our algorithm for calculating binary holograms for the DMD, and how to use these to correct for aberrations. Finally, in chapter 6, I present the main experimental results for trapping and transporting single atoms. I measure the temperature and lifetime of the trapped atoms, and demonstrate that they can be trapped in arbitrary arrangements.

Chapter 2

Theory

The following chapter describes aspects of the theory of laser cooling and trapping that are relevant to our experiments. We follow the approach and notation of D. A. Steck [80]. More detailed treatments of the subject can be found in [81, 82].

First, we describe the interaction between light and matter and derive the formula for the dipole force - the principle of operation of our atom traps. Next, we describe the principles of laser cooling, which we use to load atoms in to the dipole traps. We then provide a simple model for the electric field of a focused Gaussian light beam. Finally, we present a complete numerical diffraction calculation of the focal field, which goes beyond the paraxial approximation. We show that this gives rise to unusual effects such as spatially varying polarisation in the focal plane, and describe the implications for our experiment.

2.1 Light-atom interaction

To illustrate the basic principles of laser cooling and trapping, it is sufficient to consider a two-level atom interacting with laser light. The transition between the two levels, $|g\rangle$ and $|e\rangle$, has an associated electric dipole moment, $\mathbf{d} = \langle e| \, q\mathbf{r} \, |g\rangle$, where q is the charge and $\mathbf{r}$ is the position of the electron. The light can be described by an

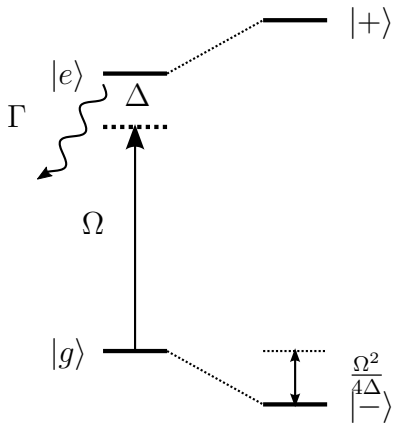

Figure 2.1: Energy levels in a two-level atom interacting with an off-resonant laser. The laser is detuned by Δ and couples the ground and excited states $|g\rangle$ and $|e\rangle$ with a Rabi frequency Ω. The excited state decays at a rate Γ due to spontaneous emission. The interaction causes the ground and excited states to form light-shifted dressed states $|-\rangle$ and $|+\rangle$ (defined in equation 2.8.

oscillating electric field, $\mathbf{E} = \mathbf{E}_0 \cos(\omega t)$, with angular frequency ω. The interaction between the atom and the field perturbs the energy of the atomic levels by an amount $\mathcal{H}_{int} = -\mathbf{d} \cdot \mathbf{E}$. The Hamiltonian for the system is

$$\mathcal{H} = \hbar\omega_g |g\rangle\langle g| + \hbar\omega_e |e\rangle\langle e| + \hbar\Omega \cos(\omega t)(|e\rangle\langle g| + |g\rangle\langle e|), \tag{2.1}$$

or in matrix form,

$$\mathcal{H} = \hbar \begin{pmatrix} \omega_g & \Omega\cos(\omega t) \\ \Omega\cos(\omega t) & \omega_e \end{pmatrix}, \tag{2.2}$$

where $\hbar\omega_g$ and $\hbar\omega_e$ are the energies of the atomic levels, and where we have introduced the Rabi frequency: $\Omega = \mathbf{d}\cdot\mathbf{E}_0/\hbar$. The magnitude of the electric field can be calculated from the laser intensity, $I = \frac{1}{2}\epsilon_0 c E_0^2$. The Hamiltonian of the two level atom can be transformed using the rotating wave approximation [81] to

$$\mathcal{H} = \frac{\hbar}{2} \begin{pmatrix} \Delta & \Omega \\ \Omega & -\Delta \end{pmatrix}, \tag{2.3}$$

where we define the detuning of the laser relative to the transition frequency, $\Delta = \omega - (\omega_e - \omega_g)$. This Hamiltonian gives rise to two kinds of interaction: on-resonant scattering in which the atom absorbs and emits photons, and the off-resonant dipole interaction in which the laser exerts a force on the atom without absorbing any photons.

2.1.1 The scattering force

When the laser frequency, ω, is resonant with the atomic transition frequency, $\omega_0 = \omega_e - \omega_g$, atoms are transferred between the ground and excited states. The Hamiltonian treatment is not sufficient to calculate the populations of the levels because it does not account for the possibility of spontaneous emission. The rate of spontaneous emission, Γ, depends on the transition dipole moment

$$\Gamma = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}|\mathbf{d}|^2. \tag{2.4}$$

The populations of the ground and excited states are obtained by solving the optical Bloch equations [80]. The equilibrium population in the excited state, ρ_{ee}, is

$$\rho_{ee} = \frac{\Omega^2}{\Gamma^2 + 4\Delta^2 + 2\Omega^2}. \tag{2.5}$$

The scattering rate of spontaneously emitted photons depends on the population of the excited state: $R_{\mathrm{sc}} = \Gamma\rho_{\mathrm{ee}}$. Also, Ω^2 is proportional to E^2 and hence the intensity of the light. By replacing $2\Omega^2/\Gamma^2 = I/I_{\mathrm{sat}}$, we can rewrite equation 2.5 as

$$R_{\mathrm{sc}} = \frac{\Gamma}{2}\frac{I/I_{\mathrm{sat}}}{1 + 4\Delta^2/\Gamma^2 + I/I_{\mathrm{sat}}}, \tag{2.6}$$

where we have defined the saturation intensity, $I_{\mathrm{sat}} = \hbar\omega_0^3\Gamma/12\pi c^2$.

It is helpful to consider some limiting cases of this formula. When the laser

intensity is high ($\Omega^2 \gg \Gamma^2, \Delta^2$), as is the case for the cooling lasers, the scattering rate saturates at a maximum value, $R_{\rm sc} \to \Gamma/2$. When the laser beam is very far detuned ($\Delta^2 \gg \Omega^2, \Gamma^2$), which is the case for the dipole trapping laser, the scattering rate can be approximated as $R_{\rm sc} \approx \Gamma\Omega^2/4\Delta^2$. Because of the $1/\Delta^2$ dependence, the scattering rate decreases rapidly as the laser is detuned further from resonance.

The scattering *force* arises from conservation of momentum of the absorbed and emitted photons. Each photon has momentum $p = \hbar k$ and so the force can be calculated as follows

$$F_{\rm sc} = \frac{dp}{dt} = \hbar k R_{\rm sc}. \tag{2.7}$$

The force acts along the directing of the laser beam. The average force is equivalent to the classical radiation pressure $P_{\rm rad} = I/c$. The quantum mechanical treatment accounts for saturation of the atom, as well as fluctuations of the force due to the fact that the light is made up of photons.

2.1.2 The dipole force

The second effect of the laser is to perturb the ground and excited state energies. This perturbation is known as the ac-Stark shift or light shift and is the underlying mechanism for our dipole traps.

To derive the formula for the light shift, we first consider the Hamiltonian (eq. 2.3) in the absence of radiation i.e. $\Omega = 0$. The eigenstates of the atom are $|g\rangle$ and $|e\rangle$ with eigenvalues (in the rotating frame) of $\frac{\hbar}{2}\Delta$ and $-\frac{\hbar}{2}\Delta$. In the presence of radiation, the eigenstates of the atom are the *dressed* states [59]

$$\begin{aligned} |+\rangle &= \sin\theta\,|g\rangle + \cos\theta\,|e\rangle \quad \text{with eigenvalue} \quad +\tfrac{\hbar}{2}\sqrt{\Omega^2+\Delta^2}, \\ |-\rangle &= \cos\theta\,|g\rangle - \sin\theta\,|e\rangle \quad \text{with eigenvalue} \quad -\tfrac{\hbar}{2}\sqrt{\Omega^2+\Delta^2}, \end{aligned} \tag{2.8}$$

where θ satisfies

$$\cos 2\theta = \frac{-\Delta}{\sqrt{\Delta^2 + \Omega^2}}, \quad \sin 2\theta = \frac{\Omega}{\sqrt{\Delta^2 + \Omega^2}}. \tag{2.9}$$

Consider the case of red-detuned light with $\Delta < 0$ (and therefore $\theta \approx 0$). The $|-\rangle$ dressed state is a superposition consisting mostly of the ground state with a small proportion of the excited state. As the intensity is increased, the energy of the $|-\rangle$ state decreases from $-\frac{\hbar}{2}|\Delta|$ to $-\frac{\hbar}{2}|\Delta|\sqrt{1 + \Omega^2/\Delta^2}$. If we assume that the light is sufficiently far detuned such that $\Delta^2 \gg \Omega^2$, we can approximate the square root by its Taylor series $1 + \Omega^2/(2\Delta^2)$. The light shift is therefore

$$U_{\text{dip}} \approx \frac{\hbar\Omega^2}{4\Delta} = \frac{\hbar\Gamma^2}{8\Delta}\frac{I}{I_{\text{sat}}}. \tag{2.10}$$

The spatially varying light shift forms a conservative potential, U_{dip}, in which atoms can be trapped. The potential is proportional to Ω^2 and hence the intensity of the light. For $\Delta < 0$, i.e. red-detuning, the potential is attractive. The dipole *force* is the gradient of this potential

$$F_{\text{dip}} = -\nabla U_{\text{dip}}. \tag{2.11}$$

This force only exists when there is a spatial variation in intensity. For example, for a laser beam with a Gaussian profile, the dipole force is strongest at the edges of the beam where the gradient is steepest. For a tightly focused red-detuned laser beam, atoms are attracted to the the focal point where the intensity of the light is highest and U_{dip} is minimised.

2.1.3 Dipole force for multi-level atoms

The simple formula for the light shift of a two-level atom (equation 2.10) can be generalised to multi-level atoms. The multiple levels will in general have different

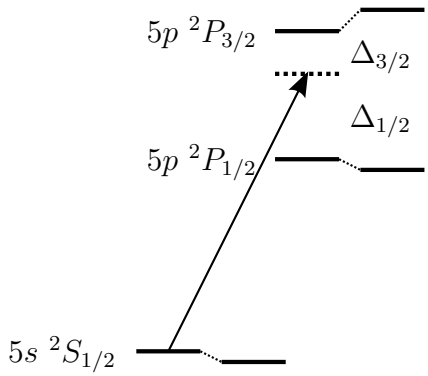

Figure 2.2: Energy levels of rubidium with a laser tuned between two excited states. The light shift due to the $^2P_{3/2}$ state and the $^2P_{1/2}$ are in opposite directions. The net result is that the ground state is shifted to lower energy.

energies and different dipole moments. The total light shift will be the sum of the contributions from each level. In the case of degenerate magnetic sublevels, the dipole moments are related by Clebsch-Gordon coefficients [80].

In the case of rubidium-87, there are two relevant transitions: the $5s\ ^2S_{1/2} \rightarrow 5p\ ^2P_{1/2}$ at 795 nm, and $5s\ ^2S_{1/2} \rightarrow 5p\ ^2P_{3/2}$ at 780 nm. Our dipole trapping laser at 785 nm is red-detuned from the $5p\ ^2P_{3/2}$ state and blue-detuned from the $5p\ ^2P_{1/2}$ state. The ground state experiences opposite light shifts from each excited state, however, the $5p\ ^2P_{3/2}$ state dominates (since the laser is tuned closer and since it has twice the number of degenerate sublevels) and the ground state is shifted to lower energy. Figure 2.3 shows the trap depth and scattering rate for a range of wavelengths.

When it comes to the light shift of the excited state, there is an additional contribution from a transition to the $5\,^2D$ state at 776 nm. This lowers the energy of the excited state by approximately 10%.

Even if the levels are degenerate, for example magnetic sublevels, then the light shift will depend on the dipole moment of each sublevel individually, which in turn depends on the orientation of the atom relative to the electric field of the light. For ground state alkali atoms, linearly polarised light, and where the detuning is much greater than the hyperfine splitting, the dipole potential is independent of the orien-

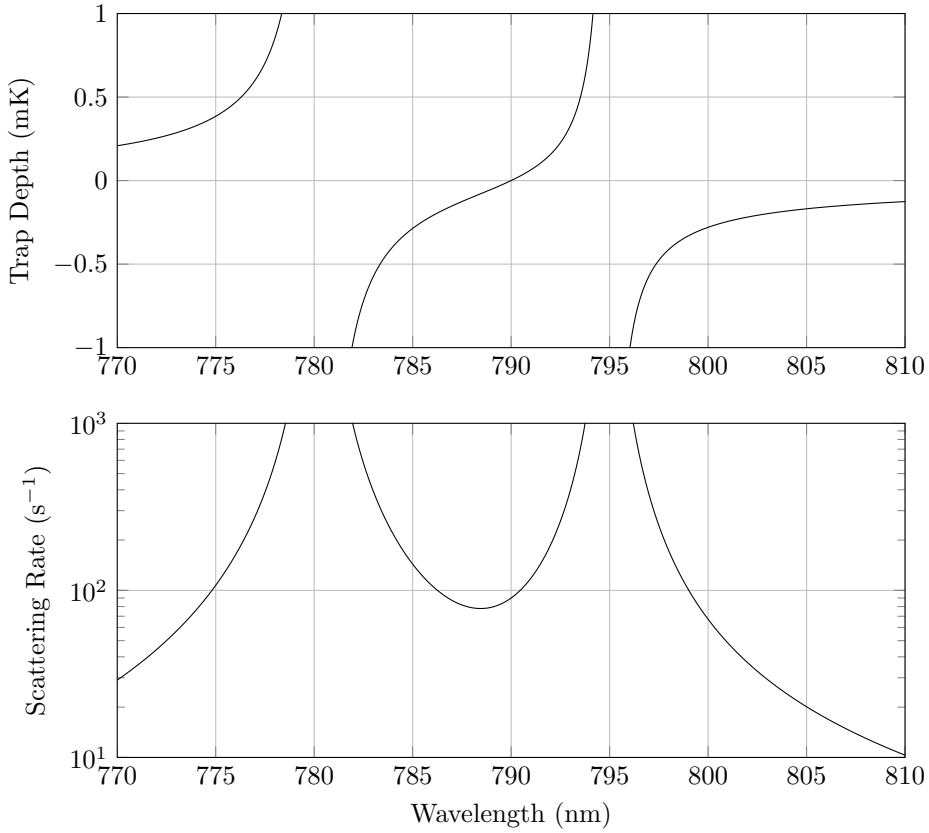

Figure 2.3: Trap depth and scattering rate for different wavelengths for an intensity of $I_0 = 10^8\,\mathrm{Wm}^{-2}$. Both the trap depth and scattering rate diverge near the transitions at 780 nm and 795 nm.

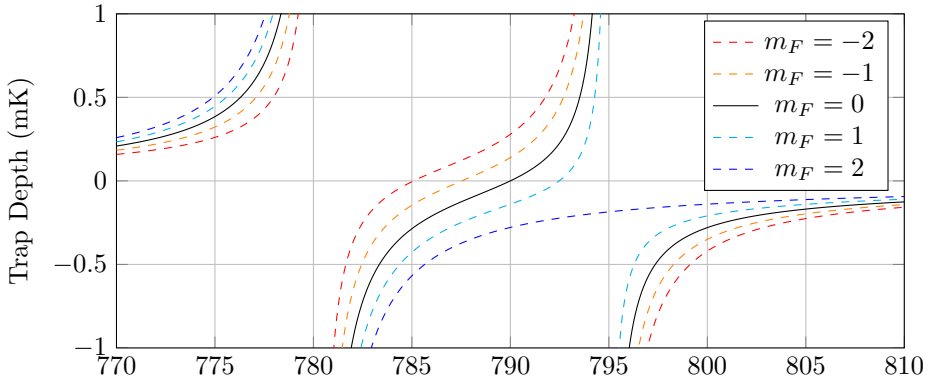

Figure 2.4: Trap depth with circularly polarised light. The different m_F states of the $F = 2$ ground state have different trap depths.

tation of the atom. This is because the $^2S_{1/2}$ state has only two possible orientations, $m_J = \pm\frac{1}{2}$, which have the same Clebsch-Gordon coefficients. However, for elliptically polarised trapping light, the dipole potential does depend on the orientation of the total angular momentum of the atom, m_F. This effect is analogous to the splitting of m_F states in a magnetic field. The complete formula [60] for the dipole potential taking into account the two transitions and the m_F dependence is

$$U_{\text{dip}} = \frac{\hbar\Omega^2}{4\Delta}\left(\left(\frac{1}{\Delta_{1/2}} + \frac{2}{\Delta_{3/2}}\right) - g_F m_F \sqrt{1-\epsilon^2}\left(\frac{1}{\Delta_{1/2}} - \frac{1}{\Delta_{3/2}}\right)\right). \tag{2.12}$$

where ϵ is the ellipticity of the light, which is 0 for circularly polarised light and 1 for linearly polarised light. Figure 2.4 shows the dipole trapping potential as a function of wavelength for circularly polarised light.

2.1.4 Doppler cooling

Doppler cooling is the primary mechanism which we use for cooling atoms into the dipole trap. The technique uses the scattering force from a laser beam to slow down atoms. An atom moving with velocity v experiences a laser frequency that is Doppler-shifted by $\Delta = kv$. If the laser is red-detuned from the atomic transition, then atoms moving towards the laser beam will be Doppler-shifted closer to resonance,

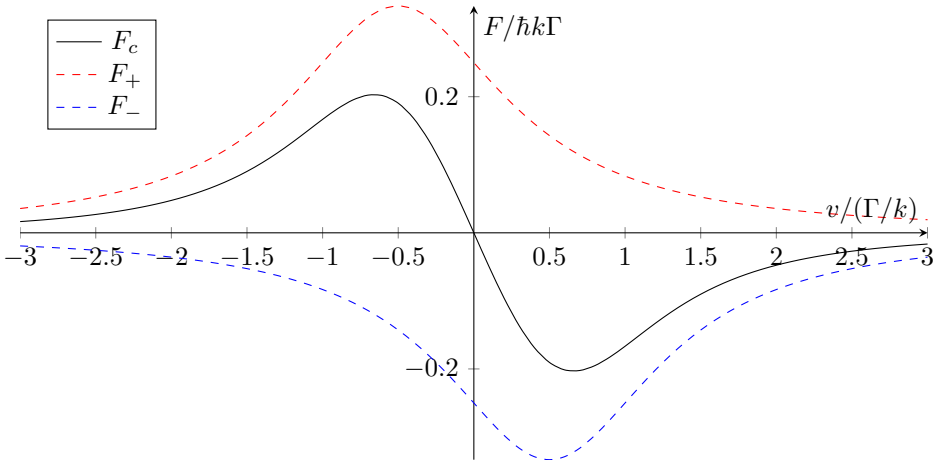

Figure 2.5: Cooling force as a function of velocity. For small velocities, the total force is approximately linear. The total cooling force is given by the sum of the left and right travelling beams, F_+ and F_-. The force and velocity are given in dimensionless units in which $\Gamma = 1$.

thus scattering more photons and being slowed down. Conversely, atoms moving away from the beam will be Doppler-shifted further away from resonance and will experience a smaller scattering force. By red-detuning the laser in this way, one can selectively slow only atoms travelling towards the laser beam. Consider a one-dimensional scenario with two counter-propagating laser beams. The force on the atom is given by

$$F_{\mathrm{c}} = \hbar k \Gamma \left(\frac{\Omega^2}{\Gamma^2 + 4(\Delta - kv)^2 + 2\Omega^2} - \frac{\Omega^2}{\Gamma^2 + 4(\Delta + kv)^2 + 2\Omega^2} \right). \tag{2.13}$$

The cooling force calculated from this equation is shown in figure 2.5. For small velocities, the force can be approximated by

$$F_{\mathrm{c}} = \frac{16\hbar k^2 \Gamma \Delta \Omega^2}{(\Gamma^2 + 4\Delta^2 + 2\Omega^2)^2} v. \tag{2.14}$$

The cooling force is proportional to velocity and therefore has the same form as a frictional force, $F = -\alpha v$. Since the detuning is negative, the force acts to reduce the velocity (and hence the temperature) of the atom. The maximum cooling force

occurs when $\Delta = -\Gamma/2$ and $\Omega = \Gamma$. Inserting these parameters

$$F_{\max} = -\frac{\hbar k^2}{2} v. \tag{2.15}$$

The cooling force is counter-balanced by recoil heating due to spontaneously emitted photons. On average, the absorbed photons reduce the velocity of the atom, however, the spontaneous photons are emitted in random directions, each one giving a small momentum kick, $p = \hbar k$, to the atom. Each kick increases the kinetic energy of the atom by $E_{\mathrm{r}} = p^2/2m = \hbar^2 k^2/2m$. There is a limit to the minimum energy of the atoms which occurs when the heating rate due to recoil from spontaneously emitted photons is equal to the cooling rate from the Doppler cooling mechanism:

$$\left(\frac{dE}{dt}\right)_{\mathrm{heat}} + \left(\frac{dE}{dt}\right)_{\mathrm{cool}} = R_{\mathrm{sc}}.E_{\mathrm{r}} + F_{\max}.v = \frac{\Gamma}{2}\frac{\hbar^2 k^2}{2m} - \frac{\hbar k^2}{2} v^2 = 0 \tag{2.16}$$

Substituting $v^2 = k_{\mathrm{B}} T/m$ gives the Doppler limit as a temperature:

$$k_{\mathrm{B}} T_{\mathrm{D}} = \frac{\hbar \Gamma}{2}. \tag{2.17}$$

For rubidium-87, $T_{\mathrm{D}} = 146\,\mu\mathrm{K}$.

2.1.5 Sub-Doppler cooling

In practice, laser-cooled atoms reach much colder temperatures than suggested by the Doppler limit. This is because of a variety of sub-Doppler cooling mechanisms known as Sisyphus cooling [82, 83]. The counter-propagating cooling beams interfere to form a standing wave, which creates a periodic dipole potential (as in section 4.2). This potential depends on whether the atom is in the ground or excited state as the corresponding light shifts have opposite sign. As ground state atoms climb the hills of this potential, they lose kinetic energy. When they reach a potential maximum,

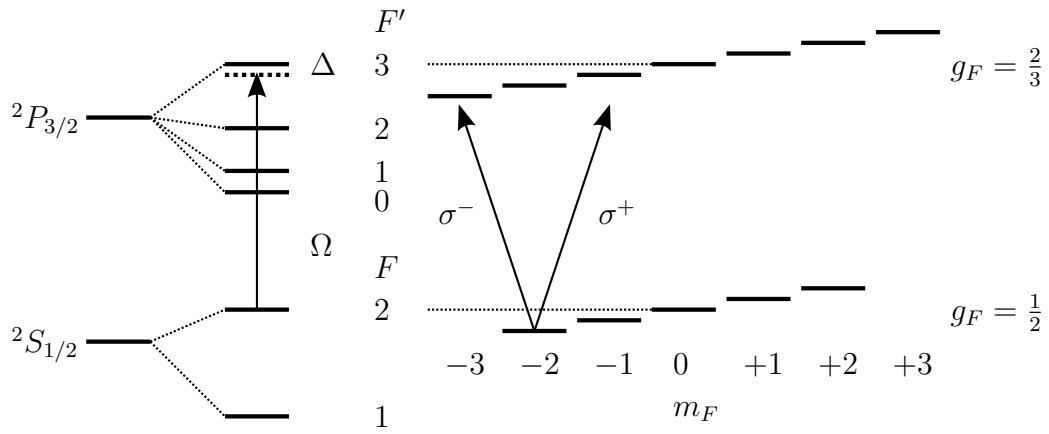

Figure 2.6: Energy levels of rubidium-87 involved in laser cooling. The ground and excited levels (left) are split in to hyperfine levels labelled by their total angular momentum, F. The cooling laser is tuned to the $F = 2 \rightarrow F' = 3$ transition. In a magnetic field, the levels are split according to their Zeeman sublevel, labelled by m_F (right). The left and right circularly polarised laser beams, σ^+ and σ^-, couple different sublevels.

the atoms have a high probability of absorbing a photon and becoming excited. The potential maximum suddenly becomes a potential minimum and the loss of kinetic energy through climbing the hill is made irreversible. The ultimate temperature limit in this case is given by the recoil limit due to the momentum of a single photon,

$$k_{\mathrm{B}} T_{\mathrm{r}} = \hbar^2 k^2 / 2m. \tag{2.18}$$

For rubidium, this limit is 362 nK [80]. The actual temperature depends on the intensity and detuning [83] of the cooling lasers, and is typically of the order of $\sim$10 µK.

2.1.6 The magneto-optical trap

Doppler cooling alone is able to cool atoms by viscous damping, but it is not sufficient to trap atoms in one place. For trapping, one requires a restoring force that is linear in position. Such a force can be created by introducing a linear magnetic field gradient, $B_z = A.z$. The energy of the atomic level varies with magnetic field (and hence

position) due to the Zeeman effect,

$$\hbar\Delta = \mu_{\mathrm{B}} B_z (g_{F'} m_{F'} - g_F m_F). \tag{2.19}$$

In the above equation, $\mu_{\mathrm{B}} = q\hbar/2m_e$ is the Bohr magneton, g_F is the Landé g-factor, and m_F is the magnetic quantum number specifying the projection of the angular momentum along the direction of the magnetic field. For the $F = 2$ to $F' = 3$ cooling transition in rubidium-87, $g_F = 1/2$ and $g_{F'} = 2/3$.

The natural quantisation axis of the atoms depends on the local value of the magnetic field. For the position dependent force to work, the atoms must interact preferentially with one laser at a time. Therefore, the cooling laser beams are chosen to be circularly polarised in opposite directions. Figure 2.6 shows the relevant Zeeman-shifted energy levels and transitions due to right and left circularly polarised cooling beams. Being continuously exposed to circularly polarised light, the atoms absorb angular momentum from the laser beams and are optically pumped in to the stretched state, $m_F = -2$. The circular polarisations are chosen so that one of the cooling lasers addresses the $m_{F'} = -3$ excited state (which is subject to the largest Zeeman shift) and the other laser addresses the $m_{F'} = -1$ state. The result is a position dependent restoring force. If we make the simplification that $g_F = g_{F'} = 1$ and $m_F = 0 \rightarrow m_{F'} = -1$, then the restoring force is given by

$$F = \frac{A\mu_{\mathrm{B}} k}{2}. \tag{2.20}$$

For atoms with different g and m_F values, the force will be proportional to this.

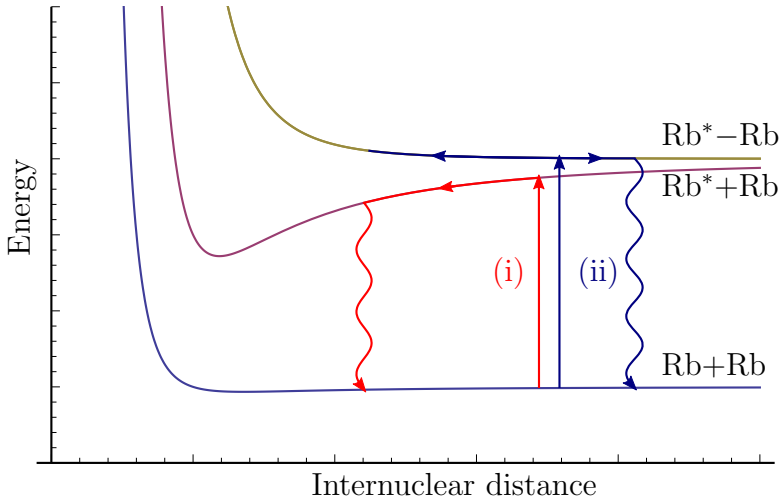

Figure 2.7: Molecular potential for ground and excited state of Rb_2. Red-detuned light (i) excites the atoms to form a bound molecule. The atoms accelerate towards each other and gain kinetic energy before returning to the ground state via spontaneous emission. Blue-detuned (ii) light excites atoms to an unbound state in which the atoms are decelerated by the molecular potential.

2.2 Collisional blockade

Ordinary collisions between rubidium atoms are elastic, that is, they involve no change in the kinetic energy of the atoms. However, if two Rubidium atoms absorb a photon during their collision, they may be excited to a molecular state. The potential energy of this molecular state depends on the distance separating the two atoms. Figure 2.7 shows the potential energy of the molecule as a function of internuclear separation. If the photon is red-detuned from the atomic level, the molecular state is a bound state. The atoms accelerate towards the minimum of the potential, gaining kinetic energy in the process. When the molecule decays via spontaneous emission, the atoms are moving too fast to remain trapped and are both lost. Pairs of atoms continue to collide in this way until there is only a single atom remaining. The result is that the atom number is projected to either zero or one depending on whether the initial number of atoms is even or odd. The phenomenon is known as the *collisional blockade* effect [61]. We exploit this phenomenon to ensure that there is at most one atom per trapping site with a loading probability of 50%.

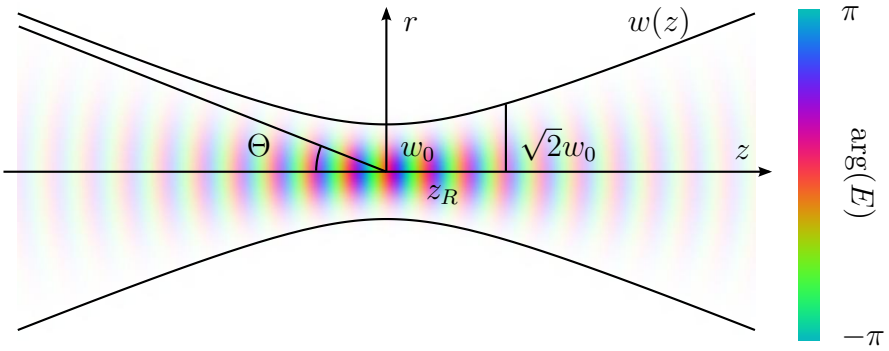

Figure 2.8: Electric field of a focused Gaussian beam. The hue represents the complex phase of the field. The waist of the beam $w(z)$ is the value of r where the intensity of the beam falls to $1/e^2$ of the maximum value. w_0 is the value of the waist at the focus, z_R is the Rayleigh length (the distance at which the waist increases by $\sqrt{2}$) and Θ is the divergence half angle.

If the photon is blue-detuned, the atoms are excited to an unbound (i.e. repulsive) molecular potential (see figure 2.7). In this case, the atoms are slowed as they climb the potential, and remain trapped when they spontaneously decay. This effect has been used to increase the probability of loading a single atom [84].

2.3 Diffraction-limited focusing

In order to reach the collisional blockade regime, atoms must be confined to a tiny volume. Therefore, the dipole trapping laser beam must be focused to as small a spot as possible. To illustrate the problem, we consider a simple model of the dipole potential due to a Gaussian laser beam. The dipole potential is proportional to the intensity of the Gaussian beam. As a function of radial and axial position, $U_\mathrm{dip}(r, z)$ is

$$U_\mathrm{dip}(r, z) = -U_0 \frac{1}{1 + z^2/z_\mathrm{R}^2} \exp(-2r^2/w(z)^2), \tag{2.21}$$

where U_0 is the value of the potential at the focus, $w(z)$ is the *waist* (the distance at which the potential drops to $1/e^2$ of the minimum value), and z_R is the Rayleigh

length. The waist and Rayleigh length are defined as

$$w(z) = w_0\sqrt{1+\frac{z^2}{z_{\mathrm{R}}^2}}, \qquad z_{\mathrm{R}} = \frac{\pi w_0^2}{\lambda}. \tag{2.22}$$

The minimum value of the waist, w_0, occurs at the focus. In terms of the divergence half-angle, Θ, the minimum waist is

$$w_0 = \frac{\lambda}{\pi \tan\Theta}. \tag{2.23}$$

The size of the waist is inversely proportional to the divergence angle of the beam, therefore, it is advantageous to use a lens with a high numerical aperture ($NA = \sin\Theta$).

In the r-direction, the trapping region is limited by the waist, while in the z-direction, the trapping region is limited by the Rayleigh length. Since the Rayleigh length depends strongly on the waist, longitudinal confinement of the atom only occurs when w_0 is of the order of λ. Near the focus, the dipole potential can be approximated as a three dimensional harmonic oscillator,

$$U_{\mathrm{dip}} = \frac{1}{2}m\omega_r^2 r^2 + \frac{1}{2}m\omega_z^2 z^2, \tag{2.24}$$

with trapping frequencies given by

$$\omega_r = \sqrt{4U_0/mw_0^2}, \qquad \omega_z = \sqrt{2U_0/mz_{\mathrm{R}}^2}. \tag{2.25}$$

2.3.1 Beyond the paraxial approximation

The Gaussian beam model is helpful for illustrating the concept of a focused beam, however, a focused Gaussian beam with uniform polarisation does not satisfy $\nabla\cdot\mathbf{E} = 0$ and hence is not a valid solution of Maxwell's equations. Such a tightly focused trap

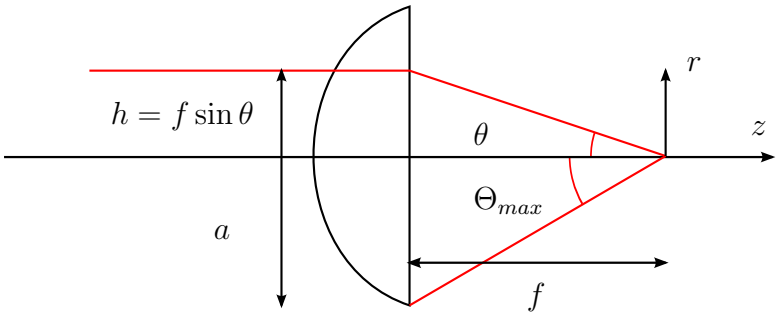

Figure 2.9: Geometry for calculation of focal field. Incoming rays a distance of h from the optical axis converge at the focus with a corresponding angle θ. a is the radius of the lens and f is the focal length.

implies that some of the incident rays must propagate at large angles with respect to the optical axis, therefore one must go beyond the paraxial approximation to obtain an accurate description of the intensity in the focal plane.

We use a model described by Novotny and Hecht [85]. The geometry of the problem is shown in figure 2.9. The incident field is assumed to be cylindrically symmetric about the optical axis, with h being the distance from the optical axis and a being the radius of the lens. The linear wavefronts of the input field are transformed by the lens in to converging circular wavefronts, with rays parameterised by θ, the angle with respect to the optical axis. The incident field, E_{inc}, is assumed to be a Gaussian beam that is linearly polarised along the x-axis,

$$E_{\text{inc}} = E_0 \exp(-h^2/w^2) = E_0 \exp(-f^2 \sin^2\theta/w^2), \tag{2.26}$$

r, ϕ, and z are cylindrical polar coordinates in the focal plane of the lens. By defining the following integrals,

$$I_{00}(r) = \int_0^{\Theta_{\max}} E_{\text{inc}}(\theta)(\cos\theta)^{1/2} \sin\theta(1+\cos\theta) J_0(kr\sin\theta) e^{ikz\cos\theta} d\theta \tag{2.27}$$

$$I_{01}(r) = \int_0^{\Theta_{\max}} E_{\text{inc}}(\theta)(\cos\theta)^{1/2} (\sin\theta)^2 J_1(kr\sin\theta) e^{ikz\cos\theta} d\theta \tag{2.28}$$

$$I_{02}(r) = \int_0^{\Theta_{\max}} E_{\text{inc}}(\theta)(\cos\theta)^{1/2} \sin\theta(1-\cos\theta) J_2(kr\sin\theta) e^{ikz\cos\theta} d\theta \tag{2.29}$$

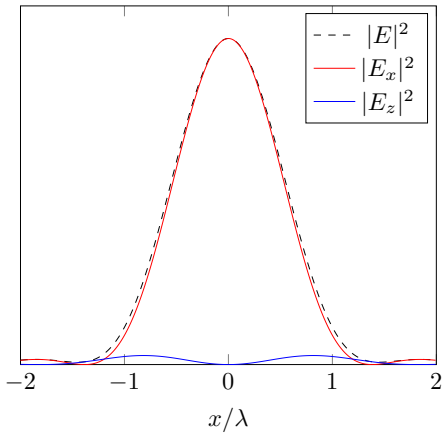

Figure 2.10: Calculated intensity of the trapping light, along with the intensity of the x and z-polarised components, in the focal plane of the aspheric lens. The y-polarised component is zero on the x axis.

we can calculate the Cartesian components of the field in the focal plane, E_x, E_y and E_z, as

$$\mathbf{E}(r,\phi) = \frac{ikf}{2} e^{-ikf} \begin{pmatrix} I_{00}(r) + I_{02}(r)\cos(2\phi) \\ I_{02}(r)\sin(2\phi) \\ -2iI_{01}(r)\cos(\phi) \end{pmatrix}. \tag{2.30}$$

We performed this calculation for our lens (with $a = 9\,\text{mm}$) in the case where the aperture of the lens is uniformly filled (i.e. where the input waist is much larger than the radius of the lens). The result is shown in figure 2.10 as the square magnitude of the various components of the field. There are two interesting features: firstly, the field is not a Gaussian but is surrounded by Airy discs due to diffraction from the circular aperture. Secondly, the focused spot has a spatially-varying polarisation. Indeed, a significant fraction of the optical power is concentrated in the longitudinal (E_z) component of the electric field.

In our experiment, the dipole trapping laser has a Gaussian waist that is comparable to the radius of the aspheric lens. We seek to maximise the peak intensity of the trapping light in the focal plane of the lens. One way to do this is to increase the

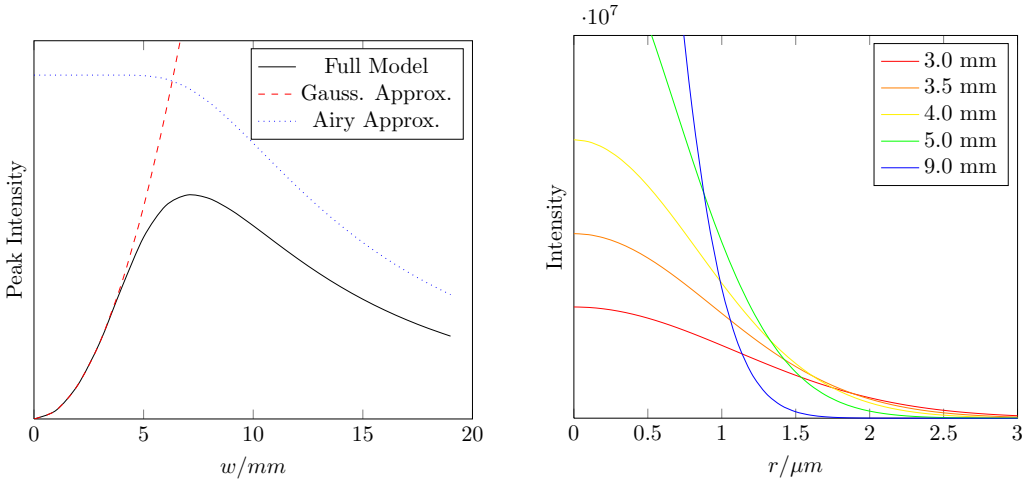

Figure 2.11: (a) Optimum waist of the trapping beam at the aspheric lens calculated from the full diffraction model. For small waists, the peak intensity is well-approximated by a Gaussian beam. For large waists, the peak intensity approaches that of an Airy disc, taking into account the fraction of the power that passes through the aperture. (b) Focal spot profiles for various aperture radii.

waist of the incident beam, in order to focus the same optical power on to a smaller diffraction-limited spot size. However, as the waist approaches the size of the lens, the power in the exponential tails does not pass through the lens aperture and is wasted. Figure 2.11a shows the peak intensity for various values of the waist. The optimum value is $w = 7.2\,\mathrm{mm}$.

Experimentally, introducing an additional aperture is a convenient way of deliberately increasing the size of the focal spot. Figure 2.11b shows the profile of the focal spot for various values of the aperture radius. These intensity profile can be approximated by fitting a Gaussian profile. This allows us to extract I_0 and w_0 parameters, which in turn allows us to estimate the trap depth and oscillation frequencies with an accuracy of $\sim 10\%$.

Now we return to the issue of additional polarisation components in the focal plane. Figure 2.12 shows the spatial intensity profile of the three components of the electric field in the focal plane. The colour scheme is used to represent the complex phase of the electric field, while the intensity is proportional to its square amplitude.

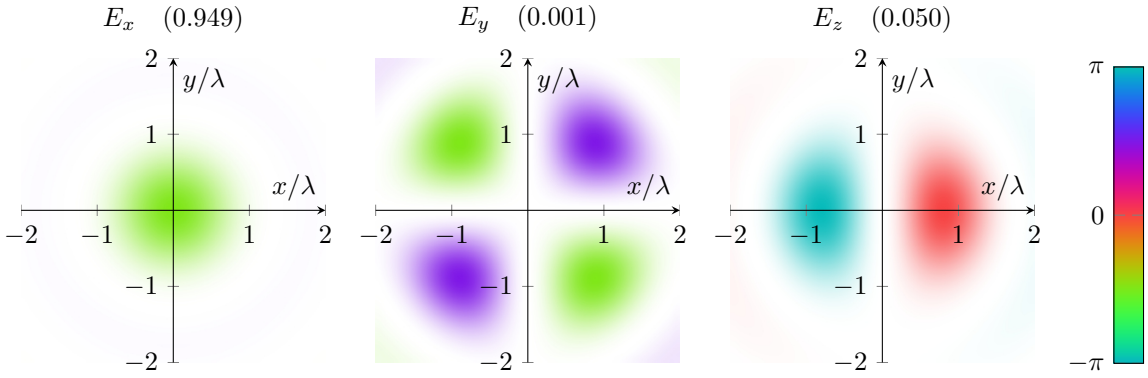

Figure 2.12: Intensity of the E_x, E_y and E_z components of the focused laser beam. The colouring represents the complex phase of the electric field. The title shows the fraction of the energy for each polarisation component.

Most of the power is concentrated in the E_x component, with 5% in the E_z component and 0.1% in E_y. Relative to the E_x component, the E_y component oscillates either in-phase or 180° out of phase, but the E_z component has a phase shift of 90°.

The sum of these linear polarisation components with various phases may lead to elliptical or even circular polarisation in some regions of the beam. For example, ignoring the very small y-component, the E_x and E_z components add together to give $E_x + iE_z$, which corresponds to circular polarisation along the y axis. In figure 2.12 the region $x < 0$ is left circularly polarised whereas $x > 0$ is right circularly polarised. Therefore, by equation 2.12, the circular polarisation dipole trapping light gives rise to an effective magnetic field B_y. In order to calculate the strength of this field, we calculate ellipticity $\epsilon = \frac{E_x^2+E_y^2-E_z^2}{E_x^2+E_y^2+E_z^2}$ and substitute this into equation 2.12.

Figure 2.13 shows the trapping potential including the effect of this non-uniform circular polarisation. The consequence is that atoms in different m_F states will experience different dipole potentials. When atoms absorb and emit photons, they will be redistributed into a random m_F state. This constant shuffling between the different potentials will heat the atoms (as documented by Kuppens et al.[60]). Furthermore, since the handedness of the circular polarisation changes from left to right along the x-axis, the trapped atoms experience an artificial magnetic field gradient. Since the

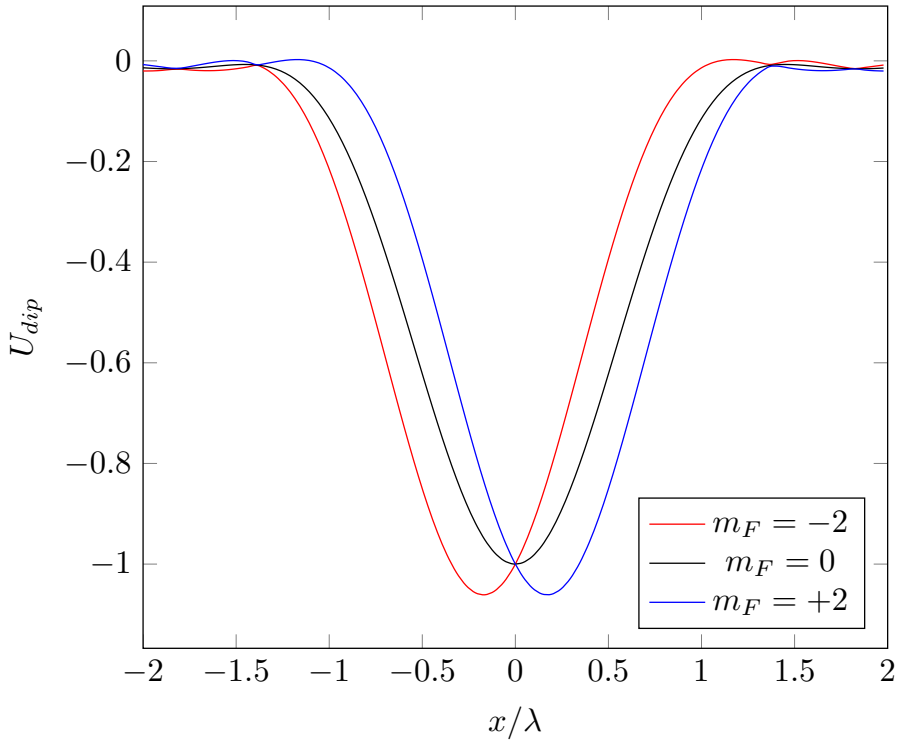

Figure 2.13: Trapping potential taking the m_F sublevels in to account.

quantum state is stored in magnetic sublevels, this artificial gradient will cause decoherence for trapped atoms. These effects become more pronounced as the light is focused to a spot that is smaller than the wavelength. For our slightly larger dipole traps, the effects are less important.

2.3.2 Spin-orbit coupling of light

As an interesting aside, it is possible to calculate the focal field of a circularly polarised input beam. A right circularly polarised beam is obtained by taking a superposition of $x-$ and $y-$polarised input beams, $\mathbf{E}^{(R)} = (\mathbf{E}^{(x)} + i\mathbf{E}^{(y)})/\sqrt{2}$. We analyse the solutions in the right/left circularly polarised basis by performing a coordinate transformation,

$$\hat{\mathbf{e}}_R = (\hat{\mathbf{e}}_x + i\hat{\mathbf{e}}_y)/\sqrt{2} \tag{2.31}$$

$$\hat{\mathbf{e}}_L = (\hat{\mathbf{e}}_x - i\hat{\mathbf{e}}_y)/\sqrt{2}. \tag{2.32}$$

The resulting solution for the focal field is given by

$$\mathbf{E}^{(x)}(r,\phi) = \frac{ikf}{2}\left(I_{00}(r)\hat{\mathbf{e}}_R + I_{02}(r)e^{-2i\phi}\hat{\mathbf{e}}_L - \sqrt{2}iI_{01}(r)e^{-i\phi}\hat{\mathbf{e}}_z\right). \tag{2.33}$$

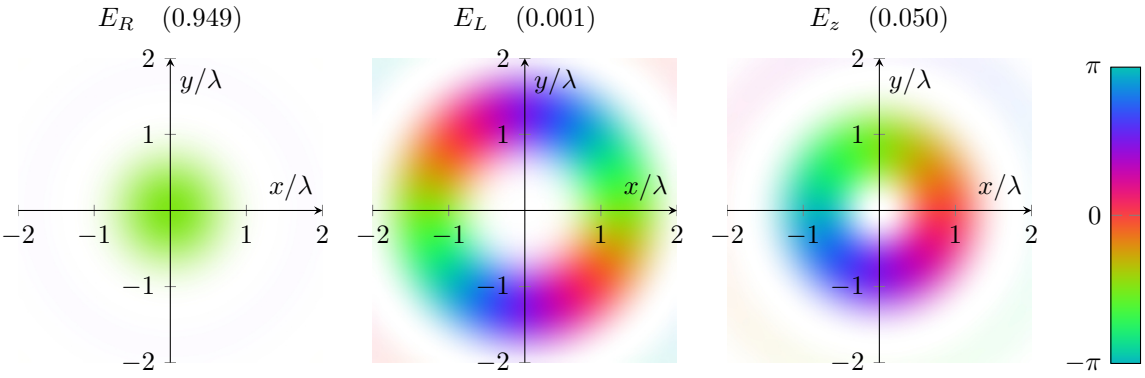

Figure 2.14: Calculated intensity of the E_R-, E_L- and E_z components for a right-circularly polarised input. The L and z components carry -2 and -1 units of orbital angular momentum, as seen by the phase winding around the optical axis. The total angular momentum is conserved for each mode.

The focal field is shown in figure 2.14. As was the case with the $x-$polarised input beam, the focus of the right-circularly polarised beam is mainly right-circularly polarised. In addition, a small amount of longitudinal and left-circularly polarised light is created. The electric field of these components contains an optical vortex in which the phase winds by 2π and 4π respectively around the optical axis. These components carry -1 and -2 units of orbital angular momentum. It is interesting to note that the sum of spin and orbital angular momentum for each component is the same as for the initial right-circularly polarised beam (i.e. -1 unit). For example the left circularly polarised component in the focal plane has -2 units of orbital angular momentum and $+1$ units of spin. The total angular momentum of -1 is conserved with respect to the input.

It may be said that the effect of focusing the light is to introduce spin-orbit coupling in to the light field. Optical spin-orbit coupling has been observed in free

space light fields [86], and been calculated theoretically in a Fabry-Perot cavity [87]. The conservation of total angular momentum in cylindrically polarised light beams has been explored by Holleczek et al. [88]. The effect has been used to create optical vortices by focusing a laser beam through a uniaxial birefringent crystal [89]. Recently, the group of Arno Rauschenbeutel has used spin-orbit coupling of light confined in an optical nano-fibre to break the mirror symmetry of the light-atom interaction [90]. Here we have shown that simply focusing a light beam is enough to introduce significant spin-orbit coupling.

2.4 Conclusions

We have covered the theory necessary for describing the cooling and trapping of single atoms. The atoms are cooled using radiation pressure in the Doppler cooling mechanism, and trapped using the dipole force. We have discussed the collisional blockade effect which is useful for obtaining single atoms. We have also presented the theory of the diffraction of a focused light beam. While this leads to unusual effects such as optical magnetism and spin-orbit coupling, these effects are only significant for very strongly focused light fields, and do not impact our current experiment.

Chapter 3

Control and Manipulation of Cold Atoms in Optical Tweezers

This chapter is adapted from the paper by Muldoon et al., *Control and Manipulation of Cold Atoms in Optical Tweezers* [1]. It describes the results from the previous experimental setup. I contributed to these experiments in the first year of my DPhil. The results are reproduced here as they provide context with which to interpret the new experimental setup.

There are several main differences between this former setup, and the new one described in chapter 4. Both use a single aspheric lens to focus the dipole trapping light and collect to the fluorescence from the atoms. Both use a DMD to control the trapping light and an EMCCD camera to image the fluorescence. However, in this former setup, the surface of the DMD was directly imaged by the aspheric lens, whereas in the new setup, the DMD is used to produce a hologram of the desired traps. Furthermore, in the former setup, a dielectric mirror was placed in the focal plane of the aspheric lens in order to retro-reflect the trapping light to form a standing wave dipole trap. This mirror was also used to create a magneto-optical surface trap (MOST) for cooling atoms. Two counter-propagating laser cooling beams were

incident on the mirror at 45°, whilst the final beam ran parallel to the mirror surface. This mirror caused a significant problem as it caused unwanted light from the cooling beams to be scattered in to the path of the EMCCD camera. This background light obscured the fluorescence signal, and prevented us from observing single atoms. In the new experiment, this mirror was removed in favour of a standard six-beam magneto-optical trap.

Despite these problems, we were able to trap clouds containing small numbers of atoms in various shaped potentials. We also managed to transport these atom clouds. The problems were eliminated in the new experimental setup.

3.1 Trap design and optical arrangement

The work presented here is the first experimental implementation of the atom trapping and manipulation scheme that we recently proposed theoretically [79]. Reconfigurable optical dipole-force traps, or optical tweezers, are produced by imaging the surface of a spatial light modulator onto a cloud of cold ^{87}Rb atoms (see Fig. 3.1). The atoms are confined in the potential caused by the resulting intensity distribution in the image plane. This potential takes the form

$$U_{dip} \approx \frac{\hbar\Omega^2}{4\delta} \tag{3.1}$$

where Ω is the Rabi frequency, and δ is the detuning of the trapping laser with respect to the closest allowed atomic transition. It results from the AC Stark shift induced by the applied field, and its gradient gives rise to a conservative force known as the dipole force [81]. If $\delta < 0$ (i.e. the trapping beam is red-detuned), this force acts in the direction of increasing intensity, and U_{dip} is attractive: this is what is referred to as a dipole-force trap. In the present work, we confine ^{87}Rb atoms in potential wells that are at most $k_B \times 100\,\mu$K deep (see Methods). To load atoms into these conservative

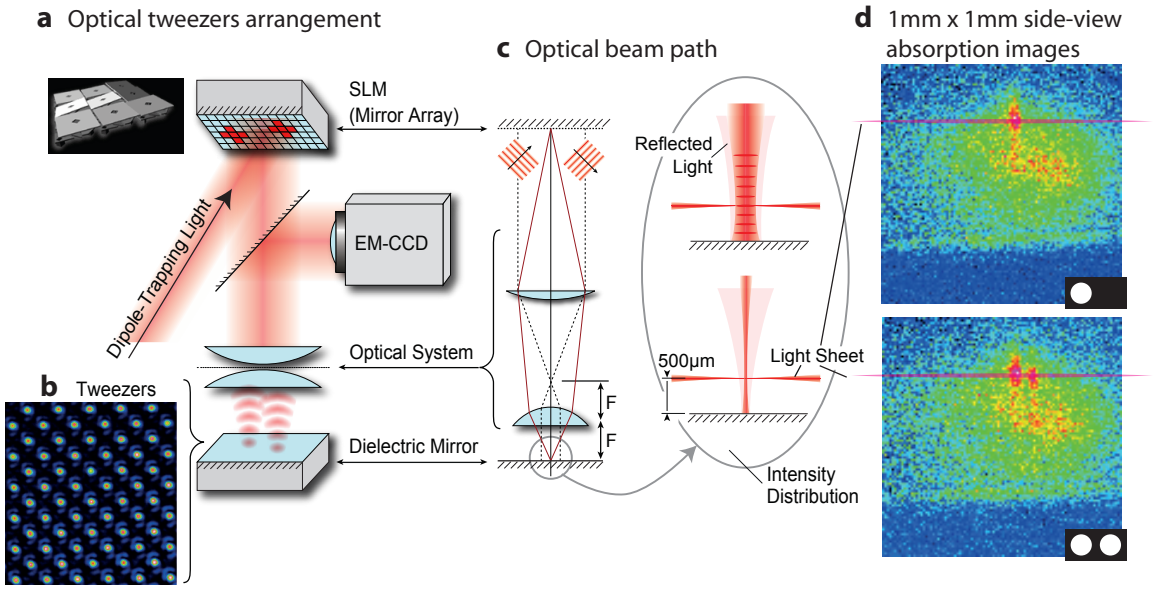

Figure 3.1: **Tweezers arrangement.** **a**, Dipole trapping light illuminates a SLM and is focused through a two lens microscope to form arbitrary potential landscapes which can be used to trap and guide atoms. A thin sheet of resonant light is used to illuminate only the atoms trapped in the image plane, and a highly sensitive EMCCD camera is used to observe the atomic fluorescence via a dichroic mirror. **b**, Tweezers array of 5 μm spacing and 1 ± 0.2 μm trap waist. Depicted is the false-colour intensity profile of the trapping light ranging from zero (dark blue) to 30 μW/μm^2 (bright red). **c** The two lens microscope used to image the surface of the SLM consists of an achromatic doublet and a high numerical aperture aspheric lens. The close-up view illustrates that interference between the incoming and reflected light leads to a standing wave, and how this effect is negligible for smaller traps. **d**, Side-view false-colour absorption images of atoms in either one or two dipole traps of ~ 30 μm radius each. The snapshots are taken immediately after loading, with a diffuse cloud of untrapped atoms from the MOST still surrounding the traps (blue: no absorption; red: maximum absorption and highest atom density).

traps, we first capture and cool them from the background gas. This is accomplished with a Magneto Optical Surface Trap (MOST) [91] providing a ~ 280 μK cold cloud of ~ 120 μm radius containing approximately 10^5 atoms at 500 μm distance from a mirror surface.

The SLM is an array of 1024×768 independently addressable micro mirrors (see Methods). Its surface is imaged onto the atoms using the two-lens microscope shown in Fig. 3.1a. For large traps, the divergence of the trapping beam is small and interference between the incoming and reflected light leads to a standing wave. This

is advantageous as in addition to axial confinement, the atoms will also experience some degree of longitudinal confinement. Atoms that are cold enough are therefore confined to anti-nodal planes stacked along the optical axis. This effect only partially applies to smaller traps (of radii $\lesssim 20\,\mu$m) because the reflected beam diverges too quickly. The microscope has a (de)magnification of 57:1, which means that individual micro mirrors are not resolved. A diffraction limited spot of $1\,\mu$m corresponds to a block of 4×4 mirrors, which implies that we can mimic up to 16 levels of intensity by dithering the mirror pattern on the otherwise digital SLM. Also, taking both the finite dimension of the SLM and the resolution limit into account, a two-dimensional array of up to 50×40 well-separated individual trapping sites can be formed. The trapping-light intensity profile depicted in Fig. 3.1b is an excerpt of such an array, with a distance between traps equal to five times the resolution limit.

Atoms are detected by fluorescence imaging through the same microscope used to produce the tweezers. The atoms in the image plane of the tweezers are illuminated by a $9.7\,\mu$m thin sheet of resonant light parallel to the mirror surface, while atoms above or below remain in the dark. The light sheet is 5.1 mm wide and is retro-reflected to avoid pushing the atoms out of the traps. It is blue-detuned by 2 MHz from the $5^2S_{1/2} \leftrightarrow 5^2P_{3/2}$ cycling transition in order to account for the tweezers-induced dynamic Stark shift. The total power in the light sheet is $7\,\mu$W, which gives rise to a scattering rate of 16 photons/μs. The fluorescence light from the atoms is separated from the reflected trapping light using a dichroic filter and directed to a sensitive camera to record fluorescence images. Additional band pass filters in front of the camera eliminate the unwanted trapping light. The camera's exposure window is synchronised with timing of the light sheet.

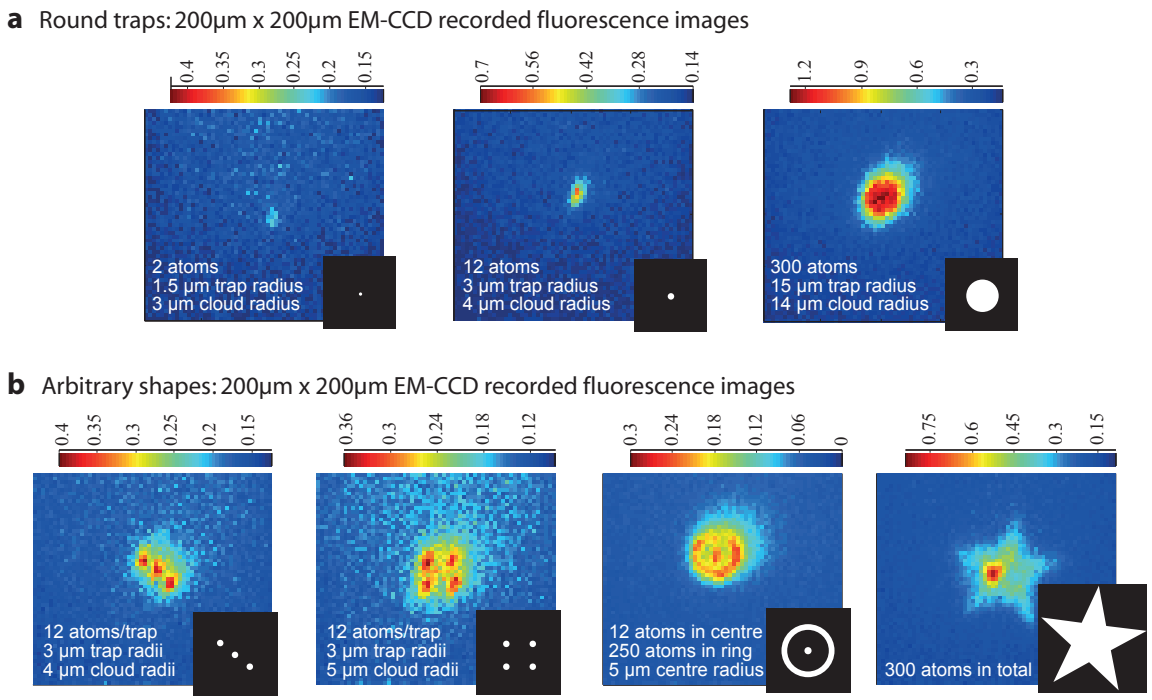

Figure 3.2: **Fluorescence images of trapped atoms. a**, Single static round traps of different sizes. **b**, Atoms in various arbitrary trapping geometries: a line, grid, bullseye and star. Each image is an average over 20 consecutive trap loading cycles. The colour scale has units of atoms per μm^2. The actual patterns on the SLM used to produce the traps are shown in the respective insets using the same length scale. The halo seen in the first three fluorescence images can be attributed to diffusion of the atoms in the light sheet.

3.2 Atoms trapped in arbitrary potential landscapes

The top view through the microscope is well-suited for arbitrarily manipulating and observing atoms in a two-dimensional plane. However, atoms are in general trapped in a vertical column perpendicular to the image plane. We monitor this from the side by absorption imaging. Fig. 3.1d shows a particular situation with atoms loaded into either one or two trapping beams of 30 μm radius. Obviously, atoms get transferred into the standing-wave dipole trapping beams at various heights above the mirror, so that observing all atoms by fluorescence imaging through the microscope would normally give rise to blurry pictures. As we are primarily interested only in those atoms trapped in or near to the image plane, we resolved that issue using the above-

mentioned light sheet to avoid excitation of atoms being out-of-focus. Independent optical characterisation of our imaging system (see Fig. 3.1b) shows that the imaging system has a resolution of $1\pm0.2\,\mu$m, which also limits the edge sharpness of any trapping potential.

The fluorescence images taken through the microscope reveal that the shape of the atom cloud adapts to the imposed trapping geometry. Fig. 3.2a shows atoms confined to a series of circular flat-bottom traps. The traps are created by displaying disks of different radii on the homogeneously illuminated SLM. Normally, one would expect the atoms to fill the trap, such that the radius of the atom cloud equals the trap radius. This is indeed found for traps of large size, which demonstrates that the atoms adapt to the imposed trapping geometries. Deviations are only found for very small traps, where some atoms trapped away from the image plane appear out of focus.

The approximate atom number in each trap was deduced from fluorescence images averaged over 20 loading cycles. During each cycle, the light sheet is flashed on for $90\,\mu$s, within which a single atom emits $\sim$ 1400 photons; of these, 21 photons are detected (see Methods for the collection efficiency). We found that the large trap ($d = 30\,\mu$m) contained 300 ± 40 atoms, the medium trap ($d = 6\,\mu$m) 12 ± 2 atoms, and the smallest observed trap ($d = 3\,\mu$m) 2.0 ± 0.3 atoms. This is very close to the single-atom regime, which one might reach by application of the collisional blockade technique [61, 60] (a widely used method for loading dipole traps with at most one atom).

The lifetime of atoms in all traps is approximately 50 ms, and does not change significantly with the size of the trap. For all traps of radius $\geq 3\,\mu$m, the atom number was sufficiently large for time-of-flight measurements, which yield a consistent temperature of $90 \pm 10\,\mu$K. This is on the same scale as the calculated trap depth.

We furthermore verified the versatility of the tweezers by trapping atoms in arbi-

trary shaped potentials. The trap arrays shown in Fig. 3.2b underline the scalability of the scheme: such arrays could be used to implement a large register of atomic qubits. In the first two images, the radius of each individual trap is $3\,\mu$m, and the average atom number in each is 12 ± 2. The third image shows atoms trapped in a bullseye-shaped trap. The $d = 6\,\mu$m centre of the bullseye contains 12 ± 2 atoms, whilst the ring contains 250 ± 33 atoms. This ring structure shows that it should be possible to transport atoms not just along straight paths, but also along curved trajectories.

The last image in Fig. 3.2b shows 300 ± 40 atoms in a star-shaped trap which is affected by an inhomogeneity in the trapping light. The intensity variation of the laser over the whole area of the SLM is about 30%. This has no significant impact on smaller traps. However, in larger flat bottomed traps like the star, it leads to a congregation of atoms in the regions of highest intensity. Notwithstanding, this large variety of different shapes highlights the flexibility of trap generation with the SLM. Trapped atoms can be arranged into arbitrary patterns using the SLM to impose any desired shape of the trapping potential.

For the larger traps depicted in Fig. 3.2, the atom number stated is always restricted to those atoms exposed to the light sheet due to diffusion of the hottest atoms out of the trap. For trap radii $\leq 3\,\mu$m, confinement along the vertical is dominated by the waist of the trapping beam, so the atom number measured equals the real atom number in the trap. The density of atoms and the small size of the traps leads to a maximum reduction of the fluorescence of 1% due to reabsorption, and hence does not lead to an underestimation of the observed atom number.

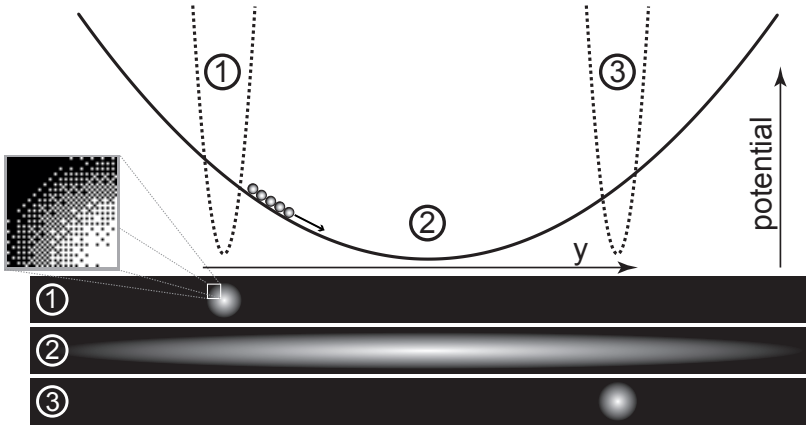

Figure 3.3: **Illustration of the concept of the atom transport** via ballistic 'release and recapture' (not to scale). The atoms are first confined in trap (1), then move along the harmonic transport trap (2) and refocus at the destination when the end trap (3) is switched on. The close-up view shows the dithering of the mirror pattern we apply to obtain a smooth trapping potential.

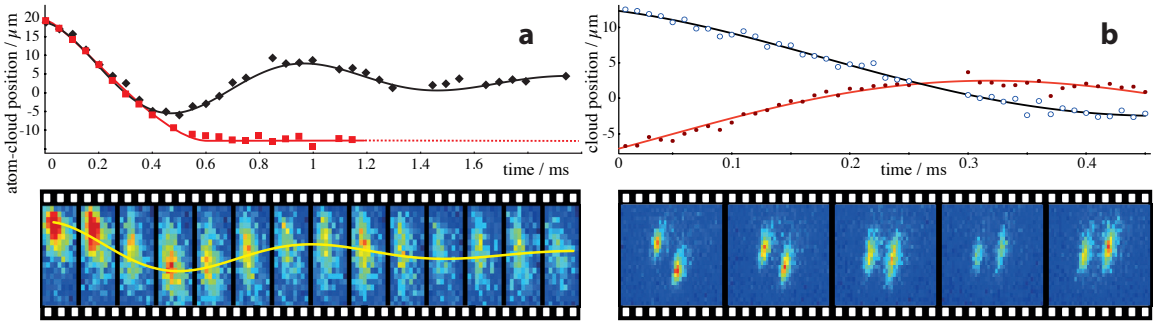

Figure 3.4: **Ballistic atom transport**. The harmonic transport traps, $U \propto -(1 - (x/w_x)^2 - (y/w_y)^2)$, are $w_x = 6\,\mu$m wide and $w_y = 26\,\mu$m wide, whilst the round harmonic start and end traps all have a waist of $w_r = 6\,\mu$m. **a**, Oscillation (black trace), and recapture (red trace) of atoms. The filmstrip below shows a sequence of fluorescence images taken at various instants of the transport. **b**, Atoms oscillating past each other in two separate harmonic wells.

3.3 Deterministic re-arrangement and controlled transport of atoms

So far, we have shown that atoms can be held in a variety of arbitrarily shaped traps. This includes regular arrays of atoms, which one could use as quantum registers in a scalable quantum processor. However, to fill such an array with preferentially one atom per site, or to arbitrarily access a random cell within a register, it is necessary to move atoms independently from one trapping site into another. Here, we show how to accomplish this task and deterministically transport atoms between two well defined positions. The most intuitive way for doing so would be to move the trap of interest by gradually changing the image on the SLM such that the centre of the trap gets displaced. However, this approach is hampered by the discrete switching of the SLM's mirrors, which results in the trap hopping along its trajectory rather than in a smooth motion. In turn, the atoms are lost due to excess heating. Our solution therefore is to let the atoms roll along a harmonic potential before recapturing them in the outer turning point. This ballistic 'release and recapture' transport effectively entails opening an elongated harmonic trapping channel between a start and end position, and timing the recapture to coincide with half the oscillation period of the atoms in this potential well. Fig. 3.3 illustrates that procedure and also shows the dithered mirror pattern used to obtain a smooth trapping potential in the image plane.

In general, one has to take into account that the atoms move in three dimensions. The SLM controls the potential in the trapping plane, while the confinement perpendicular to it is controlled by a combination of the near-field propagation and standing-wave effects resulting from the interference between the trapping light and its own reflection. We calculated the near-field intensity distribution away from the image plane using the Fresnel-Kirchhoff diffraction integral. The divergence is moderate enough to yield a standing-wave modulation visibility of 75% at the height of the

image. This is very strong, and most atoms are confined to the resulting anti-nodal planes. Therefore, we effectively have 20 identical two-dimensional trapping planes within the width of the light sheet.

Fig. 3.4a shows atoms oscillating along the transport channel (black trace) and illustrates that we are capable of recapturing them at the turning point after half a period of oscillation (red trace). The displayed atom-cloud position is the centre-of-mass of the recorded fluorescence images, which were taken at constant rate. Of the 22 atoms initially in the trap, 12 atoms are recaptured. Most surprisingly, this cannot be attributed to heating-induced losses, as time-of-flight measurements reveal a temperature rise of only $10\,\mu$K. Instead, the reduction in fluorescence is to some extent explained by atoms not exactly regrouping at the destination. The primary reason for that limitation is the small anharmonicity of the transport trap, resulting in a 1.3 ms damping time of the harmonic motion (see Fig. 3.4a, black trace). The secondary loss mechanism is a diffusion of the most energetic atoms along the standing wave, i.e. away from the image plane. Both losses could be easily reduced using a more powerful trapping laser. Hence disregarding the spreading, the successful ballistic transport in harmonic traps illustrates that the SLM is an ideal tool for positioning atoms in a deterministic way.

The manipulation of individual information carriers in a large-scale quantum network relies on the capability of re-arranging atoms sitting at different nodes of an array independently from one another. To demonstrate that this can be done, we use two transport channels to move atoms in opposite directions from two different starting positions. We apply the same release and recapture method as discussed above. Fig. 3.4b shows atoms oscillating past each other in these harmonic potentials. The atoms move to destinations $8\,\mu$m and $9\,\mu$m away from their respective starting points, with the two destinations chosen independently. This individual transport of randomly selected trapping sites shows that SLM-controlled atom traps are capable

of regrouping arrays of trapped atoms to nearly any arbitrary pattern.

3.4 Conclusion and Outlook

SLM-based optical-tweezers constitute a flexible scheme for the manipulation and transport of dipole-trapped neutral atoms. Individual trapping sites within a large array can be controlled independently, which is essential for a scalable system. The refresh rate of the SLM is fast enough for a dynamic control of trapped atoms in response to the observed fluorescence, such that a feed-back scheme could be realised to control the atom number in individual sites. With the ballistic transport applied to individual atoms, one might, e.g., move two atoms into a fibre-tip micro cavity and realise pair-wise entanglement [92], or use controlled collisions between two atoms to implement a two-qubit gate [17, 26]. Hence scalable atomic arrays for quantum computing or simulation are now in reach of current technology.

3.5 Methods

The spatial light modulator (SLM – Texas Instruments DMD Discovery 1100) consists of 1024 by 768 independently addressable micro-mirrors, 13.7 μm by 13.7 μm each, that can be digitally switched between two tilt angles corresponding to an on and off position, the on position directing the impinging light onto the desired optical axis, and the off position throwing it onto a beam dump. These mirrors can be controlled individually, allowing us to generate traps of virtually any shape. The entire optical surface covered by these mirrors has a size of 14 mm by 10.5 mm. The SLM can display either static patterns or movies made from frames of these patterns. The movie frame rate ranges between 4 and 20 KHz, making it fast enough for a dynamic control of the atoms trapped in the tweezers.

Trap depth: The SLM is illuminated using a high power diode laser (Toptica

DLX-110). This laser has a power output of 800 mW and is red-detuned by 5 nm (λ=785 nm) from the $5^2S_{1/2} \leftrightarrow 5^2P_{3/2}$ transition in ^{87}Rb. We have chosen this relatively small detuning to achieve a reasonable trap depth at light intensities below the damage threshold of the SLM. The lens distances and convergence of the illuminating beam have been chosen such that isoplanatism is ensured [93], i.e. the wavefront in the image plane is identical to the wavefront of the light in the plane of the SLM. The beam illuminating the SLM has a Gaussian profile that fills the area of interest (typically 1/3 of the SLM surface). In the image plane, the intensity of the trapping laser reaches 30 W/mm^2, which corresponds to a maximum potential depth of $k_B \times 100\,\mu$K.

Optical system: A two-lens microscope images the surface of the SLM onto a plane 500 μm above the surface of a mirror. We use a high numerical aperture aspheric lens so that our optical tweezers have the smallest possible foci, and hence the best possible axial confinement. The microscope consists of this aspheric lens, which has a focal length of 20 mm and a numerical aperture of 0.52, and an achromatic doublet as a collimator, with a focal length of 750 mm. Only the aspheric lens is mounted inside the vacuum chamber. An important feature of this system is that the same two lenses used to image the SLM are also used to observe the atoms trapped in the tweezers by laser-induced fluorescence, as illustrated in Fig. 1. A dichroic mirror is used to separate the trapping light from the collected atomic fluorescence at 780 nm. An electron-multiplying CCD camera (EMCCD) is used to observe the laser-induced fluorescence (Andor iXon 885).

Collection efficiency: The photon-detection efficiency of the system was found to be $\eta_{tot} = (1.5 \pm 0.2)\%$, where η_{tot} is the product of the collection efficiency of the lens system η_{lens}=7.6%, the losses along the optical path η_{loss}=51%, and the quantum efficiency of the camera η_{QE}=41%.

Chapter 4

Experimental Design

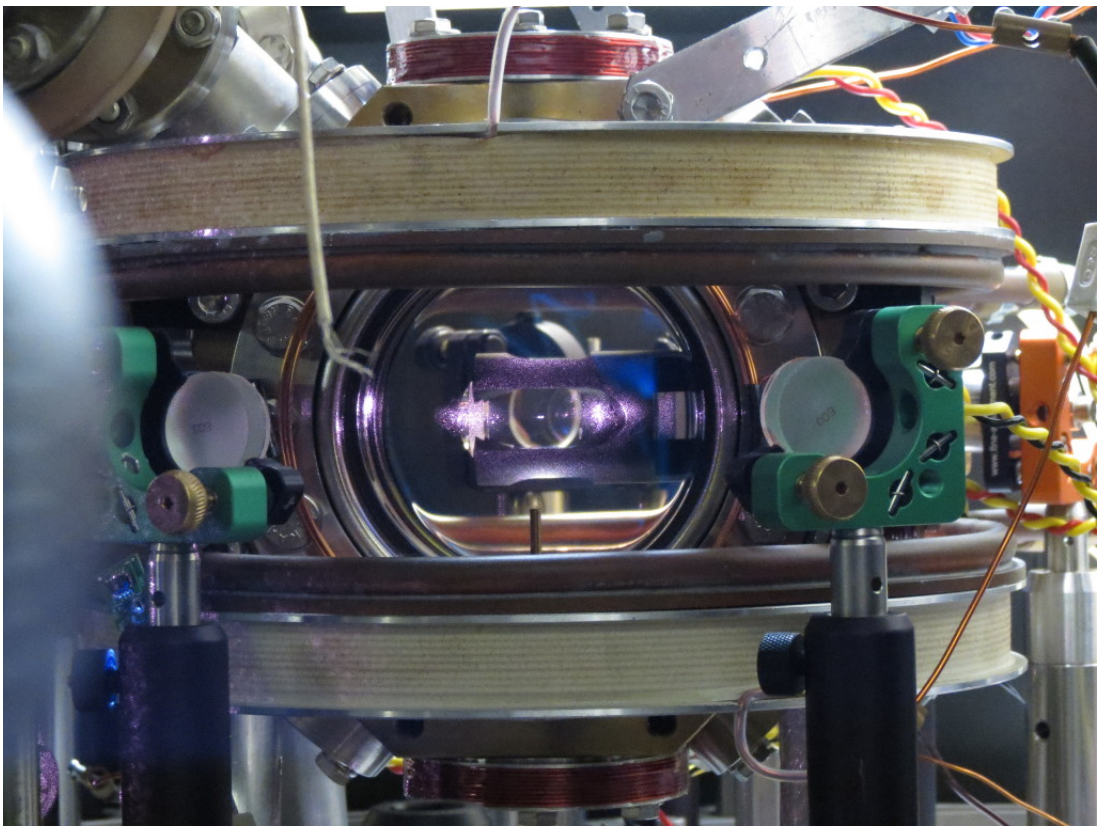

Figure 4.1: Photograph of the complete vacuum chamber taken with an infra-red camera. The laser-cooled rubidium atoms are visible in the centre of the chamber as a glowing purple cloud.

In this chapter, I describe the new apparatus for cooling and trapping atoms. The apparatus described in chapter 3 had several limitations which prevented us from

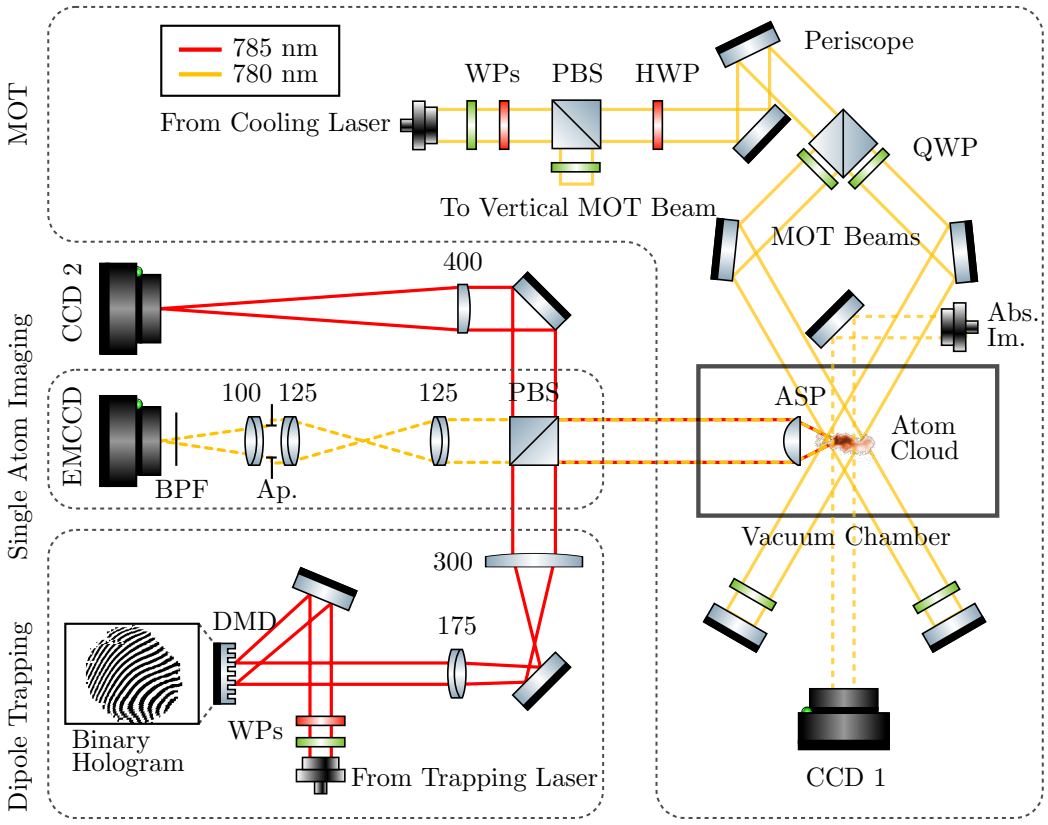

Figure 4.2: Experimental setup comprising of a magneto-optical trap (MOT), dipole trap, and single atom imaging system. The MOT (section 4.1) is formed from six counter-propagating laser beams and produces an atom cloud in the centre of the vacuum chamber. The focused dipole trapping beam (section 4.2) is superimposed on this cloud using an aspheric lens (ASP). A digital mirror device (DMD) is used to holographically generate multiple dipole traps. An electron multiplying charge coupled device (EMCCD) camera is used to image fluorescence light from single atoms (section 4.3) using the same aspheric lens.

[a] Optical diagrams were made using ComponentLibrary (http://www.gwoptics.org/ComponentLibrary/) by Alexander Franzen under a Creative Commons license (http://creativecommons.org/licenses/by-nc/3.0/)

observing single atoms. The first of these limitations was the dielectric mirror that was used to create the standing-wave dipole trap. This mirror scattered a significant amount of light from the cooling lasers into the microscope, overwhelming the fluorescence signal from single atoms. Secondly, the method of directly imaging the DMD resulted in a large amount of dipole trapping light to be wasted, as only a small number of mirrors are used for making the trap. To overcome these limitations, we redesigned the apparatus to eliminate the dielectric mirror. The apparatus (see photograph in figure 4.2, and diagram in 5.1) is essentially a combination of a magneto-optical trap (MOT) and a microscope. The purpose of the MOT is to generate a reservoir of cold atoms, and to provide the dissipative cooling force needed to load atoms into the dipole traps. The microscope has two roles: to focus the dipole trapping light to a diffraction-limited spot; and to collect the fluorescence light emitted from trapped atoms.

The apparatus can be divided into three sections: the MOT, the dipole trapping beam, and the single-atom imaging system. Each section of the experiment will be discussed individually, including the relevant design choices and steps involved in construction.

4.1 MOT

Magneto-optical traps [82] are robust and powerful tools for producing laser-cooled atoms. In our MOT, the laser beams are produced from a single-mode fibre collimator (Thorlabs F810APC-780) with an output beam waist $w = 3.75\,\mathrm{mm}$. The beam is divided using a combination of half wave plates and polarising beamsplitters to produce three beams each with a power of 3 mW. There are two horizontal and one vertical beam which all intersect at the centre of a vacuum chamber. Usually, all the MOT beams are mutually perpendicular, however, in our setup the angle between

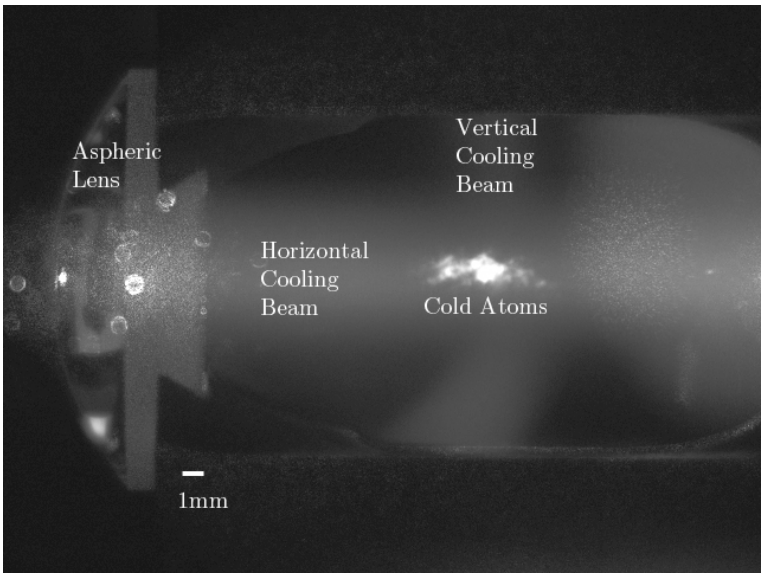

Figure 4.3: Fluorescence image of the MOT as seen from CCD1 (see figure 4.2). The horizontal and vertical MOT beams can be seen glowing faintly due to background rubidium vapour.

the horizontal MOT beams is reduced to 76°. This allows for the aspheric lens to be positioned closer to the cloud. This arrangement gives weaker confinement of the cold atoms along one axis, leading to a cigar shaped atom cloud. The beams are retro-reflected back along the same path to provide the two counter-propagating beams for Doppler cooling described in chapter 2. The incoming beams have circular polarisation (set by quarter-wave plates) while the retro-reflected beams have opposite circular polarisation.

4.1.1 Vacuum chamber

The large temperature gradient between laser cooled atoms and the atmosphere, and the need to minimise the rate of background gas collisions, the experiment must be performed under conditions of ultrahigh vacuum (UHV). The vacuum chamber is a custom designed hexagonal chamber made from non-magnetic stainless steel. The viewports are designed to fit standard CF 40 and CF 100 flanges, and are made from Kodial glass with a broadband anti-reflection coating. The chamber was designed

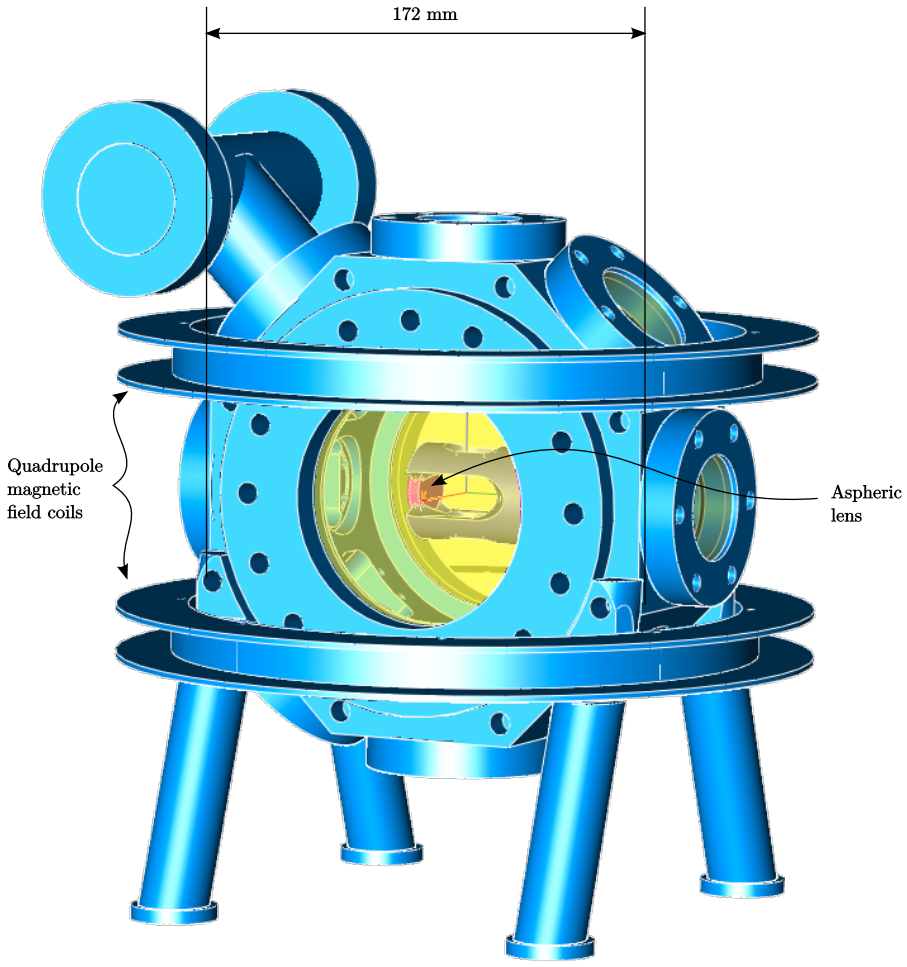

Figure 4.4: Perspective view of the vacuum chamber, showing the quadrupole magnetic field coils and aspheric lens.

with CAD software (TurboCad 2) and is shown in figure 4.4.

Rubidium vapour is introduced to the chamber using a resistively-heated dispenser (Alvasource Rb-50). Ultrahigh vacuum is maintained using an ion pump (MECA 2000 PID 25) and Titanium sublimation pump (VACOM Sublicon 51). The pressure is monitored with an ionisation gauge (Granville-Phillips 350) and is typically below 10^{-10} mbar. When the rubidium dispenser is switched on, the pressure increases to 10^{-7} mbar.

4.1.2 Magnetic field coils

A pair of anti-Helmholtz coils are used to generate the quadrupole magnetic field for the MOT. We approximate the magnetic field gradient as that of two loops of wire with opposite currents,

$$\frac{dB}{dz} = \frac{3\mu_0 INhR^2}{(h^2+R^2)^{5/2}}, \tag{4.1}$$

where μ_0 is the permeability of free space, I is the current, N is the number of turns, $2h$ is the vertical separation between the two coils, and R is the semi-diameter of the coil. Each coil consists of 110 turns of 1.6 mm diameter Copper wire with a resistivity of $\rho = 1.7 \times 10^{-8}\,\Omega\mathrm{m}^{-1}$. Taking average values of $R = 115\,\mathrm{mm}$, $h = 45\,\mathrm{mm}$, and $I = 10\,\mathrm{A}$, the expected magnetic field gradient is $8.6\,\mathrm{Gauss\,cm}^{-1}$. The total resistance of the coils is $1.5\,\Omega$, leading to a power dissipation of 150 W per coil. This large power dissipation means that when operating at full current the coils rapidly reach a temperature of 60°C. To counter this, the coils are water cooled using a computer CPU chiller (Aquatuning Koolance Exos-2.5) to keep the temperature below 30°C.

Additional x, y and z trim coils are used to compensate for Earth's magnetic field. Each of these pairs of coils consists of a few turns of copper wire. These coils allow the zero-point of the quadrupole field to be moved by several millimetres in each direction.

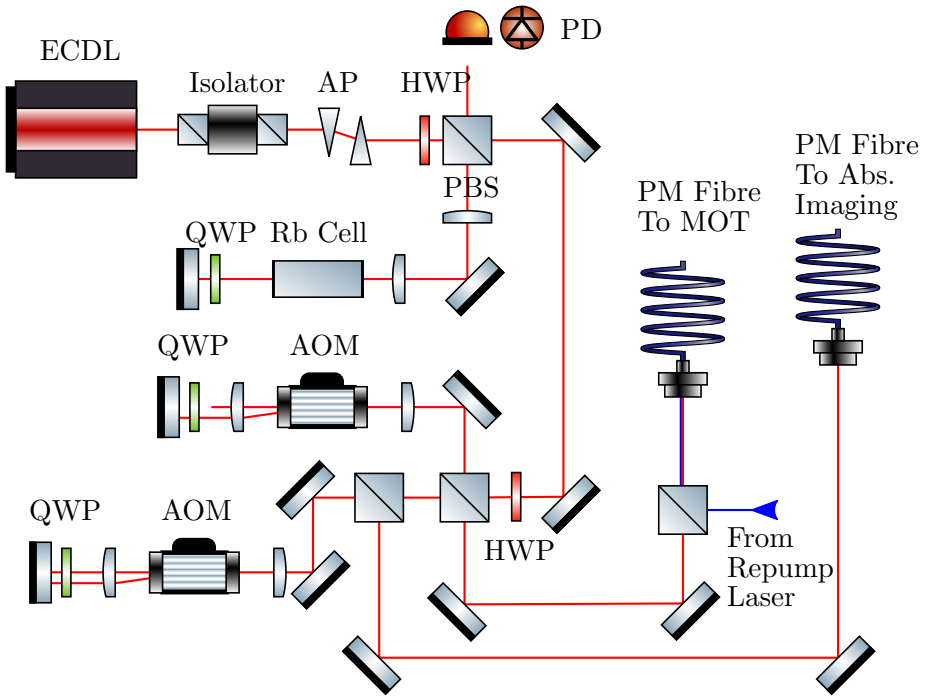

Figure 4.5: Diagram of cooling laser setup. Light from an external cavity diode laser (ECDL) passes through an optical isolator and anamorphic prism pair (AP). Approximately 1 mW of light is diverted to a rubidium cell for saturated absorption spectroscopy. The spectroscopy signal is measured with a fast photodiode (PD). The remainder of the light is frequency shifted with a double passed acousto-optic modulator (AOM) and coupled into a polarisation maintaining (PM) single-mode optical fibre. Repump light from a similar setup is combined with the cooling light in the same fibre.

4.1.3 Cooling lasers

The cooling laser provides the scattering force required for the operation of the MOT. The main requirements for this laser are a narrow linewidth ($< 1\,\mathrm{MHz}$), a stable, variable detuning relative to the cooling transition, and a power of several mW. The energy levels relevant for cooling are shown in figure 4.8. The cooling transition ($F = 2 \rightarrow F' = 3$) is known as a cycling transition, since the excited state can only decay back to the $F = 2$ ground state due to atomic selection rules. Occasionally, the cooling lasers can off-resonantly excite the $F = 2 \rightarrow F' = 2$ transition, which can spontaneously decay into the $F = 1$ ground state. As a result, a second laser is required for repumping atoms out of this dark state. Both lasers are external cavity

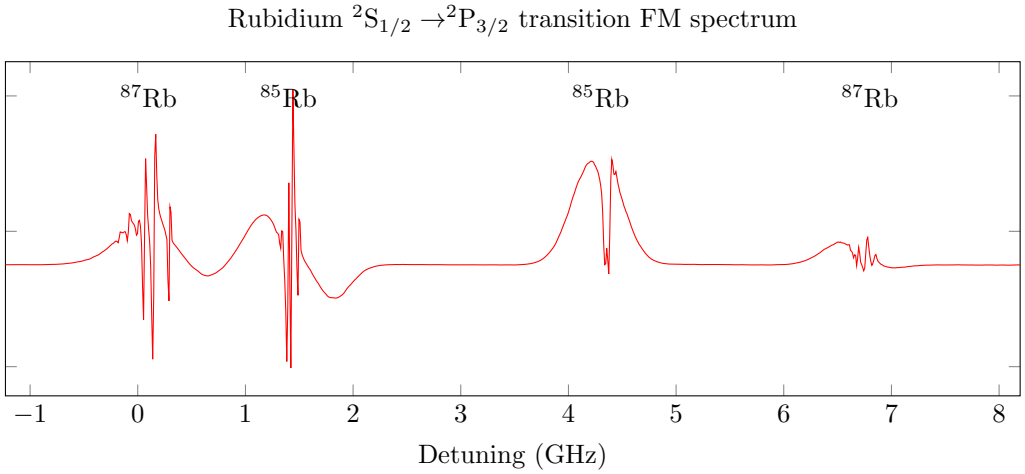

Figure 4.6: Frequency-modulated (FM) spectrum of the rubidium transitions. The cooling laser is locked to one of the lines in the leftmost group. The repump laser is locked to a line in the the rightmost group.

diode lasers (ECDL, Toptica DL 100).

The frequency of the laser is stabilised to a saturated absorption spectroscopy signal [81] using the Pound-Drever-Hall technique [94]. The optical arrangement for the cooling laser is shown in figure 4.5. The laser current is modulated with a local oscillator at 20 MHz to generate frequency modulation (FM) sidebands with a power of 1% relative to the total laser power. After passing through a rubidium vapour cell, the frequency-dependent absorption of the atoms causes the FM sidebands to be converted into amplitude modulation, which is detected with a fast photodiode. An error signal is generated using a frequency mixer to demodulate the photodiode signal with the local oscillator. The resulting FM spectrum is shown in figure 4.6. Rather than locking to the cooling transition itself, the laser is locked to the (1,3)-crossover resonance (see the level diagram in figure 4.8 and FM spectrum in figure 4.7). The beam is then frequency shifted back into resonance using a double passed acousto-optic modulator (AOM, Neos Technologies 23110). The AOM allows the both the frequency and amplitude of the cooling beam to be controlled.

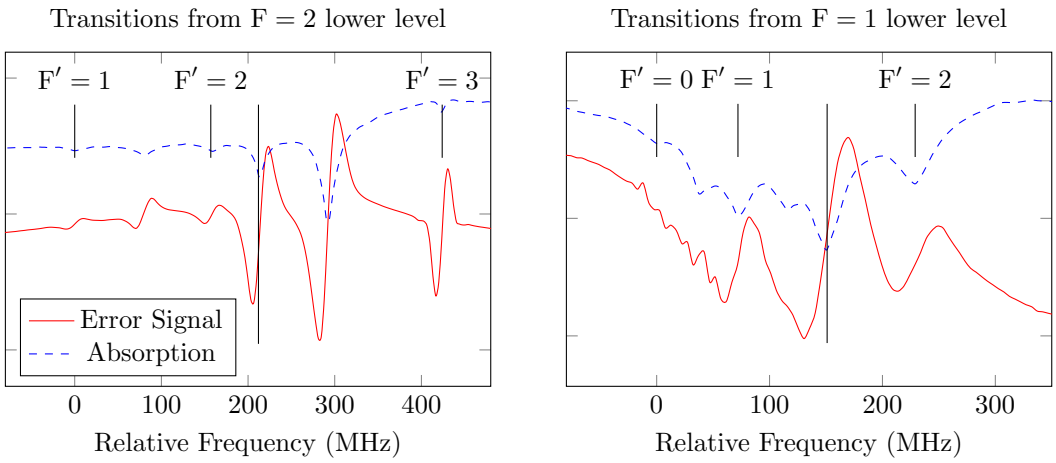

Figure 4.7: Detailed frequency modulation (FM) spectra of the cooling and repump transitions. The vertical line indicates the locking point.

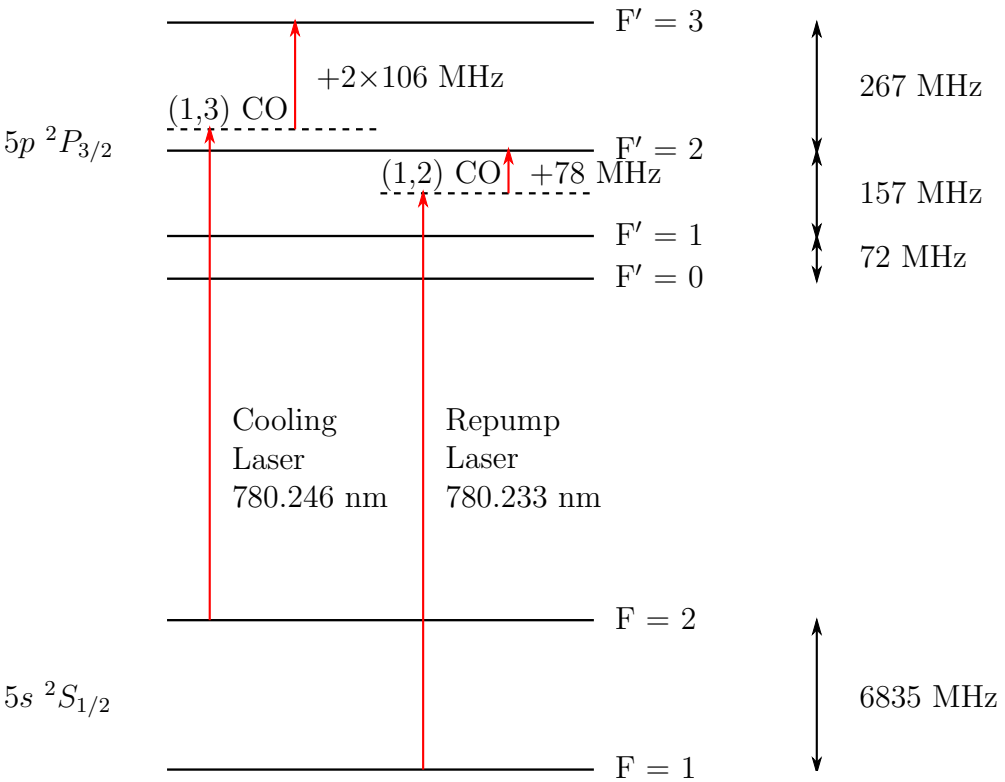

Figure 4.8: Relevant energy levels for laser locking. The cooling laser is locked to the (1,3) crossover line (CO), and frequency-shifted using a double pass acousto-optic modulator (AOM) to reach the $F = 2 \rightarrow F' = 3$ transition. The repump laser is locked to the (1,2) crossover line and frequency-shifted by 78 MHz using an AOM.

4.1.4 Sub-Doppler temperature

As discussed in chapter 2, the ultimate temperature of the MOT is determined by sub-Doppler cooling processes. We experimentally determined the lowest attainable temperature by varying the intensity and detuning of the cooling beams. The experimental procedure was as follows:

1. Load the MOT with $\Delta = -\Gamma$ and $I = I_{\text{sat}}$ for 1 second.
2. Increase the detuning to $\Delta = -3\,\Gamma$ and lower the intensity to $I = 0.2\,I_{\text{sat}}$ for a period of 20 ms.
3. Measure the temperature of atoms using the time of flight technique i.e. by measuring the rate of expansion and hence the average speed of the atoms in the cloud.

The results of the experiment are shown in figure 4.9. Larger detunings result in a lower temperature of the laser cooled atoms, as expected from sub-Doppler cooling theory [83]. The minimum temperature attained was 15 μK, at a detuning of $-6.6\,\Gamma$. Below this detuning, very few atoms were left in the MOT.

4.1.5 Control of the MOT density

The density of the MOT is especially important for loading single atoms into microscopic dipole traps. The upper limit for the density of a MOT is approximately $10^{11}\,\text{cm}^{-3}$, when spontaneously emitted photons are reabsorbed by other atoms in the MOT cloud [82]. The number of atoms in the MOT (denoted by N) behaves according to the following differential equation:

$$\dot{N} = R - \gamma N - \beta N(N-1), \tag{4.2}$$

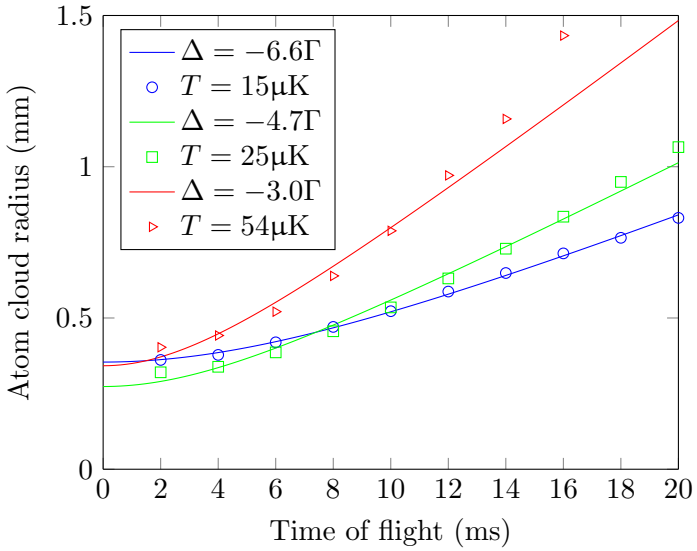

Figure 4.9: Time-of-flight measurement of the temperature of sub-Doppler cooled atoms. Larger detunings lead to colder atomic clouds. The coldest temperature of 15μK was obtained with a detuning of $\Delta = -6.6\Gamma$.

where R is the loading rate of atoms into the MOT, γ is the rate of loss due to collisions with room-temperature background gas atoms, and β is a two-body loss coefficient due to collisions between pairs of cold atoms in the MOT.

Given that our goal is to trap and observe single atoms, we rather need an extremely low density MOT. There are two reasons for this: firstly, the density of the MOT determines the rate at which atoms are loaded into the microscopic dipole traps. In the collisional blockade regime, the loss rate is the same as the loading rate, as both are caused by the arrival of another atom in the dipole trap. The timescale of atoms arriving and leaving the trap must be slower than the frame rate of the camera that is used to detect the single atoms, which in our case is 10 Hz. Secondly, atoms from the MOT also emit fluorescence light, which obscures the light from the single dipole-trapped atoms, and it is essential to minimise this unwanted background light.

One way to control the density of the MOT is to increase its size by decreasing the strength of the quadrupole magnetic field. The effect of this is to decrease the stiffness of the MOT potential, which results in a larger size for a given temperature.

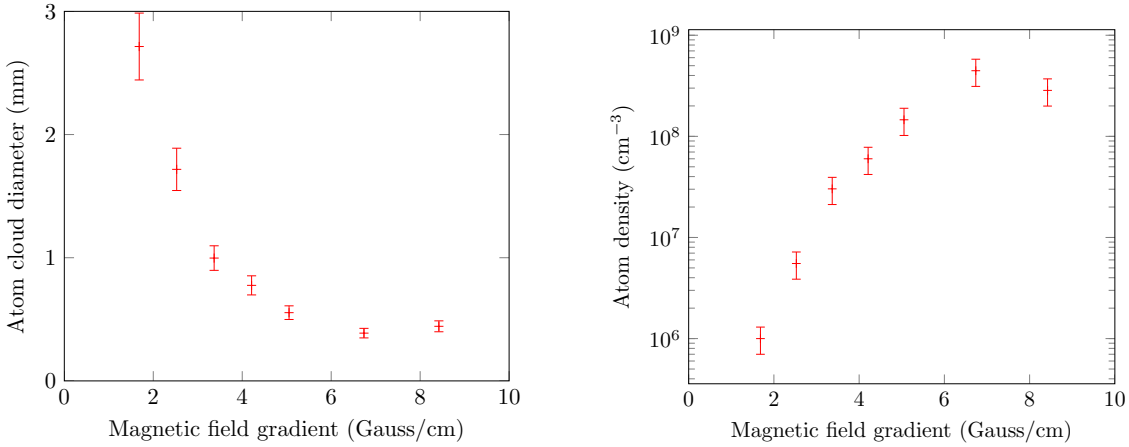

Figure 4.10: Variation of the diameter and density of atoms in the MOT with magnetic field gradient.

The number of atoms in the MOT is determined by equation 4.2, and is a constant independent of the quadrupole field. Figure 4.10 shows the variation in the diameter of the MOT with magnetic field gradient. The size of the MOT was determined by analysing fluorescence images. As shown in the figure, the density can be adjusted by three orders of magnitude using this technique.

Another method for reducing the density of the MOT is by reducing the number of atoms. The equilibrium number of atoms in the MOT, N_∞, occurs when the loading and loss rates in equation 4.2 are balanced. This equation can be solved, ignoring the two-body loss term, to give the equilibrium number of atoms in the MOT as

$$N_\infty = \frac{R}{\gamma}. \tag{4.3}$$

The loading rate, R, is proportional to the partial pressure of rubidium, p_{Rb}, whilst loss rate, γ is proportional to the total pressure, $p_{\mathrm{Rb}} + p_{\mathrm{other}}$, where p_{other} is the partial pressure of other gasses. The MOT atom number is therefore proportional to the following ratio

$$N_\infty \propto \frac{p_{\mathrm{Rb}}}{p_{\mathrm{Rb}} + p_{\mathrm{other}}}. \tag{4.4}$$

Evidently, the number of atoms in the MOT can be reduced by lowering the

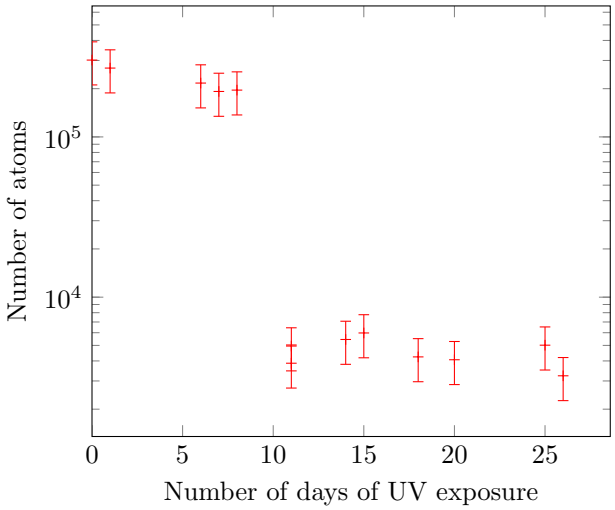

Figure 4.11: Number of atoms in the MOT measured over one month during which the rubidium dispenser was switched off and the chamber was exposed to UV light to desorb atoms from the walls of the vacuum chamber.

pressure of rubidium, however, rubidium vapour is continuously replenished by atoms adsorbed to the walls of the chamber. To eliminate these atoms, we switched off the rubidium dispenser and switched on a 1 W ultraviolet light emitting diode (LED), with a peak emission wavelength 420 nm. The effect of the LED is to remove rubidium atoms from the walls and windows of the vacuum chamber through light induced atomic desorption (LIAD) [95]. The LED was left on for a period of one month while the equilibrium number of atoms in the MOT was measured each day. The data is shown in figure 4.11. After 10 days of continuous exposure to UV light, the number of atoms in the MOT dropped by two orders of magnitude. We conclude that after this time, the majority of the Rubidium on the walls had been eliminated. For all remaining experiments in which we trap single atoms, we operate under these conditions of extremely low rubidium density.

4.2 Dipole trap

The dipole trap is used to isolate single atoms from the MOT. The dipole trapping light is generated by a tapered amplifier (TA) laser system and delivered to the ex-

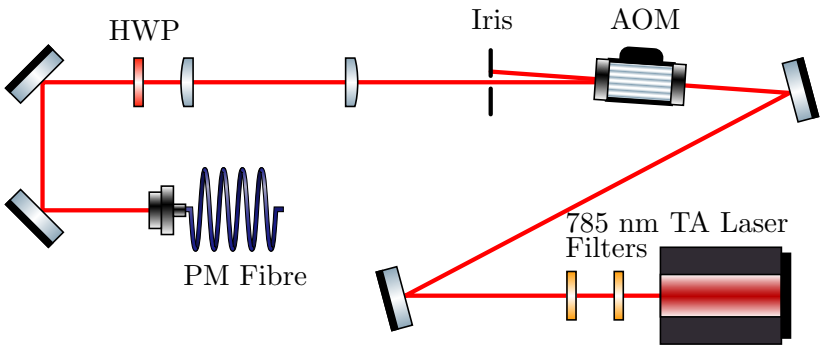

Figure 4.12: Dipole trapping laser setup. Light from the dipole trapping laser passes through two 785 nm bandpass filters to suppress amplified spontaneous emission (ASE), and an AOM for fast amplitude modulation, and coupled into an optical fibre.

periment using a single mode optical fibre. The optical layout is shown in figure 4.12. We use a DMD to split the incident beam into multiple steerable output beams for trapping and moving large numbers of single atoms.

4.2.1 Trapping laser

The trapping laser is based on a wavelength-tunable tapered amplifier laser system (Toptica DLX 110). The tapered amplifier chip (m2k-laser m2k-TAL-0785-1000) has a broad gain region of $770-800\,\mathrm{nm}$ and has a maximum output power of $1\,\mathrm{W}$. Typically such amplifiers are used in a master oscillator power amplifier (MOPA) configuration, in which they are seeded by a separate laser. However, for our chip, the front and rear facets have a reflectivity of $1\,\%$ and $0.01\,\%$ respectively, which enables it to be used as a laser in its own right. Light emitted from the rear facet is reflected back into the chip using a diffraction grating, thus the front facet and diffraction grating form an extended laser cavity. The wavelength can be selected by tuning the angle of the grating.

In addition to the wavelength set by the external cavity, the laser also emits a broad pedestal of amplified spontaneous emission (ASE), shown in figure 4.13. Any remaining ASE near the rubidium transitions at $780\,\mathrm{nm}$ and $795\,\mathrm{nm}$ will scatter

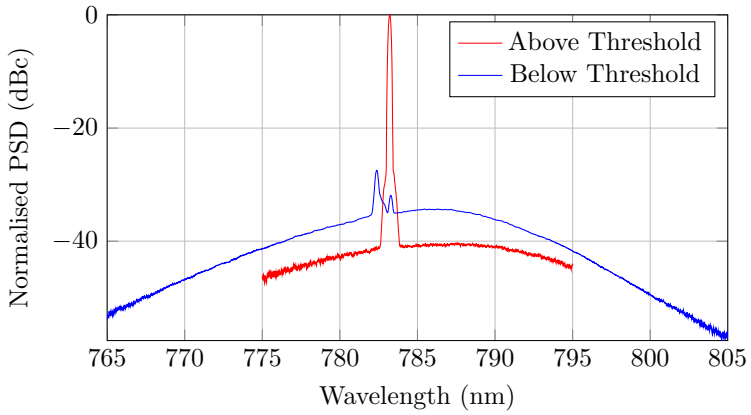

Figure 4.13: Normalised power spectral density of the dipole trapping laser, measured with an optical spectrum analyser with a resolution bandwidth of 0.05 nm (25 GHz). When the laser current is below threshold, it spontaneously emits light over a broad spectrum. Above the threshold current, laser emission occurs at 783.2 nm, along with a broad pedestal of amplified spontaneous emission (ASE).

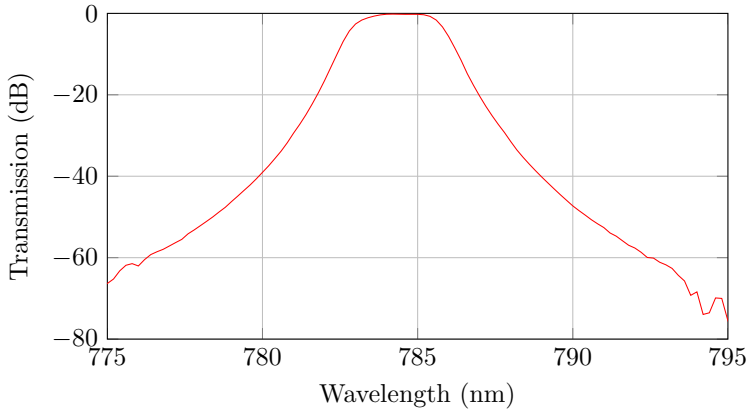

Figure 4.14: Transmission of MaxLine 785 filter. The filter suppresses light by a factor of 10^{-4} at 780 nm.

photons from the trapped atoms. These scattered photons will destroy the coherence of the quantum state, and also contribute towards recoil heating of the trapped atoms. To avoid these effects, the ASE is suppressed by a factor of 10^{-8} by using two 785 nm interference filters (Semrock MaxLine LL01-785-12.5). The arrangement of these filters and additional optics is shown in figure 4.12. The transmission profile of one of the filters is shown in figure 4.14.

The dipole trapping laser then passes through an AOM, which is used for intensity control and fast switching of the dipole trap. Finally, the laser is coupled into a single mode polarisation-maintaining optical fibre. A half-wave plate (HWP) is used to align the polarisation of the laser to the axis of the fibre. The fibre acts as an additional ASE filter, since ASE is emitted into multiple spatial modes with a wide range of angles, most of which are outside the acceptance cone of the fibre. After the AOM, bandpass filters, and optical fibre, we obtain 336 mW available for atom trapping. This power places an ultimate limit on the number of dipole traps that we can produce.

4.2.2 Testing the dipole trap

We tested the dipole trapping laser by trapping atoms in a single beam. A narrow beam from the dipole trapping laser was focused through the vacuum chamber to intersect with the centre of the MOT. This beam had a waist $w_0 = 30\,\mu\text{m}$ and a Rayleigh range $z_R = 3.6\,\text{mm}$, yielding a trap depth of 0.54 mK. The dipole trap was switched on during the loading phase of the MOT, in which atoms are transferred from the MOT to the dipole trap. The cooling lasers were switched off and the MOT was allowed to expands due to the residual thermal velocity of the atoms. As it expands, it falls under the influence of gravity, while atoms in the dipole trap remain in place. Figure 4.15 shows an image of the atom cloud 10 ms after the cooling lasers had been switched off. Most of the atoms fall away while $\sim 10\%$ are transferred to

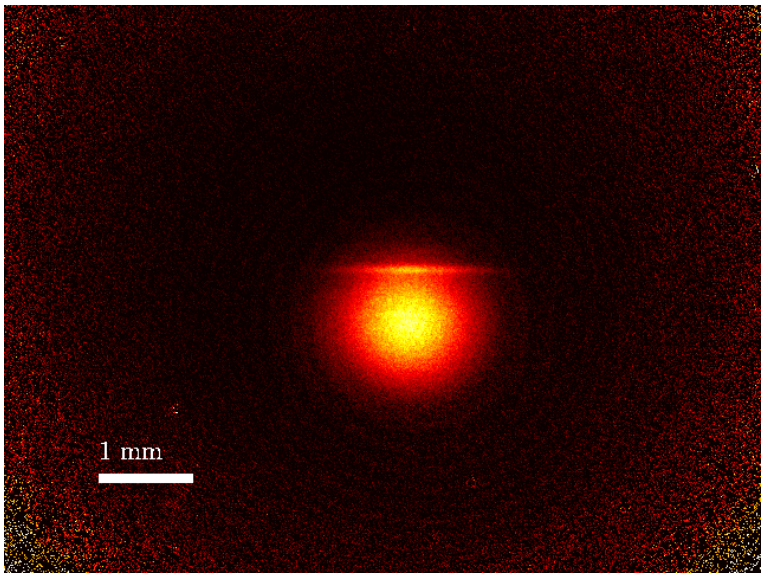

Figure 4.15: Absorption image of atoms being loaded into a dipole trap. Atoms in the MOT expand in a spherical cloud, while some are held in place in the cigar-shaped dipole trap.

the dipole trap.

4.2.3 Aspheric lens

Arguably, the most critical component of the setup is the objective lens of the microscope. We chose a high numerical (NA) aperture aspheric lens (Asphericon 20-18 HPX). This lens has an effective focal length EFL= 18 mm and a clear aperture of 18 mm, giving $NA = 0.49$. The large working distance (distance from the back face of the lens to the focal plane) of 14.1 mm means that the lens can be placed well away from the MOT beams. Furthermore, the lens has been corrected for spherical aberration for a collimated input beam. The diffraction-limited waist is 0.75 µm, allowing a tightly-focused dipole trapping beam.

An alternative to an aspheric lens would be to use a compound microscope objective, as in [28, 63]. The advantage of these lenses is that they are corrected for both spherical and chromatic aberration, which allows the trapping and fluorescence wavelengths to be far apart. Another significant advantage of compound objectives is

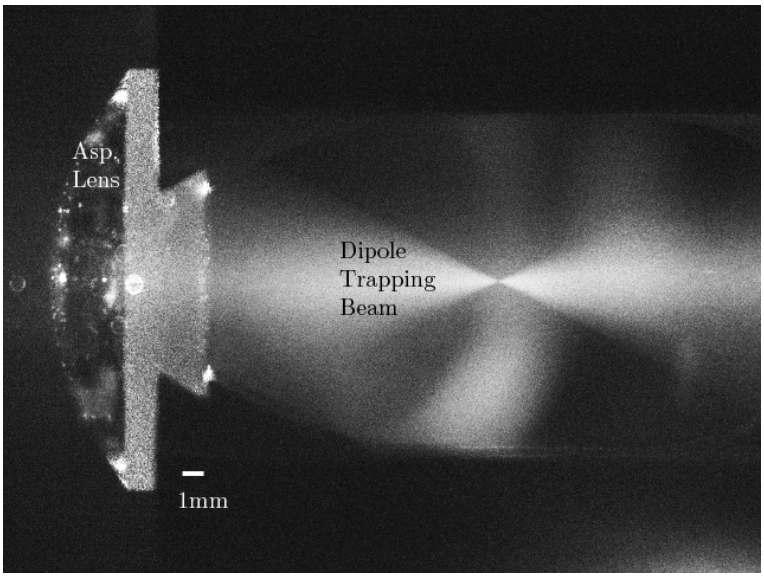

Figure 4.16: Fluorescence from the focused dipole trapping beam as seen from CCD1. 780 nm light from the cooling lasers is focused through the aspheric lens, causing the background rubidium vapour to fluoresce.

that they can be designed to allow for a window between the lens and the atom [28]. This allows the lens to be placed outside the vacuum chamber, where it can be easily realigned. The disadvantages include reduced optical transmission through multiple lenses, shorter working distances, and higher cost.

4.2.4 Lens mount

The aspheric lens was held in a custom aluminium mount, shown in figure 4.17. The mount is designed to position the lens at a fixed distance of 14.1 mm from the centre of the MOT. A black anodised coating was applied to absorb unwanted reflections from the MOT beams. A large amount of material was removed to allow the horizontal and vertical MOT beams to pass unhindered. The base of the mount is designed to allow a fibre-tip cavity to be mounted in the future. The lens was fixed in place using UHV compatible glue (EpoTek 353ND). Air holes were drilled in all of the screw threads to prevent trapped volumes of air from contaminating the vacuum.

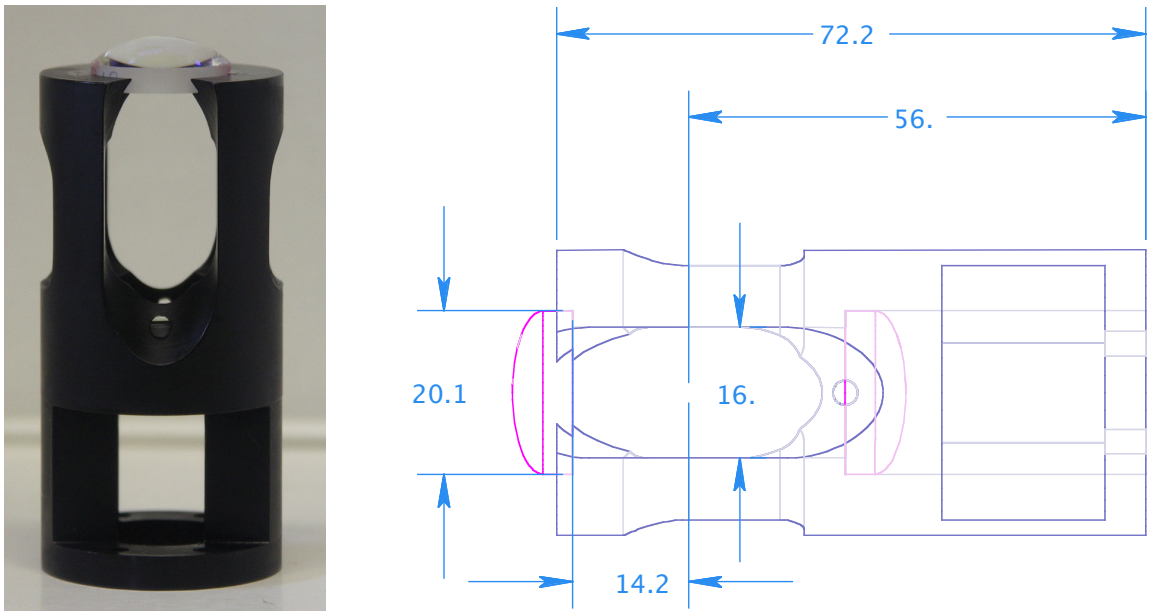

Figure 4.17: Photograph and dimensions of custom lens mount and dimensions from the CAD drawing.

4.2.5 Digital mirror device

We use a digital mirror device (DMD) to holographically generate multiple trapping beams. The DMD is a Texas Instruments DMD Discovery 1100, which consists of 1024×786 of micro-mechanical mirrors, shown in figure 4.18. Each mirror can be switched between two angles of $\pm 12°$, which we refer to as 'on' and 'off'. The 'on' position reflects light through the experiment while the 'off' position deflects light to

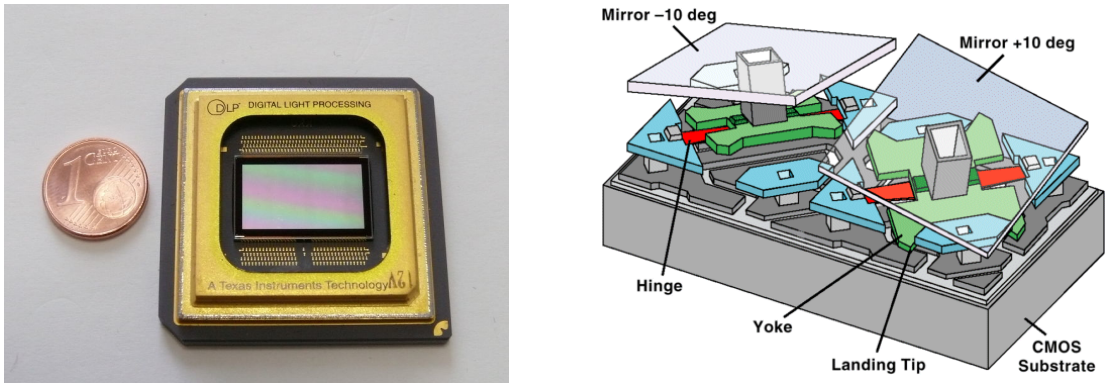

Figure 4.18: Left: photograph of the DMD chip. Right: close-up image of the DMD micro mirrors (courtesy of Texas Instruments Inc.). When voltages are applied to the CMOS substrate, the mirror tilts between two possible angles.

a beam dump, thus acting as a binary spatial light modulator.

An important market for DMDs is in digital video projectors, in which the DMD is illuminated with incoherent light from a projector bulb or LED. When the DMD is illuminated with coherent, monochromatic laser light, as in our experiment, the periodic structure of the DMD means that it acts as a blazed diffraction grating [96]. The mirrors have a spacing of $d = 13.68\,\mu\text{m}$ and a blazing angle of $\theta_b = 12°$. However, since they are blazed along the diagonal, the effective blazing angle along each direction is given by $\tan\theta_b = \tan(12°)/\sqrt{2}$. The actual angle of the reflected light θ_r is determined not by the angle of the mirrors but by the diffraction condition

$$\sin(\theta_r) = \frac{n\lambda}{d} + \sin(\theta_i), \tag{4.5}$$

where n is the order of diffraction and θ_i is the angle of incidence. For $\lambda = 785\,\text{nm}$, the blazing angle approximately coincides with the 5^{th} order of diffraction.

The DMD mirrors are made of aluminium, which has a reflectivity of 89.4%. In addition, the filling factor, diffraction efficiency and finite transmission of the DMD windows means that the overall efficiency of the device is 68%. The remaining laser power is absorbed by the DMD itself, which can eventually lead to overheating and destruction of the DMD. The damage threshold of the DMD is claimed by the manufacturer to be $10\,\text{Wcm}^{-2}$.

4.2.6 DMD speed

The main advantage of DMDs compared with other kinds of spatial light modulators are their extremely high speed. For example, our DMD has a maximum frame rate of 20 kHz, which is limited by the data transfer rate. The DMD is driven by a DDR 120 MHz bus with 64 parallel data lines (Discovery 1100 HSC). Two alternative technologies for spatial light modulation are ferroelectric liquid crystal (FLC) [97]

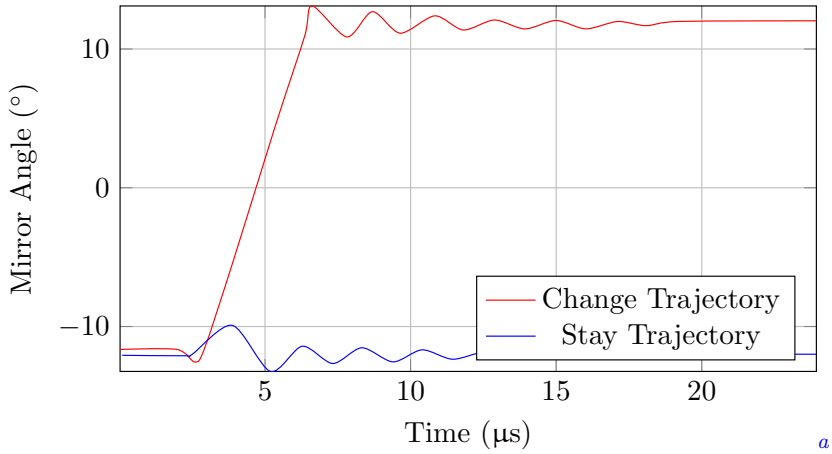

Figure 4.19: Trajectory of mirrors after a 'reset' operation. Oscillations of the mirror angles are observed regardless of whether the mirror changes state or stays in the same state.

[a]Redrawn using data from Texas Instruments Inc. http://e2e.ti.com/support/dlp__mems_micro-electro-mechanical_systems/f/94/t/213349.aspx

and nematic liquid crystal (LC) [66] modulators. Ferroelectric modulators act as binary phase modulators and have a speeds of up to 1 kHz. Liquid crystal SLMs act as continuous phase modulators, typically displaying 256 phase values at frame rates of 30 Hz. State of the art devices have 65536 phase levels and up to 204 Hz [98].

An additional limit on the maximum speed of DMDs is the mechanical settling time of the mirrors. The settling time is determined by the material properties i.e. the mass, spring constant, and damping coefficient of the mirrors. When the mirror state is reset, the mirror angle overshoots and then oscillates by $\pm 1°$ before settling at the final angle. This behaviour is shown in figure 4.19. For our DMD, the settling time is 12 µs which gives a theoretical maximum switching rate of 83 kHz. In addition, the amplitude of the mirror oscillations are large enough to deflect light out of the optical path, leading to intensity modulation of the trapping beam. Figure 4.20 shows this intensity modulation in our setup. We tried switching all 768 lines at once, as well as a smaller subset of 48 lines. This intensity modulation is a potential source of heating for dipole trapped atoms.

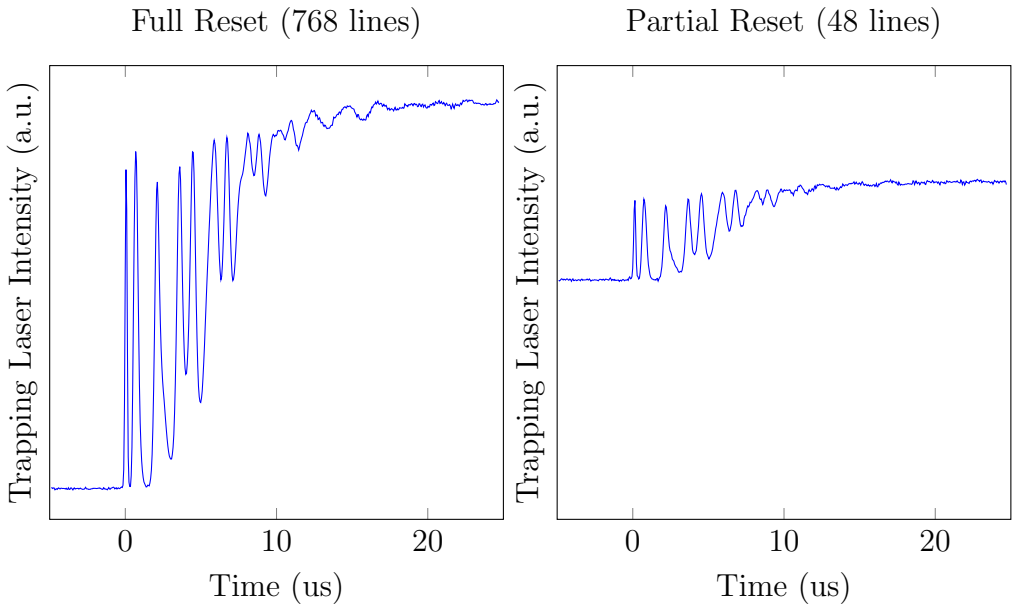

Figure 4.20: Intensity modulation due to DMD settling time. The mirror angle oscillations with amplitude $\pm 1^\circ$ cause light to be reflected out of the optical path, leading to intensity modulation. Resetting a subset of the DMD mirrors reduces the effect in proportion to the amount of light falling on those mirrors.

4.2.7 Relay optics

Rather than directly image the surface of the DMD as in chapter 3, we use the DMD to imprint a holographic pattern on to the trapping light. The aspheric lens is used to perform a Fourier transform on this hologram to reconstruct the desired arrangement of traps in the focal plane. There are two practical problems with implementing this arrangement. Firstly, the DMD is quite a large distance away from the aspheric lens. Rays leaving the DMD at different angles, after propagating over this large distance, may miss the aspheric lens entirely, causing a loss of power in the trapping beam. Secondly, the diameter of the lens, 18 mm, is not equal to the diameter of the DMD, 10.5 mm. To solve these problems, we use relay optics to form a magnified image of the DMD at the position of the aspheric lens. The relay optics consist of two lenses acting as a $2f - 2f$ telescope. A scale drawing of the lens arrangement is shown in figure 4.21.

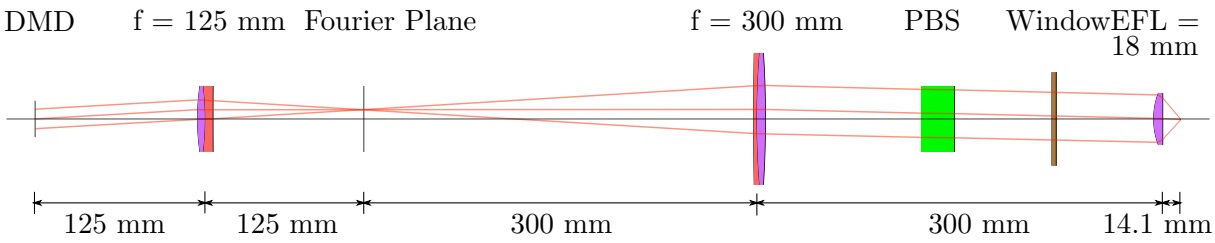

Figure 4.21: Dipole trapping optics. A holographic pattern on the DMD is magnified using a telescope formed from two achromatic doublet lenses. The telescope forms an image of the DMD on the surface of the aspheric lens. This lens performs a Fourier transform to reconstruct the image of the desired traps in the focal plane.

4.2.8 Aberrations

Careful elimination of optical aberrations is essential for producing a diffraction-limited focus. While the aspheric lens is the most critical component in the setup, the remainder of the optics must not introduce any aberrations either. The most important aberrations in our experiment are spherical aberration, astigmatism, and chromatic aberration. Spherical aberration results from the limitation of spherical lens surfaces, causing different rays to focus at different positions. This aberration is most significant for high numerical aperture lenses such as the aspheric lens and fibre collimating lens. Both of these lenses are aspheric lenses in which the surface profile has been chosen to eliminate spherical aberrations. In addition, the relay optics are made from achromatic doublets (Thorlabs AC254-125-B-ML and AC254-300-B), which are corrected for spherical aberration.

Astigmatism is when an optical system has a different focal length in the horizontal and vertical directions. This can occur if rays do not pass through a lens along its optical axis, or for cylindrical lens surfaces, which introduce a different wavefront curvature along one direction. Initially, we used a laser line filter as a dichroic mirror (Semrock MaxLine LL01-785-30), instead of the PBS in figure 5.1, to separate the dipole trapping and fluorescence light. Unfortunately, we found that the surface of the mirror was cylindrically deformed - an artifact of the manufacturing process -

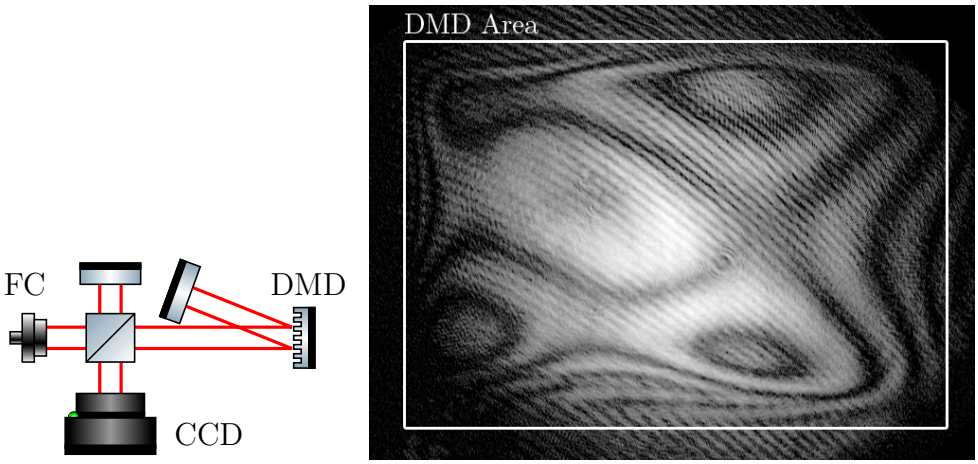

Figure 4.22: Surface of the DMD as part of a Michelson interferometer. Left: the experimental setup consisting of a beam of light from a fibre collimator (FC) incident on the DMD, which is situated in one arm of the Michelson interferometer. Right: the fringe pattern observed with the CCD is due to surface height variations of the DMD (which have an amplitude of 10λ)

which introduced an unacceptable level of astigmatism into the trapping beam.

Chromatic aberration is when there is a different focal length for different wavelengths. This is the case for the aspheric lens, which is used to focus both 780 nm fluorescence and 785 nm trapping light. The chromatic focal shift between these two wavelengths is 4.1 µm, which is smaller than the Rayleigh length of the traps. For comparison, the chromatic focal shift at 1064 nm is 170 µm. Furthermore, the additional spherical aberrations would make the lens unusable at this wavelength. Compound microscope objectives do not suffer from this problem as the multiple lenses are corrected for chromatic aberration.

Unfortunately, the largest source of aberrations in our setup is the DMD itself, as its surface of the DMD is not perfectly flat. As a result, light reflected from the DMD, when focused, forms a star-shaped spot that is much larger than the diffraction limit. We measured the wavefront error of the DMD by placing it in one arm of a Michelson interferometer. The resulting fringe pattering is shown in figure 4.22. The usual circular fringes of the Michelson interferometer are extremely distorted. By counting the fringes, we inferred that the surface varies in height by 10λ.

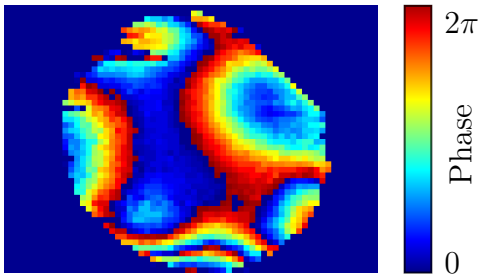

Figure 4.23: Wavefront error measured using the DMD (see section 5.2.6. Surface height variations cause a phase shift of the light. The circular aperture is due to the size of the laser beam.

Fortunately, as we use the DMD to produce artificial holograms, we are able to correct for this wavefront error by adding a spatially varying phase shift to the holograms. We developed a method for accurately measuring the wavefront error *in situ*, which is discussed in section 5.2.6. An example of the measured wavefront is shown in figure 4.23. Furthermore, this technique is able to correct for aberrations from all of the optics, not just the DMD.

4.3 Single atom imaging

One challenging aspect of the experiment is to detect the very small amount of light scattered by a single atom. To do this, we collect the scattered light using the same aspheric lens that is used to produce the dipole trap, and image it on to an electron multiplying charge coupled device (EMCCD) camera (Andor iXon 885). The principle of operation of the EMCCD camera is illustrated in figure 4.24. The camera consists of an ordinary CCD with an additional electron multiplication register. This register is capable of amplifying the electron signal by a factor of 300, large enough to overcome electronic noise in the camera, meaning that it is capable of detecting single photons.

The aim of the detection system is to collect as many photons as possible, which mainly depends on the optics that make up the microscope. These optics are shown in figure 4.25. The collected fraction of the photons is determined by the solid angle

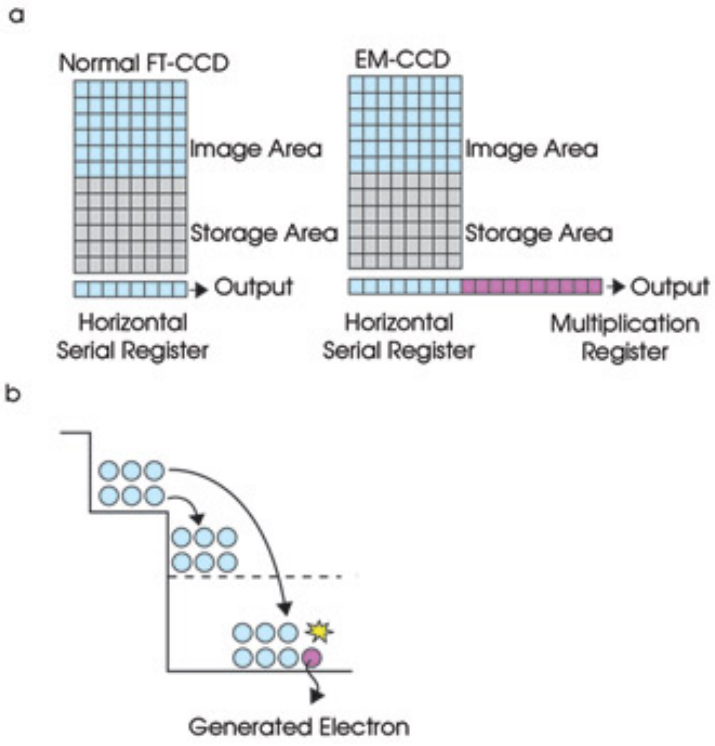

Figure 4.24: Operating principle of the EMCCD camera. (a) The camera is exposed to light producing photoelectrons in the image area. The electrons are transferred to the storage area and then read out row by row in the horizontal serial register. In an EMCCD camera, the electrons pass through a multiplication register (b) which generates additional electrons by accelerating the incident electrons through a large potential difference.

[a]Figure courtesy Andor Technology Ltd.

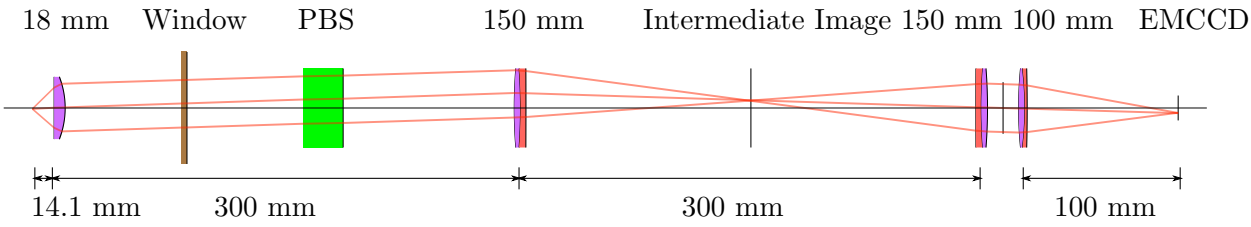

Figure 4.25: Single atom imaging optics. Light from a single atom is collected by the aspheric lens (18 mm). The collimated light passes through the vacuum window and PBS. A 100 mm lens focuses the light on to a single pixel of the EMCCD camera. A 150 mm doublet lens is used to form a 1:1 image of the aspheric lens aperture on to the 100 mm lens aperture.

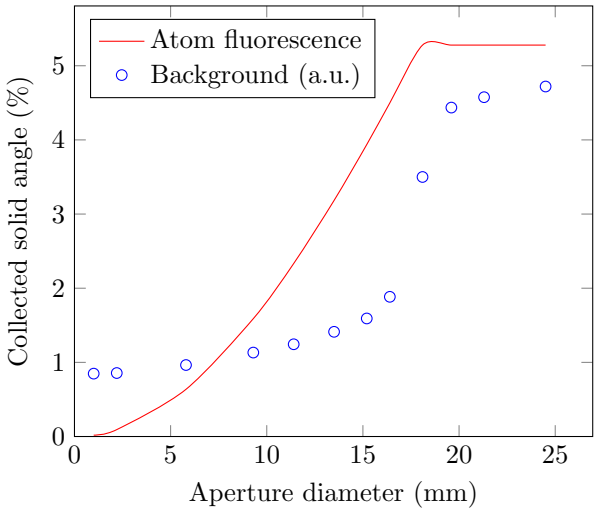

Figure 4.26: Background light variation with aperture diameter. A diameter of 16.5 mm maximises the amount of fluorescence captured while eliminating most of the background light.

encompassed by the aspheric lens, which for our NA of 0.45 is 5.3%. The light then passes through a polarising beam splitter. Since the emitted light is randomly polarised, 50% of the photons are lost. Using a dichroic mirror would avoid this problem, however, we were unable to use this because of the associated aberrations discussed in section 4.2.8. Additional lenses are used to focus the light on to a single EMCCD pixel in order to maximise the signal-to-noise ratio of the camera. A 780 nm bandpass filter is used to eliminate any residual dipole trapping light which may damage the camera. These optics have a transmission of 81%.

In addition to the light from single trapped atoms, there is some unwanted stray light from the MOT beams. This light is primarily scattered from the edges of the aspheric lens. This light is easily removed by using a circular aperture at the position of the 100 mm lens. Figure 4.26 shows the amount of background light as a function of the diameter of this aperture. An aperture diameter of 15.2 mm, i.e. slightly smaller than the diameter of the lens, is sufficient to eliminate most of the scattered light. However, the collection efficiency is reduced by a further 71% due to this aperture.

4.3.1 Calibrating the EMCCD camera

Finally, the EMCCD camera converts incident photons to electrons with a quantum efficiency of 40%. These electrons are multiplied by the chosen EM gain, amplified with a conventional amplifier, and digitised with a 16 bit analogue digital converter (ADC). The digital value output by the ADC (which we refer to as EMCCD counts and denote by C) is proportional to the number of photoelectrons, N, with a certain scaling factor, a. The mean values of these quantities, μ_C and μ_N, are also proportional,

$$\mu_C = a\mu_N, \tag{4.6}$$

while their variances are related by

$$\sigma_C^2 = a^2\sigma_N^2 + \sigma_e^2, \tag{4.7}$$

where σ_e^2 is the variance due to noise in the electronic amplifier.

It is possible to calibrate the value of a if the variance of the number of photoelectrons is limited by shot noise. In this case, $\mu_N = \sigma_N^2$, and equations 4.6 and 4.7 can be solved to relate the mean and variance of the EMCCD counts to each other as follows

$$\sigma_C^2 = a\mu_C + \sigma_e^2. \tag{4.8}$$

We can determine a by fitting a straight line to a plot of the variance versus the mean number of counts. We follow the elegant method of Tutt et al. [99] in which the EMCCD is exposed to a linear intensity gradient across its surface. The intensity gradient was produced using a red light emitting diode (LED) obscured by a triangular mask. The columns of the EMCCD were read out continuously so that the electrons from all pixels in a given column were aggregated in the horizontal serial register (see figure 4.24), a process known as full vertical binning. The number of counts in each

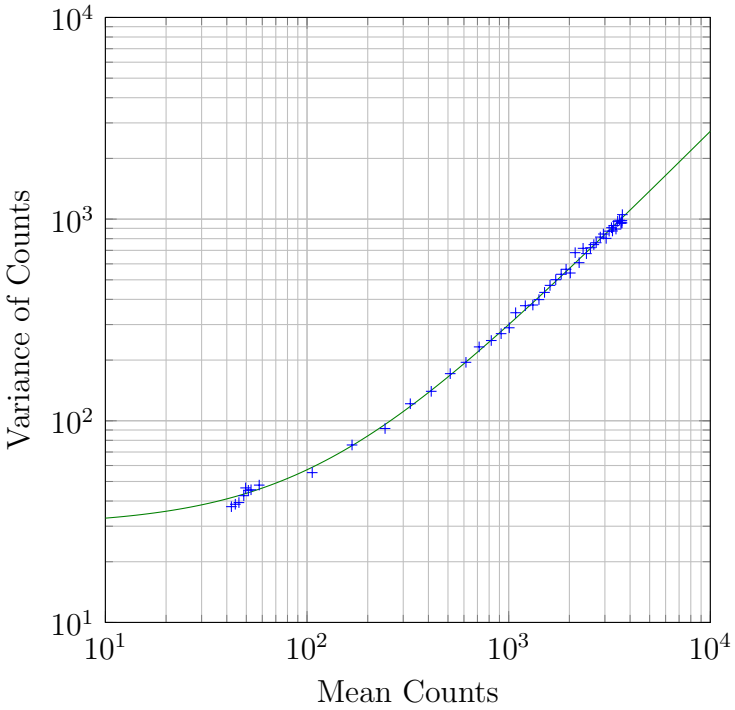

Figure 4.27: Variance versus mean of EMCCD counts. The fit parameters imply a sensitivity of $a = 0.2696 \pm 0.0005$ counts per electron and an electronic noise of $\sigma_e = 5.50 \pm 0.05$ counts.

column was recorded for 1000 frames, and the mean and variance of the counts were determined for each pixel by analysing this time series. The resulting data is shown on a logarithmic scale in figure 4.27.

The data fits extremely well to a straight line, allowing us to determine the scaling factor, a. Furthermore, the EM gain process introduces additional noise, due to the random nature of the multiplication process. When the EM gain is switched on, the variance increases by a factor of 2 [99]. We confirmed this fact for our camera (see figure 4.28).

4.4 Conclusions

I have described the experimental setup that we use for cooling and trapping atoms. This new setup eliminates the problems associated with the dielectric mirror in chap-

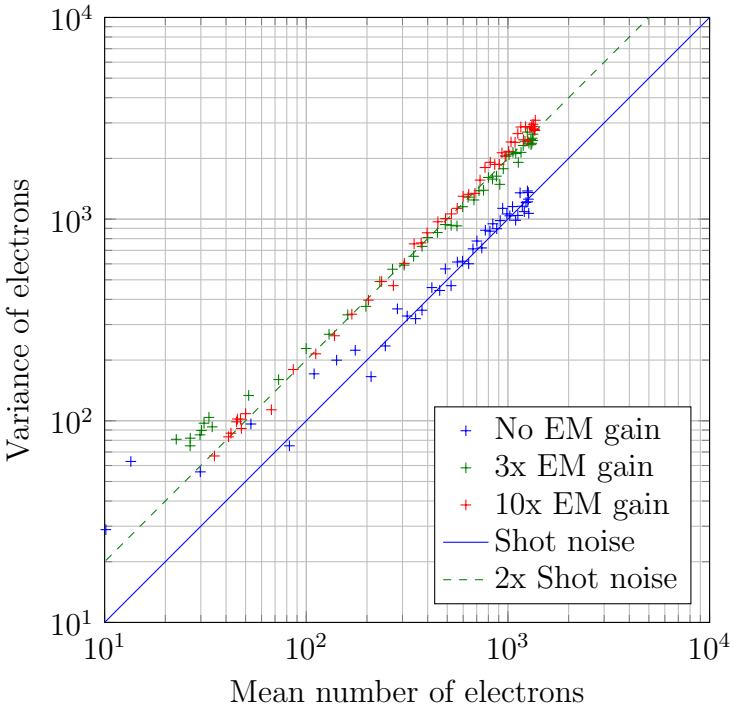

Figure 4.28: Variance versus mean of photoelectrons. Without EM gain, the variance is due to shot noise of the photoelectrons. With EM gain, the variance is increased by a factor of 2.

ter 3.

The setup, while based on standard techniques of laser cooling and trapping, has several unique features. In particular, we operate our MOT with an extremely low density in order to reduce the loading rate to and from the microscopic dipole traps. The apparatus could be further improved by including a more reliable way of controlling the rubidium vapour pressure, for example, using multiple vacuum chambers with differential pumping, and a 2D-MOT instead of a dispenser as a switchable source of cold atoms.

Chapter 5

Fast algorithms for generating binary holograms

This chapter is based on the paper *Fast algorithms for generating binary holograms* [100], which we have submitted to *Optics Express*.

We describe three algorithms for generating binary-valued holograms. Our methods are optimised for producing large arrays of tightly focused optical tweezers for trapping particles. Binary-valued holograms allow us to use a digital mirror device (DMD) as the display element, which is much faster than other alternatives. We describe how our binary amplitude holograms can be used to correct for phase errors caused by optical aberrations. Furthermore, we compare the speed and accuracy of the algorithms for both periodic and arbitrary arrangements of traps, which allows one to choose the ideal scheme depending on the circumstances.

5.1 Introduction

Since their invention by Arthur Ashkin [54], optical tweezers have had an enormous impact in diverse fields from biology to quantum physics (see [101] for a review). The underlying mechanism is the optical gradient force, which acts on polarisable particles

such as living cells [102], nanoscale tools [103] or single atoms [66], causing them to be trapped at the point of highest intensity of a tightly focused light beam. The potential energy of such a particle is proportional to the intensity of the trapping light at the position of the particle. Much effort has been devoted to designing dynamic potential landscapes for trapping and moving large numbers of particles. Regular arrays of thousands of potential wells have been created using optical lattices [64, 104], microlens arrays [75] and diffractive optical elements [105], however, these methods are limited in that the trapping sites can only be moved in unison, not individually. A second approach is to use an acousto-optic deflector (AOD) to generate a steerable trapping beam which can be used to 'paint' time-averaged potentials [106, 107]. In one demonstration [108] the authors generate 32 movable trapping beams using frequency shift key modulation. However, the time averaging is only valid when the oscillation frequency of the trapped particles is much lower than the rate of frequency shifting, which is ultimately limited by the rise time of the AOD.

A third, more flexible approach is to use a spatial light modulator (SLM) to create the desired potential landscape. The SLM can either be used as a variable attenuator to create the desired intensity pattern in the trapping light which can then be directly imaged on to the particles [1], or to display a hologram which is converted into the desired intensity landscape after propagation through a lens [103] The direct imaging technique is simple but does not make efficient use of optical power as only a small area of the SLM is active for each trap. The holographic technique concentrates a large fraction of the optical power in the active trapping sites and allows for three-dimensional positioning, but requires additional effort to calculate the hologram. There are two broad categories of SLMs: phase modulators such as ferroelectric (FLC) modulators [97] and liquid crystal displays (LCD); and amplitude modulators such as digital mirror devices DMDs [109]. Holographic optical tweezers typically use a phase modulating LCD, coupled with an iterative phase retrieval algo-

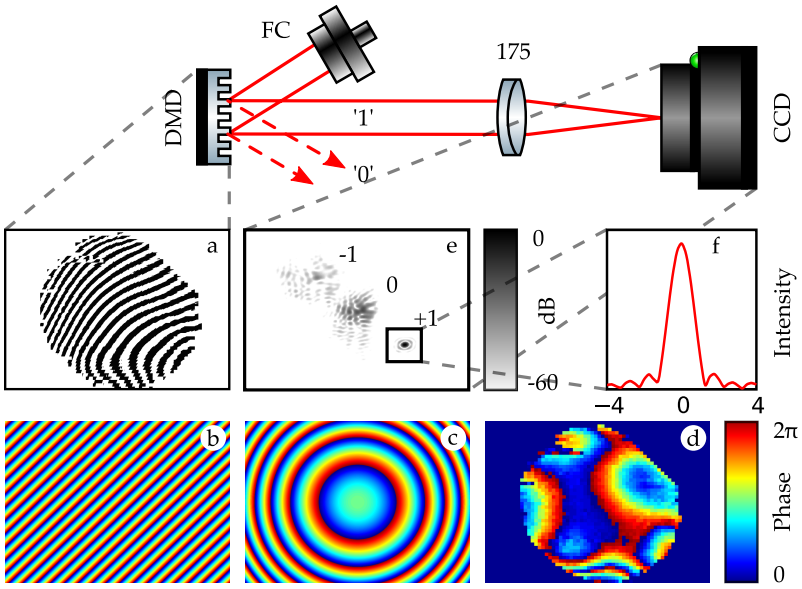

Figure 5.1: The experimental setup used for testing the holographic optical tweezers. A collimated beam of laser light at 785 nm is delivered from a fibre coupler (FC) and is incident on the DMD. The reflected light from mirrors in the '1' state is focused using an achromatic doublet lens (f = 175 mm). The traps are produced in the focal plane of this lens and imaged with a CCD camera. (a) The binary hologram on the DMD, formed from (b) a linear phase ramp, (c) a quadratic lens and (d) an arbitrary phase correction map. (e) The intensity measured by the CCD camera. (f) A line profile through the desired trap showing diffraction-limited performance (the x-axis is in dimensionless units of $\lambda D/f$ where D is the diameter of the limiting aperture).

rithm such as the Gerchberg-Saxton algorithm [110], mixed-region amplitude freedom (MRAF) [111], offset MRAF [112], or conjugate gradient minimisation [113]. Here we demonstrate the use of a DMD in a holographic optical tweezers apparatus, along with algorithms that are much faster than standard iterative phase retrieval. DMDs are extremely fast compared to all other available SLMs, with full frame rates up to 20 kHz. The DMD we are using at present in our laboratory (Texas Instruments DMD Discovery 1100) is an array of 1024×768 micro-mechanical mirrors. Each mirror can be switched between two different angles, which we refer to as '0' and '1'. A mirror in the '1' position reflects light through the remainder of the optical setup whereas a mirror in the '0' position deflects light towards a beam stop, thus acting as a binary amplitude modulator.

In this chapter, we show how to use a DMD to holographically generate large arrays of individually movable trapping sites. The chapter is divided into three sections. Firstly, we describe several different algorithms for rapidly calculating artificial holograms for an amplitude-modulating SLM, and we show how to apply these to modulators that only permit binary amplitude modulation such as a DMD. Secondly, we show how the DMD can be used as a wavefront sensor to accurately measure and correct for optical aberrations. The ability to do this is essential for producing tightly focused diffraction-limited traps. Finally, we compare the different algorithms on the basis of speed of convergence, efficiency of use of laser power and accuracy of the resulting trapping potentials.

5.2 Computation of holograms

The general problem of holography is to produce an optical field $H(x, y)$, the hologram, which after propagation through a lens results in a desired optical field $F(x', y')$, the image. The problem can be divided in to two parts [114]: the computational prob-

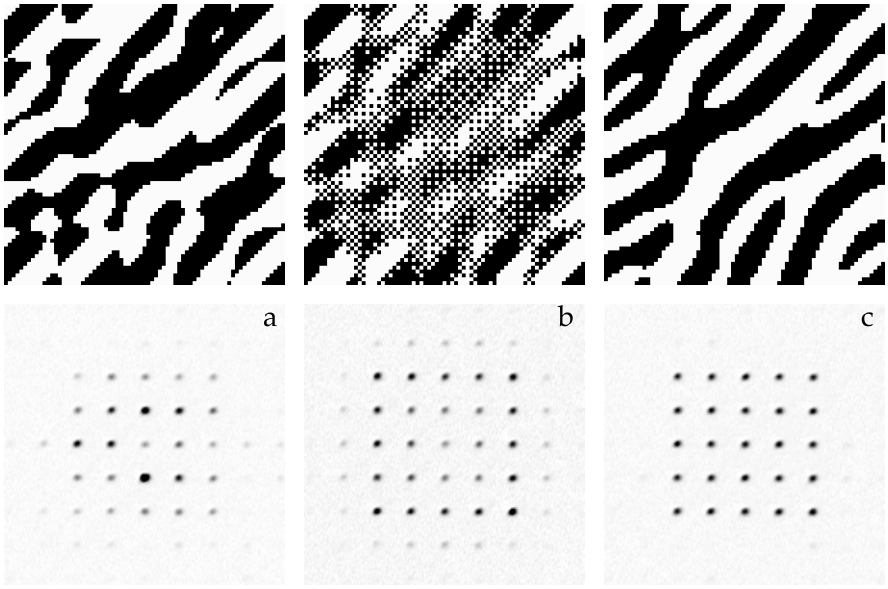

Figure 5.2: Example images illustrating three of the different hologram algorithms: (a) binary rounding, (b) ordered dithering, and (c) the Gerchberg-Saxton algorithm. Each case shows portion of the DMD (73×73 px) showing the binary hologram and an image of the resulting intensity pattern taken with a CCD camera. The distance between the traps is 173 μm.

lem of how to calculate the required optical fields; and the representational problem of how to display the complex-valued fields using a physical light modulator. We consider three different algorithms which from here on in we refer to as binary rounding, dithering, and weighted Gerchberg-Saxton. Figure 5.2 shows examples of holograms and images generated using the different techniques.

5.2.1 Hologram of a single focused trap

The common characteristics and limitations of all four algorithms are best illustrated by considering the hologram field of a single diffraction-limited trap at a precisely defined location:

$$H(x,y) = A\exp(i\frac{2\pi}{f\lambda}(x_0'x + y_0'y + \frac{z_0'}{2f}(x^2+y^2)) + i\phi(x,y)) \tag{5.1}$$

$A = 1$ encodes the amplitude and phase of the trap, x_0', y_0', z_0' are its coordinates

relative to the focal plane of the lens (with focal length f), and λ is the wavelength of the trapping light. x and y are the coordinates in the plane of the DMD and $\phi(x, y)$ is an arbitrary function that we introduce to compensate for any systematic phase variations in the actual optical system. The terms in this equation are illustrated graphically in figure 5.1. The two terms linear in x and y (fig. 5.1b) allow the trap to be displaced within the focal plane of the lens. The $(x^2 + y^2)$ term (c) introduces a small defocus which allows the trap to be moved in and out of the focal plane by a distance z_0'. Finally, the arbitrary phase correction (d) primarily compensates for the wave-front distortion due to the non-planar DMD surface we found in our demonstration setup. The image $F(x', y')$ may be calculated using the Fresnel diffraction formula [114]. If we neglect the quadratic lens term, this is essentially a Fourier transform of $H(x, y)$. The resulting $F(x', y')$ is a delta function spot at the location (x_0, y_0). In reality, the spot is an Airy pattern whose width is determined by the limiting aperture of the setup (see fig. 5.1f). Furthermore, since only the intensity of the image $|F(x', y')|^2$ is relevant for trapping, we are free to choose the phase $\arg(A)$ of the trap.

5.2.2 Binary rounding algorithm

The simplest way to represent this complex-valued field on the DMD is the binary rounding algorithm: to map all pixels whose phase is between $-\pi/2$ and $\pi/2$ to the '1' state, and all other pixels to the '0' state. This simple mapping results in the maximum amount of optical power being directed in to the desired trap, since all the pixels in state '1' interfere constructively, and all the pixels that would interfere destructively are in state '0'. The result is a top hat grating (fig. 5.1a) where the fraction of the total optical power in the n^{th} order of diffraction is given by $\frac{1}{2}\mathrm{sinc}(n\pi/2)^2$. The desired trap is produced in the $+1$ diffraction order, with a theoretical maximum power of $1/\pi^2 \approx 10.1\%$. The -1 and 0 orders are blurred because the phase correction

acts in the opposite direction for these orders - doubling the wavefront error.

5.2.3 Extension to multiple traps

It is straightforward to extend the above method to generating multiple traps by simply summing the complex-valued holograms for each trap. The A parameter is used to set the magnitude and phase of each trap. We set the magnitude of the n^{th} trap $|A_n| = 1/\sqrt{N}$ in order to evenly distribute the total power amongst the N traps. In addition, we set the phase of each trap to a random value between 0 and 2π in order to prevent the amplitude maxima of all of the individual trap holograms from accumulating at the same position.

Now we consider how to rapidly compute this sum. One approach is to explicitly evaluate equation 5.1 for each of the N traps at all M pixels and sum the results. This approach was taken by Bowman et al. [98] for a phase modulator, using a graphics processing unit (GPU) to parallelise the calculation. For $N = 25$ traps and a $M = 512 \times 512$ pixel grid, they achieve a frame rate of between 60 Hz and 500 Hz depending on the choice of hardware and accuracy of the output. A disadvantage is that the computation time scales as $N.M$ and so increases linearly with the number of traps.

The approach we have taken is to implement the sum using a fast Fourier transform (FFT). The advantages of this approach are numerous. Highly optimised FFT libraries (such as FFTW [115]) exist that run on ordinary computer hardware. Furthermore, the algorithm scales as $M \log M$, which is independent of the number of traps. Thirdly, if the number of traps is sparse such that $N \ll M$ then only the rows that actually contain traps need to be transformed, yielding a speedup by a further factor of two. The FFT method is particularly efficient if the lens and phase correction terms are the same for all traps, as these can simply be multiplied with the end result of the FFT. If different lens terms are desired, the lenses can be incorporated

efficiently via a convolution. Benchmarking on a 2GHz Intel Core 2 Duo CPU gives a frame rate of 300 Hz, which is comparable to the fastest speeds hitherto achieved but with the advantage of being independent of the number of traps and not requiring any specialised hardware.

5.2.4 Ordered dithering algorithm

Unfortunately, for the case of multiple traps, the problem of representing the complex hologram as an amplitude-only one is more difficult. The simple binary rounding algorithm is highly non-linear, which results in additional unwanted ghost traps [116] and drastic variations in intensities of the traps. The non-linearity arises from two sources: the quantisation of continuous pixel values to two binary values, and the fact that any pixels that are outside the range $(0, 1)$ are truncated. The ordered dithering algorithm seeks to minimise this non-linearity by mimicking continuous greyscales on the binary DMD [117]. Dithering works by alternating pixels between '0' and '1' with a high spatial frequency so that the average pixel value is approximately equal to a continuous value.

The first step in the algorithm is to create a 512×512 array in which to store the target image F. We set each entry corresponding to a trap to $\exp(i\theta_n)/\sqrt{N}$ where θ_n is a random phase. Next, we perform an inverse Fourier transform on the array to find the hologram H. After applying the quadratic lens and phase correction, we take the real part of the complex-valued hologram field. This step similar to Gabor's [118] original approach (of recording the interference pattern between the field and a uniform reference wavefront) since the interference term is proportional to $\Re(H(x, y))$. In addition, we scale and add an offset of $1/2$ since the real part can be both positive and negative while the DMD only accepts positive values. Finally, we truncate to $(0, 1)$ and dither the result by comparing it to a fixed threshold matrix, specifically, an 8×8 Bayer matrix [119]. Alternatively, the ordered dithering step can be replaced

with an error diffusion dithering algorithm such as Floyd-Steinberg (see [120] for a description of the algorithm). A pseudocode version is given below:

1: **function** ORDERED DITHER(x_n, y_n)
2: $F(x_n, y_n) \leftarrow \exp(i\theta_n)/\sqrt{N}$
3: $H \leftarrow$IFFT(F)
4: $H \leftarrow H.\exp(i\frac{z_0\pi}{f^2\lambda}(x^2+y^2) + i\phi(x,y))$
5: $H \leftarrow$REAL($(H+1)/2$)
6: $H \leftarrow$CLIP($H, 0, 1$)
7: $H \leftarrow H > threshold\_matrix$
8: **return** H
9: **end function**

5.2.5 Weighted Gerchberg-Saxton algorithm

Another alternative algorithm that we consider is based on an iterative phase-retrieval algorithm: the weighted Gerchberg-Saxton algorithm. This algorithm exploits the phase freedom of the individual trap holograms in equation 5.1. The algorithm is initialised with a target image of traps with equal amplitudes and random phases. One then calculates the hologram by inverse Fourier transform and binary rounding. Next, the image that would be formed from this hologram is calculated by a forward Fourier transform. After this first iteration, the resulting traps in the image may have widely varied intensities. The initial image is updated according to the newly calculated image: the new phases are adopted and the intensity of the n^{th} trap is weighted by w_n such that weaker traps are boosted and more intense traps are attenuated. After a few dozen iterations, the algorithm converges towards a result with equal power in each trap. In pseudocode:

1: **function** GERCHBERGSAXTON(x_n, y_n)
2: $F(x_n, y_n) \leftarrow \exp(i\theta_n)/\sqrt{N}$

3: **repeat**

4: $H \leftarrow \text{IFFT}(F)$

5: $H \leftarrow \text{Real}(b) > 0$

6: $F \leftarrow \text{FFT}(H)$

7: $w_n \leftarrow w_n / abs(F(x_n, y_n))$

8: $F(x_n, y_n) \leftarrow w_n . \exp(i \arg(F(x_n, y_n)))$

9: **until** convergence

10: $H \leftarrow \text{IFFT}(F)$

11: $H \leftarrow H . \exp(i \frac{z_0 \pi}{f^2 \lambda}(x^2 + y^2) + i\phi(x, y))$

12: $H \leftarrow \text{Real}(b) > 0$

13: **return** H

14: **end function**

5.2.6 Wavefront correction

Even though the DMD is an amplitude modulator, equation 5.1 allows for an arbitrary phase function $\phi(x, y)$ to be included in the hologram. This enables the correction of wavefront errors in the illuminating light. Such wavefront errors can be caused by optical aberrations in the lenses, or by the non-planarity of the DMD itself. The latter are particularly severe: we measured wavefront errors of up to 10λ over the surface of our particular DMD.

We determined the wavefront error at each point on the surface of the DMD using a technique similar to Cizmar et al. [121]. The technique is illustrated graphically in figure 5.3. The surface of the DMD is divided in to blocks of 16×16 pixels. All blocks were switched off except for two: one block of 16×16 pixels centred around the position (x, y) and another centred at the origin. We denote the distance between the two blocks as b. The light from these two regions interferes on the CCD camera to form a sinusoidal Young's slit interference pattern whose fringe spacing is proportional

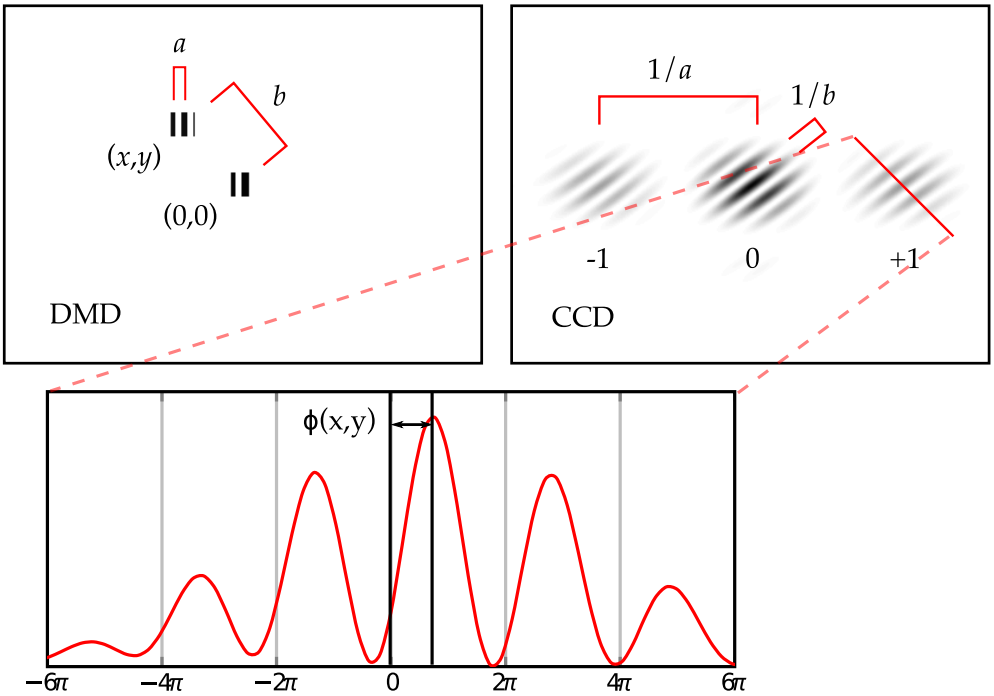

Figure 5.3: Illustration of the technique for measuring wavefront errors (see text for description).

to $1/b$. Within each block on the DMD, there is a diffraction grating with period a. This causes the -1, 0 and $+1$ diffraction orders on the CCD to be separated by $1/a$. The fringes in the $+1$ order are examined by taking a line profile. The displacement of the interference fringe maximum from zero is proportional to $\phi(x, y)$, the phase of the light due to the wavefront error. We fit the interference pattern with a least-squares algorithm to extract this phase offset. This procedure is repeated for each block on the surface of the DMD, which takes approximately one minute. The resulting phase map is the one shown in figure 5.1d. Including this phase map in the hologram calculation eliminates the wavefront errors and yields a diffraction-limited spots in the first order of diffraction.

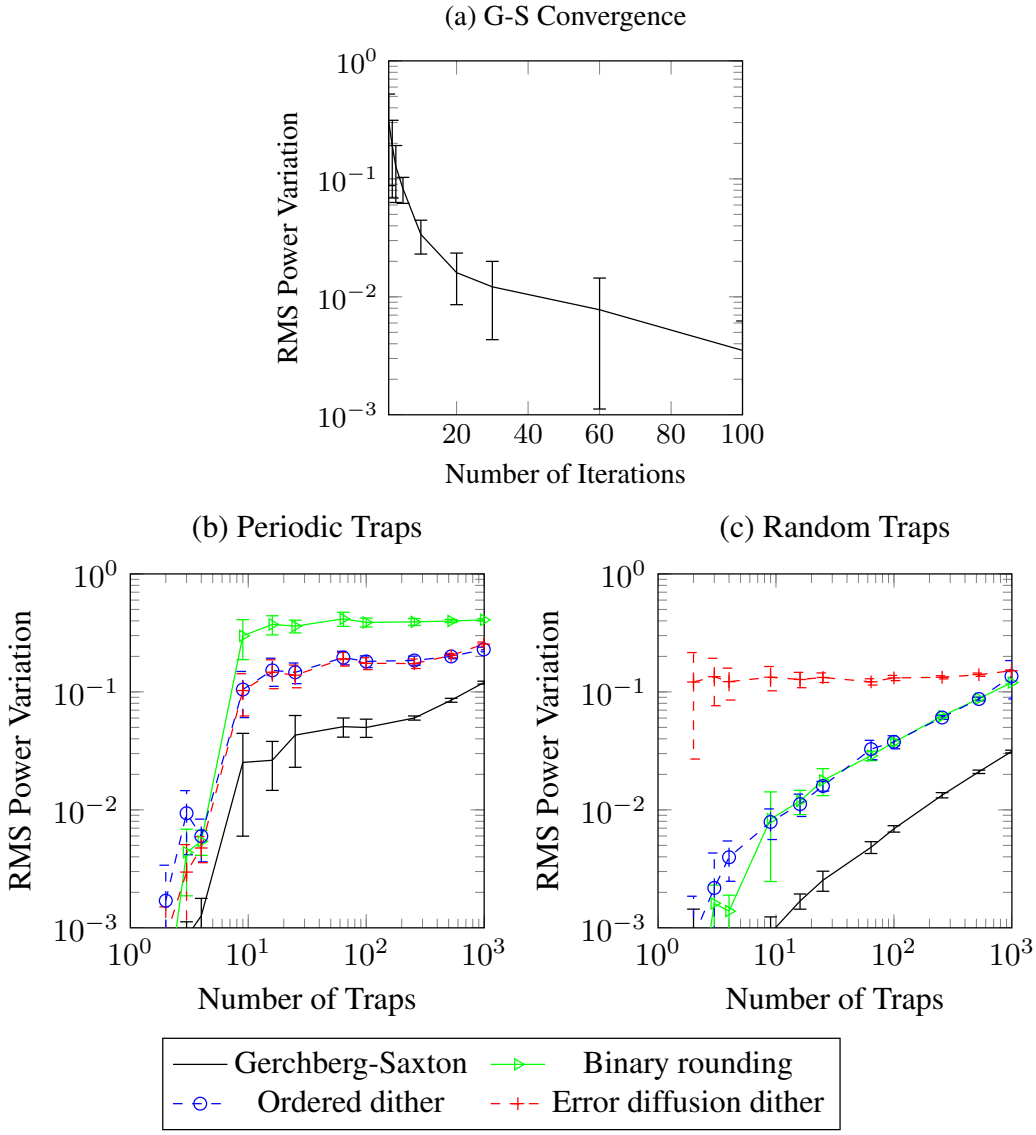

Figure 5.4: Evaluation of the performance of the different algorithms for different numbers of traps. (a) The speed of convergence of the Gerchberg-Saxton algorithm for 25 periodic traps. (b) Relative RMS power variation for different periodic arrays of traps and (c) for randomly positioned traps. The error bars are statistical from a sample of 10 runs with different initial random phases.

5.3 Comparison of different algorithms

All of the algorithms have similar capabilities. They can be used to position a large number of traps in three dimensions, and can all correct for wavefront errors. Furthermore, they all suffer from the same set of problems: loss of optical power in unwanted diffraction orders, generation of unwanted ghost traps, and variation of the power in existing traps. We evaluated the performance of each algorithm with a numerical simulation for two arrangements of traps: a periodic square lattice (as in fig. 5.2, spacing 4 pixels), and a set of random, uniformly distributed traps. The simulations were performed at a resolution of 512×512 pixels and averaged over 10 runs with different initial random phases for the traps. For the case of randomly positioned traps, their position was also randomised for each run. Figure 5.4 shows the results of how the severity of the power variation scales with the number of traps for each algorithm.

The loss of power originates from the conversion to a real-valued hologram. Both the binary rounding and Gerchberg-Saxton algorithm achieve the theoretical maximum of 10.1 %, while the power per trap for the dithering algorithms is slightly lower. This is a consequence of line 5 of the ordered dithering algorithm, which can be expanded as

$$\frac{1}{2}(\Re(H(x,y)) + 1) = \frac{1}{2} + \frac{1}{4}H^*(x,y) + \frac{1}{4}H(x,y) \tag{5.2}$$

The $H(x,y)$ term is Fourier transformed to give $F(x',y')$, resulting in the desired arrangement of traps with a relative power of $1/4^2 = 1/16 = 6.25\%$, somewhat less than the theoretical maximum. Additionally, the $1/2$ term results in the zeroth-order diffraction spot containing $1/4$ of the total power, and the $H^*(x,y)$ term becomes $F(-x',-y')$, a mirrored version of the traps containing $1/16$ of the power.

The problem of ghost traps is not especially important for optical trapping as

the ghost traps usually appear spatially separated from other traps, and are often too weak to trap particles. However, the problem of intensity variations is significant, and depends on both the arrangement of traps and the choice of algorithm. We quantify the variation by the relative RMS power variation i.e. the standard deviation of the power in each trap divided by the mean.

First, consider the general case of randomly positioned traps with random initial phases. The holograms of each trap add together incoherently. Therefore it is appropriate to add the intensities of each trap in quadrature, giving a total RMS amplitude of 1. After scaling according to line 5 of the ordered dithering algorithm, the majority of pixels in the hologram lie in the range $(0, 1)$ and so the clipping step has a negligible effect on the traps. For random traps, even the binary rounding algorithm yields acceptable results, with an RMS power variation of a few percent. The ordered dither algorithm achieves similar performance. Surprisingly, the error diffusion algorithm performs relatively poorly, especially for small numbers of traps. This is because the error diffusion dither distributes some noise to locations on the DMD which contain traps. In the case of ordered dithering, the noise is concentrated around a few specific locations corresponding to the spatial frequencies contained within the threshold matrix. Finally, the Gerchberg-Saxton algorithm is capable of reducing the RMS power variation even further. Figure 5.4a shows how the RMS power variation improves with more iterations. High numbers of iterations suffer from diminishing returns. We choose 10 iterations as a compromise between speed and accuracy, but this is still 10 times slower than the ordered dithering algorithm.

However, for certain specific arrangements of traps, the binary rounding algorithm fails utterly. In the particular case of a periodic grids of traps, the maxima of the individual trap holograms tend to coincide at a few points on the DMD even with random initial phases. As a result, these locations in the hologram have an amplitude far outside the range $(0, 1)$. The subsequent clipping step severely distorts the trap

intensities. For periodic traps, the dithering algorithm is a significant improvement compared to the binary rounding algorithm. As before, the Gerchberg-Saxton algorithm can reduce the trap variation even further, but requires many iterations in order to do so.

For periodic arrangements of traps, the ordered dithering algorithm is a good compromise between power per trap, uniformity, and speed. When a periodic lattice of traps is not required, all algorithms provide good accuracy and the simple binary rounding algorithm is sufficient.

5.4 Conclusion

In conclusion, we have shown that a DMD can be used to holographically generate arrays of diffraction-limited light spots that are ideally suited for optical tweezers. While the efficiency is low (10.1%), their fast full frame rate of 20 kHz is on the same timescale as the motion of laser-cooled atoms. Our tweezers have wide-ranging applications in the atomic physics wherever single atoms need to be manipulated accurately, for example, close to a dielectric surface, photonic crystal cavity, or superconducting circuit. In addition the high speed means that we can execute fast feedback on the atoms to precisely control their position. Our approach is not just limited to single atoms, but should enable new experiments in other fields (e.g. biological systems). Finally, the speed of each algorithm is crucial in order to utilise the full capability of the DMD. High speed algorithms such as ordered dithering are a significant simplification over iterative phase retrieval algorithms. High speed is necessary for real-time manipulation and feedback on trapped particles, and we have shown here that these can be realised using current technologies.

Chapter 6

Moving single atoms

This chapter describes the main results of the experiment and is being prepared for publication.

6.1 Introduction

Quantum information science is a very active field of research that seeks to revolutionise computation, metrology, and cryptography. Single neutral atoms are amongst the most promising candidates for physically implementing quantum bits, the fundamental unit of quantum information. Many schemes have been demonstrated for performing quantum logic gates with atoms, for example, using single photons in high-finesse cavities [20, 92], controlled collisions of cold atoms [122], or interactions between atoms in Rydberg states [27]. For a large scale quantum computer, all these schemes rely on being able to trap large arrays of atoms in which each atom can be individually addressed, positioned, and transported. Two types of transport are of interest: *bright* transport in which atoms are continuously cooled as they are repositioned (which is useful for rapidly reconfiguring atoms), and *dark* transport where the atoms do not interact with radiation (preserving their coherence).

In this chapter, we describe a new method for creating such a reconfigurable array

of traps. Our traps are formed from a tightly focused far off-resonant laser beam, which traps atoms using the optical dipole force. Several other approaches have been used to create large arrays of optical dipole traps, for example optical lattices [64], microlens arrays [75] and diffractive optical elements [105]. However, in all these systems, the trapping sites can only be moved in unison, not individually. Several other approaches have been tried for creating reconfigurable traps, for example, 2D acousto-optic deflectors (AODs) [107, 108] and liquid crystal spatial light modulators (SLMs) [65, 97]. AODs are fast and can be used to paint complicated time-averaged trap shapes, but the time-averaging is only valid for slowly moving atoms. Impressive results have been achieved by using liquid crystal SLMs to imprint a phase hologram on the trapping light - allowing for arbitrary arrangements of traps [66]. However, these devices are quite slow (typically limited to frame rates of 60 Hz), only work at a specific wavelength, and have limited power handling capability which limits the maximum number of dipole traps.

Here we demonstrate that it is feasible to use a digital mirror device (DMD) to holographically generate arrays of movable dipole traps, which overcomes most of these limitations. Our DMD (Texas Instruments DLP Discovery 1100) is a 1024×768 array of micro-mechanical mirrors. The angle of each mirror can be digitally switched between $-12°$ and $+12°$, which we refer to as 'on' and 'off'. With a full frame rate of 20 kHz, our DMD is much faster than all available liquid crystal SLMs. Furthermore, as was shown in chapter 5, the DMD can be used to correct optical aberrations in the experimental apparatus, which is essential for producing tightly focused dipole traps.

This chapter is divided into four sections. Firstly, we describe the experimental setup for cooling and trapping atoms, and characterise the properties of a single dipole trap. Secondly, we show how the DMD can be used to generate arbitrary arrangements of up to 20 atom traps. Thirdly, we demonstrate both bright and dark transport of arrays of single atoms, and finally, we discuss the future prospects for

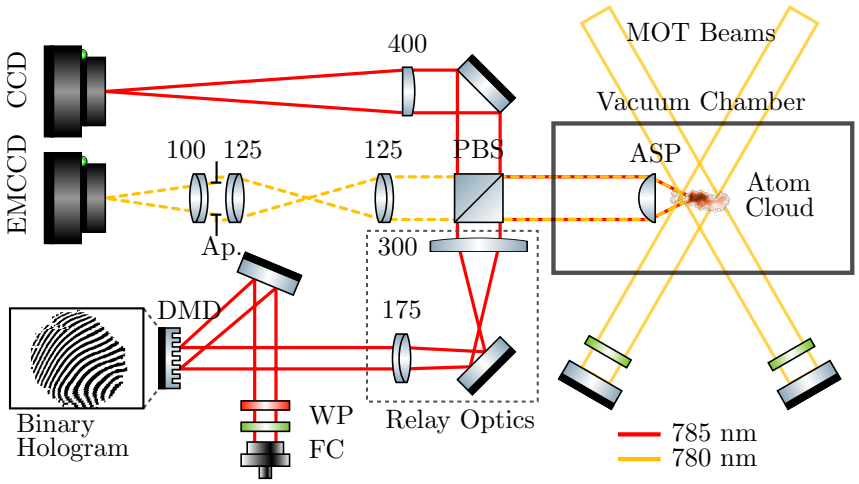

Figure 6.1: Diagram of the experimental setup. A magneto-optical trap (MOT) formed from three retro-reflected cooling laser beams confines a cloud of cold rubidium atoms at the centre of a vacuum chamber. Dipole trapping light at 785 nm is delivered from a single mode fibre collimator (FC) through waveplates (WP) to correct the polarisation, and directed on to the digital mirror device (DMD). The DMD imprints a holographic intensity pattern on to the trapping beam, which is transferred via relay optics on to an aspheric lens (ASP) inside the vacuum chamber. The microscopic dipole traps are formed in the focal plane of this lens. A polarising beam splitter (PBS) reflects most of the trapping light into the vacuum chamber, and transmits the remainder to a CCD camera for monitoring the shape of the traps. Fluorescence light at 780 nm from single trapped atoms is collected by the same aspheric lens and imaged on to an EMCCD camera. An additional aperture (Ap.) was used to eliminate background light from the MOT beams, and a bandpass filter (BPF) at 780 nm was used to eliminate any residual dipole-trapping light.

implementing a quantum information processing scheme using our setup.

6.2 Trapping and imaging single atoms

First we describe the apparatus that we use for cooling and trapping arrays of atoms (depicted in figure 6.1). The apparatus is essentially the combination of a magneto-optical trap (MOT) for cooling small clouds of atoms and a microscope for trapping and imaging single atoms. At the heart of the experiment is a single aspheric lens (Asphericon HPX 20-18) with a numerical aperture $NA = 0.49$. This lens serves both to focus the trapping light to a diffraction limited spot, and to collect the light

scattered by single atoms. The role of the DMD is to create multiple reconfigurable dipole traps from a single laser beam.

6.2.1 Dipole traps

The physical principle underlying the traps is the ac-Stark effect: in the presence of an oscillating electric field from a laser beam with a frequency close to the rubidium 5s $\rightarrow$ 5p transition, the interaction between the laser beam and the atom causes a shift of the energy of the atomic states. For red-detuned light, the ground state is shifted to a lower energy by an amount proportional to the intensity of the light. When the laser is detuned far from resonance, this spatially varying light shift forms a conservative potential of the form

$$U_{\mathrm{dip}} = \frac{\hbar\Gamma^2}{24}\frac{I}{I_{\mathrm{sat}}}\left(\frac{1}{\Delta_{1/2}} + \frac{2}{\Delta_{3/2}}\right) \tag{6.1}$$

where $\Gamma = 2\pi \times 6$ MHz is the natural linewidth of the transition, $I_{\mathrm{sat}} = 2\hbar\pi^2 c\Gamma/3\lambda^3 = 16.7\ \mathrm{Wm}^{-2}$ is the saturation intensity and $\Delta_{1/2}, \Delta_{3/2}$ are the frequency detunings from the $^2\mathrm{P}_{1/2}$ and $^2\mathrm{P}_{3/2}$ excited states respectively. Our trapping light is produced by tapered amplifier laser (Toptica DLX 110) at a wavelength of $\lambda = 785.4\,\mathrm{nm}$. This value, which lies between the $^2\mathrm{P}_{3/2}$ and $^2\mathrm{P}_{1/2}$ states, is a compromise that gives a large trap depth while minimising the scattering rate due to the trapping light. We model the focused laser as a Gaussian beam, giving a potential of the form $U_{\mathrm{dip}}(r) = U_0 \exp(-2r^2/w_0^2)$. For a laser power of 0.91 mW per trap and a waist $w_0 = 1.7$ µm we expect a trap depth $U_0 = k_{\mathrm{B}} \times 0.53\,\mathrm{mK}$.

The source of the atoms is a magneto-optical trap (MOT), which produces a cloud of $\sim 10^6$ laser-cooled rubidium-87 atoms. The cooling lasers are tuned to the $^2\mathrm{S}_{1/2}\ F = 2 \rightarrow\ ^2\mathrm{P}_{3/2}\ F' = 3$ cooling transition at 780.2 nm. In addition to forming the initial cloud, the MOT beams are used to cool the atoms into the dipole traps,

and to scatter photons which allow the trapped atoms to be detected. For loading the dipole traps, we set the intensity of the MOT beams to $I/I_{\text{sat}} = 4$ and the detuning to $\Delta = -6\,\Gamma$. We measured the temperature of the atoms in the MOT to be $T = 25\,\mu\text{K}$ which is below the Doppler limit [83].

6.2.2 Imaging system

The scattered light is collected using the same aspheric lens as is used to produce the dipole traps and is imaged on to an electron multiplying EMCCD camera (Andor iXon 885). This camera includes an EM gain register to amplify the electron signal and is capable of detecting single photons. We use an EM gain of between 30 and 100, which is sufficient to overcome the electronic noise of the EMCCD readout amplifier. The magnification was chosen so that a diffraction-limited point source is focussed on to a single camera pixel. The fluorescence image was contaminated with a small amount of light from MOT beams that was reflected from the edges of the lens. An additional aperture with $d = 15.2\,\text{mm}$ was introduced into the imaging system so as to eliminate this background light.

The expected scattering rate from a single atom is given by

$$R_{\text{sc}} = \frac{\Gamma}{2}\frac{I/I_{\text{sat}}}{1 + 4\Delta^2/\Gamma^2 + I/I_{\text{sat}}}. \tag{6.2}$$

To calculate the expected number of photons per atom, one must take into account the light shift of the transition frequency due to the dipole trap. According to equation 6.1, the ground state is shifted down by $U_0/\hbar = 2\pi \times 11\,\text{MHz}$, while the excited state is shifted up by $2\pi \times 15\,\text{MHz}$ (ignoring the $\sim 6\%$ contribution due to the 5p $\rightarrow$ 6s and 5p $\rightarrow$ 5d transitions). Adding the total light shift $\Delta_{ls} = 2\pi \times 25\,MHz$ to the laser detuning yields a photon scattering rate of $R_{\text{sc}} = 180\,\text{kHz}$ for a single atom. The overall detection efficiency is 0.61%, which is due to the fraction of the

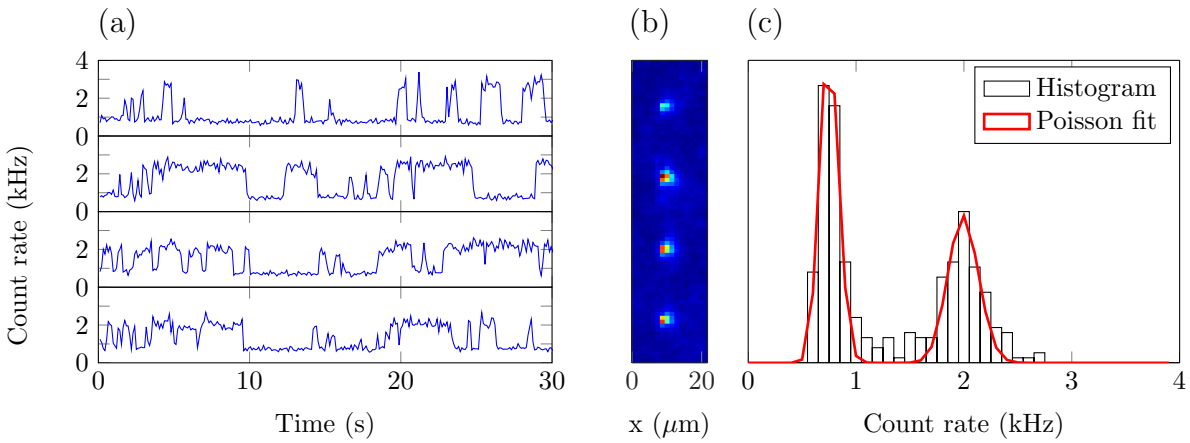

Figure 6.2: Statistical analysis of the light emitted from trapped atoms. (a) Number of photons collected from each dipole trap. (b) Image of the fluorescence light averaged for 30 seconds. (c) Histogram of the photon counts from trap 4. The peak at 0.8 kHz is due to background light, while the peak at 2 kHz corresponds to the presence of a single atom. The variance of the count rate is consistent with Poissonian photon counting statistics. We never observe two atoms in the dipole trap due to the collisional blockade effect.

solid angle collected by the lens (5.3%), splitting of the randomly-polarised light at the PBS (50%), the additional aperture (71%) the quantum efficiency of the EMCCD camera (40%) and imperfect transmission of the optics (81%), giving an expected photoelectron count rate of 1.1 kHz on the EMCCD camera.

6.2.3 Observation of single atoms

Figure 6.2 shows the photon counts from four individual dipole traps. Atoms are continuously exchanged between the MOT and the traps, meaning that only a subset of the traps are occupied at any given time. The number of counts from each trap jumps between two distinct values indicating the presence or absence of a single atom. The background count rate of 0.8 kHz is due to both fluorescence from other atoms in the MOT and unwanted reflections of the MOT beams from the aspheric lens. The average count rate from a single atom is 1.2 kHz, which is in good agreement with the calculated value in section 6.2.2. Figure 6.3 shows images that have been post-selected and averaged for all possible combinations of filled traps. The atoms

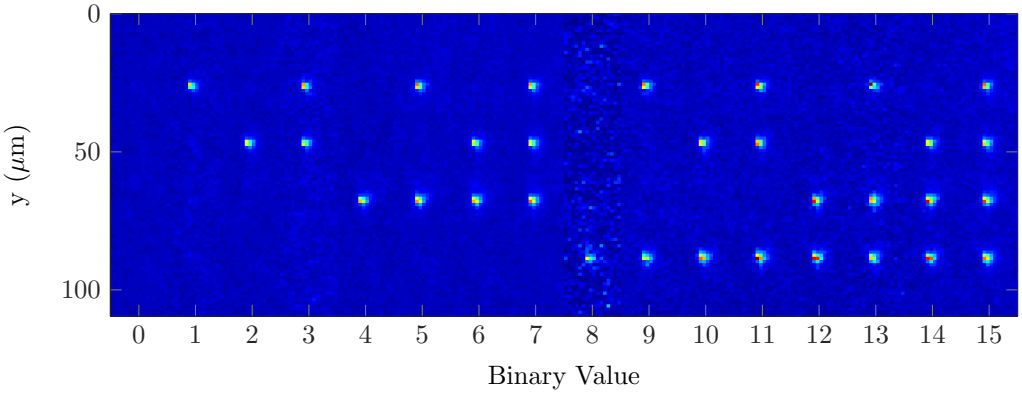

Figure 6.3: Post-selected images of all possible combinations of filled traps. Each image is averaged over many occurrences. The images are arranged in ascending order of their binary representation (from 0 to 15).

are arranged in order of their binary value, demonstrating the potential for the dipole traps to be used as a quantum register.

The number of atoms in the trap is restricted to either zero or one by the collisional blockade mechanism [61]: when pairs of rubidium atoms collide in the presence of red-detuned cooling light, they can be photo-excited to form a molecule and are lost from the trap. This explains why there is no two-atom peak in figure 6.2c. The collisional blockade mechanism depends on the interplay between R_L, the rate of atom loading from the MOT into the dipole trap; γ, the rate of one-body losses due to heating from the trapping lasers or collisions with room-temperature atoms; and β, the rate of two-body collisions between pairs of rubidium atoms within the dipole trap. This leads to the following differential equation for the number of trapped atoms:

$$\dot{N} = R_L - \gamma N - \beta N(N-1) \tag{6.3}$$

Under our experimental conditions where the trap volume is $< 10\,\mu\text{m}^3$, the two-body loss rate β is too large to be measured as the arrival of a second atom when the trap is filled causes both atoms to be lost instantaneously. However, the rates R_L and γ can be inferred from the average length of time for which an atom is trapped using the

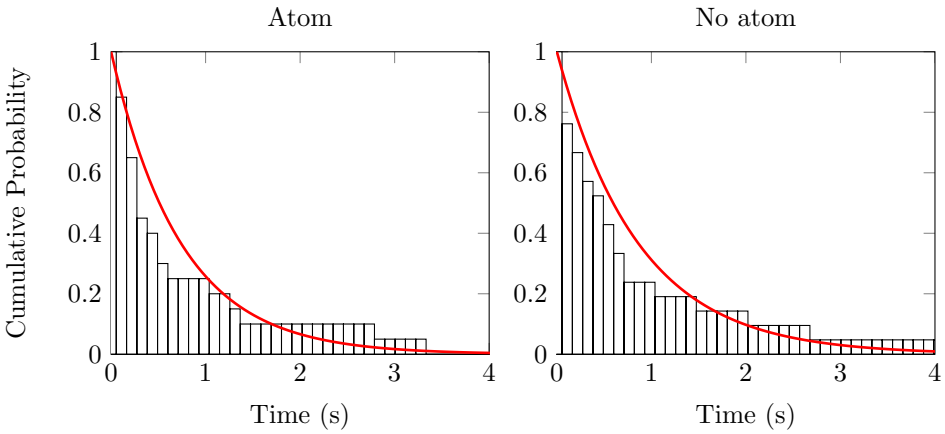

Figure 6.4: Cumulative histogram of trapping times in trap 4 (i.e. the probability of holding an atom for less than time t). The histograms agree with an exponential distribution with a mean duration of $\tau_b = 0.74 \pm 0.05\,\mathrm{s}$ for one atom and $\tau_d = 0.86 \pm 0.05\,\mathrm{s}$ for no atoms.

data in figure 6.2a. We refer to the average time for which an atom is trapped (and the trap is bright) as τ_b, and the average time for which the trap is empty (dark) as τ_d. During a dark interval, the only process that can occur is the loading of a single atom, thus, the length of the dark intervals τ_d is inversely proportional to the loading rate. On the other hand, when an atom is present, two processes can lead to atom loss: one-body collisions, and the arrival of a second atom. This can be expressed mathematically as

$$1/\tau_d = R_L, \tag{6.4}$$

$$1/\tau_b = R_L + \gamma. \tag{6.5}$$

which can be rearranged to give the one-body loss rate

$$\gamma = \frac{1}{\tau_b} - \frac{1}{\tau_d} \tag{6.6}$$

We assume that atom loading and one-body loss are statistically independent Poisson processes. In this case, we expect the times τ_b and τ_d to follow an exponential

Trap	τ_b (s)	τ_d (s)	R_L (s^{-1})	γ (s^{-1})
1	0.5 ± 0.1	1.7 ± 0.5	0.6 ± 0.2	1.4 ± 0.5
2	1.5 ± 0.6	1.2 ± 0.3	0.9 ± 0.3	-0.2 ± 0.4
3	1.2 ± 0.3	1.0 ± 0.3	1.0 ± 0.3	-0.2 ± 0.4
4	0.7 ± 0.2	0.9 ± 0.2	1.2 ± 0.3	0.2 ± 0.5

Table 6.1: Average length of bright and dark intervals, and inferred values of the loading rate R_L and one-body loss rate γ. The negative loss rates are consistent with zero within the statistical error.

probability distribution. Figure 6.4 shows an example of the cumulative histograms of the lengths of these intervals for trap 4, showing good agreement with a fitted exponential distribution. The reason that we plot the cumulative distribution (the probability $p(\tau < t)$) instead of the ordinary probability density is because it is less noisy, automatically has $p(0) = 1$ and (in the case of an exponential distribution) has the same mean.

Consider the situation where the one-body loss rate is negligible. In this case, the length of the bright and dark intervals must be equal as the same process (the arrival of a single atom) leads to both loading and loss. As one-body losses become larger, then the bright intervals will become shorter while the dark intervals will remain the same length. Table 6.1 shows the mean bright and dark times for each of the four traps. In traps 2 – 4 the length of bright and dark intervals are equal, which means that collisional blockade is the dominant process. In these traps, the one-body losses are too small to measure. However, in trap 1, the bright intervals are significantly shorter than the dark intervals, from which we infer a one-body loss rate of $\gamma = 1.4 \pm 0.5$ s^{-1}. The trap depths were not exactly uniform, and trap 1 was shallower than the other three traps. In shallower traps, the atoms are more easily heated by the MOT beams, which explains the increased loss rate for this trap.

Another important consideration is that loading rate R_L should be slow enough for the atom number to be monitored by the EMCCD camera. To this respect, the density of the MOT is very important as it determines the dynamics of both atom

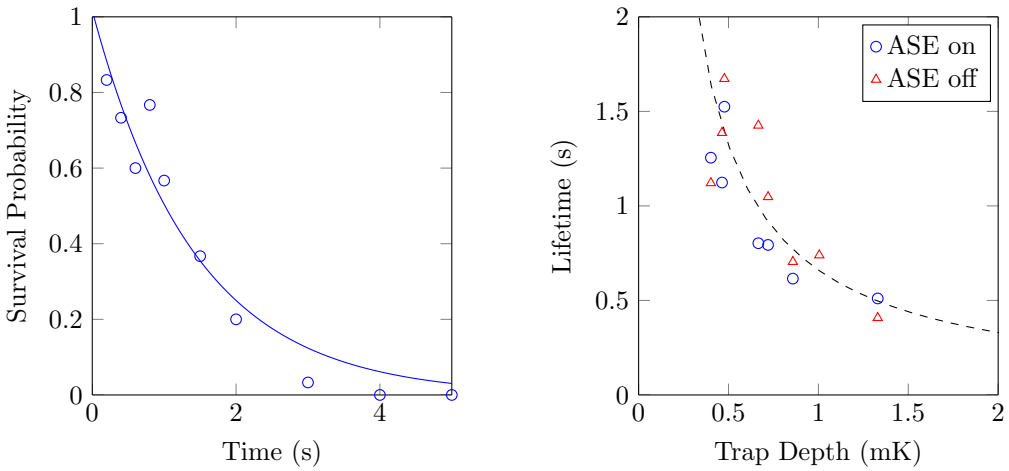

Figure 6.5: (a) The $1/e$ lifetime of a single trapped atom is 1.4 ± 0.2 s. (b) Variation in the lifetime with the trap depth. The lifetime is shorter for deeper traps because of the increased scattering rate. The lifetime is not affected by the amount of amplified spontaneous emission (see text).

loading and loss. We operate our MOT with a very low density of $< 10^6\,\mathrm{cm}^{-3}$ (as opposed to $\sim 10^{11}\,\mathrm{cm}^{-3}$ for a typical MOT [82]). From table 6.1, we infer loading rates of between 0.6 – 1.2 s^{-1}, which is much slower than our chosen frame rate of the camera of 10 Hz.

6.2.4 Lifetime of trapped atoms

To further characterise our dipole traps, we measured the lifetime of a single trapped atom. We switched off the cooling light for a variable delay and measured the survival probability the atom. After fitting an exponential decay to the data (see figure 6.5a), the maximum $1/e$ lifetime we observed was 1.4 s. The primary reason for this lifetime is heating due to off-resonant scattering from the dipole trapping light. Using equation 6.2, the scattering rate at the centre of the trap is $R_{\mathrm{sc}} = 260\,\mathrm{Hz}$, with each scattered photon increasing the kinetic energy of the atom by $E_{\mathrm{r}} = \hbar^2 k^2/2m = k_{\mathrm{B}} \times 362\,\mathrm{nK}$. From this, we estimate the lifetime as $U_{\mathrm{dip}}/E_{\mathrm{r}}/R_{\mathrm{sc}} = 5.6\,\mathrm{s}$, which is of a similar magnitude to the observed lifetime.

There are several other effects that could potentially reduce the lifetime further,

such as collisions with room temperature background gas atoms, heating due to intensity fluctuations of dipole trapping laser, and resonant scattering due to amplified spontaneous emission (ASE) of the tapered amplifier laser. The rate of background gas collisions was estimated from the MOT loading time to be ∼60 s, and thus a negligible cause of atom loss. Similarly, we measured the intensity noise spectrum of the dipole trapping laser (following [123]) and calculated that the heating rate was negligible. In addition to light at 785.4 nm, the tapered amplifier emits a broad pedestal of ASE between 770 nm and 800 nm. To isolate this effect, we measured the lifetime of a single atom both with and without a bandpass filter (Semrock MaxLine 785 nm). This filter had an attenuation of 10^{-4} for resonant light at 780 nm and 795 nm while allowing the dipole trapping light at 785 nm light to pass through. We repeated the lifetime measurements for a variety of different laser powers and trap waists, with the results shown in figure 6.5b. The lifetime is significantly shorter for deeper traps - which have a higher intensity of dipole trapping light and therefore a higher scattering rate. Furthermore, the lifetime is not affected by removing or replacing the bandpass filter, eliminating ASE as a possible source of heating.

6.2.5 Temperature of trapped atoms

We estimated the temperature of the atoms in the dipole trap by a release-and-recapture method [62, 73]. We assume that atoms are at the centre of the dipole trap, $r = 0$, and that the radial velocity, v, follows a 2D Maxwell-Boltzmann distribution, $p(v) = \frac{mv}{k_B T} \exp(-mv^2/k_B T)$. We switch off the cooling lasers and momentarily switch off the trap for a time t. During this time, the atoms move a distance of $r = vt$ away from the centre of the trap. When the trap is switched back on, the atoms will have gained potential energy. Above a certain velocity, the atoms will have gained enough energy to escape the trap. This escape velocity, v_e, can be calculated from $\frac{1}{2}mv_e^2 - U_0 \exp(-2v_e^2 t^2/w_0^2) > 0$. The probability of recapturing the atom is obtained

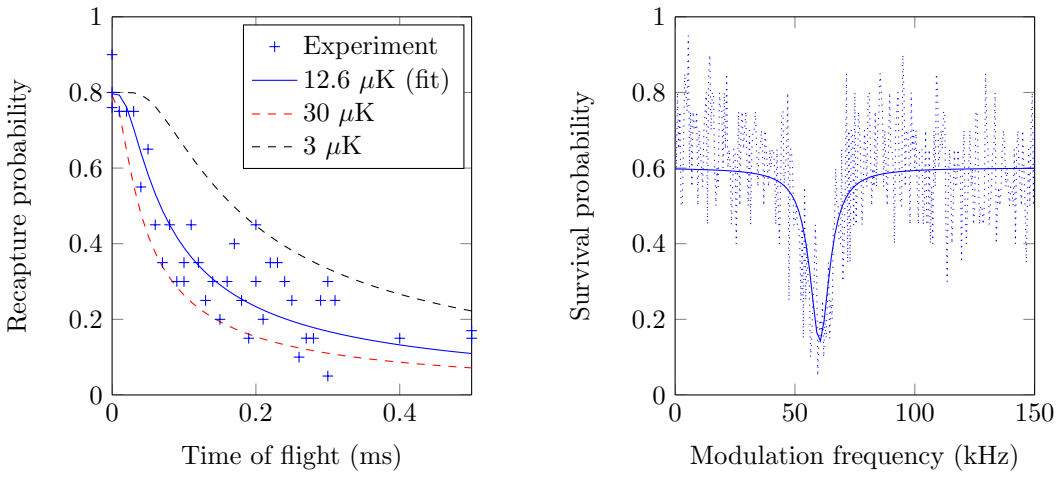

Figure 6.6: (a) Measurement of temperature by release and recapture. (b) Measurement of the trap oscillation frequencies by modulating the trap depth. There is significant loss when the trap depth is modulated at twice the radial oscillation frequency $2\omega_r = 2\pi \times 60.3$ kHz. The longitudinal oscillation frequency was too low to be observed in this experiment.

by integrating the Maxwell-Boltzmann distribution to find the fraction of atoms with $v < v_e$. By fitting this theory to the measured recapture probability (see figure 6.6a) we find a temperature of 12.6 μK. This value is on the same scale as the temperature of the MOT, and is much less than the trap depth, confirming the assumption that the atoms are cooled to the bottom of the trap.

6.2.6 Trap frequencies

Finally, we experimentally measured the trapping frequencies by modulating the trap depth. When the modulation frequency is equal to twice the oscillation frequency ω, the atoms are parametrically driven out of the trap. Using the expression

$$\omega_r = \left(\frac{4U_0}{mw_0^2}\right)^{1/2} \tag{6.7}$$

we calculated radial trapping frequency to be $\omega_r = 2\pi \times 42\,\text{kHz}$. The longitudinal trapping frequency $\omega_z = 2\pi \times 4.6\,\text{kHz}$ is expected to be much smaller due to the relatively long Rayleigh length. We modulated the intensity of the trapping laser

by 20% for 100 ms over a frequency range of $0 < f_{\mathrm{mod}} < 150$ kHz (see figure 6.6). We observed a radial trapping frequency of 30.1 kHz, significantly smaller than the expected value. This implies that the actual trap depth is lower than expected, which may be due to residual aberrations in the optical setup.

6.3 Trapping arrays of atoms

We now demonstrate that our setup can be scaled to large numbers of traps. We use the DMD to imprint a binary amplitude mask on to the trapping beam. The desired arrangement of traps (the Fourier transform of the mask) is produced in the focal plane of the aspheric lens.

As simple example, we consider a single focused trap. The complex-valued field of this trap in the plane of the DMD is given by

$$E(x, y) = A \exp(i \frac{2\pi}{f\lambda}(x_0' x + y_0' y)) \tag{6.8}$$

where f is the effective focal length of the lens (including the relay optics), x, y are the coordinates on the DMD and x_0', y_0' are the coordinates of the trap. The easiest way to convert this complex-valued field to a binary amplitude hologram is to switch all mirrors 'on' which have $\Re[E] > 0$. The hologram consists of stripes of mirrors which alternate between 'on' and 'off', forming an artificial diffraction grating. The first-order diffraction peak can be moved by changing the period and orientation of the grating, allowing for arbitrary positioning of the trap. The maximum possible displacement of the trap is 516 µm, and occurs when the period of the artificial grating is double the pitch of the micro-mirrors (analogous to the Nyquist frequency). The power in the n^{th} order of diffraction is given by $\frac{1}{2}\mathrm{sinc}^2(n\pi/2)$, which for the first order is 10.1 %.

The case for multiple traps is more complicated. Ideally, one would simply sum

the complex-valued holograms for each trap. However, the act of converting this to a binary amplitude hologram results in drastic variation of the trap depths, as well as generating unwanted ghost traps. These problems can be avoided by dithering the hologram, or by using an iterative phase retrieval algorithm such as Gerchberg-Saxton, as discussed in chapter 5. For the case of a static array of traps, we use the weighted Gerchberg-Saxton algorithm, which is slow but allows for the generation of very accurate trapping potentials.

The initial complex-valued field can be multiplied by an arbitrary phase map $\exp(i\phi(x, y))$. This allows the phase of the wavefront to be modified despite the fact that the DMD is an amplitude-only modulator. We take advantage of this feature for two purposes: firstly, we apply a quadratic phase function $\phi(x, y) = \frac{2\pi z_0'}{2f^2\lambda}(x^2 + y^2)$, which acts as an artificial Fresnel lens allowing us to adjust the longitudinal position of the traps by a distance z_0'. Secondly, we can choose $\phi(x, y)$ to correct for wavefront errors caused by aberrations in the optical setup. The largest source of aberrations is the DMD itself, whose surface varies in height by up to 10λ. This results in a spot size that is much larger than the diffraction limit. We measure these wavefront errors *in situ* by using the CCD camera to observe interference patterns between different regions of the DMD (see section 5.2.6 for details). This allows us to correct for aberrations along the entire optical path except for the vacuum window and the aspheric lens.

Examples of possible trapping arrangements of trapping sites are shown in figure 6.7. The number of trapping sites is limited by the total available laser power. From an initial power of 524 mW from the tapered amplifier, we obtain 14.4 mW available for trapping after taking into account the coupling through a single-mode fibre (64%), the reflectivity of the DMD (55%), efficiency of the holograms (10%), and transmission of the optics (76%). This limits the maximum number of traps to around 20. As in section 6.2.3, we observe occupation probabilities $\chi < 0.5$. This

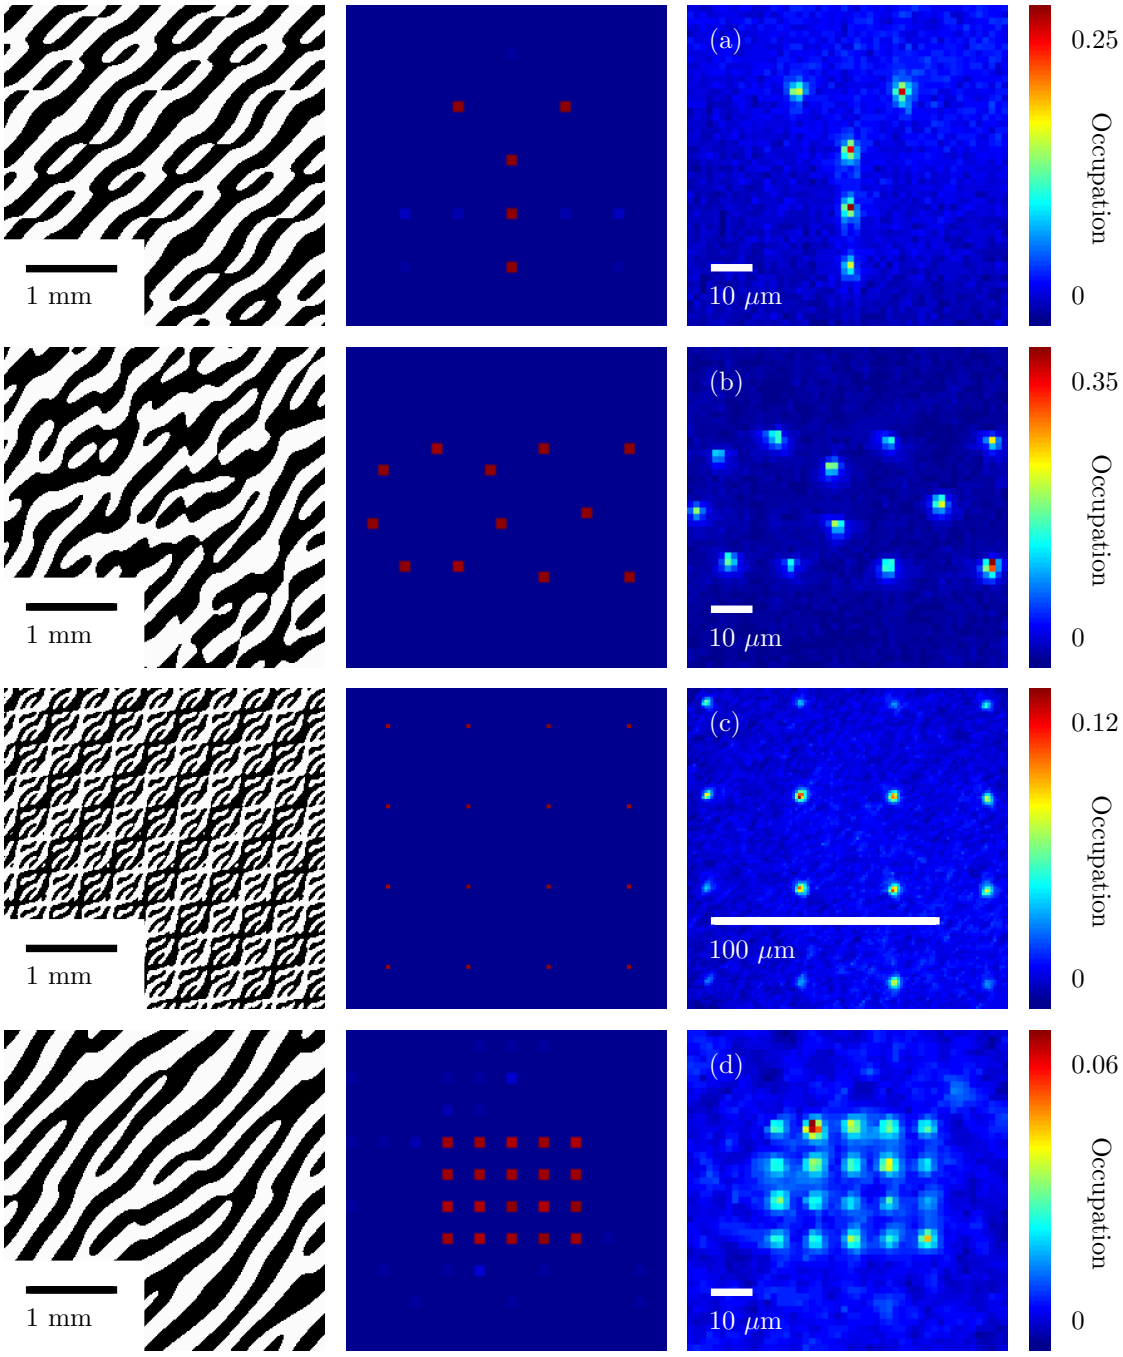

Figure 6.7: Gallery of trapped atoms with arbitrary positions. For each case, we show a portion of the binary amplitude hologram, simulated trapping potential, and fluorescence images of the trapped atoms. The images are averaged over 30 seconds, excluding frames where < 2 of the traps are filled. (a) The letter 'Y' with 5 atoms (mean occupation probability $\chi = 0.22$), (b) the letters 'OX' with 12 atoms ($\chi = 0.12$), (c) a widely-spaced grid of 16 atoms ($\chi = 0.1$), and (d) a grid of 20 atoms ($\chi = 0.044$). At any given time, each trap contains either zero or one atom. The average occupation probability is lower where the laser power is divided amongst a larger number of traps. The variation in brightness is due to differences in the trap depth from site to site.

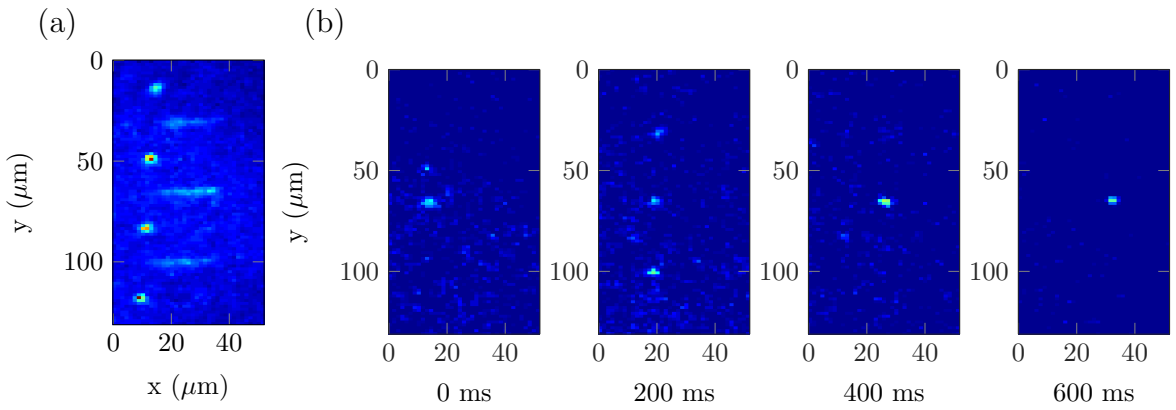

Figure 6.8: Bright transport of atoms with continuous laser cooling. A linear array of 7 traps is produced by the DMD. Odd traps are stationary while even traps are moved with a constant velocity. Atoms randomly enter and exit the trap while being transported over a distance of 14 μm. (a) Fluorescence image averaged over 30 seconds and (b) A sequence of fluorescence images showing a single atom being transported in the fourth trap.

is because the trap depth decreases as the number of traps increases, since the total power is fixed.

6.4 Transporting single atoms

We now demonstrate *bright* and *dark* transport of atoms. In our previous experiment described in chapter 3), we did this by switching from an initial circular trap to an elongated transport channel, followed by a final circular trap. This technique was well-suited to the former experiment, in which we directly imaged the DMD rather than using a holographic technique. While it is possible to generate elongated channels holographically, it is not efficient in terms of laser power. The reason for this is the reciprocal nature of Fourier transforms; making the trap larger by a factor of L in one direction reduces the extent of the hologram in that direction by $1/L$, and hence lowers the intensity of the trap by $1/L^2$.

Instead, we simply display a sequence of holograms on the DMD in which the trapping site is moved in small steps δx. This method maintains efficient, tightly-

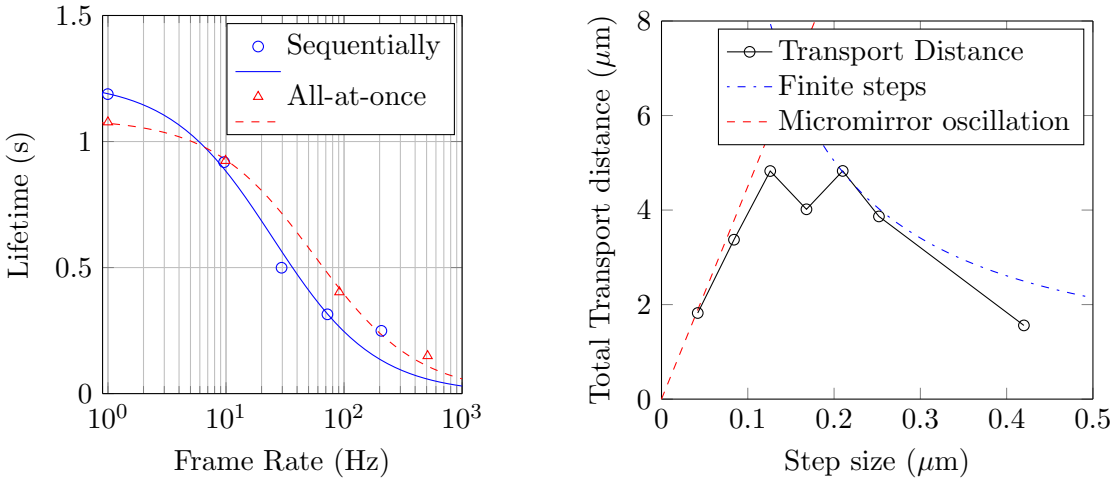

Figure 6.9: (a) Reduction in lifetime due to the DMD micromirror settling behaviour. After switching 50 full frames, the lifetime is reduced by 50%. Switching the mirrors sequentially in groups instead of all at once led to a higher heating rate. (b) Maximum transport distance versus step size. The heating rate increases quadratically with the step size, placing a limit on the maximum transport distance. For small step sizes, more frames are required for a given transport distance, and the limited is set by heating due to the mirror settling behaviour.

focussed traps, but leads to additional heating of the atoms due to the finite step size. We can estimate the heating rate as follows; assuming an atom is at rest in the centre of the trap, each step δ_x increases the energy of the atom by $E_s = m\omega_r^2 \delta x^2/2$. This heating rate can be suppressed by performing the transport in the presence of cooling light, i.e. bright transport. Figure 6.8 illustrates the transport of 7 single atoms over a distance of 14 μm using this technique. We demonstrate the individual addressing of the atoms by holding some atoms stationary while moving others.

When the micro-mirrors are switched between the 'on' and 'off' states, they have a finite settling time of 12 μs. During this time, the mirror angle oscillates about its final position with an amplitude of several degrees (see figure 4.19). The resulting angular displacement of the trapping beam means that it no longer passes through the relay optics, momentarily reducing the intensity of the trapping light. This leads to additional heating in proportion to the number of frames. The mirror oscillation occurs regardless of the previous state of the mirror. This source of heating is purely

technical, and could be overcome by using lenses with larger diameters in the relay optics to capture the rays at large angles, or by mechanically damping the DMD micro-mirrors. We measured the severity of this heating in our current setup by observing the reduction in lifetime for different frame rates of the DMD (see figure 6.9a). At a frame rate of 50 Hz, there is an equal contribution from both the off-resonant scattering and mirror vibration, and the lifetime is reduced by 50%. We attempted to reduce this source of heating by switching the mirrors sequentially in small groups, rather than all at once. Unfortunately this increased the heating rate, presumably because the increased number of switching events outweighs the benefit of switching a smaller number of mirrors.

Finally, we demonstrate dark transport of a single atom. The duration of the transport was held constant at 0.5 s, and the frame rate set to 50 Hz. Under these conditions, there is an equal contribution to the heating rate from off-resonant scattering and micro-mirror oscillation. Next we measured the transport distance as a function of step size δx. Figure 6.9b shows the total transport distance, which we define as the distance at which the survival probability falls to $1/e$. For small step sizes, the transport distance is limited by the small number of frames, which is in turn limited by the heating rate due to micro-mirror oscillation. For large step sizes, the distance is limited by the finite step size. The measured transport distance agrees well with the simple models for the two heating effects. The optimum step size of 150 nm at which the three sources of heating contribute equally, yielding a maximum transport distance of 5 µm. If the micro-mirror oscillation were eliminated, then it would be possible to perform the transport with smaller steps and the maximum transport distance would increase dramatically.

6.5 Prospects for quantum information processing

Having demonstrated the trapping and transport of single neutral atoms, we now consider the feasibility of using them as qubits in a quantum information processor. For reasons that will be discussed, we propose using the $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$ ground state sublevels as the qubit states $|0\rangle$ and $|1\rangle$. A commonly accepted list of requirements for a quantum computer are the DiVincenzo criteria [8]. We consider each of these criteria as they apply to our setup.

6.5.1 Scalability of qubits

The main advantage of our scheme compared to other approaches is scalability. While demonstrations of quantum gates have been made with one or two atoms [71, 9, 26, 27], we are able to trap up to 20 atomic qubits. This is only limited by the power of our trapping laser. Other groups have also made impressive progress towards trapping arrays of atoms [64, 75, 66]. A unique feature of our apparatus is that each atom can be moved independently, which allows for interactions between any pair of qubits, rather than only nearest neighbours. This feature could also be used to correct the unsolved problem of atom losses by replenishing empty traps from an atom reservoir.

6.5.2 Preparation fidelity

The initial qubit states can be prepared by optical pumping. For example, a laser resonant with the $F = 1 \rightarrow F' = 2$ transition will optically pump atoms to the $F = 2$ ground states. An additional π-polarised laser beam resonant with the $F = 2 \rightarrow F' = 2$ transition will depopulate atoms from all states except $|F = 2, m_F = 0\rangle$ (as the $m_F = 0 \rightarrow m'_F = 0$ transition is forbidden by selection rules). Over time, the atom accumulate in this state (the $|1\rangle$ qubit state). The preparation fidelity is limited because the same laser can off-resonantly drive $F = 2 \rightarrow F' = 1, 3$ transitions which

depopulate the $|1\rangle$ state. The probability of these unwanted transitions relative to the desired ones is given by

$$p \approx \frac{4\Gamma^2}{\Delta_{hf}^2}$$

where Δ_{hf} is the hyperfine splitting between the excited states in question. These splittings are 267 MHz and 157 MHz, which gives a preparation fidelity of $1-p \approx 0.99$ for the $|1\rangle$ state. The corresponding fidelity for the $|0\rangle$ state is 0.96. We assume that the polarisation purity is sufficient so as not to drive any unwanted transitions.

6.5.3 Coherence time

The coherence time of the $|0\rangle$ and $|1\rangle$ states is very important as it limits the fidelity of quantum operations. Usually, the major source of decoherence for atomic qubits are magnetic fields, which shift the energy of the states due to the Zeeman effect. By working with $m_F = 0$ states, the magnetic field dependence is limited to second-order Zeeman shifts, which give a coherence time $\tau > 1000\,\mathrm{s}$ for $B < 1\,\mathrm{mG}$ [124]. In our setup coherence time is limited by two main effects: off-resonant scattering from the dipole trapping laser, and differential Stark shifts of the $|0\rangle$ and $|1\rangle$ due to the thermal motion of the atoms in the traps.

As stated in section 6.2.4, the rate of off-resonant scattering from the dipole trapping laser is 260 Hz, giving a coherence time of 4 ms. Since $R_{sc} \propto 1/\Delta^2$, this could be improved by using a laser with a longer wavelength, for example 1064 nm. This would give a scattering rate of 4 Hz, or an average time between scattering events of 0.25 s. Furthermore, as this wavelength is far-detuned relative to the splitting between the resonances, the rate of inelastic scattering to different m_F states is suppressed even further due to quantum interference effects [125]. Coherence times of more than 4 s have been demonstrated experimentally at this wavelength[126].

As the atoms move within the trap, they experience a varying light shift due to the

spatial variation of the intensity of the trapping laser. Although we assume that the two ground states are shifted equally by the trapping laser, in reality, there is a small differential shift [124]. For the $F = 1$ and $F = 2$ states, this differential Stark shift is $\delta = 4 \times 10^{-5}$. The extent to which atoms move is determined by their temperature, which limits the maximum height that they can climb within the trap potential. Since the potential energy of the atoms in the trap is proportional to the light shift, our estimated temperature of 12.5 μK leads to a variable light shift of $k_B T/h = 260\,\text{kHz}$. Multiplying this by δ gives a differential shift in the qubit frequency of 11 Hz, and hence a coherence time of 90 ms.

The differential Stark shift may be compensated by introducing an additional laser beam with the same spatial profile as the dipole trap, but with a frequency tuned halfway between the hyperfine levels. Kaplan et al. observed a 50-fold increase in coherence time using this method in a dipole trap with a 50 μm waist [127]. Unfortunately, their method of mode matching the spatial profiles – by coupling both beams through a single-mode fibre – is incompatible with holographic positioning, as the two beams have different wavelengths and would be diffracted at different angles by the DMD. In any case, matching the spatial profile of diffraction-limited dipole traps at two different wavelengths would be extremely challenging due to optical aberrations. Alternatively, one could eliminate the atomic motion (and hence the differential Stark shift) by cooling the atoms to the vibrational ground state of the traps.

6.5.4 Quantum gates

Single qubit rotations of neutral atoms in microscopic dipole traps have been demonstrated previously [9, 75, 128]. In these experiments, Raman transitions between the $|0\rangle$ and $|1\rangle$ states were driven by a pair of phase-locked lasers. For example, Jones et al. [9] achieved Rabi frequencies of $\Omega = 2\pi \times 6.7\,\text{MHz}$ and coherence times of $\tau = 34\,\text{ms}$, giving an error per gate of better than 10^{-5}. Furthermore, Lengwenus et

al. [75] showed that coherence is preserved even when the dipole traps are moved. Given the shorter coherence time in our experiment it will be difficult to observe such gates without changing to a far-detuned trapping laser.

The most promising two-qubit gates for neutral atoms are based on the Rydberg blockade [26, 27]. The strength of the Rydberg interaction between neighbouring atoms depends on their separation. Zhang et al. [129] demonstrated a cNOT gate with a fidelity of 0.71 using a trap separation of 8.7 μm and a waist of 3.2 μm. Fortunately, the gate mechanism is insensitive to fluctuations in the position of the atoms as long as the separation is small-enough to ensure a strong Rydberg blockade. We have demonstrated atom separations ranging between 10 μm and 100 μm (see figure 6.7) and step sizes as low as 40 nm (figure 6.9). In principle we could achieve a minimum separation of $2w_0 = 3.4\,\mu\text{m}$, which is well-within the parameters of other demonstrations.

The largest sources of error in the experimentally-demonstrated Rydberg gates are Doppler broadening of the transition to the Rydberg state and spontaneous emission from the Rydberg state. These parameters depend on the intensity of the Rydberg excitation laser at 480 nm and the temperature of the atoms [129]. Our estimated temperature of 12.5 μK is substantially lower than other demonstrations (140 μK), which should allow us to achieve comparable or better fidelities. The intrinsic error rate for Rydberg gates has been estimated to be $E < 10^{-3}$ [77]. However, the largest error in the demonstrated gates is atom loss during state detection, which is discussed below.

6.5.5 Qubit readout

The state of the qubit can be read by using resonance fluorescence. Following the method of Kruse et al. [9, 128] the atoms in the $|F = 2\rangle$ state are first heated out of the traps by a resonant laser on the $F = 2 \rightarrow F' = 3$ transition. Next, the

remaining atoms in the $|F = 1\rangle$ state are detected by illuminating them with cooling lasers and collecting the fluorescence photons using the same technique as described in section 6.2.3. The presence of an atom is signalled when the number of photons is above a fixed threshold. The fluorescence signal is contaminated by background counts due to scattering from the laser beams and detector dark counts. Accurate assignment of the qubit state is possible when the number of photons due to a single atom, N_a, exceeds the chosen threshold by more than the shot noise of the total counts, i.e.

$$N_a > 2\sqrt{N_a + N_b} \tag{6.9}$$

where N_b is the number of background counts. The factor of 2 arises from setting the threshold to half the expected number of photons. For example, using our data presented in figure 6.2, $N_a = 120$ and $N_b = 80$ and the total shot noise is 14 (the frame rate of the EMCCD camera is 10 Hz). Theoretically, this should allow the discrimination of a single atom at the 4σ confidence level, i.e. with an error of 10^{-4}.

In practice, this qubit measurement scheme is corrupted by atom loss processes – since an atom in the $|1\rangle$ state is indistinguishable from no atom at all. In recent demonstrations of two-qubit Rydberg gates [26, 27], the probability of losing one of the atoms is 0.19 – making it the by far the largest source of error. Therefore, care must be taken to eliminate sources of atom loss, and to make the traps as robust as possible.

6.5.6 Coupling via optical cavities

Finally, our tweezers could be used to position an atom in the mode of an optical cavity. The positioning accuracy is limited by the width of the thermal distribution of the atoms

$$\sigma_x = \left(\frac{k_B T}{m\omega_r^2}\right)^{1/2} \tag{6.10}$$

which for our experimental parameters is $\sigma_x = 180\,\text{nm}$. This is less than the spacing between the antinodes of the cavity mode ($\lambda/2 = 390\,\text{nm}$), and should allow atoms to be localised at a single antinode (where the cavity coupling rate g is maximised). Better localisation could be achieved by cooling the atom to the vibrational ground state of the trap [130, 63].

There is one final obstacle to combining our tweezers with an optical cavity. The large divergence angle of the trapping beam requires significant optical access from the side of the cavity to avoid clipping the edges of the beam. This means that the diameter of the cavity mirrors must be on the order of the cavity length ($\sim 100\,\mu\text{m}$). The fibre-tip cavities [31, 44] that we are developing in our group (see figure 1.2) have mirror diameters of $100\,\mu\text{m}$ and are ideally suited for this purpose.

6.5.7 Concluding remarks

In conclusion, our optical tweezers setup is suitable in principle for a scalable implementation of quantum gates. Our results for trap size, depth, separation, and temperature are comparable to other demonstrations of one- and two-qubit gates. The predicted error rates for such gates are promising, however, the fidelity of state preparation and measurement is worse than has been achieved with trapped ions [25]. The most urgent area for improvement is the trapping laser, as currently this will limit the coherence time of the atoms to 4 ms. Others have shown that the coherence of internal atomic state is unaffected by transporting the traps. Related to this is the issue of atom loss. The fundamental lifetime due to background gas collisions is $> 30\,\text{s}$ in our experiment, however, the lifetime due to heating from the trapping laser is only 1.4 s, which needs to be improved.

Chapter 7

Conclusion

In conclusion, we have demonstrated the feasibility of cooling, trapping and transporting individual atoms using a digital mirror device. A wide variety of trapping geometries are possible. The lifetime of the trapped atoms is limited by scattering from the dipole trapping laser due to its proximity to resonance. The transport distance is limited by vibration of the micro-mirrors of the DMD.

In this thesis we have overcome the two main technical issues with using a DMD, namely, rapidly calculating the holograms and correcting for optical aberrations. Our techniques enable a DMD to be used as a replacement for a liquid crystal SLM to, for example, perform dipole trapping of other particles such as cells [102] and nanoscale tools [103]. Furthermore, the fast switching speed of the DMD may be used to perform feedback on the dipole traps in response to measurements in real-time.

There has been significant progress in the use of optical tweezers for trapping cold atoms [66]. To this end, we anticipate the application of our tweezers to extend the demonstrations of Rydberg gates [27, 26], matter wave interference [76], and cavity coupling [30] to large arrays of atoms. The tweezers would allow neighbouring trapping sites to be merged and separated, allowing one to create extended many-body quantum states. They could also be used to position multiple atoms in the mode

of an optical cavity, where their quantum state can be entangled through engineered dissipation [92].

Further improvements to our results are possible with some simple modifications to the apparatus. Firstly, large-diameter relay lenses would reduce the intensity modulation due to the micro-mirror oscillation. Alternatively, we could use a different type of spatial light modulator that doesn't suffer from these problems. A more powerful dipole trapping laser with a larger detuning, for example a fibre laser at 1064 nm, would reduce off-resonant scattering and increase the lifetime. In addition, a more controlled method of loading atoms in to the dipole traps, such as a 2D-MOT, would be enormously beneficial. We are currently designing the next generation trapping apparatus which will overcome these technical issues. In the meantime, combining these traps with optical cavities to allow the precise control of atoms close to the dielectric mirrors is of particular interest to our group. However, our approach has wide-ranging applications in the atomic physics community wherever single atoms need to be manipulated accurately.

List of publications

- D. Stuart and A. Kuhn. Moving single atoms, *(in preparation)*
- D. Stuart, O. Barter and A. Kuhn. Fast algorithms for generating binary holograms, *ArXiv e-prints*, 1409.1841, http://arxiv.org/abs/1409.1841, *(submitted 15 August 2014)*
- C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsworth, A. Kuhn, Control and manipulation of cold atoms in optical tweezers, *New J. Phys.*, **14** 073051, http://arxiv.org/abs/1109.0657, July 2012

Bibliography

[1] C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsworth, and A. Kuhn, "Control and manipulation of cold atoms in optical tweezers," New J. Phys. **14**, 073051 (2012).

[2] R. Feynman, "Simulating physics with computers," Int. J. Theor. Phys. **21**, 467–488 (1982).

[3] L. K. Grover, "From Schrödinger's equation to the quantum search algorithm," Am. J. Phys. **69**, 769–777 (2001).

[4] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, SFCS '94 (IEEE Computer Society, Washington, DC, USA, 1994) pp. 124–134.

[5] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," Commun. ACM **21**, 120–126 (1978).

[6] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information (Cambridge Series on Information and the Natural Sciences)*, 1st ed. (Cambridge University Press, 2004).

[7] A. M. Steane, "Error correcting codes in quantum theory," Phys. Rev. Lett. **77**, 793–797 (1996).

[8] D. P. DiVincenzo, "The Physical Implementation of Quantum Computation," Fortschritte der Physik **48**, 771–783 (2000), quant-ph/0002077 .

[9] M. P. A. Jones, J. Beugnon, A. Gaëtan, J. Zhang, G. Messin, A. Browaeys, and P. Grangier, "Fast quantum state control of a single trapped neutral atom," Phys. Rev. A **75**, 040301 (2007).

[10] E. Knill, R. Laflamme, and G. J. Milburn, "A scheme for efficient quantum computation with linear optics," Nature **409**, 46–52 (2001).

[11] C. J. Ballance, T. P. Harty, N. M. Linke, and D. M. Lucas, "High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions," (2014), arXiv:1406.5473 [quant-ph] .

[12] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, "Qubit architecture with high coherence and fast tunable coupling," (2014), arXiv:1402.7367 [cond-mat.supr-con] .

[13] J. Chan, T. P. M. Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Groblacher, M. Aspelmeyer, and O. Painter, "Laser cooling of a nanomechanical oscillator into its quantum ground state," Nature **478**, 89–92 (2011).

[14] J. P. Reithmaier, G. Sek, A. Loffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel, "Strong coupling in a single quantum dot-semiconductor microcavity system," Nature **432**, 197–200 (2004).

[15] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish,

M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, "14-qubit entanglement: creation and coherence," Phys. Rev. Lett. **106**, 130506 (2011).

[16] J. I. Cirac and P. Zoller, "Quantum computations with cold trapped ions," Phys. Rev. Lett. **74**, 4091–4094 (1995).

[17] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller, "Entanglement of Atoms via Cold Controlled Collisions," Phys. Rev. Lett. **82**, 1975–1978 (1999), quant-ph/9810087 .

[18] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, "Fast quantum gates for neutral atoms," Phys. Rev. Lett. **85**, 2208–2211 (2000).

[19] K. Mølmer and A. Sørensen, "Multiparticle entanglement of hot trapped ions," Phys. Rev. Lett. **82**, 1835–1838 (1999).

[20] T. Pellizzari, S. A. Gardiner, J. I. Cirac, and P. Zoller, "Decoherence, continuous observation, and quantum computing: A cavity qed model," Phys. Rev. Lett. **75**, 3788–3791 (1995).

[21] A. G. Fowler, "Proof of finite surface code threshold for matching," Phys. Rev. Lett. **109**, 180502 (2012).

[22] A. G. Fowler and J. M. Martinis, "Quantifying the effects of local many-qubit errors and nonlocal two-qubit errors on the surface code," Phys. Rev. A **89**, 032316 (2014).

[23] S. D. Barrett and T. M. Stace, "Fault tolerant quantum computation with very high threshold for loss errors," Phys. Rev. Lett. **105**, 200502 (2010).

[24] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell,

M. B. Ketchen, and M. Steffen, "Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms," Phys. Rev. B **86**, 100506 (2012).

[25] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, "High-fidelity preparation, gates, memory and readout of a trapped-ion quantum bit," (2014), arXiv:1403.1524 [quant-ph] .

[26] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, "Demonstration of a neutral atom controlled-not quantum gate," Phys. Rev. Lett. **104**, 010503 (2010).

[27] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, "Entanglement of two individual neutral atoms using Rydberg blockade," Phys. Rev. Lett. **104**, 010502 (2010).

[28] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, "Single-atom-resolved fluorescence imaging of an atomic Mott insulator," Nature **467**, 68–72 (2010).

[29] S. Ritter, C. Nolleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mucke, E. Figueroa, J. Bochmann, and G. Rempe, "An elementary quantum network of single atoms in optical cavities," Nature **484**, 195–200 (2012).

[30] J. D. Thompson, T. G. Tiecke, N. P. de Leon, J. Feist, A. V. Akimov, M. Gullans, A. S. Zibrov, V. Vuletić, and M. D. Lukin, "Coupling a single trapped atom to a nanoscale optical cavity," Science **340**, 1202–1205 (2013).

[31] D. Hunger, T. Steinmetz, Y. Colombe, C. Deutsch, T. W. Hänsch, and J. Reichel, "A fiber Fabry-Perot cavity with high finesse," New J. Phys. **12**, 065038 (2010).

[32] J. M. Raimond, M. Brune, and S. Haroche, "Manipulating quantum entanglement with atoms and photons in a cavity," Rev. Mod. Phys. **73**, 565–582 (2001).

[33] A. D. O'Connell, M. Hofheinz, M. Ansmann, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, D. Sank, H. Wang, M. Weides, J. Wenner, J. M. Martinis, and A. N. Cleland, "Quantum ground state and single-phonon control of a mechanical resonator," Nature **464**, 697–703 (2010).

[34] H. J. Kimble, "The quantum internet," Nature **453**, 1023–1030 (2008).

[35] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, "Quantum state transfer and entanglement distribution among distant nodes in a quantum network," Phys. Rev. Lett. **78**, 3221–3224 (1997).

[36] L. M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," Nature **414**, 413–418 (2001).

[37] N. H. Nickerson, Y. Li, and S. C. Benjamin, "Topological quantum computing with a very noisy network and local error rates approaching one percent," Nat. Commun. **4**, 1756 (2013).

[38] R. Raussendorf and H. J. Briegel, "A one-way quantum computer," Phys. Rev. Lett. **86**, 5188–5191 (2001).

[39] M. V. Balabas, T. Karaulanov, M. P. Ledbetter, and D. Budker, "Polarized alkali-metal vapor with minute-long transverse spin-relaxation time," Phys. Rev. Lett. **105**, 070801 (2010).

[40] C. K. Law and H. J. Kimble, "Deterministic generation of a bit-stream of single-photon pulses," J. Mod. Opt. **44**, 2067–2074 (1997).

[41] A. Kuhn, M. Hennrich, and G. Rempe, "Deterministic single-photon source for distributed quantum networking," Phys. Rev. Lett. **89**, 067901 (2002).

[42] P. B. R. Nisbet-Jones, J. Dilley, D. Ljunggren, and A. Kuhn, "Highly efficient source for indistinguishable single photons of controlled shape," New J. Phys. **13**, 103036 (2011).

[43] J. Dilley, P. Nisbet-Jones, B. W. Shore, and A. Kuhn, "Single-photon absorption in coupled atom-cavity systems," Phys. Rev. A **85**, 023834 (2012).

[44] Jian Dong, *Laser Trapping of Atoms and Cavity Quantum Electrodynamics In Fibre-Tip Microcavities*, Ph.D. thesis, University of Oxford (2014).

[45] Calum E. MacAulay and Andrew L. P. Dlugan, "Use of digital micromirror devices in quantitative microscopy," Proc. SPIE **3260**, 201–206 (1998).

[46] Hanley, Verveer, Gemkow, Arndt-Jovin, and Jovin, "An optical sectioning programmable array microscope implemented with a digital micromirror device," Journal of Microscopy **196**, 317–331 (1999).

[47] J. R. Hui, *Optical Tweezers Using the Texas Instruments' Digital Micromirror Device*, Master's thesis, Massachusetts Institute of Technology (2001).

[48] W. Neuhauser, M. Hohenstatt, P. E. Toschek, and H. Dehmelt, "Localized visible Ba^+ mono-ion oscillator," Phys. Rev. A **22**, 1137–1140 (1980).

[49] Z. Hu and H. J. Kimble, "Observation of a single atom in a magneto-optical trap," Opt. Lett. **19**, 1888–1890 (1994).

[50] D. Haubrich, H. Schadwinkel, F. Strauch, B. Ueberholz, R. Wynands, and D. Meschede, "Observation of individual neutral atoms in magnetic and magneto-optical traps," Europhys. Lett. **34**, 663 (1996).

[51] M. Steiner, H. M. Meyer, C. Deutsch, J. Reichel, and M. Köhl, "Single ion coupled to an optical fiber cavity," Phys. Rev. Lett. **110**, 043003 (2013).

[52] G. R. Guthohrlein, M. Keller, K. Hayasaka, W. Lange, and H. Walther, "A single ion as a nanoscopic probe of an optical field," Nature **414**, 49–51 (2001).

[53] A. Ashkin, "Acceleration and trapping of particles by radiation pressure," Phys. Rev. Lett. **24**, 156–159 (1970).

[54] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm, and Steven Chu, "Observation of a single-beam gradient force optical trap for dielectric particles," Opt. Lett. **11**, 288–290 (1986).

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

[56] Steven Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable, "Experimental observation of optically trapped atoms," Phys. Rev. Lett. **57**, 314–317 (1986).

[57] S. Chu, "Laser manipulation of atoms and particles," Science **253**, 861–866 (1991).

[58] V I Balykin, V G Minogin, and V S Letokhov, "Electromagnetic trapping of cold atoms," Rep. Prog. Phys. **63**, 1429 (2000).

[59] J. Dalibard and C. Cohen-Tannoudji, "Dressed-atom approach to atomic motion in laser light: the dipole force revisited," J. Opt. Soc. Am. B **2**, 1707–1720 (1985).

[60] S. J. M. Kuppens, K. L. Corwin, K. W. Miller, T. E. Chupp, and C. E. Wieman, "Loading an optical dipole trap," Phys. Rev. A **62**, 013406 (2000).

[61] N. Schlosser, G. Reymond, and P. Grangier, "Collisional blockade in microscopic optical dipole traps," Phys. Rev. Lett. **89**, 023005 (2002).

[62] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, "Energy distribution and cooling of a single atom in an optical tweezer," Phys. Rev. A **78**, 033425 (2008).

[63] A. M. Kaufman, B. J. Lester, and C. A. Regal, "Cooling a single atom in an optical tweezer to its quantum ground state," Phys. Rev. X **2**, 041014 (2012).

[64] K. D. Nelson, X. Li, and D. S. Weiss, "Imaging single atoms in a three-dimensional array," Nat. Phys. **3**, 556–560 (2007).

[65] S. Bergamini, B. Darquié, M. Jones, L. Jacubowiez, A. Browaeys, and P. Grangier, "Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator," J. Opt. Soc. Am. B **21**, 1889–1894 (2004).

[66] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, "Single-atom trapping in holographic 2d arrays of microtraps with arbitrary geometries," Phys. Rev. X **4**, 021034 (2014).

[67] Philip P. J. Zupancic, *Dynamic Holography and Beamshaping using Digital Micromirror Devices*, Master's thesis, Ludwig-Maximilians-Universitat (2013).

[68] R. Dumke, M. Volk, T. Müther, F. B. J. Buchkremer, G. Birkl, and W. Ertmer, "Micro-optical realization of arrays of selectively addressable dipole traps: A scalable configuration for quantum computation with atomic qubits," Phys. Rev. Lett. **89**, 097903 (2002).

[69] G. Li, S. Zhang, L. Isenhower, K. Maller, and M. Saffman, "Crossed vortex bottle beam trap for single-atom qubits," Opt. Lett. **37**, 851–853 (2012).

[70] J. Arlt and M. J. Padgett, "Generation of a beam with a dark focus surrounded by regions of higher intensity: the optical bottle beam," Opt. Lett. **25**, 191–193 (2000).

[71] J. Beugnon, C. Tuchendler, H. Marion, A. Gaetan, Y. Miroshnychenko, Y. R. P. Sortais, A. M. Lance, M. P. A. Jones, G. Messin, A. Browaeys, and P. Grangier, "Two-dimensional transport and transfer of a single atomic qubit in optical tweezers," Nat. Phys. **3**, 696–699 (2007).

[72] S. Kuhr, W. Alt, D. Schrader, M. Müller, V. Gomer, and D. Meschede, "Deterministic delivery of a single atom," Science **293**, 278–280 (2001).

[73] W. Alt, S. Schrader, S. Kuhr, M. Müller, V. Gomer, and D. Meschede, "Single atoms in a standing-wave dipole trap," Phys. Rev. A **67**, 033403 (2003).

[74] Y. Miroshnychenko, D. Schrader, S. Kuhr, W. Alt, I Dotsenko, M. Khudaverdyan, A. Rauschenbeutel, and D. Meschede, "Continued imaging of the transport of a single neutral atom," Opt. Express **11**, 3498–3502 (2003).

[75] A. Lengwenus, J. Kruse, M. Schlosser, S. Tichelmann, and G. Birkl, "Coherent transport of atomic quantum states in a scalable shift register," Phys. Rev. Lett. **105**, 170502 (2010).

[76] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, "Two-particle quantum interference in tunnel-coupled optical tweezers," Science **345**, 306–309 (2014).

[77] M. Saffman, T. G. Walker, and K. Mølmer, "Quantum information with Rydberg atoms," Rev. Mod. Phys. **82**, 2313–2363 (2010).

[78] M. M. Müller, M. Murphy, S. Montangero, T. Calarco, P. Grangier, and A. Browaeys, "Implementation of an experimentally feasible controlled-phase gate on two blockaded Rydberg atoms," Phys. Rev. A **89**, 032334 (2014).

[79] L. Brandt, C. Muldoon, T. Thiele, J. Dong, E. Brainis, and A. Kuhn, "Spatial

light modulators for the manipulation of individual atoms," App. Phys. B **102**, 443–450 (2011).

[80] Daniel Steck, "Rubidium 87 D line data," available online at http://steck.us/alkalidata (2010).

[81] C. J. Foot, *Atomic physics*, Oxford master series in physics (Oxford University Press, 2005).

[82] H.J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Graduate Texts in Contemporary Physics (Springer New York, 1999).

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

[84] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen, "Near-deterministic preparation of a single atom in an optical microtrap," Nat. Phys. **6**, 951–954 (2010).

[85] L. Novotny and B. Hecht, *Principles of Nano-Optics* (Cambridge University Press, 2012).

[86] L. Allen, V. E. Lembessis, and M. Babiker, "Spin-orbit coupling in free-space Laguerre-Gaussian light beams," Phys. Rev. A **53**, R2937–R2939 (1996).

[87] M Zeppenfeld, "Solutions to maxwell's equations using spheroidal coordinates," New J. Phys. **11**, 073007 (2009).

[88] A. Holleczek, A. Aiello, C. Gabriel, C. Marquardt, and G. Leuchs, "Classical and quantum properties of cylindrically polarized states of light," Opt. Express **19**, 9714–9736 (2011).

[89] A. V. Volyar and T. A. Fadeeva, "Generation of singular beams in uniaxial crystals," Optics and Spectroscopy **94**, 235–244 (2003).

[90] Jan Petersen, Jürgen Volz, and Arno Rauschenbeutel, "Chiral nanophotonic waveguide interface based on spin-orbit interaction of light," Science **346**, 67–71 (2014), http://www.sciencemag.org/content/346/6205/67.full.pdf .

[91] S. Wildermuth, P. Krüger, C. Becker, M. Brajdic, S. Haupt, A. Kasper, R. Folman, and J. Schmiedmayer, "Optimized magneto-optical trap for experiments with ultracold atoms near surfaces," Phys. Rev. A **69**, 030901 (2004).

[92] C. Marr, A. Beige, and G. Rempe, "Entangled-state preparation via dissipation-assisted adiabatic passages," Phys. Rev. A **68**, 033817 (2003).

[93] E. Brainis, C. Muldoon, L. Brandt, and A. Kuhn, "Coherent imaging of extended objects," Optics Communications **282**, 465 – 472 (2009).

[94] E. D. Black, "An introduction to Pound–Drever–Hall laser frequency stabilization," Am. J. Phys. **69**, 79–87 (2001).

[95] B. P. Anderson and M. A. Kasevich, "Loading a vapor-cell magneto-optic trap using light-induced atom desorption," Phys. Rev. A **63**, 023404 (2001).

[96] *Using Lasers with DLP DMD technology*, Tech. Rep. TI DN 2509927 (Texas Instruments, 2008).

[97] V. Boyer, C. M.M Chandrashekar, C. J. Foot, and Z. J. Laczik, "Dynamic optical trap generation using FLC SLMs for the manipulation of cold atoms," J. Mod. Opt. **51**, 2235–2240 (2004).

[98] R. W. Bowman, G. M. Gibson, A. Linnenberger, D. B. Phillips, J. A. Grieve, D. M. Carberry, S. Serati, M. J. Miles, and M. J. Padgett, "Red tweezers:

Fast, customisable hologram generation for optical tweezers," Comput. Phys. Commun. **185**, 268 – 273 (2014).

[99] James H. Tutt, Andrew D. Holland, David J. Hall, Richard D. Harriss, and Neil J. Murray, "The noise performance of electron-multiplying charge-coupled devices at x-ray energies," IEEE Transactions on Electron Devices **59**, 167–175 (2012).

[100] D. Stuart, O. Barter, and A. Kuhn, "Fast algorithms for generating binary holograms," (2014), arXiv:1409.1841 [physics.optics] .

[101] A. Ashkin, "History of optical trapping and manipulation of small-neutral particle, atoms, and molecules," IEEE J. Sel. Topics Quantum Electron. , 841–856 (2000).

[102] M. Dienerowitz, L. V. Cowan, G. M. Gibson, R. Hay, M. J. Padgett, and V. R. Phoenix, "Optically trapped bacteria pairs reveal discrete motile response to control aggregation upon cell-cell approach," Curr. Microbiol. (2014).

[103] G. M. Gibson, R. W. Bowman, A. Linnenberger, M. Dienerowitz, D. B. Phillips, D. M. Carberry, M. J. Miles, and M. J. Padgett, "A compact holographic optical tweezers instrument," Rev. Sci. Instrum. **83**, 113107 (2012).

[104] W. S. Bakr, J. I. Gillen, A. Peng, S. Folling, and M. Greiner, "A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice," Nature **462**, 74–77 (2009).

[105] M. J. Piotrowicz, M. Lichtman, K. Maller, G. Li, S. Zhang, L. Isenhower, and M. Saffman, "Two-dimensional lattice of blue-detuned atom traps using a projected gaussian beam array," Phys. Rev. A **88**, 013420 (2013).

[106] K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier, "Experimental demonstration of painting arbitrary and dynamic potentials for Bose–Einstein condensates," New J. Phys. **11**, 043030 (2009).

[107] D. Trypogeorgos, T. Harte, A. Bonnin, and C. Foot, "Precise shaping of laser light by an acousto-optic deflector," Opt. Express **21**, 24837–24846 (2013).

[108] K. O. Roberts, T. McKellar, J. Fekete, A. Rakonjac, A. B. Deb, and N. Kjærgaard, "Steerable optical tweezers for ultracold atom studies," Opt. Lett. **39**, 2012–2015 (2014).

[109] M. Mirhosseini, O. S. Magaña-Loaiza, C. Chen, B. Rodenburg, M. Malik, and R. W. Boyd, "Rapid generation of light beams carrying orbital angular momentum," Opt. Express **21**, 30196–30203 (2013).

[110] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," Optik **35**, 237–246 (1972).

[111] M. Pasienski and B. DeMarco, "A high-accuracy algorithm for designing arbitrary holographic atom traps," Opt. Express **16**, 2176–2190 (2008).

[112] A. L. Gaunt and Z. Hadzibabic, "Robust digital holography for ultracold atom trapping." Scientific reports **2**, 721 (2012).

[113] T. Harte, G. D. Bruce, J. Keeling, and D. Cassettari, "A conjugate gradient minimisation approach to generating holographic traps for ultracold atoms," (2014), arXiv:1408.0188 [cond-mat.quant-gas] .

[114] J.W. Goodman, *Introduction to Fourier optics* (McGraw-Hill, 1996).

[115] M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," Proc. IEEE **93**, 216–231 (2005).

[116] C. Hesseling, M. Woerdemann, A. Hermerschmidt, and C. Denz, "Controlling ghost traps in holographic optical tweezers," Opt. Lett. **36**, 3657–3659 (2011).

[117] C. Dorrer and J. D. Zuegel, "Design and analysis of binary beam shapers using error diffusion," J. Opt. Soc. Am. B **24**, 1268–1275 (2007).

[118] D. Gabor, "A New Microscopic Principle," Nature **161**, 777–778 (1948).

[119] B. Bayer, "An optimum method for two-level rendition of continuous-tone pictures," in *Proceedings of IEEE International Conference on Communications* (IEEE, 1973) pp. 26–11 to 26–15.

[120] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grey scale," in *Proceedings of the Society of Information Display*, Vol. 17 (1976) pp. 75–77.

[121] K. Dholakia T. Čižmár, M. Mazilu, "In situ wavefront correction and its application to micromanipulation," Nat. Photon. **4**, 388–394 (2010).

[122] U. Dorner, T. Calarco, P. Zoller, A. Browaeys, and P. Grangier, "Quantum logic via optimal control in holographic dipole traps," J. Opt. B **7**, S341 (2005).

[123] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede, "Analysis of dephasing mechanisms in a standing-wave dipole trap," Phys. Rev. A **72**, 023406 (2005).

[124] M. Saffman and T. G. Walker, "Analysis of a quantum logic device based on dipole-dipole interactions of optically trapped Rydberg atoms," Phys. Rev. A **72**, 022347 (2005).

[125] R. A. Cline, J. D. Miller, M. R. Matthews, and D. J. Heinzen, "Spin relaxation of optically trapped atoms by light scattering," Opt. Lett. **19**, 207–209 (1994).

[126] D. Frese, B. Ueberholz, S. Kuhr, W. Alt, D. Schrader, V. Gomer, and D. Meschede, "Single atoms in an optical dipole trap: Towards a deterministic source of cold atoms," Phys. Rev. Lett. **85**, 3777–3780 (2000).

[127] A. Kaplan, M. Fredslund Andersen, and N. Davidson, "Suppression of inhomogeneous broadening in rf spectroscopy of optically trapped atoms," Phys. Rev. A **66**, 045401 (2002).

[128] J. Kruse, C. Gierl, M. Schlosser, and G. Birkl, "Reconfigurable site-selective manipulation of atomic quantum systems in two-dimensional arrays of dipole traps," Phys. Rev. A **81**, 060308 (2010).

[129] X. L. Zhang, L. Isenhower, A. T. Gill, T. G. Walker, and M. Saffman, "Deterministic entanglement of two neutral atoms via Rydberg blockade," Phys. Rev. A **82**, 030306 (2010).

[130] A. Reiserer, C. Nölleke, S. Ritter, and G. Rempe, "Ground-state cooling of a single atom at the center of an optical cavity," Phys. Rev. Lett. **110**, 223003 (2013).